Frequency Division Duplexing Massive MIMO:
Novel Theoretical Results and Efficient Algorithms

vorgelegt von
M. Sc.
Mahdi Barzegar Khalilsarai
ORCID: 0000-0003-0579-0200

von der Fakultät IV - Elektrotechnik und Informatik
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Ingenieurwissenschaften
- Dr.-Ing. -

genehmigte Dissertation

Promotionsausschuss:

Vorsitzender: Prof. Dr. Rafael Schaefer
Gutachter: Prof. Dr. Giuseppe Caire
Gutachter: Prof. Dr. Gerhard Wunder
Gutachter: Prof. Dr. Andreas Molisch

Tag der wissenschaftlichen Aussprache: 19.01.2021

Berlin 2021

For my dear parents, Roqaye and Mousa,

and my lovely sisters, Saeedeh and Samira

Abstract

Massive Multiple-Input Multiple-Output (MIMO) systems are envisioned for employment in the upcoming generations of wireless communication networks. The characteristic of these systems is the use of a large number of antenna elements ($M \gg 1$) in the Base Station (BS) array, that enables data transmission to several users via spatial multiplexing. These systems enjoy favorable communication properties including energy efficiency due to channel hardening and large beamforming gain, suppression of the inter-cell interference, and a simplified user scheduling and rate adaptation.

In allocating time-frequency resources for Uplink-Downlink (UL-DL) transmission, a massive MIMO BS can operate in two principal modes: Time Division Duplexing (TDD), and Frequency Division Duplexing (FDD), where in the former UL-DL data transmission occurs over disjoint time intervals and in the latter it occurs over disjoint frequency bands. This dissertation is dedicated to the study of various signal processing aspects in realizing an FDD massive MIMO system. These systems comprise an interesting topic of research, not only due to their favorability in symmetric-traffic and delay-sensitive communication scenarios, but also due to their widespread implementation in the existing cellular networks. The benefits of FDD massive MIMO systems, however, come at the cost of a more challenging signal processing task as compared to TDD systems. Due to a lack of UL-DL channel reciprocity, it is more difficult for the BS to obtain fresh channel state information (CSI) for effective DL transmission, which is conventionally achieved via training of the DL channels by transmitting pilot sequences to the users and receiving their closed-loop feedback. In a massive MIMO scenario where the signal dimension is large, this burdens the system with a sizeable training and feedback overhead, which may exhaust the time-frequency resources, resulting in a poor DL data rate.

In this dissertation, we address several aspects of the mentioned issues, by analysing relevant theoretical problems as well as proposing efficient practical algorithms to solve them. In particular, first we study the problem of channel covariance estimation from limited pilot samples (either in UL or in DL), in which the goal is to enable an efficient estimation of the high-dimensional channel covariance with the relatively small pilot dimension, where we propose two efficient estimators. Second, we study the problem of estimating the DL channel covariance from the available UL channel covariance. We de-

rive universal minimax error bounds for this UL-DL covariance "transformation" and we propose an algorithm for performing the transformation. Finally, we address the problem of multi-user DL precoding in an FDD massive MIMO system, in which the goal is to train the high-dimensional user channels with a given, limited DL pilot dimension. Our suggested technique, coined as *active channel sparsification* does *not* rely on the typical channel sparsity assumptions and is shown to outperform the state-of-the-art methods in the literature.

Zusammenfassung

Massive Multiple-Input Multiple-Output (MIMO) Systeme sind für den Einsatz in den zukünftigen Generationen von drahtlosen Kommunikationsnetzen vorgesehen. Charakteristisch für diese Systeme ist die Verwendung einer großen Anzahl von Antennenelementen ($M \gg 1$) im Basisstation (BS) Array, das die Datenübertragung an mehrere Benutzer über räumliches Multiplexing ermöglicht. Diese Systeme weisen günstige Kommunikationseigenschaften auf, darunter eine Energieeffizienz durch Kanalhärtung und großen Strahlformungsgewinn, Unterdrückung der Interzellularinterferenz und eine vereinfachte Benutzerplanung und Ratenanpassung.

Bei der Zuweisung von Zeit-Frequenz Ressourcen für die Uplink-Downlink (UL-DL) Übertragung kann ein massive MIMO BS in zwei Hauptmodi arbeiten: Time Division Duplexing (TDD) und Frequency Division Duplexing (FDD), wobei im ersteren Fall die UL-DL Datenübertragung über disjunkte Zeitintervalle und im letzteren Fall über disjunkte Frequenzbänder erfolgt. Diese Dissertation ist der Untersuchung verschiedener Aspekte der Signalverarbeitung bei der Realisierung eines FDD massive MIMO System gewidmet. Diese Systeme stellen ein interessantes Forschungsthema dar, nicht nur wegen ihrer Nutzlichkeit bei symmetrischem Verkehr und verzögerungsempfindlicher Kommunikationszenarien, sondern auch aufgrund ihrer weit verbreiteten Implementierung in den bereits bestehenden zellularen Netzwerken. Die Vorteile von FDD massivs MIMO Systemen gehen jedoch auf Kosten einer anspruchsvolleren Signalverarbeitung im Vergleich zu TDD Systemen. Aufgrund mangelnder Reziprozität der UL-DL Kanäle ist es für die BS schwieriger, aktuelle Kanalstatusinformationen (CSI) für eine effektive DL Übertragung zu erhalten, welche konventionell durch Training der DL Kanäle erreicht wird, indem Pilotsequenzen an die Benutzer übertragen und deren Feedback in geschlossener Schleife empfangen werden. In einem massiven MIMO-Szenario, in dem die Signaldimension umfangreich ist, belastet dies das System mit einem beträchtlichen Trainings- und Feedback-Overhead, der die Zeit-Frequenz Ressourcen erschöpfen kann, was zu einer schlechten DL-Datenrate führt.

In dieser Dissertation untersuchen wir mehrere Aspekte der zuvor genannten Themen, indem wir sowohl relevante theoretische Probleme analysieren als auch effiziente praktische Algorithmen zu deren Lösung vorschlagen. Insbesondere untersuchen wir zunächst das Problem der Kanalkovarianzschätzung aus kleinen Pilotstichproben (entweder in UL

oder in DL), wobei das Ziel darin besteht, eine effiziente Schätzung der hochdimensionalen Kanalkovarianz mit einer relativ kleinen Pilotdimension zu ermöglichen, wobei wir zwei effiziente Methoden vorschlagen. Zweitens untersuchen wir das Problem der Schätzung der DL-Kanal-Kovarianz aus der verfügbaren UL-Kanal-Kovarianz. Wir universelle Minimax-Fehlergrenzen für diese UL-DL-Kovarianz "Transformation" ab und schlagen einen Algorithmus zur Durchführung der Transformation vor. Schließlich befassen wir uns mit dem Problem der Mehrbenutzer-DL-Precodierung in einem FDD-Massiv-MIMO-System, bei dem das Ziel darin besteht, die hochdimensionalen Nutzerkanäle mit einer begrenzten DL-Pilotdimension zu trainieren. Die von uns vorgeschlagene Technik, die als *active channel sparsification* bezeichnet wird, beruht *nicht* auf den typischen Kanalsparsamkeitsannahmen und übertrifft nachweislich die modernen Methoden der Literatur.

Acknowledgments

This dissertation is the summary of my work over about four and a half years of studies as a Ph.D. student. It would not have been accomplished without the help of several important people in my life.

I am forever thankful to my parents for their support and unconditional love, not only during these years, but all throughout my life. The same goes for my sisters, who have always believed in me and who never stopped to motivate me in this journey.

I would like to thank my supervisor, professor Giuseppe Caire, for his support and mentoring. It has been a joy to work with him, to have been constantly inspired by his scientific curiosity and ambition, and to have learned so much from his vast knowledge and experience. In addition, I am thankful to professor Gerhard Wunder, for his guidance, support and encouragement and for the plenty of productive collaborations we have had.

I have been truly lucky to have worked with Dr. Saeid Haghighatshoar for the most part of my Ph.D. studies. We have had numerous discussions on various topics in which, for the most part, he has been a patient teacher and I have been a bothersome student. Multiple parts of this dissertation are the outcomes of such exchanges. I am indebted to Saeid for all these, and for the scientific passion and curiosity he provoked in me.

Discussions with my colleagues Tianyu Yang and Yi Song on some topics have been quite useful and here I express my gratitude for their help and contribution. In addition, I would like to thank Daniyal Amir Awan, Saeed Dehkordi, and Rick Fritschek for their comments and feedback about parts of this dissertation.

Ever since starting to work at the Communications and Information Theory (CommIT) group, I have felt at home because of the friendly and constructive vibe among the members. I am thankful to all of them. Their company has left so many good memories in my mind. I refrain from naming the group members, being afraid that I will be embarrassed by forgetting to mention every single one of them.

My sincere apologies go to all those who have contributed to my work, and whose names I have unintentionally failed to mention.

<div align="right">

Mahdi Barzegar Khalilsarai

Berlin, November 2020

</div>

Contents

List of Figures

Note: All figures are generated using the Ti*k*Z package and LaTeX typesetting [119].

1 Introduction

Multi-user *multiple-input multiple-output* (MIMO) consists of employing multiple antennas at the Base Station (BS), in order to multiplex several data streams over the spatial domain to multiple users sharing the same time-frequency transmission resource (time slots and frequency bands). For a block-fading channel with spatially independent fading, the fading channel coefficients can be considered constant over a time-frequency *coherence block* of T symbols [123]. For such a channel, the high-SNR sum-capacity behaves according to

$$C(\text{SNR}) = M^*(1 - M^*/T) \log \text{SNR} + O(1) \tag{1.1}$$

where $M^* = \min\{M, K, T/2\}$, M denotes the number of BS antennas, and K denotes the number of single-antenna users [1, 82, 141]. The pre-log factor in this expression indicates the number of spatial-domain data streams supported by the system, such that each stream has a spectral efficiency that behaves as an interference-free Gaussian channel, i.e., $\log \text{SNR} + \mathcal{O}(1)$. In practice, although the system may be interference-limited (e.g., due to inter-cell interference in multi-cell wireless systems), a well-designed system would exhibit a regime of practically relevant SNR for which its sum-rate behaves as an affine function of $\log \text{SNR}$ [78].

When M and the number of users are potentially very large, the system pre-log factor is maximized by serving $K = T/2$ data streams (users). While any number $M \geqslant K$ of BS antennas yields the same (optimal) pre-log factor, a key observation made in [83] is that, when training a very large number of antennas comes at no additional overhead cost, it is indeed convenient to use $M \gg K$ antennas at the BS. In this way, at the cost of some additional *hardware complexity*, very significant benefits at the system level can be achieved. These include: i) energy efficiency (due to the large beamforming gain); ii) inter-cell interference reduction; iii) a dramatic simplification of user scheduling and rate adaptation, due to the inherent large-dimensional channel hardening [75]. Systems for which the number of BS antennas (M) is much larger than the number of DL data streams (K) are generally referred to as *massive MIMO* (see [75, 83, 84] and references therein). Massive MIMO has been the object of intense research investigation and development and is expected to be a cornerstone of the forthcoming 5th generation of wireless/cellular systems [13].

In order to achieve the benefits of massive MIMO, the BS must learn the Downlink (DL) channel coefficients for K users and $M \gg K$ BS antennas. For *Time Division Duplexing* (TDD) systems, due to the inherent Uplink-Downlink (UL-DL) channel reciprocity [82], this can be obtained from K mutually orthogonal UL pilots transmitted by the users. Unfortunately, the UL-DL channel reciprocity does not hold for *Frequency Division Duplexing* (FDD) systems, since the UL and DL channels are separated in frequency by much more than the channel coherence bandwidth [123]. Hence, unlike TDD systems, in FDD the BS must actively probe the DL channel by sending a common DL pilot signal, and request the users to feed their channel state back.

In order to obtain a "fresh" channel estimate for each coherence block, T_{dl} out of T symbols per coherence block must be dedicated to the common DL pilot transmission. Assuming (for simplicity of exposition) a delay-free channel state feedback, the resulting DL pre-log factor is given by $K \times \max\{0, 1 - T_{\mathsf{dl}}/T\}$, where K is the number of served users, and $\max\{0, 1 - T_{\mathsf{dl}}/T\}$ is the penalty factor incurred by DL channel training. Conventional DL training consists of sending orthogonal pilot signals from each BS antenna. Thus, in order to train M antennas, the minimum required training dimension is $T_{\mathsf{dl}} = M$. Hence, with such scheme, the number of BS antennas M cannot grow arbitrarily large. For example, consider a typical scenario in an LTE system [111], where the BS schedules groups of users over resource blocks spanning 14 OFDM symbols × 12 subcarriers, for a total dimension of $T = 168$ symbols in the time-frequency plane. Consider a typical massive MIMO configuration serving $K \sim 20$ users with $M \geqslant 200$ antennas (e.g., see [79]). In this case, the entire resource block dimension would be consumed by the DL pilots, leaving no room for data communication. Furthermore, feeding back the M-dimensional measurements (or estimated/quantized channel vectors) represents also a significant feedback overhead for the UL [19, 63, 69, 77, 137].

While the argument above is kept informal on purpose, it can be made information-theoretically rigorous. The main issue is that, if one insists to estimate the $K \times M$ channel matrix in an "agnostic" way, i.e., without exploiting the channel's fine structure, a hard dimensionality bottleneck kicks in and fundamentally limits the number of data streams that can be supported in the DL by FDD systems. It follows that gathering "massive MIMO gains" in FDD systems is a challenging problem. Despite this fact, FDD massive MIMO may be preferred to its TDD counterpart due to the following reasons:

- Current wireless networks are mostly based on FDD. Such systems are easier to operate and more effective than TDD systems in situations with symmetric traffic and delay-sensitive applications [23, 62, 98]. Besides, converting current FDD systems to TDD would represent a non-trivial cost for wireless operators.

- TDD massive MIMO systems are prone to system imperfections such as calibration errors in UL and DL RF chains. These result in the issue that estimating the DL channel from the UL channel by leveraging channel reciprocity may not be accurate [11, 49, 96], and therefore it is too optimistic to expect the results predicted by theory for these systems to hold in practice.

- In a noise-limited transmission scenario, an FDD system shows a 3-dB SNR advantage compared to a TDD system [84]. Let BW denote the total system bandwidth (sum of UL and DL bandwidths), P the received power at the user side and N_0 the noise spectral density. In TDD, the DL sum-rate is given by $\frac{\text{BW}}{2} \log_2 \left(1 + \frac{P}{\text{BW} \, N_0}\right)$, where the 1/2 factor appears in the pre-log term, since transmission takes place over the full bandwidth but only half of the time. In contrast, the DL sum-rate in a FDD system is given by $\frac{\text{BW}}{2} \log_2 \left(1 + \frac{P}{\frac{\text{BW}}{2} \, N_0}\right)$, where the 1/2 factor appears both in and outside the logarithm, since transmission is performed continuously in time but only over half of the bandwidth. This results in the mentioned 3-dB advantage of FDD over TDD.

With these motivations in mind, researchers have devoted a significant effort to reduce the common DL training dimension and feedback overhead in order to materialize significant massive MIMO gains also for FDD systems.

1.1 Outline and Contributions

This dissertation addresses three major problems in realizing an FDD massive MIMO system, with the eventual goal that the BS can enjoy the advantages of massive MIMO in FDD mode with as small training and feedback overhead as possible. The three elements of our work are listed as follows.

I. Channel Covariance Estimation

Channel covariance knowledge is crucial for minimum mean squared error (MMSE) instantaneous channel estimation as well as for designing efficient channel training schemes and, in a multi-user setup, for precoder design [90, 139]. In order to obtain covariance information, the BS typically probes the instantaneous channel across multiple coherence blocks by sending/receiving pilots to/from the user. With low or moderate array size, the number of available pilots exceeds the channel dimension, and this enables the BS to reliably estimate the covariance using the simple unbiased sample covariance [12]. However, with a massive array, the channel dimension is large ($M \gg 1$) and may be comparable

with the pilot dimension $(M \sim N)$. It is well-known that in such a regime, the sample covariance can be substantially improved by exploiting the covariance structure which narrows down the set of admissible *estimands* [15]. We propose two covariance estimation methods that exploit the structure of a MIMO covariance to estimate it from a noisy channel sample set of a size that is comparable to the channel dimension. The first method is based on solving a (convex) non-negative least-squares (NNLS) problem [113], and the second method is based on maximum likelihood (ML) estimation of the coefficients of a carefully-designed parametric approximation of the covariance. These estimators rely on the key observation that the channel covariance can be seen as the inner product of the *array manifold* with a non-negative measure over the angle domain, which we call the *angular scattering function* (ASF) that is independent from the array geometry. Given a sample set, one can solve an inverse problem to estimate the ASF and then use it for covariance estimation. The proposed methods can be used either for UL or DL channels, and both in FDD and TDD modes. We develop the mentioned methods in **Chapter 2** and empirically evaluate their performances.

II. Uplink-Downlink Channel Covariance Transformation

Downlink channel covariance estimation in FDD mode is particularly challenging, since unlike TDD the BS does not naturally receive full-dimensional channel samples, as this must be done via DL probing (pilot transmission) and closed-loop UL feedback. Collecting enough channel samples for a reliable DL covariance estimation in the massive MIMO regime imposes a large overhead on the system [26, 46]. However, the BS receives UL pilots during UL, enabling an accurate estimation of the UL covariance. It is then desirable to seek a way to deduce DL covariance information from its UL counterpart. Not that, even the UL and DL covariance matrices are *not* identical in FDD mode, since the array response varies from the UL band to the DL band.

A variety of recent works have proposed techniques for *UL-DL covariance transformation*[1], by employing specific channel models [32, 35, 54, 56]. In most of these works the transformation algorithm is heuristic and relies on restrictive structural assumptions such as channel sparsity. Also a principled analysis of the limitations of covariance transformation for a *general* class of MIMO covariances is missing.[2] In **Chapter 3** we derive *minimax* bounds on the error of estimating the DL covariance from the observation of

---

[1]The term "transformation" appears in different sources by different synonyms, such as conversion, interpolation, and extrapolation [22, 34].

[2]After the publication of our paper based on the material of this chapter, a number of works addressed this problem in similar generality and with different error metrics (see [22, 86, 87]). These works will be discussed on the side of the core material of this chapter.

the UL covariance, where the bound is given per covariance entry (for a large portion of the entries), and holds for any channel angular scattering function for all estimators that satisfy minimal properties (namely, non-negativity and data-consistency constraints). Our result relies on a standard assumption of *channel angular reciprocity* [40, 59], such that the angular scattering function is considered to be invariant over UL and DL frequency bands. We also propose a transformation method based on NNLS that first estimates the angular scattering function from the UL covariance and, combined with a change of basis, uses it to estimate the DL covariance. We provide extensive empirical results to compare our covariance transformation method to those suggested in the literature.

III. Downlink Precoding

In a multi-user MIMO system, the BS needs to probe instantaneous user channels (via pilot transmission) and design a precoder that mitigates inter-user interference [3, 139]. As mentioned earlier, in a massive MIMO regime the channel dimension can potentially be larger than the dimension of the fraction of the channel coherence block dedicated to pilot transmission ($M > T_{\mathsf{dl}}$). In this case, the BS has to estimate the channel $\mathbf{h}$ from a set of (noisy) linear feedback pilot measurements $\mathbf{y} = \mathbf{\Phi}\mathbf{h} + \mathbf{z}$, where $\mathbf{\Phi} \in \mathbb{C}^{T_{\mathsf{dl}} \times M}$ is the pilot matrix [98]. In order to solve this under-determined system of equations, several works have proposed exploiting the *hypothetical* [12] sparse structure of the massive MIMO channel and using *compressed sensing* (CS) methods to recover the channel [71, 98, 112]. However, when the channel degrees of freedom s (number of non-zero coefficients in the sparsity domain) is larger than a certain threshold this strategy results in poor channel estimates. Standard CS theory states that stable sparse signal reconstruction is possible using $T_{\mathsf{dl}} = \mathcal{O}(s \log M)$ measurements, so when $T_{\mathsf{dl}} < s \log M$, one can not stably recover $\mathbf{h}$ from $\mathbf{y}$. In other words, CS-based methods are at the mercy of nature, in the sense that they would fail under circumstances in which the channel is *not* sparse enough, e.g. in the case of rich scattering environments.

In order to allow stable channel estimation with a given pilot dimension T_{dl}, in **Chapter 4** we propose *active channel sparsification* (ACS), a novel method that uses the DL covariance information of all users to design an optimal *sparsifying precoder*. This is a linear transformation that depends only on the channel second order statistics (estimated DL covariances) that imposes that the effective channel matrix (including the precoder) has large rank and yet each column has sparsity not larger than T_{dl}. In this way, our method is not at the mercy of nature, i.e. it is flexible with respect to various types of environments and channel sparsity orders. We cast the optimization of the sparsifying precoder as a

matching-size maximization problem over a user-virtual beam bipartite graph and solve it using a *mixed-integer linear program* (MILP) [91].

To extend the ACS to arbitrary array geometries (other than the ULA), in **Chapter 4** we also propose a method of MIMO virtual beam design that can serve as a *common* spatial eigenbasis for all user channels. This method is based on maximizing the likelihood of a unitary matrix that (approximately) jointly diagonalizes the covariances of all users, given a set of noisy pilot samples– i.e., without accurate covariance information. We implement the optimization problem via a projected gradient descent algorithm and we show that it converges to a stationary point of the likelihood function. The common eigenbasis in turns enables modeling of the user-virtual beam bipartite graph and application of the ACS.

1.2 Notation and Abbreviations

We denote scalars, vectors, matrices and sets by lower case letters, lower case bold letters, upper case bold letters and calligraphic letters, i.e., x, $\mathbf{x}$, $\mathbf{X}$, $\mathcal{X}$, respectively. Real and complex scalars are denoted by $\mathbb{R}$ and $\mathbb{C}$, respectively. The natural numbers are denoted by $\mathbb{N} = \{1, 2, \ldots\}$ and the non-negative integers by $\mathbb{Z}_+$. For a non-negative integer n, the set $\{0, 1, \ldots, n-1\}$ is represented by $\mathbb{Z}_n$. We denote the d-dimensional non-negative orthant by $\mathbb{R}_+^d$, dropping the exponent for $d = 1$. The cone of d-dimensional positive semidefinite (PSD) matrices, Toeplitz Hermitian matrices, and Toeplitz PSD matrices are denoted by $\mathbb{S}_+^d$, $\mathbb{T}^M$, and $\mathbb{T}_+^M$, respectively. Unless otherwise defined, we refer to the i-th element of a vector $\mathbf{x}$ by $[\mathbf{x}]_i$, and to the (i, j)-th element of a matrix $\mathbf{X}$ by $[\mathbf{X}]_{i,j}$. The i-th row of $\mathbf{X}$ is denoted by $\mathbf{X}_{i,\cdot}$ and its j-th column by $\mathbf{X}_{\cdot,j}$. Superscripts $(\cdot)^*$, $(\cdot)^\mathsf{T}$, $(\cdot)^\mathsf{H}$, $(\cdot)^{-1}$, and $(\cdot)^\dagger$ represent the complex conjugate, transpose, conjugate transpose, inverse and Moore-Penrose pseudo-inverse, respectively. The Kronecker product is denoted by $\otimes$. For a matrix $\mathbf{X}$ of dimension $m \times n$, the vector $\text{vec}(\mathbf{X})$ of dimension $mn \times 1$ is composed by stacking the columns of $\mathbf{X}$ on top of each other. For a vector vector $\mathbf{x}$, the symbol $\text{diag}(\mathbf{x})$ denotes a matrix with $\mathbf{x}$ as its main diagonal. For the same vector, given that its first element is a real scalar, $\mathcal{T}(\mathbf{x})$ denotes a Hermitian Toeplitz matrix with $\mathbf{x}$ as its first column. The trace and determinant of a square matrix $\mathbf{X}$, are denoted by $\text{tr}(\mathbf{X})$ and $\det(\mathbf{X})$, respectively. The ℓ_p norm of a vector $\mathbf{x}$ is referred to as $\|\mathbf{x}\|_p$, where for simplicity we drop the subscript for the case of $p = 2$. The Frobenius norm of a matrix $\mathbf{X}$ is denoted by $\|\mathbf{X}\|_\mathsf{F}$.

Abbreviations

3GPP	3rd Generation Partnership Project
ACS	active channel sparsification
AoA	angle-of-arrival
AWGN	additive white Gaussian noise
BS	base station
CSI	channel state information
CSIT	channel state information at the transmitter
DL	Downlink
FDD	frequency division duplexing
iff	if and only if
JSDM	joint spatial division and multiplexing
LoS	Line of Sight
MIMO	Multiple-Input Multiple-Output
ML	maximum likelihood
MUSIC	MUltiple Signal Classification method
MSE	mean square error
NNLS	non-negative least-squares
OFDM	Orthogonal Frequency Division Multiplexing
PGD	projected gradient descent
PSD	positive semi-definite
SDP	semi-definite program
SINR	signal to sum of noise and interference ratio
SNR	signal to noise ratio
TDD	time division duplexing
UL	Uplink
w.l.o.g	without loss of generality

1.3 Publications and Copyright Disclaimer

I have re-used material from my previously published work, including journal articles and conference papers. All of these works are protected by the IEEE intellectual property rights ©2017-2020 IEEE. Regarding each individual publication, all copyrights are reserved by the authors and re-use or distribution of all or some parts of the publication by third-party individuals and organizations without appropriate citation is prohibited.

The following list includes my publications during my Ph.D. studies.

Published

1. **Barzegar Khalilsarai, Mahdi**, Saeid Haghighatshoar, and Giuseppe Caire. "Joint Approximate Covariance Diagonalization with Applications in MIMO Virtual Beam Design." proceedings of the IEEE Global Communications Conference (Globecom), 2020.

2. **Barzegar Khalilsarai, Mahdi**, Benedikt Gross, Stelios Stefanatos, Gerhard Wunder, Giuseppe Caire. "WiFi-Based Channel Impulse Response Estimation and Localization via Multi-Band Splicing." proceedings of the IEEE Global Communications Conference (Globecom), 2020.

3. **Barzegar Khalilsarai, Mahdi**, Tianyu Yang, Saeid Haghighatshoar, and Giuseppe Caire. "Structured Channel Covariance Estimation for Dual-Polarized Massive MIMO Arrays." In WSA 2020; 24th International ITG Workshop on Smart Antennas, pp. 1-6. VDE, 2020.

4. **Barzegar Khalilsarai, Mahdi**, Tianyu Yang, Saeid Haghighatshoar, Xinping Yi, and Giuseppe Caire. "Active Channel Sparsification and Precoding for Dual-Polarized FDD Massive MIMO." In WSA 2020; 24th International ITG Workshop on Smart Antennas, pp. 1-6. VDE, 2020.

5. **Barzegar Khalilsarai, Mahdi**, Tianyu Yang, Saeid Haghighatshoar, and Giuseppe Caire. "Structured channel covariance estimation from limited samples in massive MIMO." In ICC 2020-2020 IEEE International Conference on Communications (ICC), pp. 1-7. IEEE, 2020.

6. **Barzegar Khalilsarai, Mahdi**, Stelios Stefanatos, Gerhard Wunder, and Giuseppe Caire. "WiFi-based indoor localization via multi-band splicing and phase retrieval." In 2019 IEEE 20th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC), pp. 1-5. IEEE, 2019.

7. **Barzegar Khalilsarai, Mahdi**, Yi Song, Tianyu Yang, Saeid Haghighatshoar, and Giuseppe Caire. "Uplink-Downlink Channel Covariance Transformations and Precoding Design for FDD Massive MIMO." In 2019 53rd Asilomar Conference on Signals, Systems, and Computers, pp. 199-206. IEEE, 2019.

8. **Barzegar Khalilsarai, Mahdi**, Saeid Haghighatshoar, Xinping Yi, and Giuseppe Caire. "FDD massive MIMO via UL/DL channel covariance extrapolation and active channel sparsification." IEEE Transactions on Wireless Communications 18, no. 1 (2018): 121-135.

9. **Barzegar Khalilsarai, Mahdi**, Saeid Haghighatshoar, and Giuseppe Caire. "How to achieve massive MIMO gains in FDD systems?." In 2018 IEEE 19th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC), pp. 1-5. IEEE, 2018.

10. **Barzegar Khalilsarai, Mahdi**, Saeid Haghighatshoar, Xinping Yi, and Giuseppe Caire. "FDD massive MIMO: Efficient downlink probing and uplink feedback via active channel sparsification." In 2018 IEEE International Conference on Communications (ICC), pp. 1-6. IEEE, 2018.

11. **Barzegar Khalilsarai, Mahdi**, Saeid Haghighatshoar, Giuseppe Caire, and Gerhard Wunder. "Compressive estimation of a stochastic process with unknown autocorrelation function." In 2017 IEEE International Symposium on Information Theory (ISIT), pp. 1282-1286. IEEE, 2017.

12. Song, Yi, **Mahdi Barzegar Khalilsarai**, Saeid Haghighatshoar, and Giuseppe Caire. "Deep Learning for Geometrically-Consistent Angular Spread Function Estimation in Massive MIMO." proceedings of the IEEE Global Communications Conference (Globecom), 2020.

13. Haghighatshoar, Saeid, **Mahdi Barzegar Khalilsarai**, and Giuseppe Caire. "Multiband covariance interpolation with applications in massive MIMO." In 2018 IEEE International Symposium on Information Theory (ISIT), pp. 386-390. IEEE, 2018.

14. Wunder, Gerhard, Ingo Roth, **Mahdi Barzegar Khalilsarai**, Axel Flinth, Saeid Haghighatshoar, Giuseppe Caire, and Gitta Kutyniok. "Hierarchical sparse channel estimation for massive MIMO." In WSA 2018; 22nd International ITG Workshop on Smart Antennas, pp. 1-8. VDE, 2018.

15. Stefanatos, Stelios, **Mahdi Barzegar Khalilsarai**, and Gerhard Wunder. "Wideband massive MIMO channel estimation via sequential atomic norm minimization."

In 2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP), pp. 998-1002. IEEE, 2018.

16. Tian, Chang, An Liu, **Mahdi Barzegar Khalilsarai**, Giuseppe Caire, Wu Luo, and Min-Jian Zhao. "Randomized Channel Sparsifying Hybrid Precoding for FDD Massive MIMO Systems." IEEE Transactions on Wireless Communications (2020).

2 Channel Covariance Estimation

2.1 Problem Statement

Consider a MIMO system in which a BS is equipped with a ULA consisting of M equispaced antennas and communicates with a single-antenna user. During UL transmission, the user sends a number N of pilot symbols to the BS, one per time-frequency coherence block. The signal received at the BS at each of these transmissions can be expressed as [123]

$$\mathbf{y}[i] = \mathbf{h}[i]s[i] + \mathbf{z}[i] \in \mathbb{C}^M, \ i \in \mathbb{Z}_N, \tag{2.1}$$

where $\mathbf{h}[i] \in \mathbb{C}^M$ is a realization of the narrow-band random channel vector $\mathbf{h}$ which contains the random channel gains associated with the M antenna elements, $s[i] \in \mathbb{C}$ is a symbol chosen from a constellation, and $\mathbf{z}[i] \in \mathbb{C}^M$ is the zero-mean additive white Gaussian noise (AWGN) vector with $\mathbb{E}[\mathbf{z}\mathbf{z}^{\mathsf{H}}] = N_0\mathbf{I}_M$. Since the BS knows the pilot symbol a priori, we can assume without loss of generality (w.l.o.g), that $s[i] = 1$ for all $i \in \mathbb{Z}_N$. Note that, taken at distinct channel coherence blocks, $\{\mathbf{h}[i]\}_i$ represent independent realizations. Assuming zero mean and the existence of the second-order statistics, the BS aims at estimating the channel covariance matrix,

$$\mathbf{\Sigma} = \mathbb{E}\left[\mathbf{h}\mathbf{h}^{\mathsf{H}}\right]. \tag{2.2}$$

When the number of samples is much larger than the channel dimension ($N \gg M$), a simple estimator such as the sample covariance

$$\hat{\mathbf{\Sigma}} = \frac{1}{N} \sum_{i \in \mathbb{Z}_N} \mathbf{y}[i]\mathbf{y}[i]^{\mathsf{H}} - N_0\mathbf{I}, \tag{2.3}$$

provides a reliable estimate of the covariance. However, with massive antenna arrays, the channel dimension is much larger than one ($M \gg 1$) and is typically of the same order as the sample size ($M \sim N$); in an extreme case we may even have $N < M$. In these regimes, it is well-known that the sample covariance yields a poor estimate (unless the channel is extremely correlated and has too few degrees of freedom).

The purpose of this chapter is to propose practical methods for estimating the channel covariance from a noisy sample set with cardinality that is comparable to the number of antennas.

2.2 Related Work

The problem of covariance estimation from a random sample set is a classical problem in statistics and has applications in various branches including wireless communications. The older view on covariance estimation considers a case in which the number of samples is much larger than the signal dimension ($N \gg M$), and perhaps tends to infinity, while the signal dimension is kept constant. In a communication scenario, typically the number of observed (noisy) channel samples is limited by the pilot dimension, which is itself limited by the dimension of the time-frequency tile over which the channel covariance remains (approximately) constant. Therefore, for a massive MIMO channel with $M \gg 1$ antennas, it is more reasonable to consider a case in which the sample size is comparable to the channel dimension ($N \sim M$). In this case, the behavior of standard covariance estimators such as the sample covariance changes. An asymptotic study of this behavior dates back to the sixties and seventies [81, 133]. More recent research has focused on the non-asymptotic regime, in which both the signal dimension and the sample size are assumed finite [130, 131]. It is widely accepted that in the case of structured covariance matrices, such as sparse or low-rank covariances, or covariances with banded decay, one can achieve faster rates of convergence (in terms of the sample size), either with the simple sample covariance or more sophisticated estimators, such as thresholding and *tapering* estimators [10, 16, 42].

On the algorithmic side, several recent works have proposed decent channel covariance estimation methods for massive MIMO. These works typically exploit an assumed structure on the multipath channel to derive their results: super-resolution techniques have been developed by assuming the channel to consist of discrete, separable multipath components [27]. A different type of estimators assumes the channel to be sparse in the angle domain, that is to say, the number of contributing multipath components is much less than the channel dimension. These works suggest sparse recovery methods to estimate the angle-of-arrival (AoA) and power parameters of the multipath components via, for example, greedy methods such as *orthogonal matching pursuit* (OMP) [76, 92, 95], ℓ_1-norm regularization [17, 99], the multiple measurement vectors technique [24, 30], etc. Besides, low-rank matrix recovery methods have been proposed to exploit the low-rank covariance structure in certain communication scenarios [25, 51]

2.2.1 Contribution

In this chapter, we propose two channel covariance estimators. Unlike the works mentioned above, we do not assume any of the standard structures on the covariance. Instead we recognize that a MIMO covariance of Gaussian channels is associated with the channel *angular scattering function* (ASF) through an inner-product with the array manifold. We exploit this structure of the channel covariance to propose better estimators. We decompose the ASF to discrete (or specular) and continuous (or diffuse) components. We approximate the continuous component with the linear combination of the elements of a family of limited-support density functions, and estimate the finite-dimensional support of the discrete component via the well-known *multiple signal classification* (MUSIC) method [117]. Then, we propose two methods to estimate the coefficients of a parametric expansion of the covariance. The first method is based on the convex non-negative least squares (NNLS) estimator and the second method solves a maximum-likelihood (ML) optimization problem.

2.3 Channel Model

We consider the standard block-fading model for the wireless channel (see, e.g., [123] and the 3GPP channel model [111]), which assumes the fading to be approximately constant over time-frequency blocks of fixed size. We denote by $\mathbf{h}[i]$ the channel vector of a generic user over a time-frequency resource block indexed i. One can assume that the pilots are transmitted over resource blocks that are sufficiently separated in time and/or frequency such that the resulting channel samples are i.i.d. A popular (yet simplistic) view of the wireless channel holds that the channel gain per antenna element can be seen as a weighted superposition of the element's response to signals impinging on the array from separable AoAs [32,40,98,135]. Formally, for a realization $\mathbf{h}[i]$, we can express the channel according to this view as

$$\mathbf{h}[i] = \sum_{\ell \in \mathbb{Z}_L} w_\ell[i]\mathbf{a}(\theta_\ell), \tag{2.4}$$

where L is the number of separable wave-fronts, $\{\theta_\ell\}_\ell$ are the associated AoAs with each path, and $w_\ell[i]$ is the i-th realization of the random channel coefficient of path ℓ. A standard assumption is to take $w_\ell[i]$ to be zero-mean, complex Gaussian with variance σ_ℓ^2, i.e., $w_\ell \sim \mathcal{CN}(0, \sigma_\ell^2)$. The AoA θ_ℓ belongs to the angular interval $[-\theta_{\max}, \theta_{\max})$, where $\theta_{\max} \in [0, \frac{\pi}{2}]$ is the maximum array angular aperture. The vector $\mathbf{a}(\theta)$ is the *array response vector*, mapping θ to an M-dimensional vector of complex exponentials

$$\mathbf{a}(\theta) = [1, e^{j\frac{2\pi d}{\lambda}\sin(\theta)}, \dots, e^{j\frac{2\pi d}{\lambda}(M-1)\sin(\theta)}]^\mathsf{T}, \tag{2.5}$$
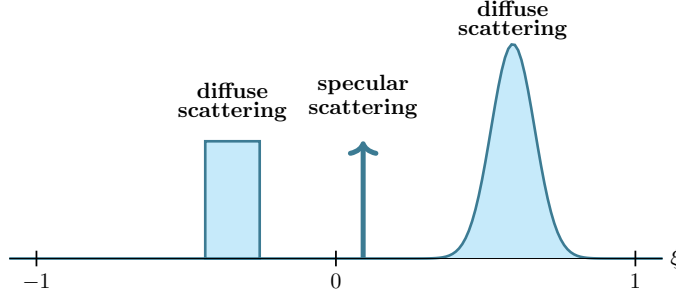
Figure 2.1: specular scattering vs diffuse scattering

where d is the distance between consecutive antenna elements, and $\lambda = c/f$ is the electromagnetic wavelength, with f the carrier frequency of the narrow-band transmitted signal and $c \approx 3 \times 10^8$ the speed of light. For convenience, in this chapter we take the antenna spacing to be $d = \frac{\lambda}{2\sin(\theta_{\max})}$, and we *normalize* the AoA and define the new variable

$$\xi := \frac{\sin(\theta)}{\sin(\theta_{\max})} \in [-1, 1). \tag{2.6}$$

Here and below we use the term AoA to refer to this normalized quantity. Then, the array response can be written in terms of ξ as

$$\mathbf{a}(\xi) = [1, e^{j\pi\xi}, \ldots, e^{j\pi(M-1)\xi}]^{\mathsf{T}}. \tag{2.7}$$

The channel model in (2.4) is inconsistent with the physical evidence provided by some measurement campaigns [93, 102, 106, 115]. These studies show that the paths contributing to the channel can be considered discrete and separable only if they correspond to line-of-sight (LoS) and *specular* scattering elements in the propagation environment. But these are not the only scattering elements: some scatterers induce *diffuse* patterns corresponding to a group of a large number of closely located components. Fig. 2.1 illustrates the difference by distinguishing between two different types of scattering. In our opinion, this distinction is fundamental and we will refer to it repeatedly throughout this dissertation. We propose a more general model of the channel as follows. The *wide-sense stationary uncorrelated scattering* (WSSUS) channel model asserts that scattering over disjoint angular intervals induces uncorrelated scattering coefficients [105]. Focusing on Gaussian channels, these coefficients can be best described as a Lévy process $W : [-1, 1] \to \mathbb{C}$ with stationary, uncorrelated complex Gaussian (hence independent) increments $dW(\xi) = W(\xi + d\xi) - W(\xi)$ [66]:

$$\mathbb{E}\left[dW(\xi)\right] = 0, \tag{2.8a}$$

$$\mathbb{E}\left[dW(\xi)dW(\xi')^*\right] = d\Gamma(\xi)\delta(\xi - \xi'), \tag{2.8b}$$

where $\delta(\cdot)$ denotes Dirac's delta. The autocorrelation Γ is a non-negative, non-increasing, right-continuous function whose increment $d\Gamma(\xi) = \Gamma(\xi + d\xi) - \Gamma(\xi) \geqslant 0$ gives the variance of $dW(\xi)$. The instantaneous channel vector at the ULA can be expressed via the following stochastic integral [28]:

$$\mathbf{h}[i] = \int_{-1}^{1} dW(\xi)[i]\, \mathbf{a}(\xi) \in \mathbb{C}^M,\ i \in \mathbb{Z}_N, \tag{2.9}$$

where $dW(\xi)[i]$ is an increment corresponding to the the i-th realization of the channel gains stochastic process and $\mathbf{h}[i]$ is the i-th realization of the channel vector $\mathbf{h}$. Obviously $\mathbf{h}$ is complex Gaussian distributed. Its mean is given by

$$\bar{\mathbf{h}} = \mathbb{E}\left[\mathbf{h}\right] = \int_{-1}^{1} \mathbb{E}\left[dW(\xi)\right]\mathbf{a}(\xi) = \mathbf{0}, \tag{2.10}$$

and its covariance is given by

$$\begin{aligned}
\boldsymbol{\Sigma} = \mathbb{E}\left[\mathbf{h}\mathbf{h}^{\mathsf{H}}\right] &= \int_{-1}^{1}\int_{-1}^{1} \mathbb{E}\left[dW(\xi)dW(\xi')^*\right]\mathbf{a}(\xi)\mathbf{a}(\xi')^{\mathsf{H}} \\
&= \int_{-1}^{1} d\Gamma(\xi)\mathbf{a}(\xi)\mathbf{a}(\xi)^{\mathsf{H}} \in \mathbb{T}_{+}^{M},
\end{aligned} \tag{2.11}$$

where $\mathbb{T}_{+}^{M}$ denotes the space of $M \times M$ positive semi-definite (PSD) Toeplitz matrices.

2.3.1 Decomposing the Angular Gain Process and the ASF

From a physics viewpoint, the channel angular coefficients process W should represent the contribution of LoS and specular scattering as well as that of diffuse scattering. This is naturally modeled by considering the Lévy-Itô decomposition of a Lévy process [72]. The related theorem states that W can be decomposed as a sum of three independent Lévy processes as

$$W = W_c + W_d + W_{sc} \tag{2.12}$$

where W_c is a Wiener process, W_d is a compound Poisson process (in this case with complex Gaussian jumps) and W_{sc} is a square-integrable pure jump martingale. Replacing (2.12) in (2.8b), one can show that the autocorrelation increment has the following decomposition:

$$d\Gamma = d\Gamma_c + d\Gamma_d + d\Gamma_{sc}, \tag{2.13}$$

where

- $d\Gamma_c$ is absolutely continuous.

- $d\Gamma_d$ is a *pure point* process, also known as a discrete measure, which is supported on a countable set of points.

- $d\Gamma_{sc}$ is a singular continuous measure, that is supported on a Lebesgue null-set.

The relation in (2.13) also follows from the Lebesgue decomposition theorem [55]. Throughout this dissertation we deal with Lebesgue integrals involving the channel angular autocorrelation. Since the contribution of $d\Gamma_{sc}$ vanishes in the result of a Lebesgue integral, we simplify the decomposition (2.13) by ignoring dW_{sc} and recognizing the fact that, any estimate of $d\Gamma$ is valid up to an additive singular-continuous measure. This will not have any effect on the result of the practical algorithms that follow. Furthermore, the discrete measure $d\Gamma_d$ can be represented using a train of Dirac deltas:

$$d\Gamma_d(\xi) = \gamma_d(\xi)d\xi = \sum_{k \in \mathbb{Z}_K} c_k \delta(\xi - \xi_k)d\xi, \, \xi \in [-1, 1], \tag{2.14}$$

where $\{c_k \in \mathbb{R}_+\}_{k=0}^{K-1}$ are non-negative coefficients, and $\{\xi_k\}_{k=0}^{K-1}$ denotes the support of γ_d. For simplicity, here we assume that γ_d is supported over a finite rather than a countable set of cardinality K.[1] Note that (2.14) involves an abuse of notation, since the delta does not have a Radon-Nikodym derivative. Nonetheless we use this notation for convenience and define the integral of a function f with respect to the delta measure as $\int_{-1}^{1} f(\xi)\delta(\xi-a)d\xi = f(a)$ for $a \in [-1, 1]$.

For the absolutely continuous component $d\Gamma_c$ we can use the Radon-Nikodym theorem to write

$$d\Gamma_c(\xi) = \gamma_c(\xi)d\xi, \, \xi \in [-1, 1], \tag{2.15}$$

where γ_c is a non-negative measure over. Analogous to (2.14), we can represent γ_c as the following linear combination:

$$\gamma_c(\xi) = \sum_{k \in \mathbb{Z}_{K'}} \widetilde{c}_k g_k(\xi). \tag{2.16}$$

Here K' is the number of diffuse scattering elements, $\widetilde{c}_k \in \mathbb{R}_+$ is a scalar, and g_k is an absolutely continuous, non-negative measure over $[-1, 1]$ for $k = 0, \ldots, K' - 1$. As an example, g_k can be modeled as a rectangular, a truncated Gaussian function or the density function of a von Mises distribution [100, 106].

From a physics standpoint, the discrete measure γ_d represents the channel power density received from LoS and specular scattering associated with the discrete AoAs ξ_k, $k =$

---

[1]This does not incur a loss of generality, since eventually we only consider recovery of a finite support for γ_d for the purpose of covariance estimation.
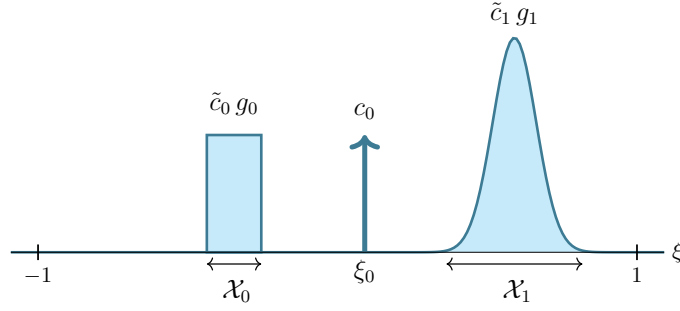
Figure 2.2: an example ASF with $K = 1$ specular component and $K' = 2$ diffuse components

$0, \ldots, K - 1$. The coefficient c_k represents the channel power along the k-th discrete AoA. In contrast, $\gamma_c(\xi)$ represents the channel power density over the diffuse scatterer associated with the interval $[\xi, \xi + d\xi]$. An element g_k, $k \in \mathbb{Z}_{K'}$ in (2.16) represents the power density associated with the k-th diffuse scatterer. Typically, the power received from such a scatterer occupies a limited subset of the angular interval, which we denote by $\mathcal{X}_k$ and call it the *effective* support of g_k.

Overall we are in a position to define the channel *angular scattering function* (ASF) $\gamma = d\Gamma$ as

$$\gamma = \gamma_c + \gamma_d, \tag{2.17}$$

where we call γ_c and γ_d the "continuous" and "discrete" ASF components. Fig. 2.2 illustrates an example of the ASF and its discrete and continuous components.

Remark 2.1. Failure to recognize the difference between the discrete channel model in (2.4) and the continuous model in (2.9) can lead to an unrealistic model of practical channels and wrong judgments about the performance of a channel processing algorithm. The diffuse ASF component significantly contributes to the MIMO channel total power and capacity [102]. Besides, it has been shown that neglecting the diffuse ASF component can lead to complications in identifying the dominant specular components of the channel [115]. Therefore, a more general characterization of the channel as developed above seems necessary.

Using Eq. (2.11) we can relate the ASF to the channel covariance as

$$\mathbf{\Sigma} = \int \mathbf{a}(\xi)\mathbf{a}(\xi)^{\mathsf{H}}\gamma(\xi)d\xi = \int \mathbf{a}(\xi)\mathbf{a}(\xi)^{\mathsf{H}}\gamma_c(\xi)d\xi + \sum_{k=0}^{K-1} c_k\mathbf{a}(\xi_k)\mathbf{a}(\xi_k)^{\mathsf{H}}. \tag{2.18}$$

With this we can also define the *class* of ULA channel covariances as

$$\mathcal{C}_{\mathsf{ULA}} = \left\{ \mathbf{\Sigma} \in \mathbb{C}^{M \times M} : \mathbf{\Sigma} = \int_{-1}^{1} \mathbf{a}(\xi)\mathbf{a}(\xi)^{\mathsf{H}} \gamma(\xi) d\xi, \ \gamma : [-1, 1] \to \mathbb{R}_{+}, \ \int_{-1}^{1} \gamma(\xi) d\xi = 1 \right\}, \tag{2.19}$$

where the the constraint on γ to be normalized such that $\int_{-1}^{1} \gamma(\xi) d\xi = 1$ is added for simplicity and w.l.o.g. It can be easily observed that $\mathbf{\Sigma}$ is an element in the convex hull of the continuous dictionary $\mathcal{A} = \left\{ \mathbf{a}(\xi)\mathbf{a}(\xi)^{\mathsf{H}}, \ \xi \in [-1, 1] \right\}$, and

$$\mathcal{C}_{\mathsf{ULA}} = \mathrm{Conv}\left(\mathcal{A} \right). \tag{2.20}$$

We exploit this structure in developing our proposed covariance estimators.

2.4 Proposed Covariance Estimation Methods

Our methods for channel covariance estimation rely on a finite- but high-dimensional approximation of the ASF, estimating the approximation components and using them to obtain the channel covariance. First, note that the discrete part of the ASF has a finite-dimensional parametric representation, with parameters $\{c_k\}_k$ and $\{\xi_k\}_k$. However, the constituents of γ_c are the infinite-dimensional measures g_k, $k \in \mathbb{Z}_{K'}$. We propose to approximate γ_c as a linear combination of the elements of a family of *limited-support* densities

$$\Psi = \{\psi_i \ : \ i \in \mathbb{Z}_{G_c}\}, \tag{2.21}$$

where ψ_i, is a density supported on a small subset $\mathrm{supp}(\psi_i)^2$ of the angular interval such that $|\mathrm{supp}(\psi_i)| \ll 1$, and the cardinality of Ψ is taken to be large ($G_c \gg M$) (typical values are $G_c = 2M$ and $G_c = 3M$). The approximation is then written as

$$\gamma_c \approx \sum_{i \in \mathbb{Z}_{G_c}} c_i' \psi_i. \tag{2.22}$$

Note that this relation only states that the approximation lies in the span of the columns of Ψ (span(Ψ)), but does not specify the properties of the approximation coefficients.

The density family Ψ can be specified in various ways. We propose two simple constructions as follows. Let $\psi^{\star} : [-1, 1] \to \mathbb{R}_{+}$ be a real, non-negative function whose support is limited to $[0, \frac{2}{G_c}]$, and define the density family to be consisting of shifted and non-overlapping versions of $\psi^{\star}$, i.e. $\psi_i(\xi) = \psi^{\star}(\xi + 1 - \frac{2i}{G_c})$, $i \in \mathbb{Z}_{G_c}$, $\xi \in [-1, 1)$. Examples of simple density families include the following:

---

[2]Here and in the rest of this dissertation we define the *support* of a measure f as $\mathrm{supp}(f) = \{x \ : \ f(x) \neq 0\}$.
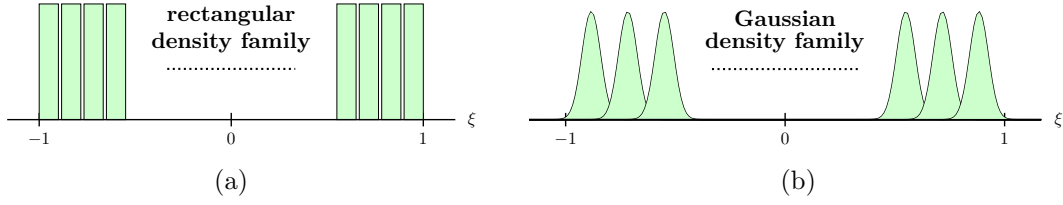
Figure 2.3: examples of density families used for approximating γ_c: (a) rectangular density family and (b) Gaussian density family.

- *impulse densities:* here we define $\psi^\star$ as a Dirac impulse at the origin, i.e.

$$\psi^\star(\xi) = \delta(\xi), \ \xi \in [-1, 1]. \tag{2.23}$$

- *rectangular densities:* in this case, we define $\psi^\star$ as a rectangular pulse over $[0, \frac{2}{G_c}]$, that is

$$\psi^\star = \frac{G_c}{2}\mathsf{rect}_{[0, \frac{2}{G_c}]}. \tag{2.24}$$

- *Gaussian denisites:* here $\psi^\star$ is defined as a Gaussian function over $[0, \frac{2}{G_c}]$:

$$\psi^\star(\xi) = \begin{cases} \frac{1}{a_0\sigma\sqrt{2\pi}}e^{-\frac{(\xi-\mu)^2}{2\sigma^2}}, & \xi \in [0, \frac{2}{G_c}], \\ 0, & \xi \notin [0, \frac{2}{G_c}], \end{cases} \tag{2.25}$$

where $\mu = 1/G_c$ and $\sigma = \frac{1}{G_c\sqrt{2\log(10)}}$ to ensure that the value of the Gaussian function drops to 0.1 of its peak at the boundary points of $(0, \frac{2}{G_c})$, and a_0 is a normalization scalar such that $\int_{-1}^{1}\psi^\star(\xi)d\xi = 1$.

Fig. 2.3 illustrates these two density families, and Fig. 2.4 illustrates an example of approximating the continuous ASF component with the rectangular density family. The limited support property and the large number of density elements helps with a high-resolution approximation of the ASF by capturing its *local* angular properties. Using (2.11), (2.17), and (2.22) we can assume an approximation of the channel covariance as

$$\boldsymbol{\Sigma} := \boldsymbol{\Sigma}_d + \boldsymbol{\Sigma}_c \approx \sum_{k\in\mathbb{Z}_K} c_k\mathbf{a}(\xi_k)\mathbf{a}(\xi_k)^\mathsf{H} + \sum_{i\in\mathbb{Z}_{G_c}} c_i'\mathbf{A}_i', \tag{2.26}$$

where we have defined the matrices $\mathbf{A}_i' = \int_{-1}^{1}\psi_i(\xi)\mathbf{a}(\xi)\mathbf{a}(\xi)^\mathsf{H}d\xi$ for $i \in \mathbb{Z}_{G_c}$, and the discrete and continuous covariance components as $\boldsymbol{\Sigma}_d$ and $\boldsymbol{\Sigma}_c$, respectively. One can easily show that with $\mathbf{a}$ as in (2.7), $\mathbf{A}_i'$ is Hermitian Toeplitz ($\mathbf{A}_i' \in \mathbb{T}^M$) for all i and therefore is fully characterized by its first column $\mathbf{a}_i' = \int_{-1}^{1}\psi_i(\xi)\mathbf{a}(\xi)d\xi$. Since $\boldsymbol{\Sigma}$ also belongs to $\mathbb{T}_+^M$, we
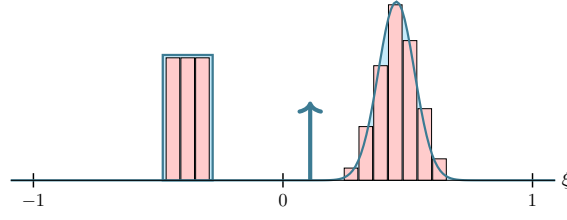
Figure 2.4: An example of approximating the continuous ASF with rectangular densities.

can explain it by its first column $\boldsymbol{\sigma}$ given by

$$\boldsymbol{\sigma} \approx \sum_{i \in \mathbb{Z}_{G_c}} c_i' \mathbf{a}_i' + \sum_{k \in \mathbb{Z}_K} c_k \mathbf{a}(\xi_k). \tag{2.27}$$

Defining $\mathbf{A} := [\mathbf{a}_0', \ldots, \mathbf{a}_{G_c-1}', \mathbf{a}(\xi_0), \ldots, \mathbf{a}(\xi_{K-1})]$ and $\boldsymbol{\gamma} := [c_0', \ldots, c_{G_c-1}', c_0, \ldots, c_{K-1}]^{\mathsf{T}} \in \mathbb{R}_+^d$ where $d := G_c + K$, we can write (2.27) as

$$\boldsymbol{\sigma} \approx \mathbf{A}\boldsymbol{\gamma}, \tag{2.28}$$

and in addition $\boldsymbol{\Sigma}(\boldsymbol{\gamma}) \approx \mathcal{T}(\mathbf{A}\boldsymbol{\gamma})$. Note that the discrete AoAs $\{\xi_k\}_{k \in \mathbb{Z}_K}$ are unknown and belong to the continuum $[-1, 1]$, and therefore the rank-1 matrices $\{\mathbf{a}(\xi_k)\mathbf{a}(\xi_k)^{\mathsf{H}}\}_{k=0}^{K-1}$ are not known a priori (in contrast to the known matrices $\mathbf{A}_i'$, $i \in \mathbb{Z}_{G_c}$). As a first step in our covariance estimation algorithm, we propose to estimate the support of γ_d, that is the discrete AoAs $\{\xi_k\}_{k \in \mathbb{Z}_K}$ from the set of noisy pilot measurements.

2.4.1 Estimating the Discrete AoAs

An accurate estimation of the discrete AoAs is crucial, not only because in many practical scenarios they may carry a significant power of the channel, but also that we are modeling the ASF as to include Dirac impulses. While a piece-wise continuous function can be well-approximated by a family of "basis" densities, these measures are not well represented by the same family, in the sense that, the approximation error

$$\min_{\{c_i \in \mathbb{R}\}_i} \int_{-1}^{1} |\delta(\xi) - \sum_i c_i \psi_i(\xi)|^2 d\xi, \tag{2.29}$$

is unbounded. This of course does not mean that if one neglects estimating the discrete AoA the covariance estimation error will be unbounded, but gives an intuitive hint that estimating discrete AoAs can result in higher accuracy.

In order to estimate the discrete AoAs from the pilot measurements $\{\mathbf{y}[i]\}_{i \in \mathbb{Z}_N}$, we employ the well-known Multiple Signal Classification (MUSIC) method [109, 117]. As a super-resolution method, MUSIC is typically used for estimating the frequencies of

multiple sinusoids from their (possibly noisy) mixture. Similarly, here we use MUSIC to estimate the angular frequencies $\{\xi_k\}_{k \in \mathbb{Z}_K}$. Recall the pilot sample covariance matrix $\widehat{\boldsymbol{\Sigma}}_\mathbf{y} = \frac{1}{N} \sum_i \mathbf{y}[i] \mathbf{y}[i]^\mathsf{H}$ and define its eigendecomposition as $\widehat{\boldsymbol{\Sigma}}_\mathbf{y} = \widehat{\mathbf{U}} \widehat{\boldsymbol{\Lambda}} \widehat{\mathbf{U}}^\mathsf{H}$, where $\widehat{\mathbf{U}} = [\widehat{\mathbf{u}}_1, \ldots, \widehat{\mathbf{u}}_M]$ denotes the eigenvectors matrix, and $\widehat{\boldsymbol{\Lambda}}$ is the diagonal eigenvalues matrix, with its diagonal elements ordered as $\widehat{\lambda}_1 \geqslant \widehat{\lambda}_2 \geqslant \ldots \geqslant \widehat{\lambda}_M$. Assume that an estimate of the number of discrete AoAs is given and denote it by $\widehat{K}$. Then, we can define the so-called "noise subspace" as the subspace spanned by the columns of the $(M - \widehat{K})$ eigenvectors in $\widehat{\mathbf{U}}$ corresponding to the smallest $(M - \widehat{K})$ eigenvalues in $\widehat{\boldsymbol{\Lambda}}$, namely the columns of $\widehat{\mathbf{U}}_{\text{noise}} = \left[\widehat{\mathbf{u}}_{\widehat{K}+1}, \ldots, \widehat{\mathbf{u}}_M \right]$ MUSIC estimates the discrete AoAs by finding the $\widehat{K}$ smallest minimizers of the *pseudo-spectrum* function

$$\eta(\xi) = \|\widehat{\mathbf{U}}_{\text{noise}}^\mathsf{H} \mathbf{a}(\xi)\|^2 = \sum_{\ell = \widehat{K}+1}^{M} \left| \widehat{\mathbf{u}}_\ell^\mathsf{H} \mathbf{a}(\xi) \right|^2. \tag{2.30}$$

We denote the estimated discrete AoAs by $\{\widehat{\xi}_k\}_{k \in \mathbb{Z}_{\widehat{K}}}$. When the observations are generated by a noisy superposition of a finite number of weighted tones, MUSIC asymptotically gives consistent estimates of the tones. In the context of our problem, this scenario translates to the case in which the ASF consists of only a discrete component and the tones are the discrete AoAs. However, in general the channel is not only a product of the discrete ASF component, but a mixture of discrete and continuous scattering components. Fortunately, a recent result has shown that even in this case, under some mild conditions on the energy distribution of the discrete and continuous parts as well as the signal dimension, MUSIC is able to consistently estimate the discrete AoAs [89]. The following theorem states a slightly modified version of this result to justify the expected success of MUSIC in identifying discrete ASF AoAs from the noisy pilot observations.

Theorem 2.1. *Consider an ASF $\gamma(\xi) = \gamma_d(\xi) + \gamma_c(\xi) = \sum_{k=1}^{K} c_k^{(M)} \delta(\xi - \xi_\ell) + \gamma_c(\xi)$, $\xi \in [-1, 1]$ and assume that the weights $\{c_k^{(M)}\}_{k \in \mathbb{Z}_K}$ may depend on (M). Consider a scaling regime where the number of antennas M and the sample size N both approach infinity such that $\frac{N}{M} \to \zeta > 0$. Then, MUSIC is asymptotically consistent, i.e., $\max_{k \in \mathbb{Z}_K} M|\widehat{\xi}_k - \xi_k| \to 0$ provided that $\limsup_{M \to \infty} \min_{k \in \mathbb{Z}_K} M c_k^{(M)} \geqslant \omega_0(\zeta, \gamma_c)$, where $\omega_0(\zeta, \gamma_c)$ is a finite parameter that depends on ζ and the continuous component γ_c.* □

Proof. See Appendix 2.7.1. □

Given an estimate of the discrete AoAs as $\{\widehat{\xi}_k\}_{k \in \mathbb{Z}_K}$, we have a model that relates the unknown $(G_c + \widehat{K})$-dimensional vector of real, positive coefficients to the covariance matrix via (2.26).

Estimating the Number of Discrete AoAs

The proof of Theorem 2.1 involves a step which shows that, for large M and under some conditions on the amplitude of the discrete impulses c_k, $k = 0, \ldots, K - 1$ and the supremum of the continuous measure γ_c, the K largest eigenvalues of $\widehat{\boldsymbol{\Sigma}}_{\mathbf{y}}$ "escape" from the rest $(M - K)$ eigenvalues, meaning that a large gap between the largest K eigenvalues and the rest occurs. This suggests a way to estimate the number of discrete AoAs. Let $\beta_i = \frac{\widehat{\lambda}_i}{\widehat{\lambda}_0}$, $i = 0, \ldots, M - 1$ denote the normalized eigenvalues of $\widehat{\boldsymbol{\Sigma}}_{\mathbf{y}}$ in descending order $\beta_0 \geqslant \beta_1 \geqslant \ldots \geqslant \beta_{M-1}$. Define the sequence $\{d_i = \beta_i - \beta_{i+1} : i = 0, \ldots, M - 2\}$ by taking differences between consecutive elements of $\{\beta_i\}_i$. Let $K_{\max} \in \mathbb{N}$ denote a pre-defined bound on the maximum number of discrete AoAs. Then, we estimate K as the index for which the largest gap between consecutive ordered eigenvalues occurs. In order to avoid poor estimates of the number of discrete impulses under challenging conditions (e.g. when the impulse amplitudes are small compared to noise and the continuous part), we assume K to be no larger than the pre-defined integer $K_{\max}$. This is mathematically formulated as

$$\widehat{K} = 1 + \underset{k \leqslant K_{\max} - 1}{\arg\max} \; d_k. \tag{2.31}$$

Remark 2.2. Note that the precise estimation of the number of discrete AoAs is *not* critical, and in particular, it is better to over-estimate the number of discrete AoAs, than to under-estimate it. If we over-estimate the number of discrete AoAs, there will be "fake" spikes identified in the support of the discrete ASF component. However, the estimation method (as will be introduced shortly) will assign small coefficients to the fake spikes, which practically means that there is no spike. Under-estimating the number of discrete AoAs can be more harmful, since some of the existing spikes will not be represented in the parametric expression and its contribution will not be fully compensated by the rest of the approximating functions. Nevertheless, even in this case these methods assign a non-zero coefficient to an element of the density dictionary Ψ that has the highest correlation with the "missing" spike. This is obviously sub-optimal, due to the poor approximation of the delta with a continuous density, but the induced error will be controlled.

2.4.2 Covariance Estimation via NNLS

Our first method estimates the non-negative coefficients $\{c_k\}_{k \in \mathbb{Z}_{\widehat{K}}}$ and $\{c'_i\}_{i \in \mathbb{Z}_{G_c}}$ by minimizing the norm of the difference between the parametric form on the right-hand-side (RHS) of (2.27) and a *coarse estimate* of the ULA covariance. Since we know a priori that a ULA covariance is Hermitian Toeplitz, the coarse estimate is obtained by projecting the sample covariance $\widehat{\boldsymbol{\Sigma}}$ in (2.3) onto the space of Hermitian Toeplitz matrices $\mathbb{T}^M$. In other

words, the coarse estimate $\widetilde{\boldsymbol{\Sigma}}$ is given as

$$\widetilde{\boldsymbol{\Sigma}} = \mathcal{P}_{\mathbb{T}}\left(\widehat{\boldsymbol{\Sigma}}\right), \tag{2.32}$$

where $\mathcal{P}_{\mathbb{T}} : \mathbb{C}^{M \times M} \to \mathbb{C}^{M \times M}$ is the orthogonal projector to the space of Hermitian Toeplitz matrices. This projection is given by averaging the diagonals of $\widehat{\boldsymbol{\Sigma}}$, and replacing the diagonal elements by the corresponding average. To be more precise, define $\widetilde{\boldsymbol{\sigma}}$ as an M-dimensional complex-valued vector such that

$$[\widetilde{\boldsymbol{\sigma}}]_i := \frac{1}{2M}\left(\sum_{m=0}^{M-(i+1)}[\widehat{\boldsymbol{\Sigma}}]_{m+i,m} + \sum_{m=0}^{M-(i+1)}[\widehat{\boldsymbol{\Sigma}}]_{m,m+i}^*\right), \ i \in \mathbb{Z}_M. \tag{2.33}$$

Then $\widetilde{\boldsymbol{\Sigma}}$ is defined as the Hermitian Toeplitz matrix with $\widetilde{\boldsymbol{\sigma}}$ as its first column, i.e. $\widetilde{\boldsymbol{\Sigma}} = \mathcal{T}(\widetilde{\boldsymbol{\sigma}})$. The non-negative least-squares (NNLS) estimator of $\boldsymbol{\gamma}$ is given by solving

$$\begin{aligned}
&\underset{\{c_k\}_k, \{c_i'\}_i}{\text{minimize}} \quad \left\|\widetilde{\boldsymbol{\Sigma}} - \sum_{k \in \mathbb{Z}_{\widehat{K}}} c_k \mathbf{a}(\xi_k)\mathbf{a}(\xi_k)^{\mathsf{H}} - \sum_{i=0}^{G_c-1} c_i' \mathbf{A}_i'\right\|_{\mathsf{F}} \\
&\text{subject to} \quad c_k \geqslant 0, \ k \in \mathbb{Z}_{\widehat{K}}, \\
&\phantom{\text{subject to} \quad} c_i' \geqslant 0, \ i \in \mathbb{Z}_{G_c}.
\end{aligned} \tag{2.34}$$

Since the matrix inside the norm expression is Hermitian Toeplitz, we can use (2.28) and write the NNLS estimate of the coefficients as

$$\boldsymbol{\gamma}_{\mathsf{NNLS}}^{\star} = \arg\min_{\boldsymbol{\gamma} \geqslant \mathbf{0}} \left\|\mathbf{M}\left(\widehat{\mathbf{A}}\boldsymbol{\gamma} - \widetilde{\boldsymbol{\sigma}}\right)\right\|, \tag{$\mathcal{P}_{\mathsf{NNLS}}$}$$

where we have defined $\widehat{\mathbf{A}} := [\mathbf{a}_0', \ldots, \mathbf{a}_{G_c-1}', \mathbf{a}(\widehat{\xi}_0), \ldots, \mathbf{a}(\widehat{\xi}_{\widehat{K}-1})]$, $\boldsymbol{\gamma}$ is a $\widehat{d}$-dimensional vector constrained to be non-negative, and

$$\mathbf{M} = \text{diag}\left(\left[\sqrt{M}, \sqrt{2(M-1)}, \sqrt{2(M-2)}, \ldots, \sqrt{2}\right]^{\mathsf{T}}\right).$$

The matrix $\mathbf{M}$ assigns suitable factors to each element to compensate for the number of times an element is repeated in a Hermitian Toeplitz matrix. The covariance estimate in this case is given by

$$\boldsymbol{\Sigma}_{\mathsf{NNLS}}^{\star} = \mathcal{T}\left(\widehat{\mathbf{A}}\boldsymbol{\gamma}_{\mathsf{NNLS}}^{\star}\right). \tag{2.35}$$

To clarify why the NNLS estimator yields a good estimate of the covariance, note that the cost in ($\mathcal{P}_{\mathsf{NNLS}}$) is equal to the Euclidean distance between the coarse covariance estimate (given by the projection onto $\mathbb{T}^M$) and the estimated covariance, by constraining γ to be

a non-negative measure. The constraint enforces the covariance to lie in a subset $\mathcal{C}_{\mathsf{ULA}}^d$ of the set of MIMO covariances $\mathcal{C}_{\mathsf{ULA}}$ in (2.19). This subset is given as

$$
\mathcal{C}_{\mathsf{ULA}}^d = \left\{ \boldsymbol{\Sigma} \in \mathbb{C}^{M \times M} : \boldsymbol{\Sigma} = \int_{-1}^{1} \gamma(\xi) \mathbf{a}(\xi) \mathbf{a}(\xi)^{\mathsf{H}} d\xi, \right.
$$
$$
\left. \gamma(\xi) = \sum_{i=0}^{G_c-1} [\boldsymbol{\gamma}]_i \psi_i(\xi) + \sum_{k=0}^{\widehat{K}-1} [\boldsymbol{\gamma}]_{G_c+k} \delta(\xi - \widehat{\xi}_k), \; \xi \in [-1, 1], \; \boldsymbol{\gamma} \in \mathbb{R}_+^{\widehat{d}} \right\}, \tag{2.36}
$$

noting that $\mathcal{C}_{\mathsf{ULA}}^d \subset \mathcal{C}_{\mathsf{ULA}}$. Therefore, $\mathcal{P}_{\mathsf{NNLS}}$ can be seen as the orthogonal projection of the coarse covariance estimate onto the $\mathcal{C}_{\mathsf{ULA}}^d$, simply described as

$$
\boldsymbol{\Sigma}_{\mathsf{NNLS}}^{\star} = \mathcal{P}_{\mathcal{C}_{\mathsf{ULA}}^d} \left(\widetilde{\boldsymbol{\Sigma}} \right). \tag{2.37}
$$

Remark 2.3. Despite the fact that the coarse covariance estimate $\widetilde{\boldsymbol{\Sigma}}$ is not necessarily PSD, the final covariance estimate after applying NNLS is always PSD due to the non-negativity of the solution coefficients.

2.4.3 Covariance Estimation via Likelihood Maximization

The NNLS estimator is given by minimizing the distance between the parametric form of the covariance and a coarse estimate of it. An alternative approach to estimate the unknown parameters vector $\boldsymbol{\gamma}$ is to maximize the likelihood of the set of observed pilot signals $\{\mathbf{y}[i]\}_{i \in \mathbb{Z}_N}$ as a function of the parametric covariance (equivalently, a function of $\boldsymbol{\gamma}$). Define $\mathbf{Y} = [\mathbf{y}[0], \dots, \mathbf{y}[N-1]] \in \mathbb{C}^{M \times N}$ as the matrix collecting pilot measurements as its columns. Assuming zero-mean Gaussian channels, we can write the likelihood function of $\mathbf{Y}$ given $\boldsymbol{\gamma}$ as (we always assume N_0 to be known)

$$
p(\mathbf{Y}|\boldsymbol{\gamma}) = \prod_{i \in \mathbb{Z}_N} p(\mathbf{y}[i] \,|\, \boldsymbol{\gamma}) = \prod_{i \in \mathbb{Z}_N} \frac{\exp\left(-\frac{1}{2} \mathbf{y}[i]^{\mathsf{H}} \left(\boldsymbol{\Sigma}(\boldsymbol{\gamma}) + N_0 \mathbf{I} \right)^{-1} \mathbf{y}[i] \right)}{\sqrt{(2\pi)^M \det\left(\boldsymbol{\Sigma}(\boldsymbol{\gamma}) + N_0 \mathbf{I} \right)}}
$$
$$
= \frac{\exp\left(-\frac{1}{2} \mathsf{tr}\left(\left(\boldsymbol{\Sigma}(\boldsymbol{\gamma}) + N_0 \mathbf{I} \right)^{-1} \mathbf{Y}\mathbf{Y}^{\mathsf{H}} \right) \right)}{(2\pi)^{\frac{MN}{2}} \left(\det(\boldsymbol{\Sigma}(\boldsymbol{\gamma}) + N_0 \mathbf{I}) \right)^{\frac{N}{2}}}, \tag{2.38}
$$

where $\boldsymbol{\Sigma}(\boldsymbol{\gamma}) = \mathcal{T}(\widehat{\mathbf{A}}\boldsymbol{\gamma})$. Using (2.38) we can form the log-likelihood function $\log p(\mathbf{Y}|\boldsymbol{\gamma})$, which after omitting constants and re-scaling amounts to

$$
f_{\mathsf{ML}}(\boldsymbol{\gamma}) := -\mathsf{tr}\left(\boldsymbol{\Sigma}_{\mathbf{y}}^{-1}(\boldsymbol{\gamma}) \widehat{\boldsymbol{\Sigma}}_{\mathbf{y}} \right) - \log \det\left(\boldsymbol{\Sigma}_{\mathbf{y}}(\boldsymbol{\gamma}) \right) \tag{2.39}
$$

where $\mathbf{\Sigma_y}(\boldsymbol{\gamma}) = \mathbf{\Sigma}(\boldsymbol{\gamma}) + N_0\mathbf{I}$, and $\widehat{\mathbf{\Sigma}}_{\mathbf{y}} := \frac{1}{N}\mathbf{Y}\mathbf{Y}^{\mathsf{H}}$ is the observations sample covariance. Then, we can formulate the maximum likelihood (ML) optimization problem as:

$$
\begin{aligned}
\underset{\boldsymbol{\gamma}}{\text{maximize}} \quad & f_{\mathsf{ML}}(\boldsymbol{\gamma}) \\
\text{subject to} \quad & \boldsymbol{\gamma} \geqslant \mathbf{0},
\end{aligned}
\qquad (\mathcal{P}_{\mathsf{ML}})
$$

Unfortunately, the objective f_{ML} is not concave. In fact, it entails a clear decomposition into the sum of a convex and a concave function $f_{\mathsf{ML}} = f_{\text{cav}} + f_{\text{vex}}$ given as

$$
\begin{aligned}
f_{\text{cav}}(\boldsymbol{\gamma}) &= -\mathsf{tr}\left(\mathbf{\Sigma}_{\mathbf{y}}^{-1}(\boldsymbol{\gamma})\widehat{\mathbf{\Sigma}}_{\mathbf{y}}\right), \\
f_{\text{vex}}(\boldsymbol{\gamma}) &= -\log\det\left(\mathbf{\Sigma}_{\mathbf{y}}(\boldsymbol{\gamma})\right),
\end{aligned}
\qquad (2.40)
$$

for $\boldsymbol{\gamma} \in \mathbb{R}^{\widehat{d}}$. As is well-known, it is in general prohibitively difficult to maximize a non-concave function, due to the existence of local maxima. A standard approach in such cases is to adopt a *minorization-maximization* (aka majorization-minimization) (MM) algorithm. MM is an umbrella term for a variety of iterative optimization algorithms, that consist of two steps [61, 118]: in the minorization step, one finds a surrogate function that is a lower-bound for the original function, and locally approximates it such that the difference between the two is minimized at the current estimation point. The surrogate typically has favorable optimization properties, such as convexity or even that its maximizer can be derived in closed form. In the second MM step we obtain the maximizer and update the current estimate and re-iterate this process until a convergence criterion is met. Examples of an MM algorithm include the expectation maximization (EM) method [37], cyclic minimization [64, 116], and the concave-convex procedure [138]. Here we choose the EM algorithm to find a solution for the ML optimization problem $\mathcal{P}_{\mathsf{ML}}$. As we will see, this algorithm yields a computationally cheap update rule (even for high channel dimensions) and excellent empirical results for the task of estimating the parametric ASF coefficients.

ML via Expectation Maximization

Before laying out the EM algorithm for solving $\mathcal{P}_{\mathsf{ML}}$, let us make a few simplifying assumptions that make the exposition clear and straightforward. Assume that we know the exact location of the discrete AoAs $\{\xi_k\}_{k=0}^{K-1}$ (in the end, we replace these with their estimated values $\{\xi_k\}_{k=0}^{\widehat{K}-1}$ obtained from the MUSIC step). We introduce a *latent variable* vector $\mathbf{x} = [x_0, \ldots, x_{d-1}]^{\mathsf{T}} \in \mathbb{C}^d$, where $d = G_c + K$. We define the components of this variable as follows. Recall the approximation of the continuous part of the ASF γ_c with a family of density functions $\{\psi_i\}_{i \in \mathbb{Z}_{G_c}}$ in (2.22), and assume that each density ψ_i is supported on an interval Ξ_i and the supports of distinct densities are non-overlapping,

i.e. $\Xi_i \cap \Xi_j = \varnothing$ for $i \neq j$. Now, consider an approximation of the continuous part of the angular gain process W_c (see (2.12)) in terms of the densities $\{\psi_i\}_{i \in \mathbb{Z}_{G_c}}$ as $W_c(\xi) \approx \sum_{i=0}^{G_c-1} x_i \psi_i$. Considering in addition the discrete process W_d to be expressed as $W_d(\xi) = \sum_{k=0}^{K-1} x_{G_c+k} \delta(\xi - \xi_k)$, we conclude the following approximation for the overall angular gain process W:

$$W(\xi) \approx \sum_{i=0}^{G_c-1} x_i \psi_i(\xi) + \sum_{k=0}^{K-1} x_{G_c+k} \delta(\xi - \xi_k). \tag{2.41}$$

This is clearly an approximation, since in general W is infinite-dimensional, and any expansion of it in terms of a finite set of functions entails a non-trivial error. Using (2.41) and (2.9), we can approximate the channel vector as

$$\mathbf{h} = \int_{-1}^{1} \mathbf{a}(\xi) dW(\xi) \approx \sum_{i=0}^{G_c-1} x_i \int_{-1}^{1} \psi_i(\xi) \mathbf{a}(\xi) d\xi + \sum_{k=0}^{K-1} x_{G_c+k} \mathbf{a}(\widehat{\xi}_k) = \mathbf{A}\mathbf{x}, \tag{2.42}$$

where $\mathbf{x} := [x_0, \ldots, x_{d-1}]$, $\mathbf{A}_{\cdot,i} = \int_{-1}^{1} \psi_i(\xi) \mathbf{a}(\xi) d\xi$ for $i = 0, \ldots, G_c - 1$, and $\mathbf{A}_{\cdot,G_c+k} = \mathbf{a}(\xi_k)$, for $k = 0, \ldots, K - 1$. Considering N samples of the channel, we have

$$\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{Z}, \tag{2.43}$$

where $\mathbf{Y} = [\mathbf{y}[0], \ldots, \mathbf{y}[N-1]]$, $\mathbf{X} = [\mathbf{x}[0], \ldots, \mathbf{x}[N-1]]$, and $\mathbf{Z} = [\mathbf{z}[0], \ldots, \mathbf{z}[N-1]]$. Note that $\mathbf{x}[i]$ is the i-th realization of the random vector $\mathbf{x}$, and similarly for $\mathbf{Y}$ and $\mathbf{Z}$. Vectorizing both sides of (2.43), we can rewrite it as

$$\widetilde{\mathbf{y}} = \widetilde{\mathbf{A}}\widetilde{\mathbf{x}} + \widetilde{\mathbf{z}} \in \mathbb{C}^{MN}, \tag{2.44}$$

where $\widetilde{\mathbf{y}} := \mathrm{vec}(\mathbf{Y})$, $\widetilde{\mathbf{x}} := \mathrm{vec}(\mathbf{X}) \in \mathbb{C}^{Nd}$, $\widetilde{\mathbf{z}} := \mathrm{vec}(\mathbf{Z})$, and $\widetilde{\mathbf{A}} := \mathbf{I} \otimes \mathbf{A} \in \mathbb{C}^{MN \times Nd}$. Using the new variables, we can write the log-likelihood function as

$$L(\boldsymbol{\gamma}) = \log p(\widetilde{\mathbf{y}}|\boldsymbol{\gamma}) = \log \mathbb{E}_{\widetilde{\mathbf{x}}|\boldsymbol{\gamma}} [p(\widetilde{\mathbf{y}}|\widetilde{\mathbf{x}}, \boldsymbol{\gamma})], \tag{2.45}$$

where we have used conditioning on $\widetilde{\mathbf{x}}$ to derive the RHS. The function L is equivalent to f_{ML} up to additive and multiplicative constants. Given the current estimate of the parameters vector as $\boldsymbol{\gamma}^{(t)}$, in the E-step of the EM algorithm, we compute the following function:

$$\begin{aligned} g(\boldsymbol{\gamma}|\boldsymbol{\gamma}^{(t)}) &= \mathbb{E}_{\widetilde{\mathbf{x}}|\widetilde{\mathbf{y}}, \boldsymbol{\gamma}^{(t)}} [\log p(\widetilde{\mathbf{y}}, \widetilde{\mathbf{x}}|\boldsymbol{\gamma})] \\ &= \mathbb{E}_{\widetilde{\mathbf{x}}|\widetilde{\mathbf{y}}, \boldsymbol{\gamma}^{(t)}} [\log p(\widetilde{\mathbf{y}}|\widetilde{\mathbf{x}}) + \log p(\widetilde{\mathbf{x}}|\boldsymbol{\gamma})] \qquad \text{(E-Step)} \\ &= \mathbb{E}_{\widetilde{\mathbf{x}}|\widetilde{\mathbf{y}}, \boldsymbol{\gamma}^{(t)}} [\log p(\widetilde{\mathbf{x}}|\boldsymbol{\gamma})] + \mathsf{const}. \end{aligned}$$

Using Jensen's inequality, one can show that $g(\boldsymbol{\gamma}|\boldsymbol{\gamma}^{(t)})$ is a lower bound for $L(\boldsymbol{\gamma})$, i.e. $g(\boldsymbol{\gamma}|\boldsymbol{\gamma}^{(t)}) \leqslant L(\boldsymbol{\gamma})$, $\boldsymbol{\gamma} \in \mathbb{R}_+^d$ [118]. In order to compute $g(\boldsymbol{\gamma}|\boldsymbol{\gamma}^{(t)})$, we need to first compute the posterior distribution $p(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}, \boldsymbol{\gamma}^{(t)})$. Using the Bayes rule, we have

$$p(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}, \boldsymbol{\gamma}^{(t)}) = \frac{p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}, \boldsymbol{\gamma}^{(t)})p(\tilde{\mathbf{x}}|\boldsymbol{\gamma}^{(t)})}{p(\tilde{\mathbf{y}}|\boldsymbol{\gamma}^{(t)})}. \tag{2.46}$$

Here, the term $p(\tilde{\mathbf{x}}|\boldsymbol{\gamma}^{(t)})$ represents the prior on $\tilde{\mathbf{x}}$. One can decide on this prior in various ways, for example by considering structural assumptions such as sparsity [114]. Here we propose the following construction for a prior on $\tilde{\mathbf{x}}$. Consider $\{x_i\}_{i\in\mathbb{Z}_d}$ in (2.41) to be the coefficients that give the minimum mean-squared approximation error defined as

$$
\begin{aligned}
E(\mathbf{q}) &= \mathbb{E}_W \left[\left\| W - \sum_{i=0}^{G_c-1} q_i\psi_i - \sum_{k=0}^{K-1} q_{G_c+k}\delta(\cdot - \xi_k) \right\|^2 \right] \\
&\overset{(a)}{=} \mathbb{E}_W \left[\left\| W_c - \sum_{i=0}^{G_c-1} q_i\psi_i \right\|^2 + \left\| W_d - \sum_{k=0}^{K-1} q_{G_c+k}\delta(\cdot - \xi_k) \right\|^2 \right]
\end{aligned}
\tag{2.47}
$$

where (a) follows from the independence of W_c and W_d. Therefore, we define $\mathbf{x}$ as

$$\mathbf{x} = \arg\min_{\mathbf{q}\in\mathbb{C}^d} E(\mathbf{q}). \tag{2.48}$$

Taking the derivative of E and setting it to zero, we can compute the elements of $\mathbf{x}$ as

$$x_i = \frac{1}{\|\psi_i\|^2}\langle\psi_i, W_c\rangle = \frac{2}{G_c}\langle\psi_i, W_c\rangle = \int_{\Xi_i} dW_c(\xi), \ i \in \mathbb{Z}_{G_c}, \tag{2.49a}$$

$$x_{G_c+k} = dW_d(\xi_k), \ k \in \mathbb{Z}_K, \tag{2.49b}$$

where we have used the mutual orthogonality of the densities, i.e. $\langle\psi_i, \psi_j\rangle = 0$, $i \neq j$, which follows from our assumption on the non-overlapping density supports. From (2.49a) we can infer the following properties about the elements of $\mathbf{x}$:

- Each element x_j with $j \in \mathbb{Z}_{G_c}$ is a complex Gaussian random variable, with zero mean and variance

$$
\begin{aligned}
\mathrm{Var}(x_j) &= \mathbb{E}\left[\int_{\Xi_i}\int_{\Xi_j} dW_c(\xi)dW_c(\xi')^* \right] = \int_{\Xi_j} \gamma_c(\xi)d\xi \\
&= [\boldsymbol{\gamma}]_j, \ j \in \mathbb{Z}_{G_c}.
\end{aligned}
\tag{2.50}
$$

- Each element x_{G_c+k} with $k \in \mathbb{Z}_K$ is a complex Gaussian random variable, with zero mean and variance

$$\mathrm{Var}(x_{G_c+k}) = \mathbb{E}\left[dW_d(\xi_k)dW_d(\xi_k)^*\right] = [\boldsymbol{\gamma}]_{G_c+k}, \ k \in \mathbb{Z}_K. \tag{2.51}$$

- The elements of $\mathbf{x}$ are mutually independent.

This shows that $\mathbf{x} \sim \mathcal{CN}(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{x}}(\boldsymbol{\gamma}))$, where

$$\boldsymbol{\Sigma}_{\mathbf{x}}(\boldsymbol{\gamma}) = \mathrm{diag}(\boldsymbol{\gamma}), \tag{2.52}$$

Since the columns of $\mathbf{X}$ are independent, we have $\widetilde{\mathbf{x}} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I} \otimes \boldsymbol{\Sigma}_{\mathbf{x}}(\boldsymbol{\gamma}))$, and the prior is given as

$$p(\widetilde{\mathbf{x}}|\boldsymbol{\gamma}) \sim \mathcal{CN}(\mathbf{0}, \mathbf{I} \otimes \boldsymbol{\Sigma}_{\mathbf{x}}(\boldsymbol{\gamma})). \tag{2.53}$$

Furthermore, we can write the posterior distribution of the observation as

$$p(\widetilde{\mathbf{y}}|\widetilde{\mathbf{x}}, \boldsymbol{\gamma}^{(t)}) = p(\widetilde{\mathbf{y}}|\widetilde{\mathbf{x}}) \sim \mathcal{CN}(\widetilde{\mathbf{A}}\widetilde{\mathbf{x}}, N_0\mathbf{I}). \tag{2.54}$$

Then, from the Bayes rule (2.46), we can compute the posterior $p(\widetilde{\mathbf{x}}|\widetilde{\mathbf{y}}, \boldsymbol{\gamma}^{(t)})$ as

$$p(\widetilde{\mathbf{x}}|\widetilde{\mathbf{y}}, \boldsymbol{\gamma}^{(t)}) \sim \mathcal{CN}\left(\boldsymbol{\mu}_t, \mathbf{C}_t\right), \tag{2.55}$$

where the posterior covariance and mean are explicitly given, respectively, as

$$\mathbf{C}_t = \left(\frac{1}{N_0}\widetilde{\mathbf{A}}^{\mathsf{H}}\widetilde{\mathbf{A}} + N_0(\mathbf{I} \otimes \boldsymbol{\Sigma}_{\mathbf{x}}(\boldsymbol{\gamma}^{(t)}))^{-1}\right)^{-1}$$
$$= \left(\frac{1}{N_0}\mathbf{I} \otimes (\mathbf{A}^{\mathsf{H}}\mathbf{A}) + \mathbf{I} \otimes \boldsymbol{\Sigma}_{\mathbf{x}}^{-1}(\boldsymbol{\gamma}^{(t)})\right)^{-1} \tag{2.56a}$$
$$= \mathbf{I} \otimes \left(\frac{1}{N_0}\mathbf{A}^{\mathsf{H}}\mathbf{A} + \boldsymbol{\Sigma}_{\mathbf{x}}^{-1}(\boldsymbol{\gamma}^{(t)})\right)^{-1} \in \mathbb{C}^{Nd \times Nd}$$

$$\boldsymbol{\mu}_t = \frac{1}{N_0}\mathbf{C}_t\widetilde{\mathbf{A}}^{\mathsf{H}}\widetilde{\mathbf{y}} \in \mathbb{C}^{Nd}. \tag{2.56b}$$

Remark 2.4. Note that the posterior is computed by writing the identity

$$p(\widetilde{\mathbf{y}}|\boldsymbol{\gamma}^{(t)})p(\widetilde{\mathbf{x}}|\widetilde{\mathbf{y}}, \boldsymbol{\gamma}^{(t)}) = p(\widetilde{\mathbf{y}}|\widetilde{\mathbf{x}}, \boldsymbol{\gamma}^{(t)})p(\widetilde{\mathbf{x}}|\boldsymbol{\gamma}^{(t)}),$$

expanding the RHS, and collecting the terms involving $\widetilde{\mathbf{x}}$ within the exponential. Comparing the resulting expression with the probability density function (pdf) of a Gaussian distribution gives (2.55) [122].

Now, recalling the (E-Step), we need to derive an explicit expression for $g(\boldsymbol{\gamma}|\boldsymbol{\gamma}^{(t)})$. A series of long calculations shows that the term within the expectation can be written as

$$\log p(\tilde{\mathbf{x}}|\boldsymbol{\gamma}) = -\frac{1}{2}\left(N\sum_{i=0}^{d-1}\log[\boldsymbol{\gamma}]_i + \sum_{i=0}^{d-1}\frac{1}{[\boldsymbol{\gamma}]_i}\left(\sum_{n=0}^{N-1}|[x[n]]_i|^2\right)\right) + \mathsf{const.} \qquad (2.57)$$

Consequently, we have

$$g(\boldsymbol{\gamma}|\boldsymbol{\gamma}^{(t)}) = -\frac{1}{2}\left(N\sum_{i=0}^{d-1}\log[\boldsymbol{\gamma}]_i + \sum_{i=0}^{d-1}\frac{[\mathbf{p}_t]_i}{[\boldsymbol{\gamma}]_i}\right) + \mathsf{const}, \qquad (2.58)$$

where we have defined the vector $\mathbf{p}_t \in \mathbb{R}^d$ such that

$$[\mathbf{p}_t]_j = \frac{1}{N}\sum_{n=0}^{N-1}\left([\mathbf{C}_t]_{nd+j,nd+j} + |[\boldsymbol{\mu}_t]_{nd+j}|^2\right), \qquad (2.59)$$

and where the dependence on $\boldsymbol{\gamma}^{(t)}$ is implicit from the calculation of posterior covariance and mean from (2.56).

The M-Step of the EM algorithm updates the current estimate of the parameters vector $\boldsymbol{\gamma}^{(t+1)}$ and can be written as

$$\boldsymbol{\gamma}^{(t+1)} = \arg\max_{\boldsymbol{\gamma} \geqslant \mathbf{0}} g(\boldsymbol{\gamma}|\boldsymbol{\gamma}^{(t)}). \qquad \text{(M-Step)}$$

Note that the cost of the minimization problem above is decoupled in terms of the elements of $\boldsymbol{\gamma}$, which enables calculation of the minimizer in closed form as

$$[\boldsymbol{\gamma}^{(t+1)}]_i = [\mathbf{p}_t]_i \quad i = 0,\dots,d-1, \qquad (2.60)$$

This concludes the EM algorithm for estimating the vector of parameters $\boldsymbol{\gamma}$. We stop the algorithm once the normalized distance between two consecutive points in the solutions sequence is less than a threshold, i.e. at an index $t+1$ for which the following condition is met

$$\frac{\|\boldsymbol{\gamma}^{(t+1)} - \boldsymbol{\gamma}^{(t)}\|}{\|\boldsymbol{\gamma}^{(t+1)}\|} \leqslant \epsilon_{\mathrm{stop}}. \qquad (2.61)$$

In the empirical results at the end of this chapter, we pick a typical value of $\epsilon_{\mathrm{stop}} = 10^{-4}$ for the threshold. The point reached at the final iteration is an estimate of the ASF coefficients vector and is denoted by $\boldsymbol{\gamma}^{\star}_{\mathsf{ML}}$. Once $\boldsymbol{\gamma}^{\star}_{\mathsf{ML}}$ is computed, we have an estimate of the channel covariance as:

$$\boldsymbol{\Sigma}^{\star}_{\mathsf{ML}} = \mathcal{T}\left(\hat{\mathbf{A}}\,\boldsymbol{\gamma}^{\star}_{\mathsf{ML}}\right). \qquad (2.62)$$

Algorithm 1 Covariance Estimation via EM

1: **Input: Y**, $\widehat{\mathbf{A}}$, $\boldsymbol{\gamma}^{(0)}$, N_0, G_c, and $\widehat{K}$
2: $\widetilde{\mathbf{A}} \leftarrow \mathbf{I} \otimes \widehat{\mathbf{A}}$
3: $\widetilde{\mathbf{y}} \leftarrow \mathrm{vec}(\mathbf{Y})$
4: $t \leftarrow 0$
5: **while** not converged **do**
6:      Compute $\boldsymbol{\Sigma}_{\mathbf{x}}(\boldsymbol{\gamma}^{(t)})$ from (2.52).
7:      Compute $\mathbf{C}_t$ and $\boldsymbol{\mu}_t$ from (2.56) by replacing $\mathbf{A}$ with $\widehat{\mathbf{A}}$.
8:      Compute $\mathbf{p}_t$ from (2.59).
9:      Obtain $\boldsymbol{\gamma}^{(t+1)}$ using the update rules (2.60).
10:      $t \leftarrow t + 1$
11: **end while**
12: $\boldsymbol{\gamma}^{\star}_{\mathsf{EM}} = \boldsymbol{\gamma}^{(t)}$
13: $\boldsymbol{\Sigma}^{\star}_{\mathsf{EM}} = \mathcal{T}\left(\widehat{\mathbf{A}}\,\boldsymbol{\gamma}^{\star}_{\mathsf{EM}}\right)$.
14: **Output:** $\boldsymbol{\Sigma}^{\star}_{\mathsf{EM}}$

The whole process is summarized in Algorithm 1.

Remark 2.5. In computing the posterior mean and covariance, we need to use the dictionary matrix $\mathbf{A}$. The last K columns of this matrix are not known a priori, and are computed after performing the MUSIC step. Therefore, we replace K with $\widehat{K}$ and $\mathbf{A}$ with its estimate $\widehat{\mathbf{A}}$.

2.5 Extension to Dual-Polarized Arrays

In order to enjoy a larger benefit from massive arrays, many network developers consider using dual-polarized (DP) antenna elements in the array, since it offers a doubling of the number of inputs with a less-than-proportional increase in array size [36, 136]. The effect of adopting DP antennas at the array on performance metrics such as the multiplexing gain depends on the degree of co-polarization (co-pol) and cross-polarization (X-pol) between the two polarization states (namely, *horizontal* (H) and *vertical* (V) polarizations). While specular reflection components lead to a low degree of X-pol (hence an approximate decoupling of the polarizations), diffuse scattering results in relatively high X-pol [36]. In order to study these effects, we assign a pair of (correlated) channel coefficients to each element of the array and introduce a statistical model to represent the co-pol and X-pol properties of a particular environment. This highlights a difference between DP arrays and single-polarized arrays, as studied earlier in this chapter, in that the the channel dimension is doubled by assigning a pair of coefficients to each element.

2.5.1 Dual-Polarized Channel Model

Similar to the model introduced at the beginning of this chapter, we consider a ULA of M dual-polarized antenna elements that communicates with a single-antenna, single-polarized user. The channel between antenna m of the array and the user antenna consists of two elements, corresponding to H and V polarization coefficients, respectively denoted as $h_{m,\mathsf{H}}$, $h_{m,\mathsf{V}} \in \mathbb{C}$. The channel gain for either polarization is a superposition of random gains along a continuum of AoAs, weighted by the antenna element response which, assuming as before an antenna spacing of $d = \frac{\lambda}{2\sin(\theta_{\max})}$, for antenna $m \in \mathbb{Z}_M$ is given by $a_m = e^{jm\pi\xi}$, $m \in \mathbb{Z}_M$ for the normalized AoA $\xi \in [-1, 1]$. Then, one can express H and V channel coefficients as

$$h_{m,\mathsf{H}} = \int_{-1}^{1} dW_{\mathsf{H}}(\xi)e^{j\pi m\xi} \tag{2.63a}$$

$$h_{m,\mathsf{V}} = \int_{-1}^{1} dW_{\mathsf{V}}(\xi)e^{j\pi m\xi} \tag{2.63b}$$

for $m \in \mathbb{Z}_M$, where W_{H} and W_{V} are random processes representing the random gains along each AoA for H and V polarizations, respectively. We assume both of these to be zero-mean, complex Gaussian processes $\mathbb{E}\left[dW_{\mathsf{V}}(\xi)\right] = \mathbb{E}\left[dW_{\mathsf{H}}(\xi)\right] = 0$ with the following autocorrelations:[3]

$$\mathbb{E}\left[dW_{\mathsf{H}}(\xi)dW_{\mathsf{H}}^*(\xi')\right] = \gamma_{\mathsf{H}}(\xi)\delta(\xi - \xi')d\xi, \tag{2.64a}$$

$$\mathbb{E}\left[dW_{\mathsf{V}}(\xi)dW_{\mathsf{V}}^*(\xi')\right] = \gamma_{\mathsf{V}}(\xi)\delta(\xi - \xi')d\xi. \tag{2.64b}$$

The measures γ_{H} and γ_{V} are both real and non-negative, representing the channel power density received along each AoA for H and V polarizations, respectively. We call these the horizontal and vertical ASF's (see Fig. 2.5).

In practice, the H and V links *cannot* be entirely isolated from each other and therefore, there exists a leakage of channel power between the two. This implies that, for each AoA, the random gains $dW_{\mathsf{H}}(\xi)$ and $dW_{\mathsf{V}}(\xi)$ are correlated such that we have

$$\mathbb{E}\left[dW_{\mathsf{H}}(\xi)dW_{\mathsf{V}}^*(\xi')\right] = \rho(\xi)\delta(\xi - \xi')d\xi \tag{2.65}$$

where ρ is a generally complex-valued function. As such, the dual-polarized channel can be more conveniently expressed as follows. Denote M-dimensional horizontal and vertical channel vectors $\mathbf{h}_{\mathsf{H}} = [h_{0,\mathsf{H}}, \ldots, h_{M-1,\mathsf{H}}]^\mathsf{T}$, $\mathbf{h}_{\mathsf{V}} = [h_{0,\mathsf{V}}, \ldots, h_{M-1,\mathsf{V}}]^\mathsf{T}$ and note that using

---

[3]For simplicity, we have dropped the technicalities in defining the Gaussian processes W_{H} and W_{V} and their corresponding autocorrelation functions. For a more involved treatment, we refer to Section 2.3 and we note that the same notions as presented in the single-polarized case extend to the present dual-polarized case.
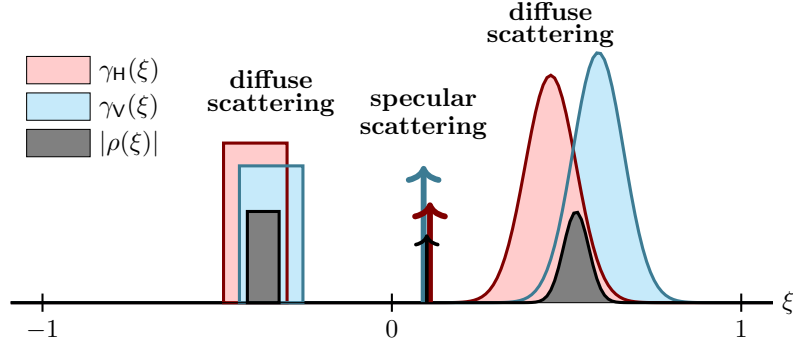
Figure 2.5: An example of H and V ASF's as well as the H-V cross-correlation modulus. The blue shaded area highlights γ_V, the red one highlights γ_H and the black one highlights $|\rho|$.

(2.63) we have

$$\mathbf{h}_H = \int_{-1}^{1} dW_H(\xi)\mathbf{a}(\xi), \tag{2.66a}$$

$$\mathbf{h}_V = \int_{-1}^{1} dW_V(\xi)\mathbf{a}(\xi). \tag{2.66b}$$

where $\mathbf{a}(\xi) = [1, e^{j\pi\xi}, \dots, e^{j\pi(M-1)\xi}]^\mathsf{T} \in \mathbb{C}^M$ is the array response vector. Note that, while we have assumed the array response to be the same for H and V polarizations, the responses need not be the same. The assumption is made for convenience, and a clearer exposition of the idea; if the array responses are different, we can duly adapt the proposed method with a few straightforward modifications. Finally, we define the dual-polarized channel as the $2M$-dimensional vector $\mathbf{h} = [\mathbf{h}_H^\mathsf{T}, \mathbf{h}_V^\mathsf{T}]^\mathsf{T}$ which can be written as

$$\mathbf{h} = \int_{-1}^{1} \begin{bmatrix} \mathbf{a}(\xi) & \mathbf{0} \\ \mathbf{0} & \mathbf{a}(\xi) \end{bmatrix} \begin{bmatrix} dW_H(\xi) \\ dW_V(\xi) \end{bmatrix} = \int_{-1}^{1} (\mathbf{I}_2 \otimes \mathbf{a}(\xi))\, d\mathbf{w}(\xi), \tag{2.67}$$

where $\otimes$ denotes Kronecker product, and $d\mathbf{w}(\xi) := [dW_H(\xi), dW_V(\xi)]^\mathsf{T}$ is the *vector* of random H and V gains, with $\mathbb{E}[d\mathbf{w}(\xi)] = \mathbf{0}$ and a covariance

$$\mathbb{E}\left[d\mathbf{w}(\xi)d\mathbf{w}(\xi')^\mathsf{H}\right] = \begin{bmatrix} \gamma_H(\xi) & \rho(\xi) \\ \rho(\xi)^* & \gamma_V(\xi) \end{bmatrix} \delta(\xi - \xi')d\xi := \mathbf{\Gamma}(\xi)\delta(\xi - \xi')d\xi \in \mathbb{C}^{2\times 2}. \tag{2.68}$$

We call the 2×2 matrix-valued measure $\mathbf{\Gamma}$ the *dual-polarized angular scattering function* (DASF), and we note that it is PSD for all $\xi \in [-1, 1]$.[4] The channel covariance can be

---

[4]Note that $\mathbf{\Gamma}$ (with bold-face font) is the 2×2 DASF matrix and should not be confused with the autocorrelation Γ introduced in (2.8b) of Section 2.3.

computed according to (2.67) and using (2.68) as

$$
\begin{aligned}
\boldsymbol{\Sigma} &= \mathbb{E}\left[\mathbf{h}\mathbf{h}^{\mathsf{H}}\right] \\
&= \int_{-1}^{1}\int_{-1}^{1} \left(\mathbf{I}_2 \otimes \mathbf{a}(\xi)\right) \mathbb{E}\left[d\mathbf{w}(\xi)d\mathbf{w}(\xi')^{\mathsf{H}}\right] \left(\mathbf{I}_2 \otimes \mathbf{a}(\xi')\right)^{\mathsf{H}} \\
&= \int_{-1}^{1} \boldsymbol{\Gamma}(\xi) \otimes \left(\mathbf{a}(\xi)\mathbf{a}(\xi)^{\mathsf{H}}d\xi\right),
\end{aligned}
\tag{2.69}
$$

Similar to the role played by the ASF in a single-polarized array, the DASF captures the angular spectral properties of the channel, i.e. the power density along H and V links and the power leakage density between the two. Note that since $\boldsymbol{\Gamma}$ is PSD, we have $|\rho(\xi)|^2 \leqslant \gamma_{\mathsf{H}}(\xi)\gamma_{\mathsf{V}}(\xi)$, for all $\xi \in [-1,1]$, putting a bound on the modulus of ρ. In particular, if for some ξ we have $\gamma_{\mathsf{H}}(\xi) = 0$ or $\gamma_{\mathsf{V}}(\xi) = 0$, then necessarily $\rho(\xi) = 0$, which shows that the support of ρ is limited to the intersection of the supports of γ_{V} and γ_{H}.

Similar to the scenario described at the beginning of this chapter, we consider a case in which the BS observes N noisy pilot measurements of the dual-polarized channel as $\mathbf{y}[i] = \mathbf{h}[i] + \mathbf{z}[i]$, $i \in \mathbb{Z}_N$, where $\{\mathbf{h}[i]\}_i$ are independent channel realizations and $\mathbf{z}[i] \sim \mathcal{CN}(\mathbf{0}, N_0\mathbf{I}_{2M})$, $i \in \mathbb{Z}_N$ are independent AWGN vectors. The goal is to estimate the channel covariance $\boldsymbol{\Sigma}$ given the pilot observations $\{\mathbf{y}[i]\}_i$.

2.5.2 Decomposition of the DASF

The DASF of a channel models the received power density over each AoA. This power density in turn depends on the scattering properties of the environment: partly it comes from line of sight (LoS) propagation and specular reflection in the environment, that occupy narrow angular intervals, while the rest of the power comes from diffuse scattering, occupying wide angular intervals [36] (see Fig. 2.5). Similar to the decomposition of the ASF, we decompose the DASF into discrete and continuous components:

$$
\boldsymbol{\Gamma} = \boldsymbol{\Gamma}_d + \boldsymbol{\Gamma}_c,
\tag{2.70}
$$

where $\boldsymbol{\Gamma}_c$ is the continuous component and $\boldsymbol{\Gamma}_d$ is the discrete component. For the discrete part, the parametric form is simply given by a train of weighted delta functions:

$$
\boldsymbol{\Gamma}_d(\xi) = \sum_{k \in \mathbb{Z}_K} \mathbf{C}_k \delta(\xi - \xi_k) d\xi, \ \xi \in [-1,1],
\tag{2.71}
$$

where $\{\mathbf{C}_k\}_k$ are 2×2 PSD matrices and $\{\xi_k\}_k$ are discrete AoAs. In contrast, we cannot assume a parametric description of $\boldsymbol{\Gamma}_c$ in terms of delta functions. Instead, we introduce an

approximation of $\mathbf{\Gamma}_c$ using a family of limited-support densities as discussed in Section 2.4. We define this density family as the one introduced in (2.21), namely $\Psi = \{\psi_i : i \in \mathbb{Z}_{G_c}\}$. Using this family we can approximate $\mathbf{\Gamma}_c$ as

$$\mathbf{\Gamma}_c \approx \sum_{i \in \mathbb{Z}_{G_c}} \mathbf{C}'_i \psi_i \tag{2.72}$$

where similar to (2.71) $\{\mathbf{C}'_i\}_i$ are 2×2 PSD matrices. If G_c is large enough ($G_c \gg 1$), then one can find the coefficients $\mathbf{C}'_i$ such that the approximation error in (2.72) is negligible. Using (2.69), (2.71) and (2.72), we can write the dual-polarized channel covariance in parametric form as

$$
\begin{aligned}
\mathbf{\Sigma} = \mathbf{\Sigma}_d + \mathbf{\Sigma}_c &:= \sum_{k \in \mathbb{Z}_K} \mathbf{C}_k \otimes \left(\mathbf{a}(\xi_k)\mathbf{a}(\xi_k)^{\mathsf{H}} \right) + \int_{-1}^{1} \Gamma_c(\xi) \otimes \left(\mathbf{a}(\xi_k)\mathbf{a}(\xi_k)^{\mathsf{H}} \right) d\xi \\
&\approx \sum_{k \in \mathbb{Z}_K} \mathbf{C}_k \otimes \left(\mathbf{a}(\xi_k)\mathbf{a}(\xi_k)^{\mathsf{H}} \right) + \sum_{i \in \mathbb{Z}_{G_c}} \mathbf{C}'_i \otimes \mathbf{A}'_i,
\end{aligned}
\tag{2.73}
$$

where we have defined $\mathbf{A}'_i = \int_{-1}^{1} \psi_i(\xi)\mathbf{a}(\xi)\mathbf{a}(\xi)^{\mathsf{H}} d\xi$ for $i \in \mathbb{Z}_{G_c}$. If the discrete AoAs $\{\xi_k\}_{k \in \mathbb{Z}_K}$ were known, we could claim via Eq. (2.73) that estimating $\mathbf{\Sigma}$ is equivalent to estimating the coefficient matrices $\{\mathbf{C}_k\}_{k \in \mathbb{Z}_K}$ and $\{\mathbf{C}'_i\}_{i \in \mathbb{Z}_{G_c}}$. Similar to the case of single-polarized antennas, here we propose to first estimate the discrete AoAs $\{\xi_i\}_{i=1}^{r}$ from pilot measurements.

2.5.3 Estimating Discrete AoAs

We use MUSIC to estimate the discrete AoAs present in the DASF. Suppose we have an estimate of the number of discrete AoAs as $\widehat{K}$, using the heuristic model-order estimation method in Section 2.4.1. This implies that the discrete covariance component $\mathbf{\Sigma}_d = \sum_{k \in \mathbb{Z}_K} \mathbf{C}_k \otimes \left(\mathbf{a}(\xi_k)\mathbf{a}(\xi_k)^{\mathsf{H}} \right)$ is of maximum rank $2\widehat{K}$. Define the eigendecomposition of the sample covariance $\widehat{\mathbf{\Sigma}}$ as $\widehat{\mathbf{\Sigma}} = \widehat{\mathbf{U}}\widehat{\mathbf{\Lambda}}\widehat{\mathbf{U}}^{\mathsf{H}}$, where $\widehat{\mathbf{U}} \in \mathbb{C}^{2M \times 2M} = [\widehat{\mathbf{u}}_1, \ldots, \widehat{\mathbf{u}}_{2M}]$ is the matrix of eigenvectors and $\widehat{\mathbf{\Lambda}} \in \mathbb{R}_+^{2M \times 2M}$ is the diagonal eigenvalues matrix, with its diagonal elements in descending order. We call the space spanned by the $2M - 2\widehat{K}$ eigenvectors of $\widehat{\mathbf{\Sigma}}$ corresponding to its $2M - 2\widehat{K}$ smallest eigenvalues as the noise subspace. The vectors spanning the noise subspace are collected in the matrix $\mathbf{U}_{\mathrm{noi}} = [\widehat{\mathbf{u}}_{2\widehat{K}+1}, \ldots, \widehat{\mathbf{u}}_{2M}]$. Then we form the pseudo-spectrum function $\eta(\xi) = \left\| \mathbf{U}_{\mathrm{noi}}^{\mathsf{H}} \left(\mathbf{I}_2 \otimes \mathbf{a}(\xi) \right) \right\|_{\mathsf{F}}^2$ and estimate the discrete AoAs as the $\widehat{K}$ smallest minimizers of η. Intuitively, in this way we find a number of $\widehat{K}$ AoAs that the $2M \times 2$ dual-polarized array response $\mathbf{I}_2 \otimes \mathbf{a}(\xi)$ along them, has the smallest norm when projected to the noise subspace. This heuristic follows the same rationale as the one explained in Section 2.4.1. We denote the estimated discrete AoAs by $\{\widehat{\xi}_k\}_{k \in \mathbb{Z}_{\widehat{K}}}$.

Recalling (2.73), now we can say that estimating $\boldsymbol{\Sigma}$ is equivalent to estimating the $\widehat{d}$ coefficient matrices, namely $\{\mathbf{C}_k\}_{k \in \mathbb{Z}_{\widehat{K}}}$ and $\{\mathbf{C}'_i\}_{i \in \mathbb{Z}_{G_c}}$.

2.5.4 Estimating the Coefficients

Let us first reformulate the channel covariance parametric description in a simpler form. Define the known $M \times M$ matrices $\mathbf{S}_i := \mathbf{A}'_i$ for $i = 0, \ldots, G_c - 1$ and $\mathbf{S}_i := \mathbf{a}(\widehat{\xi}_{i-G_c})\mathbf{a}(\widehat{\xi}_{i-G_c})^{\mathsf{H}}$ for $i = G_c, \ldots, \widehat{d} - 1$. Also define their associated unknown coefficients as $\mathbf{W}_i := \mathbf{C}'_i$ for $i = 0, \ldots, G_c - 1$ and $\mathbf{W}_i := \mathbf{C}_{i-G_c}$ for $i = G_c, \ldots, \widehat{d} - 1$. Then (2.73) can be reformulated as

$$\boldsymbol{\Sigma}(\{\mathbf{W}_i\}_{i=1}^{\widehat{d}}) \approx \sum_{i \in \mathbb{Z}_{G_c}} \mathbf{W}_i \otimes \mathbf{S}_i. \tag{2.74}$$

Now, the problem is to estimate the coefficient matrices $\{\mathbf{W}_i \in \mathbb{S}_+^2\}_{i \in \mathbb{Z}_{\widehat{d}}}$, given the noisy pilot measurements $\{\mathbf{y}[i]\}_{i \in \mathbb{Z}_N}$. Our proposition for performing this task is based on minimizing the difference between a *coarse* estimate of the channel covariance and its parametric form as a function of the matrix coefficients. It is easy to see that the dual-polarized channel covariance is a blocky matrix, consisting of four Hermitian Toeplitz blocks:

$$\boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_{\mathsf{H}} & \boldsymbol{\Sigma}_{\mathsf{HV}} \\ \boldsymbol{\Sigma}_{\mathsf{VH}} & \boldsymbol{\Sigma}_{\mathsf{V}} \end{bmatrix}, \tag{2.75}$$

where $\boldsymbol{\Sigma}_{\mathsf{H}} = \mathbb{E}\left[\mathbf{h}_{\mathsf{H}}\mathbf{h}_{\mathsf{H}}^{\mathsf{H}}\right]$, $\boldsymbol{\Sigma}_{\mathsf{V}} = \mathbb{E}\left[\mathbf{h}_{\mathsf{V}}\mathbf{h}_{\mathsf{V}}^{\mathsf{H}}\right]$, and $\boldsymbol{\Sigma}_{\mathsf{HV}} = \mathbb{E}\left[\mathbf{h}_{\mathsf{H}}\mathbf{h}_{\mathsf{V}}^{\mathsf{H}}\right] = \boldsymbol{\Sigma}_{\mathsf{HV}}^{\mathsf{H}}$. This blocky structure specifies a linear space, and we define the coarse covariance estimate as the orthogonal projection of the sample covariance $\widehat{\boldsymbol{\Sigma}}$ onto this space. The sample covariance itself has a blocky structure as

$$\widehat{\boldsymbol{\Sigma}} = \begin{bmatrix} \widehat{\boldsymbol{\Sigma}}_{\mathsf{H}} & \widehat{\boldsymbol{\Sigma}}_{\mathsf{HV}} \\ \widehat{\boldsymbol{\Sigma}}_{\mathsf{VH}} & \widehat{\boldsymbol{\Sigma}}_{\mathsf{V},} \end{bmatrix}$$

where the blocks are not necessarily Toeplitz. One can simply show that the mentioned projection is given by projecting each block onto the space of $M \times M$ Hermitian Toeplitz matrices $\mathbb{T}^M$. Therefore, the coarse estimate is given as

$$\widetilde{\boldsymbol{\Sigma}} = \begin{bmatrix} \mathcal{P}_{\mathbb{T}^M}(\widehat{\boldsymbol{\Sigma}}_{\mathsf{H}}) & \mathcal{P}_{\mathbb{T}^M}(\widehat{\boldsymbol{\Sigma}}_{\mathsf{HV}}) \\ \mathcal{P}_{\mathbb{T}^M}(\widehat{\boldsymbol{\Sigma}}_{\mathsf{VH}}) & \mathcal{P}_{\mathbb{T}^M}(\widehat{\boldsymbol{\Sigma}}_{\mathsf{V}}) \end{bmatrix} \tag{2.76}$$

After computing $\widetilde{\boldsymbol{\Sigma}}$, we estimate the matrix coefficients $\{\mathbf{W}_i\}_{i \in \mathbb{Z}_{\widehat{d}}}$ by solving the following optimization problem:

$$
\begin{aligned}
\{\mathbf{W}_i^\star\}_i &= \underset{\{\mathbf{W}_i\}_i}{\arg\min} \quad \|\widetilde{\boldsymbol{\Sigma}} - \sum_{i \in \mathbb{Z}_{\widehat{d}}} \mathbf{W}_i \otimes \mathbf{S}_i\|_{\mathsf{F}}^2 \\
&\text{subject to} \quad \mathbf{W}_i \succeq \mathbf{0}, \ i \in \mathbb{Z}_{\widehat{d}}.
\end{aligned}
\tag{2.77}
$$

We call this problem a *positive semidefinite least-squares* (PSD-LS) program. The PSD-LS is convex and can be solved using standard convex optimization algorithms. Finally, we obtain the covariance estimate as

$$
\boldsymbol{\Sigma}^\star = \sum_{i \in \mathbb{Z}_{\widehat{d}}} \mathbf{W}_i^\star \otimes \mathbf{S}_i.
\tag{2.78}
$$

2.6 Simulation Results

We construct an ASF using expressions (2.14), (2.16), and (2.17). The discrete component of the ASF is generated by specifying a set of discrete impulse locations $\{\xi_k : k = 0, \ldots, K-1\}$ and their corresponding positive coefficients $\{c_k : k = 0, \ldots, K-1\}$, where K is specified later for each experiment, the spike locations are randomly generated as $\xi_k \sim \text{Unif}([-1, 1]), \ k \in \mathbb{Z}_K$, and their coefficients as $c_k \sim 1 + |\mathcal{CN}(0, 1)|, \ k \in \mathbb{Z}_K$.

The continuous component consists of $\widetilde{K}$ rectangular functions

$$
g_k = \frac{1}{|\mathcal{X}_k|} \mathsf{rect}_{\mathcal{X}_k},
$$

where $\mathcal{X}_k = \text{supp}(g_k)$, $|\mathcal{X}_k| = w \, \forall \, k$ denotes the angular support of the diffuse scatterer k, where the width of the scatterer is set to $w = 0.3$. The support starting point $\inf_\xi \mathcal{X}_k$ is located, uniformly at random, over the line segment $[\frac{2k}{\widetilde{K}} - 1, \frac{2(k+1)}{\widetilde{K}} - w - 1]$. The coefficient of each function is given as $\{\widetilde{c}_k : k = 0, \ldots, \widetilde{K} - 1\}$, where $\widetilde{c}_k \sim 1 + |\mathcal{CN}(0, 1)|, \ k \in \mathbb{Z}_{\widetilde{K}}$. The overall ASF is expressed as

$$
\gamma(\xi) = \frac{\alpha}{Z_d} \gamma_d(\xi) + \frac{(1-\alpha)}{Z_c} \gamma_c(\xi) = \frac{\alpha}{Z_d} \sum_{k=0}^{K-1} c_k \delta(\xi - \xi_k) + \frac{(1-\alpha)}{Z_c} \sum_{k=0}^{\widetilde{K}-1} \widetilde{c}_k g_k(\xi), \ \xi \in [-1, 1],
\tag{2.79}
$$

where $Z_d = \int_{-1}^{1} \gamma_d(\xi) d\xi$, and $Z_c = \int_{-1}^{1} \gamma_c(\xi) d\xi$ (normalizing each component of the ASF), and α controls the contribution of each part to the overall ASF γ, i.e. if $\alpha = 0$, we have a purely continuous ASF, while if $\alpha = 1$, we have a purely discrete one.

Once the ASF is specified, we compute the channel covariance $\boldsymbol{\Sigma}$ using the synthesis formula in (2.11). Denote the eigendecomposition of $\boldsymbol{\Sigma}$ as $\boldsymbol{\Sigma} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^{\mathsf{H}}$. Then, a random

realization of the channel can be drawn as

$$\mathbf{h} = \mathbf{U}\boldsymbol{\Lambda}^{\frac{1}{2}}\mathbf{g}, \tag{2.80}$$

where $\mathbf{g} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_M)$ is a white, complex Gaussian random vector. A random sample set of the channel $\mathbf{h}$ is obtained by generating N such random realizations. We add white Gaussian noise $\mathbf{z} \sim \mathcal{CN}(\mathbf{0}, N_0\mathbf{I})$ to each sample to generate the noisy pilot samples in (2.1), noting that the noise variance per component is given as $N_0 = 1/\text{SNR}$. The noise variance is supposed to be known to all methods.

Error Metrics

There are different ways to compare covariance matrices in terms of error metrics. Among these, we choose the following two metrics.

1. **Normalized Euclidean Distance.** Let $\boldsymbol{\Sigma} \in \mathbb{S}_+^M$ be the true channel covariance and let $\widehat{\boldsymbol{\Sigma}}$ denote its estimate. A standard error metric is given by the normalized Euclidean norm of the difference, given as

$$E_{\text{euc}} = \frac{\|\boldsymbol{\Sigma} - \widehat{\boldsymbol{\Sigma}}\|_{\text{F}}}{\|\boldsymbol{\Sigma}\|_{\text{F}}}. \tag{2.81}$$

2. **Relative Efficiency.** In wireless communication it is sometimes crucial to know, how similar their associated "dominant" subspaces. Let $p \in \{1, \ldots, M\}$ denote a subspace dimension parameter, $\mathbf{U}_p = [\mathbf{U}_{\cdot,1}, \ldots, \mathbf{U}_{\cdot,p}] \in \mathbb{C}^{M \times p}$ the p eigenvectors of $\boldsymbol{\Sigma}$ corresponding to its largest p eigenvalues. Similarly, define $\widehat{\mathbf{U}}_p$ as the matrix of the p eigenvectors of $\widehat{\boldsymbol{\Sigma}}$ corresponding to its p largest eigenvalues. We define the relative efficiency parameter as

$$\varepsilon(p) = 1 - \frac{\langle \boldsymbol{\Sigma}, \widehat{\mathbf{U}}_p \widehat{\mathbf{U}}_p^{\mathsf{H}} \rangle}{\langle \boldsymbol{\Sigma}, \mathbf{U}_p \mathbf{U}_p^{\mathsf{H}} \rangle}. \tag{2.82}$$

It is easy to see that $\varepsilon(p) \in [0, 1]$, for all $p = 1, \ldots, M$. Disregarding the estimated eigenvalues, this metric compares the power captured by the estimated p-dominant subspace to that captured by the true p-dominant subspace. The close $\varepsilon(p)$ is to 0, the better the p-dominant subspace is estimated.

Averages of these metrics $\mathbb{E}[E_{\text{euc}}]$ and $\mathbb{E}[\varepsilon(p)]$ are empirically calculated, in each case for 100 Monte Carlo simulations. We compare the performance of the following estimators with regards to these metrics. We compare the results for the sample covariance estimator given as $\widehat{\boldsymbol{\Sigma}} = \frac{1}{N}\mathbf{Y}\mathbf{Y}^{\mathsf{H}} - N_0\mathbf{I}$, the NNLS and the ML estimators. In the two latter cases, we

Figure 2.6: covariance estimation error curves for the normalized Euclidean distance and relative efficiency metrics against the sampling ratio N/M for three different channel dimensions M. Here $\alpha = 0.5$ and SNR $= 20$ dB.

use a rectangular density family introduced in (2.24) to approximate the continuous ASF component, where the number of density functions is set to $G_c = 3M$. In each case, the MUSIC step is performed to estimate the discrete spike locations, consequently adding the corresponding array response vectors at those locations to obtain the overall dictionary $\widehat{\mathbf{A}} \in \mathbb{C}^{M \times \widehat{d}}$.

Fig. 2.6 illustrates the comparison results for the two error metrics against the *sampling ratio* N/M. Here we have considered channels whose ASF is composed of $K = 2$ spikes and $\widetilde{K} = 2$ diffuse (rectangular) components. The discrete-continuous power-splitting ratio is set to $\alpha = 0.5$ and the SNR is 20 dB. The results are considered for three different channel dimensions $M = 50$, $M = 100$, and $M = 200$. We observe that, in all scenarios, the NNLS and ML estimators perform much better than the sample covariance estimator, with the the normalized Euclidean error decreasing with an increase of the sampling ratio (Figs. (a) to (c)). This is also the case for relative efficiency metrics in Figs. (d) to (f), where we plot the average of $\varepsilon(p)$ for $p = 1, \ldots, 20$. Note that, we do *not* expect the metric to necessarily decrease with p increasing. The important point is rather, that the NNLS and ML estimators perform much better than the sample covariance for all p. Also, the sharp minimum observed at $p = 2$, in all figures for the ML and NNLS relative efficiency

(a) $M = 50$ (b) $M = 100$ (c) $M = 200$

(d) $M = 50$, $s = 4$ (e) $M = 100$, $s = 4$ (f) $M = 200$, $s = 4$
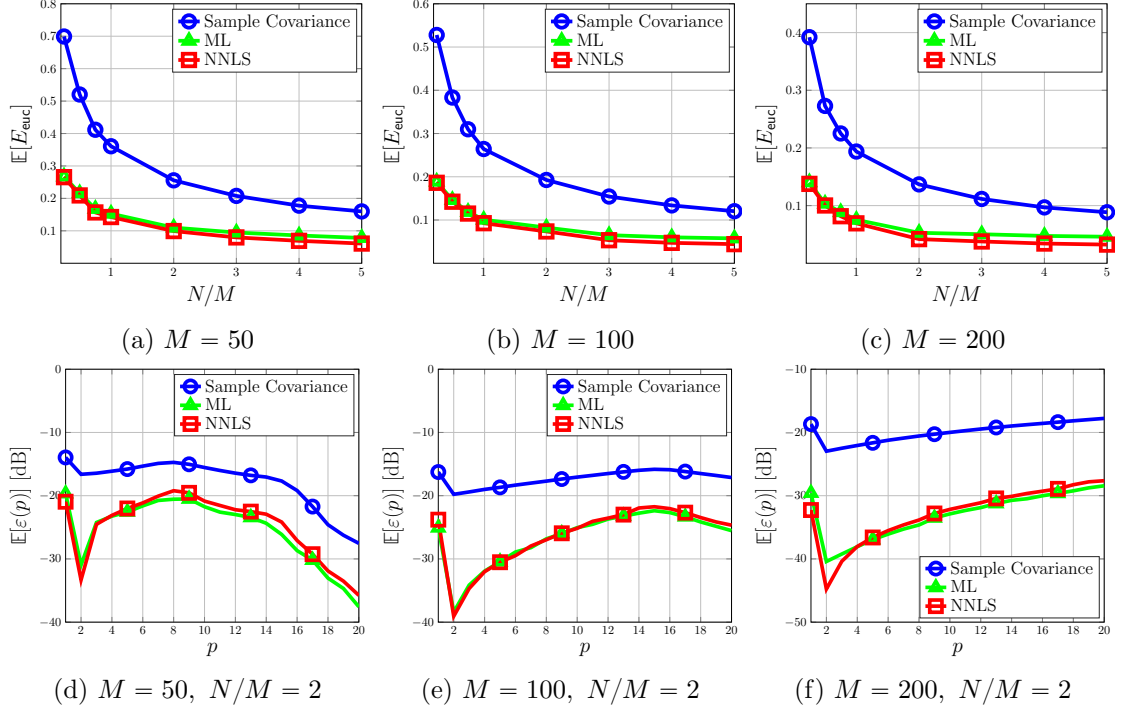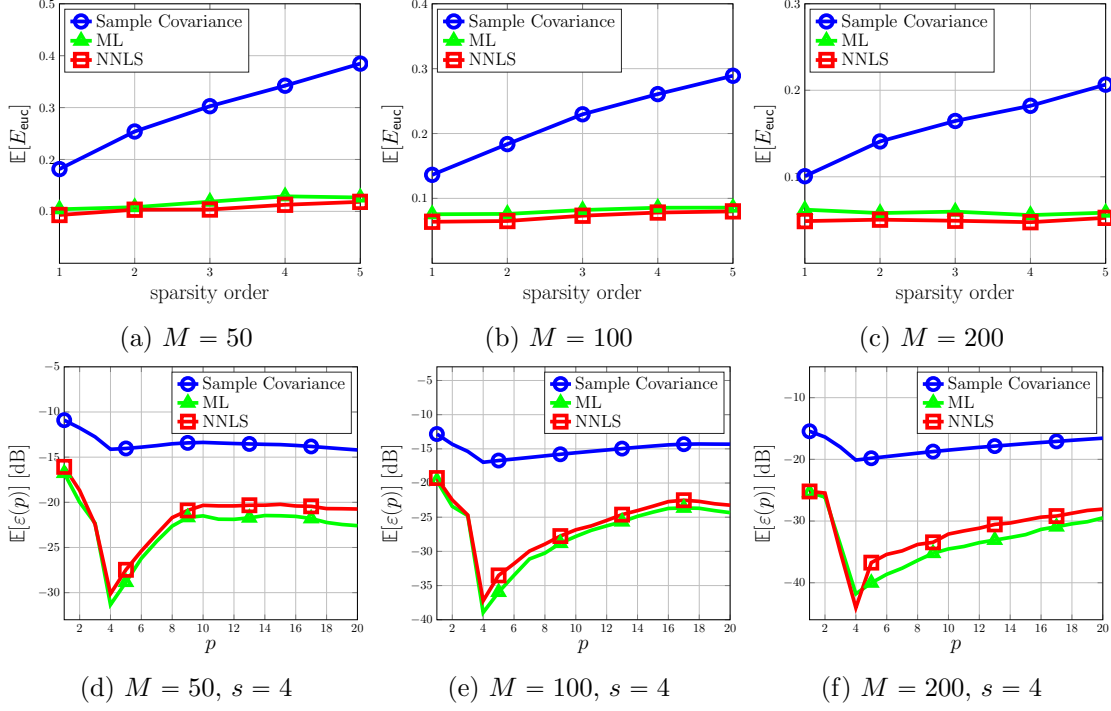
Figure 2.7: covariance estimation error curves for the normalized Euclidean distance and relative efficiency metrics against the sparsity order for three different channel dimensions M. Here $\alpha = 0.5$, $N/M = 2$, and SNR = 20 dB.

curves is justified as follows. The ASFs generated in this simulation contain two spikes and we have considered an equal splitting of channel power over discrete and continuous components. Therefore, for large M the span of the $p = 2$ dominant eigenvectors of $\boldsymbol{\Sigma}$ highly overlap with the span of the two array response vectors at the spike locations. Since the ML and NNLS methods initially obtain the spike locations via MUSIC, they estimate the dominant 2-dimensional channel subspace with high accuracy.

In addition to this, the results for NNLS and ML are quite similar, with NNLS showing a slightly better performance in terms of the normalized Euclidean error. This is understandable, since as seen from the formulation of $\mathcal{P}_{\mathsf{NNLS}}$, NNLS directly minimizes the Frobenius norm of the error between the parametric covariance and (the Toeplitz projection of) the sample covariance, while ML maximizes the likelihood function in $\mathcal{P}_{\mathsf{ML}}$. Therefore, this advantage of NNLS compared to ML in terms of a Euclidean distance measure should not be surprising. On the other hand, the ML method performs better in some cases in terms of the relative efficiency metric. Admittedly, this performance difference between the two methods is so small that one can safely ignore it.

As explained in the beginning of this section, in our numerical experiments we consider the channel ASF to consist of K spikes and $\widetilde{K}$ diffuse components. The larger these

parameters are, the less sparse the channel is. This means that with an increase of K and $\widetilde{K}$, the number of angular directions with a non-zero power (variance) increases. It is interesting to observe the behavior of the covariance estimation methods for different orders of channel sparsity. Many channel estimation methods in the literature rely on the assumption of the channel to be sparse in the angular domain. The performance of these methods naturally degrades when the sparsity assumption is violated. Therefore, it is desirable to have estimators that perform well for a variety of channel sparsity orders. Fig. 3.11 illustrates the comparison of NNLS and ML estimators with the sample covariance estimator for different channel sparsity orders. Here, for each sparsity order $s \in \{1, 2, 3, 4, 5\}$, we set $K = \widetilde{K} = s$, and we select the location of the scatterers and their width as explained earlier. As we can see, by increasing s, the sample covariance estimator shows an increasing error. The reason is that, for a fixed dimension M and sampling ratio $N/M = 2$, as s grows large, the degrees of freedom of the channel covariance increases, making it harder for the sample covariance to estimate yield an accurate estimate. In contrast, the NNLS and ML estimators seem to be robust with respect to an increase of the channel sparsity order, as their error either does not increase or increases only slightly.

Dual-Polarized Covariance Estimation

Finally, we provide simulation results for the task of dual-polarized channel covariance estimation. We consider a ULA of size $M = 32$. In order to produce (semi-)random Horizontal and Vertical ASFs we consider the following generative model for the horizontal ASF:

$$\gamma_{\mathsf{H}}(\xi) = \frac{1}{|\mathcal{I}_1| + |\mathcal{I}_2|} \left(\mathsf{rect}_{\mathcal{I}_1}(\xi) + \mathsf{rect}_{\mathcal{I}_2}(\xi) \right) + \frac{1}{2} \left(\delta(\xi - \xi_1) + \delta(\xi - \xi_2) \right), \qquad (2.83)$$

where $\mathcal{I}_1, \mathcal{I}_2 \subset [-1, 1]$ where the the interval lengths are chosen uniformly at random between 0.1 and 0.4, i.e. $|\mathcal{I}_j| \sim \mathrm{Unif}([0.1, 0.4])$, independently for $j = 1$ and $j = 2$. Besides, $\xi_1, \xi_2 \in$ denote discrete AoAs, generated independently and uniformly at random over $[-1, 1]$. Similarly, we generate the vertical ASF as:

$$\gamma_{\mathsf{V}}(\xi) = \frac{\alpha}{|\mathcal{I}_1'| + |\mathcal{I}_2'|} \left(\mathsf{rect}_{\mathcal{I}_1'}(\xi) + \mathsf{rect}_{\mathcal{I}_2'}(\xi) \right) + \frac{1 - \alpha}{2} \left(\delta(\xi - \xi_1') + \delta(\xi - \xi_2') \right), \qquad (2.84)$$

Since it is natural for the horizontal and vertical ASFs to overlap in their support, we assume the discrete AoAs to be the same, i.e. $\xi_1' = \xi_1$ and $\xi_2' = \xi_2$, and we assume $\mathcal{I}_1'$ and $\mathcal{I}_2'$ to be slightly shifted versions of $\mathcal{I}_1$ and $\mathcal{I}_2$ as $\mathcal{I}_1' = \mathcal{I}_1 + 0.1$ and $\mathcal{I}_2' = \mathcal{I}_2 + 0.1$. Finally, we assume the cross-correlation function $\rho(\xi)$ to take on the form $\rho(\xi) = \frac{1}{2} \sqrt{\gamma_{\mathsf{H}}(\xi) \, \gamma_{\mathsf{V}}(\xi)}$. This is a simplifying assumption on the form of $\rho(\xi)$, which does not undermine the generality
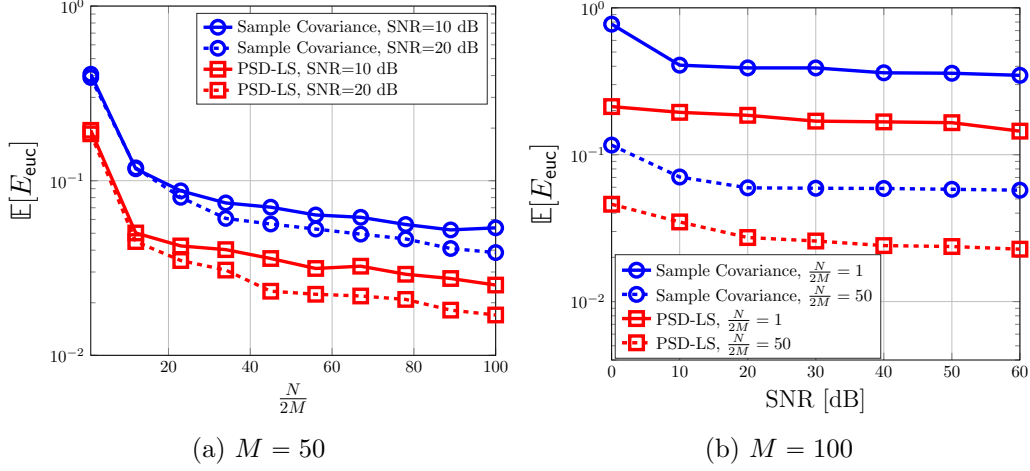
Figure 2.8: dual-polarized channel covariance estimation error in terms of normalized Euclidean distance vs the sampling ratio $\frac{N}{2M}$ (left) and vs SNR (right). Here $M = 32$.

of the DASF, and satisfies the necessary condition $|\rho(\xi)|^2 \leqslant \gamma_{\mathsf{H}}(\xi)\gamma_{\mathsf{V}}(\xi)$ for the DASF $\mathbf{\Gamma}(\xi)$ to be a PSD matrix-valued function for all $\xi \in [-1, 1]$. To perform a Monte-Carlo simulation, we generate 100 random DASFs according to the model explained earlier. For each random DASF, we generate N independent samples of the channel and AWGN vectors to generate the noisy pilot signals. Fig. 2.8b compares the normalized Euclidean error as a function of the sampling ratio (left figure) as well as the SNR (right figure). The error figures show that the method based on PSD-LS considerably improves estimation accuracy in comparison to the sample covariance estimator. The main reason is that, the PSD-LS program (2.77) captures the structure of the dual-polarized covariance: it enforces the Kronecker structure by adopting the parametric covariance form $\sum_{i=0}^{\hat{d}-1} \mathbf{W}_i \otimes \mathbf{S}_i$ and it constraints the coefficients $\mathbf{W}_i, \ i = 0, \ldots, \hat{d} - 1$ to be PSD in accordance with the DASF being a PSD matrix-valued function.

2.7 Appendices

2.7.1 Proof of Theorem 2.1

We prove Theorem 2.1 by first referring to the result in [89]. Consider a ULA with M elements. The channel ASF is expressed as $\gamma(\xi) = \sum_{k \in \mathbb{Z}_K} c_k \delta(\xi - \xi_k) + \gamma_c(\xi)$ as in (2.17) for $\xi \in [-1, 1]$. The asymptotic scaling law studied in [89] assumes a challenging case in which the spike coefficients $\{c_k\}_{k \in \mathbb{Z}_K}$ scale with M according to $c_k^{(M)} = \frac{\kappa_k}{M}$ (hence, the newly introduced superscript (M)), where $\kappa_k, \ k \in \mathbb{Z}_K$ are positive constants that represent the relative energy distribution over the support of γ_d. Note that, with this assumption, identifying the support set $\{\xi_k\}_{k \in \mathbb{Z}_K}$ becomes more challenging with increasing the array

dimension, since the associated coefficients keep decreasing. In contrast, in our setup the spike coefficients remain constant regardless of the number of antennas due to the physical fact that the power received from a specific AoA is independent of the array dimension. It can be seen that, we can adopt the analysis in [89] by assuming that the parameters κ_k follow a scaling law as $\kappa_k = M\kappa_{k,0}$ for positive constants $\kappa_{k,0}$, so that dividing them by M results in constant, dimension-independent spike coefficients.

Consider the noisy pilot observations $\mathbf{y}[i] = \mathbf{h}[i] + \mathbf{z}[i]$, $i \in \mathbb{Z}_N$, which is a zero-mean complex Gaussian vector with a covariance $\mathbf{\Sigma_y} = \mathbf{\Sigma} + N_0\mathbf{I}$, where we have

$$\mathbf{\Sigma} = \sum_{k\in\mathbb{Z}_K} \rho_k^{(M)}\mathbf{a}(\xi_k)\mathbf{a}(\xi_k)^{\mathsf{H}} + \int_{-1}^{1} \gamma_c(\xi)\mathbf{a}(\xi)\mathbf{a}(\xi)^{\mathsf{H}}d\xi, \qquad (2.85)$$

with $\mathbf{a}(\xi) = [1, e^{j\pi\xi}, \ldots, e^{j\pi(M-1)\xi}]^{\mathsf{T}}$ being the array response vector. Let $\lambda_{0,M} \geqslant \ldots \geqslant \lambda_{M-1,M}$ denote the eigenvalues (also singular values) of $\mathbf{\Sigma_y}$. The following proposition shows that as M grows, the K largest singular values of $\mathbf{\Sigma_y}$ "escape" from the rest of the singular values and converge to fixed values as $M \to \infty$.

Proposition 2.1 (escape of the K largest singular values of $\mathbf{\Sigma_y}$)**.** *Let $\gamma(\xi)$ denote an ASF with a discrete part $\gamma_d(\xi) = \sum_{k\in\mathbb{Z}_K} c_k^{(M)}\delta(\xi - \xi_k)$ with $\rho_k^{(M)} = \frac{\kappa_k}{M}$ and a continuous part $\gamma_c(\xi)$, and let $\mathbf{\Sigma}$ be the covariance matrix generated by $\gamma(\xi)$ and define $\mathbf{\Sigma_y} = \mathbf{\Sigma} + N_0\mathbf{I}$. If $\kappa_k + \gamma_c(\xi_k) > \|\gamma_c\|_\infty$ for all $k \in 0, \ldots, K-1$, then*

$$\lambda_{k,M} \xrightarrow[M\to\infty]{} \lambda_k = \kappa_k + \gamma_c(\xi_k) + N_0, \qquad (2.86)$$

for $k = 0, \ldots, K-1$, while $\limsup_{m\to\infty} \lambda_{K,M} \leqslant \|\gamma_c\|_\infty + N_0$ $\left(\|\gamma_c\|_\infty := \sup_{\xi\in[-1,1]} \gamma_c(\xi)\right)$.
□

Proof. See [89], Proposition 1. □

This proposition has the following implication: if for all $k \in \mathbb{Z}_K$, the spike coefficient plus the value of the continuous component γ_c at the spike location is greater than the supremum of the continuous part γ_c over the whole set of AoAs, then the K largest singular values converge according to (2.86) to a value larger than $\|\gamma_c\|_\infty + N_0$, while the rest of the singular values are upper-bounded by $\|\gamma_c\|_\infty + N_0$. As a result, by increasing M, we observe a separation between the first K singular values and the remaining $M - K$ ones, which can be exploited to identify the number of spikes K (see Section 2.4.1). This result holds for the true covariance matrix $\mathbf{\Sigma_y}$. However, it can be modified to also apply to the sample covariance $\widehat{\mathbf{\Sigma}}$, provided that the number of available signal samples for covariance estimation N is sufficiently large.

To characterize this rigorously, let us consider an asymptotic regime, in which the number of pilots N grows proportionally to the number of antennas M such that $\zeta_M = \frac{M}{N} \to \zeta > 0$ as $M \to \infty$. Although in practice the dimension M is finite, the ensuing asymptotic result gives a flavor of conditions under which support recovery is feasible for the discrete spikes.

Let $\eta_M(\lambda) = \frac{1}{M} \sum_{k \in \mathbb{Z}_M} \delta(\lambda - \lambda_{k,M})$ be the empirical measure that gives the spectrum of $\mathbf{\Sigma}_{\mathbf{y},M}$ for every M, where we have made the dimension (M) of the observation covariance $\mathbf{\Sigma}_{\mathbf{y}}$ explicit by adding a subscript. One can show that this empirical measure converges to a deterministic one. Call this deterministic measure η. Also, let $\widehat{\lambda}_{0,M}, \ldots, \widehat{\lambda}_{M-1,M}$ denote the eigenvalues of $\widehat{\mathbf{\Sigma}}_{\mathbf{y}}$ and define the random probability measure $\widehat{\mu}_M(\lambda) = \frac{1}{M} \sum_{k \in \mathbb{Z}_M} \delta(\lambda - \widehat{\lambda}_{k,M})$ as the empirical distribution of these eigenvalues. It is known that almost surely (a.s.) as $M \to \infty$ we have $\widehat{\mu}_M \overset{\text{weakly}}{\longrightarrow} \mu,$, where μ is a deterministic density with a Stieltjes transform as [4] (see [132] for the definition of the Stieltjes transform)

$$m(z) = \int_{\mathbb{R}} \frac{d\mu(\lambda)}{\lambda - z}, \tag{2.87}$$

where $m(z)$ is a function that satisfies the following fixed-point equation

$$m(z) = \int_{\mathbb{R}} \frac{d\eta(\lambda)}{\lambda(1 - \zeta - \zeta\, z\, m(z)) - z}, \tag{2.88}$$

for all $z \in \mathbb{C}\backslash\text{supp}(\mu)$, where $\text{supp}(\mu)$ denotes the support μ. Note that, although the eigenvalues of $\mathbf{\Sigma}_{\mathbf{y}}$ and $\widehat{\mathbf{\Sigma}}_{\mathbf{y}}$ have a well-defined limit as η and μ, these two limit distributions are different from each other for arbitrary $\zeta > 0$ and approach each for $\zeta \to 0$, namely, when the number of samples N becomes much larger than M, where in that case $\widehat{\mathbf{\Sigma}}_{\mathbf{y}}$ also converges to $\mathbf{\Sigma}_{\mathbf{y}}$.

To extend the separation condition proved in Proposition 2.1 for the true covariance $\mathbf{\Sigma}_{\mathbf{y}}$ to the sample covariance $\widehat{\mathbf{\Sigma}}_{\mathbf{y}}$, we need to study η further. From Szegö's theorem [48], it is well-known that η is given by the distribution of the random variable $\gamma_c(d\bar{\xi}) + N_0$, where $\bar{\xi}$ is a uniformly distributed random variable in $[-1, 1]$. Note that since the random variable $\gamma_c(d\bar{\xi}) + N_0$ is upper bounded by $\|\gamma_c\|_\infty + N_0$, the support of η always lies in the interval $\left[0, \|\gamma_c\|_\infty + N_0\right]$, and in particular $\max(\text{supp}(\zeta)) = \|\gamma_c\|_\infty + N_0$. This implies that the function

$$\phi(\omega) = \omega \left(1 - \zeta \int_{\mathbb{R}} \frac{\lambda}{\lambda - \omega} d\eta(\lambda) \right), \tag{2.89}$$

is well-defined for all $\omega \in (\|\gamma_c\|_\infty + N_0, +\infty)$. Note that $\phi(\omega)$ is a continuous and differentiable (of any order) function in this interval. Moreover, $\phi(\omega) \to \infty$ as $\omega \to \|\gamma_c\|_\infty + N_0$

and $\lim_{\omega \to \infty} \phi(\omega) = \infty$. Thus, $\phi(\omega)$ should have a local minimum $\omega_0 \in (\|\gamma_c\|_\infty + N_0, \infty)$. A direct computation shows that

$$\phi''(\omega) = \int \frac{2\zeta\lambda^2}{(\omega - \lambda)^3} d\eta(\lambda), \tag{2.90}$$

which is always positive in the interval $(\|\gamma_c\|_\infty + N_0, \infty)$. Hence, $\phi(\omega)$ is a convex function in this interval and a point ω_0 is the unique minimizer of $\phi(\omega)$. Note that ω_0 is the unique minimizer of $\phi(\omega)$, thus, it satisfies $\phi'(\omega_0) = 0$. Taking the derivative of $\phi(\omega)$, we can write the condition $\phi'(\omega_0) = 0$ as

$$\int \frac{\lambda^2}{(\omega_0 - \lambda)^2} d\eta(\lambda) = \frac{1}{\zeta}. \tag{2.91}$$

We can simply check that $\omega_0(\zeta)$ is an increasing function of ζ. In particular, by increasing the number of samples N, thus, letting $\zeta \to 0$, we have $\frac{1}{\zeta} \to \infty$, which is satisfied provided that ω_0 approaches the boundary value $\|\gamma_c\|_\infty + N_0$. Similarly, we can check that by decreasing the number of samples N in a scaling regime where $\zeta \to \infty$, we obtain $\omega_0 = \infty$. In brief, ω_0 ranges monotonically in the interval $(\|\gamma_c\|_\infty + N_0, +\infty)$ for $\zeta \in (0, +\infty)$.

The following theorem, proved in [89], shows that, similar to the escape of the K largest singular values in the spectrum of the true covariance $\boldsymbol{\Sigma_y}$ illustrated in Proposition 2.1, the K largest singular values of the sample covariance $\widehat{\boldsymbol{\Sigma}}_{\mathbf{y}}$ escape from the rest of its spectrum if a separation condition is satisfied. This separation condition can be formulated in terms of ω_0 as follows.

Theorem 2.2. *Let $\{\lambda_{k,M} : k \in \mathbb{Z}_K\}$ denote the set of singular values of $\boldsymbol{\Sigma_y}$ as before and suppose that*

$$\lambda_{k,M} > \omega_0. \tag{2.92}$$

Then, for $k = 0, \dots, K - 1$, with probability one as $M \to \infty$ we have

$$\widehat{\lambda}_{k,M} \to \phi(\lambda_k), \tag{2.93}$$

whereas $\widehat{\lambda}_{K+1,m} \to \phi(\omega_0) < \phi(\lambda_{K,M})$ (as ω_0 is the maximizer of $\phi(\omega)$). □

The separation condition (2.92) implicitly depends on the parameters of the spike elements $\{\kappa_k : k \in \mathbb{Z}_K\}$ as well as the continuous part of the ASF γ_c (through the function ϕ defined in (2.89)) and in particular on the asymptotic sampling ratio ζ. As a sanity check by increasing the number of samples $\zeta \to 0$, and $\omega_0(\zeta) \to \|\gamma_c\|_\infty + N_0$, and the separation condition in Theorem 2.2 becomes the same as that in Proposition 2.1, which makes sense since for large number of samples the sample covariance matrix $\widehat{\boldsymbol{\Sigma}}_{\mathbf{y}}$ converges to the

original covariance matrix $\mathbf{\Sigma_y}$. Moreover, as $\omega_0(\zeta) > \|\gamma_c\|_\infty + N_0$ for all ζ, Theorem 2.2 requires a stronger separation condition than Proposition 2.1, which is the cost one needs to pay for not having the original covariance matrix but the sample covariance matrix. Overall, if the separation condition (2.92) is satisfied, we are able to consistently detect the number of spikes by identifying the gap between the singular values. In particular,

$$\widehat{r}_M = \max\{k : \widehat{\lambda}_{k,m} > \phi(\omega_0) + \epsilon\} \xrightarrow{a.s.} K,$$

as $M \to \infty$ for any $\epsilon \in \big(0, \phi(\lambda_{K-1,M}) - \phi(\omega_0)\big)$. Eventually, the following theorem proves the consistency of the MUSIC estimator in the present context.

Theorem 2.3 (Consistency of MUSIC). *If the separation condition* (2.92) *holds, then*

$$M\left(\widehat{\xi}_k - \xi_k\right) \xrightarrow{a.s.} 0, \;\; k = 0, \ldots, K-1, \tag{2.94}$$

as $M \to \infty$, where $\widehat{\xi}_k$ denotes the K-th dominant minimizer of (2.30). □

It is worthwhile here to pose these results in the semi-rigorous setting we already discussed. More specifically, since in our case the amplitude of the spikes $\{\rho_k\}_{k \in \mathbb{Z}_K}$ remain constant (rather than decreasing with M), we can mimic this by assuming that the coefficients υ_k grow proportionally to M as $\kappa_k = M\rho_k$.

An important point of analysis in [89] summarized in this section is that for any asymptotic sampling ratio $\frac{M}{N} \to \zeta \in (0, \infty)$, no matter how small ζ may be, we can make the detection of all K spikes, namely, their number K and also their support, feasible by increasing the coefficients $\{\kappa_k : k \in \mathbb{Z}_K\}$ until the separability condition in (2.92) is fulfilled. Let us first illustrate this point step by step. First note that the measure $\eta(\lambda)$ depends only on the continuous part γ_c, thus, is not affected by changing the weights $\{\kappa_k : k \in \mathbb{Z}_K\}$. Therefore, for a fixed $\zeta \in (0, \infty)$, the function $\phi(\omega)$ and as a result the parameter ω_0 are not affected by changing $\{\kappa_k\}_{k \in \mathbb{Z}_K}$. Second, by dropping the contribution of the continuous part γ_c from $\mathbf{\Sigma_y}$, we can easily check that the first K singular values of $\mathbf{\Sigma_y}$ are larger than the first K singular values of the matrix

$$\mathbf{\Sigma_h} + N_0\mathbf{I} = \sum_{k \in \mathbb{Z}_K} \frac{\kappa_k}{M} \mathbf{a}(\xi_k)\mathbf{a}(\xi_k)^{\mathsf{H}} + N_0\mathbf{I}. \tag{2.95}$$

A direct calculation shows that as $M \to \infty$, where as a result $\frac{\langle \mathbf{b}(\xi_k), \mathbf{b}(\xi_{k'}) \rangle}{M} \to 0$ for $k \neq k'$, the K largest singular values approach to $\{\kappa_k + N_0 : k \in \mathbb{Z}_K\}$ which would satisfy (2.92) by increasing $\{\kappa_k : k \in \mathbb{Z}_K\}$ (as ω_0 is not affected by changing $\{\kappa_k : k \in \mathbb{Z}_K\}$). In brief, we

can say that the separability condition (2.92) is fulfilled if

$$\min\{\kappa_k : k \in \mathbb{Z}_K\} \geqslant \tau(\zeta, \gamma_c), \tag{2.96}$$

where $\tau(\zeta, \gamma_c)$ is a finite threshold that depends on the sampling ratio ζ and γ_c. As a result, assuming that $\kappa_k = \rho_k M$ grows proportionally to M, this condition would be satisfied for any finite $\zeta \in (0, \infty)$ and for any practically relevant γ_c provided that M is sufficiently large.

3 Uplink-Downlink Channel Covariance Transformation

3.1 Problem Statement

In a multi-user MIMO system, the base station (BS) needs to estimate instantaneous Downlink (DL) channels of all users to transmit precoded data that mitigates interference and ensures reliable multi-user communication. Knowledge of the second-order channel statistics at the BS is crucial, not only to obtain minimum mean squared error (MMSE) estimates of the instantaneous DL channel, but also in user grouping and precoder design (see [90] and references therein). In a TDD system obtaining the DL channel covariance is straightforward: the Uplink (UL) and DL channels are identical (due to channel reciprocity) and therefore the UL and DL channel covariances are the same. During UL transmission, the BS collects a number of independent pilot signals and estimates the UL channel covariance using, for example, the covariance estimation methods developed in Chapter 2. Since UL and DL covariances are identical, the UL covariance estimate is also an estimate for the DL covariance. In contrast, in an FDD system, channel reciprocity does not hold, since UL and DL signaling is performed over two disjoint frequency bands. Therefore, a conventional approach suggests that the BS transmits DL pilot signals, receives pilot measurements from the user via a closed-loop UL feedback mechanism, and uses these pilot measurements to estimate the DL instantaneous channel and eventually its covariance. This is an unfeasible strategy for a massive MIMO system, in which the channel dimension (= the number of BS antennas) is large ($M \gg 1$) and DL channel training requires a high pilot dimension, imposing a large training and feedback overhead on the system. Note that, UL channel covariance estimation is similar to the TDD case. The UL pilots are naturally sent from the user to the BS, which enables the BS to obtain a reliable estimate of the UL channel covariance. Is it possible to obtain a reliable estimate of the DL covariance $\boldsymbol{\Sigma}_{\mathsf{dl}}$ from the UL covariance $\boldsymbol{\Sigma}_{\mathsf{ul}}$, without DL pilot transmission and channel training?

In this chapter, we will show that the answer to the question above is affirmative: one can estimate $\boldsymbol{\Sigma}_{\mathsf{dl}}$ from $\boldsymbol{\Sigma}_{\mathsf{ul}}$ by solving convex programs and, under mild technical conditions,

most DL covariance entries can be estimated with sufficiently small error. Our results rely on the following two widely-accepted assumptions [59, 87, 134].

Assumption 3.1 (**Channel Angular Reciprocity**). *Let W_{ul} and W_{dl} be the zero-mean, complex Gaussian stochastic processes representing the angular channel gains in UL and DL, respectively. Although instantaneous realizations of W_{ul} and W_{dl} may be different, we assume that their increments have identical autocorrelation, i.e.*

$$\mathbb{E}[dW_{\mathsf{ul}}(\xi)dW_{\mathsf{ul}}(\xi')] = \mathbb{E}[dW_{\mathsf{dl}}(\xi)dW_{\mathsf{dl}}(\xi')] = \gamma(\xi)\delta(\xi - \xi')d\xi, \tag{3.1}$$

where $\delta(\cdot)$ is Dirac's delta, and γ is a non-negative measure that may include Deltas. The equality above implies that we assume the angular scattering function (ASF) γ to be an invariant of the channel during UL and DL transmission.

This assumption can be physically justified, by considering the fact that the UL-DL frequency bands occupy a fairly small part of the spectrum and they are typically closely spaced compared to their carrier frequencies (UL-DL band separation of the order of 100 MHz, for carrier frequencies ranging between 2 and 6 GHz) [111]. Therefore, being a function of the physical properties of scatterers and reflectors, the ASF remains the same over UL and DL bands.

Assumption 3.2 (**Stationarity**). *We assume that the ASF γ is piece-wise constant across time and remains unchanged over many time slots (channel coherence times). In other words, the ASF remains unchanged over a time range much larger than the (instantaneous) channel coherence time.*

Note that, this assumption is necessary for the BS to reliably estimate the UL covariance from UL pilots, that are transmitted by the user over multiple time slots. The assumption is justified as follows. The channel instantaneous fading coefficients are typically modeled as Gaussian by appealing to the central limit theorem [105, 123]. The variance of each coefficient depends on the scattering properties of the objects in the environment. Therefore, while the fading coefficients may change quickly in time and frequency, their second order statistics changes at a much slower rate.

3.2 System Setup

Consider a BS equipped with a uniform linear array (ULA) consisting of M equi-spaced antennas that operates in FDD mode over two bands with corresponding carrier frequencies f_{ul} for UL and f_{dl} for DL. We assume the same channel model as in Chapter 2. Similar

to (2.9) we can express UL and DL channels of an arbitrary user respectively as

$$\mathbf{h}_{\mathsf{ul}} = \int_{-1}^{1} dW_{\mathsf{ul}}(\xi)\mathbf{a}_{\mathsf{ul}}(\xi), \tag{3.2a}$$

$$\mathbf{h}_{\mathsf{dl}} = \int_{-1}^{1} dW_{\mathsf{dl}}(\xi)\mathbf{a}_{\mathsf{dl}}(\xi), \tag{3.2b}$$

where the integrals are understood as stochastic integrals [68], and where the UL and DL array response vectors are distinctly defined: the response of element $m \in \mathbb{Z}_M$ in UL and DL are given as

$$[\mathbf{a}_{\mathsf{ul}}(\xi)]_m = e^{j\frac{2\pi d}{\lambda_{\mathsf{ul}}}m\xi} \tag{3.3a}$$

$$[\mathbf{a}_{\mathsf{dl}}(\xi)]_m = e^{j\frac{2\pi d}{\lambda_{\mathsf{dl}}}m\xi} \tag{3.3b}$$

for $\xi \in [-1,1]$, where $\lambda_{\mathsf{ul}} = \frac{c}{f_{\mathsf{ul}}}$ and $\lambda_{\mathsf{ul}} = \frac{c}{f_{\mathsf{dl}}}$ are UL and DL wavelengths, respectively. Assume an antenna spacing of $d = \frac{\varrho\lambda_{\mathsf{ul}}}{2\sin(\theta_{\max})}$, where we have introduced the scalar $\varrho \in (0,1]$ as the *spatial oversampling factor*. This parameter is set to 1, for simplicity, in the other chapters of this dissertation. Note that since the normalized angular range $[-1,1]$ scanned by the array has an angular span of 2, the antenna spacing with $\varrho = 1$ and given by $d = \frac{\lambda_{\mathsf{ul}}}{2\sin(\theta_{\max})}$ is the maximum one required to avoid spatial aliasing or grating lobes, which is the reason we call ϱ the spatial oversampling factor. In array processing applications, where one deals with only a single frequency band, say, $\mathcal{F}_{\mathsf{ul}}$ or $\mathcal{F}_{\mathsf{dl}}$, it is conventional to set $\varrho = 1$ since this allows maximum physical span, thus, maximum angular resolution for the array. In this chapter, we study communication at two disjoint frequency bands $\mathcal{F} = \mathcal{F}_{\mathsf{ul}} \bigcup \mathcal{F}_{\mathsf{dl}}$ and we will assume $\varrho < \frac{f_{\mathsf{ul}}}{f_{\mathsf{dl}}}$ to avoid grating lobes [70,128] in both bands $\mathcal{F}_{\mathsf{ul}}$ and $\mathcal{F}_{\mathsf{dl}}$. We will see the role played by ϱ in our analysis shortly. With these conventions, the UL and DL array response vectors are given as

$$\mathbf{a}_{\mathsf{ul}} = \left[1, e^{j\pi\varrho\xi}, \ldots, e^{j(M-1)\pi\varrho\xi}\right]^{\mathsf{T}} \in \mathbb{C}^M, \tag{3.4a}$$

$$\mathbf{a}_{\mathsf{dl}} = \left[1, e^{j\pi\varrho\xi/\nu}, \ldots, e^{j(M-1)\pi\varrho\xi/\nu}\right]^{\mathsf{T}} \in \mathbb{C}^M, \tag{3.4b}$$

where $\nu := \frac{f_{\mathsf{ul}}}{f_{\mathsf{dl}}}$ is the UL to DL carrier ratio. In current FDD systems, the UL carrier is typically smaller than the DL carrier and therefore we have $\nu < 1$.

Remark 3.1. In general, the BS uses OFDM signaling to send and receive symbols to and from the user. The array response from one subcarrier to the next changes in UL as well as in DL. Here, we compare the channel corresponding to the *carrier frequency* in the UL and its counterpart in the DL. For other subcarriers, the same comparison applies

with the slight modification that the ratio ν defined above changes to the ratio between a subcarrier in UL to its counterpart in DL.

Remark 3.2. In (3.2) we have modeled UL/DL channels as M-dimensional random *Gaussian* vectors. This model represents a standard mathematical approximation of the *small-scale* channel fading, which is different across antenna elements since the elements have different responses to an incoming wave. However, this model assumes that *large-scale* channel coefficients due to path-loss and shadowing are identical across antenna elements. This assumption is justified when the array dimension is much smaller compared to the distance between the user and the BS. Later on, we will study the transformation problem in large channel dimensions ($M \to \infty$), which seems to be violating the aforementioned assumption. However, this is not a fundamental limitation, since in all practical implementations of massive MIMO, the array dimension will be large, but not so large as to be comparable to the user-BS distance. The large-dimension analysis is simply a theoretical tool for predicting the performance behavior of practical, large arrays.

The angular random processes W_{ul} and W_{dl} in (3.2) represent zero-mean, complex Gaussian channel gains in UL and DL, respectively. These processes can be seen as realizations of the small-scale fading along the AoAs as seen at the BS. The gain variance (average fading power) is captured by the respective autocorrelation function, that is the by the ASF. Based on Assumption 3.1, the ASF is identical in UL and DL, and therefore, the UL and DL covariances can be expressed in terms of the ASF, respectively, as:

$$\boldsymbol{\Sigma}_{\mathsf{ul}} = \int_{-1}^{1} \gamma(\xi) \mathbf{a}_{\mathsf{ul}}(\xi) \mathbf{a}_{\mathsf{ul}}(\xi)^{\mathsf{H}} d\xi, \tag{3.5a}$$

$$\boldsymbol{\Sigma}_{\mathsf{dl}} = \int_{-1}^{1} \gamma(\xi) \mathbf{a}_{\mathsf{dl}}(\xi) \mathbf{a}_{\mathsf{dl}}(\xi)^{\mathsf{H}} d\xi. \tag{3.5b}$$

Despite the fact that the ASF is an invariant of the UL and DL channels, the covariances differ since the array response is different in UL and DL. Being Hermitian Toeplitz matrices, $\boldsymbol{\Sigma}_{\mathsf{ul}}$ and $\boldsymbol{\Sigma}_{\mathsf{dl}}$ are characterized by their first columns:

$$\boldsymbol{\sigma}_{\mathsf{ul}} = \int_{-1}^{1} \gamma(\xi) \mathbf{a}_{\mathsf{ul}}(\xi) d\xi, \tag{3.6a}$$

$$\boldsymbol{\sigma}_{\mathsf{dl}} = \int_{-1}^{1} \gamma(\xi) \mathbf{a}_{\mathsf{dl}}(\xi) d\xi. \tag{3.6b}$$

Now, suppose we are given the UL channel covariance $\boldsymbol{\Sigma}_{\mathsf{ul}}$. Estimates of the UL covariance can be obtained via UL pilot transmission and using, for example, the covariance estimation techniques proposed in Chapter 2. Assumption 3.2 enables the BS to receive

sufficient UL pilots while the channel covariance remains the same. This chapter provides answers for the following two questions.

- What is a minimax estimation error bound on the estimated DL covariance $\hat{\boldsymbol{\Sigma}}_{\mathsf{dl}}$ obtained from $\boldsymbol{\Sigma}_{\mathsf{ul}}$ with respect to the true DL covariance $\boldsymbol{\Sigma}_{\mathsf{dl}}$?

- How can one efficiently estimate the DL covariance given the UL covariance?

3.3 Related Work and Contribution

The problem of UL-DL covariance transformation has been studied now for about two decades [56, 60]. It has gained more attention in the last five years due to the advent of massive MIMO and the need for adapting these systems to the FDD operation mode [32, 35, 54, 107, 129]. These works typically rely on heuristic techniques and barely discuss performance limits in terms of, for example, transformation error bounds. Compared with these works, we provide the following major contributions:

1. We mathematically prove that under the assumptions 3.1 and 3.2 and mild conditions on the array spatial oversampling factor (ϱ) and the UL-DL carrier ratio ($\nu = \frac{f_{\mathsf{ul}}}{f_{\mathsf{dl}}}$) the DL covariance matrix $\boldsymbol{\Sigma}_{\mathsf{dl}}$ can be *stably* estimated from the UL covariance $\boldsymbol{\Sigma}_{\mathsf{ul}}$.

2. Our analysis of the UL-DL covariance transformation problem implies a robustness guarantee for a special class of estimators. In short, all estimators that satisfy minimal structural assumptions on the underlying channel ASF, namely non-negativity and data-consistency, result in fairly small estimation error for most DL covariance entries.

3. As a by-product of our analysis, we prove that for any given pair of parameters (ϱ, ν), there is always a subset of elements of $\boldsymbol{\Sigma}_{\mathsf{dl}}$ that may not be well estimated from the observation of $\boldsymbol{\Sigma}_{\mathsf{ul}}$. This effect is totally neglected in all the ad-hoc schemes previously proposed in the literature.

Perhaps the most relevant results to the topic of this chapter were published in a series of papers [22], [86], and [87].[1] Using set-theoretic results, the authors of [22] derive general error bounds for estimating a DL covariance entry, given the UL covariance. The strong aspect of this result lies in its generality for arbitrary arrays, and a wide range of admissible angular power spectra, treating the ASF as an infinite-dimensional element of a Hilbert space without too restrictive structural assumptions (such as sparsity and/or consisting of

---

[1]All three papers were published after our work first appeared on the online repository arXiv [52], and later published at the IEEE International Symposium on Information Theory (ISIT), 2018 [53].

only line spectra). However, one caveat of this work is that, explicit error bounds are given only for the case in which the non-negative angular scattering function γ is bounded in the induced norm as $\|\gamma\|_{\mathcal{H}}$, where $\|\cdot\|_{\mathcal{H}}$ is the norm defined by the inner-product associated with the Hilbert space $\mathcal{H}$. Particularly, when $\mathcal{H} = L_2$ the space of square-integrable functions, one uses the standard Euclidean norm which necessitates $\|\gamma\|_2 \leqslant B$ for some $B \in \mathbb{R}_+$. But when γ includes Delta impulses, this condition does not hold, since the Delta is not an element in L_2, and therefore the error bound is violated. This however is not a negligible example. It is quite natural for the ASF to consist of Delta impulses that are associated with line-of-sight (LoS) and specular scattering components. In fact the majority of related works in the literature consider the ASF to be consisting of *only* Dirac impulses. It is not clear how one can resolve this issue by using a different inner product and norm. For example, resorting to the ℓ_1 norm enables bounding the norm of γ by assuming $\int_{-\infty}^{\infty} \delta = 1$, but then it is not clear what inner product induces the ℓ_1 norm. Since the bounds depend heavily on the definition of the inner product, we can not evaluate the proposed bounds in [22] when considering $\gamma \in L_1$, where L_1 is the space of Lebesgue-integrable functions. In contrast, our error bounds hold also for the case in which the spectrum involves Delta impulses.

Apart from theoretical results, strong UL-DL covariance transformation algorithms were proposed in the papers [86], and [87] companion to [22]. These methods are based on solving convex feasibility problems [6], and are proposed in two variants which differ in the constraints they impose on the ASF, namely, non-negativity and data consistency constraints. We discuss the relation between our proposed covariance transformation method and those suggested in [86, 87] and provide simulation results to compare their performance.

3.4 Minimax Error Bounds

A natural way to study the limitations and potentials of UL-DL covariance transformation is to derive *minimax* estimation error bounds. We rigorously define this worst-case error and specify regimes in which a reliable estimation can be performed. We hasten to emphasize that, unlike most works in the literature, we do not impose any structural assumption on the spectrum, such as for it to be consisting of only line spectra or to be a sparse measure. This generality is crucial in order to cover a wider class of possible ASF forms in a communication scenario.

We assume w.l.o.g that γ is a normalized non-negative measure over $[-1, 1]$ with $\gamma([-1, 1]) = 1$. Define the continuous Fourier transform of γ as $\breve{\gamma} : \mathbb{R} \to \mathbb{C}$:

$$\breve{\gamma}(x) = \int_{-1}^{1} \gamma(\xi) e^{j\pi\xi x} d\xi, \ x \in \mathbb{R}. \tag{3.7}$$

Since γ is a normalized measure with a bounded support $[-1, 1]$, $\breve{\gamma}(x)$ is a continuous function of x due to

$$\lim_{h \to 0} |\breve{\gamma}(x + h) - \breve{\gamma}(x)| \leqslant \lim_{h \to 0} \int_{-1}^{1} \gamma(\xi)|e^{j\pi\xi x} - e^{j\pi\xi(x+h)}|d\xi$$

$$\leqslant \lim_{h \to 0} \int_{-1}^{1} \gamma(\xi)|1 - e^{j\pi\xi h}|d\xi = 0,$$

where the last equality follows from the dominated convergence theorem as $|1 - e^{j\pi\xi h}| \leqslant 2$ is a bounded function approaching 0 for all $\xi \in [-1, 1]$ as $h \to 0$. Similarly, one can verify that $\breve{\gamma}(0) = 1$ and $|\breve{\gamma}(x)| \leqslant |\breve{\gamma}(0)|$ for all $x \in \mathbb{R}_+$, and that $\breve{\gamma}(x)$ has conjugate symmetry, namely, $\breve{\gamma}(-x)^* = \breve{\gamma}(x)$ for any $x \in \mathbb{R}$. Moreover, being the Fourier transform of a positive measure, it is also a *positive definite* function, i.e., $\sum_{i,j=1}^{l} c_i c_j^* \breve{\gamma}(x_i - x_j) \geqslant 0$ for any l, any $\{x_i\}_{i \in \mathbb{Z}_l} \subset \mathbb{R}$, and any sequence of complex numbers $\{c_i\}_{i \in \mathbb{Z}_l}$; We refer to [126] for a comprehensive introduction to positive definite functions and their applications, and to [8] for the connection with *reproducing kernel Hilbert spaces* (RKHS). Also, seen as a function of $x \in \mathbb{R}$, $\breve{\gamma}(x)$ is band-limited with a spectrum bounded in $[-1, 1]$.

Now let us consider $\boldsymbol{\sigma}_{\mathsf{ul}}$. From (3.6a) and (3.7) it is seen that

$$[\boldsymbol{\sigma}_{\mathsf{ul}}]_m = \int_{-1}^{1} \gamma(\xi) e^{jm\pi\varrho\xi} d\xi = \breve{\gamma}(m\varrho), \ m \in \mathbb{Z}_M. \tag{3.8}$$

Hence, the "samples" $\{\breve{\gamma}(m\varrho) : m \in \mathbb{Z}_M\}$ of $\breve{\gamma}$ at the lattice sampling points $\{m\varrho : m \in \mathbb{Z}_M\}$ correspond to the M elements of $\boldsymbol{\sigma}_{\mathsf{ul}}$. Similarly, it is not difficult to check that

$$[\boldsymbol{\sigma}_{\mathsf{dl}}]_m = \breve{\gamma}(\frac{m\varrho}{\nu}), \ m \in \mathbb{Z}_M. \tag{3.9}$$

This implies that $\boldsymbol{\sigma}_{\mathsf{dl}}$, and as a results $\boldsymbol{\Sigma}_{\mathsf{dl}}$, can be obtained from the samples of $\breve{\gamma}$ at positions $\{\frac{m\varrho}{\nu} : m \in \mathbb{Z}_M\}$. With this explanation, we can pose the problem of UL-DL covariance transformation as follows.

Problem 1. *Given the set of M UL samples $\{\breve{\gamma}(m\varrho) : m \in \mathbb{Z}_M\}$ of the band-limited function $\breve{\gamma}$, with an unknown non-negative spectrum $\gamma(\xi)$ supported over $[-1, 1]$, find the values of the corresponding DL samples $\{\breve{\gamma}(\frac{m\varrho}{\nu}) : m \in \mathbb{Z}_M\}$.*
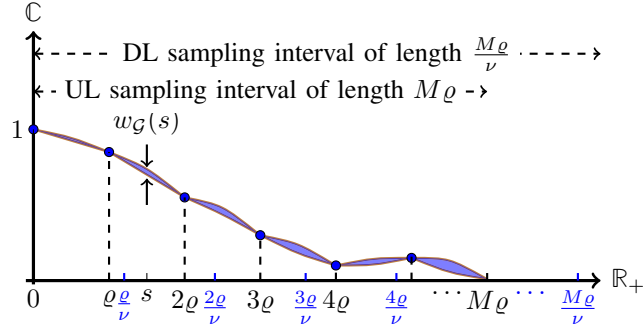
Figure 3.1: Illustration of the graph $\mathcal{G} \subset [0, M\varrho] \times \mathbb{C}$ associated to a specific measure and its width at a specific probing point s. Note that the vertical axis corresponds to the complex plane or more precisely to the interior of the complex unit sphere in the complex plane as $|\breve{\gamma}(x)| \leqslant |\breve{\gamma}(0)| = 1$, and due to the conjugate symmetry we have plotted $\breve{\gamma}(x)$ only for $x \in \mathbb{R}_+$.

This problem is schematically visualized in Fig. 3.1. Note that since by our assumption, $\breve{\gamma}(x)$ is band-limited, from the sampling theorem[2] [125], we should be able to recover γ from the samples $\{\breve{\gamma}(m\varrho) : m \in \mathbb{Z}_+\}$ (even without any spatial oversampling, i.e., for $\varrho = 1$), thus, to estimate $\breve{\gamma}(x)$ at any arbitrary $x \in \mathbb{R}_+$. When we have only finitely many samples $\{\breve{\gamma}(m\varrho) : m \in \mathbb{Z}_M\}$, given the band-limitedness and smoothness of $\breve{\gamma}$, we may still expect to estimate $\breve{\gamma}(x)$ for those x inside the UL sampling interval $[0, M\varrho]$ with a moderately small error that vanishes as $M \to \infty$. However, in UL-DL covariance transformation, there is always a subset of DL sampling interval $[0, \frac{M\varrho}{\nu}]$ that lies near the boundary of the UL sampling interval $[0, M\varrho]$ (see UL/DL sampling intervals in Fig. 3.1). In fact, one can argue that no matter how large M is, for most non-negative measures γ, those boundary points suffer from some interpolation/extrapolation error and cannot be approximated very well from the UL samples $\{\breve{\gamma}(m\varrho) : m \in \mathbb{Z}_M\}$. However, when $\breve{\gamma}(x)$ decays sufficiently fast in terms of x, and M is large, we expect that the samples of $\breve{\gamma}$ close to the boundary of $[0, M\varrho]$ have a small amplitude (energy) and their contribution to the DL covariance matrix is negligible.

It is also important to note that our explanation in this section confirms that for a large number of antennas M, the UL and DL covariance matrices can differ significantly from each other since they are obtained by sampling $\breve{\gamma}(x)$ at quite different sampling intervals (see Fig. 3.1), so an appropriate covariance transformation from UL to DL in relevant applications is inevitable.

---

[2]As a brief note, we would like to mention that here, for convenience, we defined the Fourier transform in (3.7) by $\pi\xi x$ rather than the conventional $2\pi\xi x$, thus, the bandwidth of $\breve{\gamma}$ in the conventional notation is $\frac{1}{2}$ (rather than 1). Therefore, samples of $\breve{\gamma}$ at $\mathbb{Z}_+$ have a sampling rate equal to (more than when $\varrho < 1$) twice the bandwidth of γ and are sufficient for its recovery according to the sampling theorem.

3.4.1 Main Result

As mentioned above, UL-DL transformation problem is equivalent to Problem 1, and here we focus on a minimax approach to this transformation (or *resampling*) problem.

We first define Γ_γ as the set of all positive normalized measures μ that are supported on $[-1, 1]$ and yield the same UL covariance matrix as γ, or more specifically:

$$\Gamma_\gamma := \left\{ \mu : \mu([-1, 1]) = 1, \breve{\mu}(m\varrho) = \breve{\gamma}(m\varrho), m \in \mathbb{Z}_M \right\}, \tag{3.10}$$

where $\breve{\mu}$ denotes the Fourier transform of μ as in (3.7). We consider the UL probing window $[0, M\varrho] \subset \mathbb{R}_+$ (see Fig. 3.1) and define the image of the set Γ_γ over the probing window $[0, M\varrho]$ under the Fourier transform as

$$\mathcal{G} := \bigcup_{\mu \in \Gamma_\gamma} \left\{ (x, \breve{\mu}(x)) : x \in [0, M\varrho] \right\} \subset [0, M\varrho] \times \mathbb{C}, \tag{3.11}$$

which is given as the union of the graph of the Fourier transforms of all $\mu \in \Gamma_\gamma$. Define the section of the graph $\mathcal{G}$ at a probing point $s \in [0, M\varrho]$ by $\mathcal{G}_s = \{z \in \mathbb{C} : (s, z) \in \mathcal{G}\}$. Also define the width of $\mathcal{G}$ at a point $s \in [0, M\varrho]$ as the *diameter* of $\mathcal{G}_s$:

$$w_{\mathcal{G}}(s) := \sup_{a,b \in \mathcal{G}_s} |a - b|. \tag{3.12}$$

Fig. 3.1 illustrates the graph associated with a specific measure γ and its width at a specific point $s \in [0, M\varrho]$. It is important to note that since all the measures in Γ_γ have the same Fourier transform at the sampling points $\{k\varrho : k \in \mathbb{Z}_M\}$, we have that $w_{\mathcal{G}}(s) = 0$ for $s \in \{k\varrho : k \in \mathbb{Z}_M\}$. Also, note that $w_{\mathcal{G}}(s)$ measures the variation in the Fourier transform of the densities in Γ_γ at a specific point s. As a result, by controlling $w_{\mathcal{G}}(s)$, we can obtain an estimate of the worst-case error of an algorithm that estimates $\breve{\gamma}(s)$ by merely observing $\breve{\gamma}$ at the sampling points $\{k\varrho : k \in \mathbb{Z}_M\}$. The following theorem is the main result of this section.

Theorem 3.1. *Let γ be an arbitrary non-negative and normalized measure supported in $[-1, 1]$. Let $[0, M\varrho]$ be the UL probing window and let $s \in [0, M\varrho]$ be an arbitrary point. Let $\mathcal{G}$ be the graph corresponding to γ and let $w_{\mathcal{G}}(s)$ be the width of $\mathcal{G}$ at $s \in [0, M\varrho]$. Then, there exists a universal constant $C \in \mathbb{R}_+$ such that*

$$w_{\mathcal{G}}(s) \leqslant \min \left\{ C\left(\sin(\frac{\pi\varrho}{2}) g(\frac{s}{M\varrho}) \right)^{2M}, 2 \right\}, \tag{3.13}$$
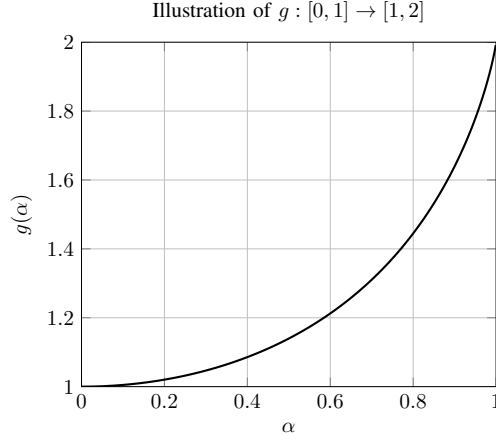
Figure 3.2: illustration of the function g over $[0, 1]$.

where $g : [0, 1] \rightarrow [1, 2]$ is a function that is independent of γ, M, and s, and given explicitly by $g(\alpha) = e^{f(\alpha)}$ for $\alpha \in [0, 1]$ where

$$f(\alpha) = \left(1 - h_2(\frac{1 + \alpha}{2})\right) \log(2), \tag{3.14}$$

and where $h_2(x) = -x \log_2(x) - (1 - x) \log_2(1 - x)$ is the binary entropy function [31] for $x \in [0, 1]$. □

Fig. 3.2 illustrates the function g in Theorem 3.1 over the interval $[0, 1]$.

3.4.2 Roadmap to the Proof

In the sequel we develop the proof of Theorem 3.1. But first, let us discuss some of the algorithmic implications of Theorem 3.1 for UL-DL covariance transformation. Consider an arbitrary algorithm that produces an estimate μ of the true ASF from the observations $\{\breve{\gamma}(m\varrho) : m \in \mathbb{Z}_M\}$ that satisfies $\{\breve{\mu}(m\varrho) = \breve{\gamma}(m\varrho) : k \in \mathbb{Z}_M\}$. Let $s \in [0, M\varrho]$ be an arbitrary probing point, and let $\breve{\gamma}(s)$ and $\breve{\mu}(s)$ be the Fourier transform of γ and that of the estimate μ at s, respectively. Theorem 3.1 implies that

$$|\breve{\gamma}(s) - \breve{\mu}(s)| \leqslant w_{\mathcal{G}}(s) \leqslant \min\left\{C\left(\sin(\frac{\pi\varrho}{2})g(\frac{s}{M\varrho})\right)^{2M}, 2\right\}$$

decays exponentially fast to 0 as M tends to infinity, provided that $\sin(\frac{\pi\varrho}{2})g(\frac{s}{M\varrho}) < 1$. Moreover, since g is an increasing function in $[0, 1]$, we also have that for any $\mathcal{W}_0 \subseteq [0, M\varrho]$

$$\sup_{s \in \mathcal{W}_0} |\breve{\gamma}(s) - \breve{\mu}(s)| \leqslant \min\left\{C \sup_{s \in \mathcal{W}_0} \left(\sin(\frac{\pi\varrho}{2})g(\frac{s}{M\varrho})\right)^{2M}, 2\right\}$$

Figure 3.3: Illustration of the variation in sampling by changing the spatial oversampling factor $\varrho \in [0,1]$. It is seen that for a fixed M, the UL and the DL observation intervals grow proportional to ϱ. For any ϱ, only a fraction of DL covariance entries in a region of size $M\varrho\alpha(\varrho)$ denoted by "S" are guaranteed to be robustly estimated from the UL samples, where a robust estimation might generally fail for the rest of the DL coefficients lying in a window denoted by "F".

$$= \min\left\{ C\left(\sin(\frac{\pi\varrho}{2})g(\frac{s_{\max}}{M\varrho}) \right)^{2M}, 2 \right\},$$

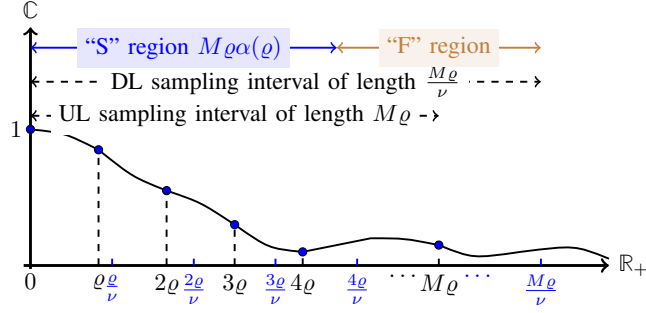where $s_{\max} = \sup\{s : s \in \mathcal{W}_0\}$. Thus, the worst case error over any probing window $\mathcal{W}_0$ can be controlled by the largest element $\sup\{s : s \in \mathcal{W}_0\}$ of $\mathcal{W}_0$. In particular, since $g(\alpha) \in [1,2]$, the estimation is precise over the whole window $[0, M\varrho]$ for $M \to \infty$, when $\sin(\frac{\pi\varrho}{2}) \leqslant \frac{1}{2}$ or equivalently when $\varrho \leqslant \frac{1}{3}$. For example, in a practical case, where the antenna scans the angular range $\Theta = [-\theta_{\max}, \theta_{\max}]$ with a $\theta_{\max} = 60$ degrees, this would require an antenna spacing of $d = \frac{\lambda_{\mathrm{ul}}}{3\sqrt{3}}$ where λ_{ul} denotes the wavelength at the UL carrier frequency f_{ul}.

Overall, since the elements of $\boldsymbol{\sigma}_{\mathsf{dl}}$ correspond to the samples of $\breve{\gamma}$ at locations $s \in \{\frac{m\varrho}{\nu} : m \in \mathbb{Z}_M\}$, from the condition $\sin(\frac{\pi\varrho}{2})g(\frac{s}{M\varrho}) < 1$ needed for $w_{\mathcal{G}}(s) \to 0$ asymptotically as $M \to \infty$, it immediately results that for any $\varrho < 1$, one can stably estimate from $\boldsymbol{\sigma}_{\mathsf{ul}}$ those components of $\boldsymbol{\sigma}_{\mathsf{dl}}$ with indices belonging to

$$\mathcal{I}_{\mathrm{dl}}(\varrho) := \{m \in \mathbb{Z}_M : m \leqslant M\nu, \sin(\frac{\pi\varrho}{2})g(\frac{m}{M\nu}) < 1\}. \tag{3.15}$$

Since $|\mathcal{I}_{\mathrm{dl}}(\varrho)| \leqslant M\nu$, the result of Theorem 3.1 guarantees the stable recovery of only a fraction of components of $\boldsymbol{\sigma}_{\mathsf{dl}}$ consisting of the first $|\mathcal{I}_{\mathrm{dl}}(\varrho)|$ elements (see Fig. 3.1 and Fig. 3.3). As a result, there are always a subset of components of $\boldsymbol{\sigma}_{\mathsf{dl}}$ consisting of $M - |\mathcal{I}_{\mathrm{dl}}(\varrho)|$ elements that are not guaranteed to be reliably estimated from $\boldsymbol{\sigma}_{\mathsf{ul}}$.

Underlying Trade-offs

The result of Theorem 3.1 might be misleading since it seems to suggest that one needs to select a smaller ϱ to obtain a better transformation performance. However, one should

note that selecting a small ϱ creates a significant spatial correlation among the antennas and reduces the spatial *degrees-of-freedom* (DoF) of the array. In words, for a given ϱ, the spatial resolution, thus, the number of effective spatial eigen-functions of the array scales as $O(M\varrho)$. For example, in the extreme case of $\varrho \to 0$, all the antennas are collocated and it is as if having only a single antenna. We refer to [97] for a more rigorous explanation of the spatial DoF of the array and related information-theoretic trade-offs. Here, we only provide an intuitive explanation and mainly focus on the non-asymptotic regime with a finite number of antennas M. Consider a specific angular ASF γ supported on $[-1,1]$. Using the error bound in Theorem 3.1, we see that we can robustly estimate, from UL observations $\{\breve{\gamma}(m\varrho) : m \in \mathbb{Z}_M\}$ obtained via UL covariance estimation, only those DL coefficients $\{\breve{\gamma}(m\frac{\varrho}{\nu}) : m \in \mathbb{Z}_M\}$ that lie inside the window $[0, M\varrho\alpha]$ provided that $\alpha \in [0,1]$ satisfies $\sin(\frac{\pi\varrho}{2})g(\alpha) < 1$. We define the largest such α by $\alpha(\varrho) = \sup\{\alpha : \sin(\frac{\pi\varrho}{2})g(\alpha) < 1\} = g^{-1}(\frac{1}{\sin(\frac{\pi\varrho}{2})})$. Thus, for any ϱ we can only guarantee the robust estimation of those DL coefficients that are obtained by sampling $\breve{\gamma}$ inside the window $[0, M\varrho\alpha(\varrho)]$ consisting of the first

$$N(\varrho) = \frac{M\varrho\alpha(\varrho)}{\frac{\varrho}{\nu}} = M\nu g^{-1}(\frac{1}{\sin(\frac{\pi\varrho}{2})}), \tag{3.16}$$

sampling points of the DL coefficients $\{\breve{\gamma}(m\frac{\varrho}{\nu}) : m \in \mathbb{Z}_M\}$. This has been illustrated in Fig. 3.3. Intuitively speaking, for any ϱ, we have $M\varrho$ spatial DoF in the UL among which only a fraction of $\alpha(\varrho) \in [0,1]$ can be robustly estimated and used for the DL, thus, a total of $MD(\varrho)$ robust DoF for the DL where $D(\varrho) = \varrho\alpha(\varrho)$. Fig. 3.4 illustrates $D(\varrho)$ for $\varrho \in [0.5, 1]$, where it is seen that the maximum $D(\varrho)$ for $\varrho \in (0,1)$ is achieved at $\varrho \approx 0.5$. For a ULA with $\theta_{\max} = 60$ degrees, this corresponds to an antenna spacing of $d = \frac{\sqrt{3}}{3}\frac{\lambda_{\mathrm{ul}}}{2}$.

We should also point out that the result stated in Theorem 3.1 and also our explanation in this section follow from the minimax analysis of the transformation problem where we consider all possible angular ASF's and even the worst-case ones. The situation is much better in practice when one deals with a much more structured class of angular ASF's γ. For example, consider the class of all ASF's γ whose Fourier transform $\breve{\gamma}(x)$ has a negligible energy beyond $|x| > \Delta$ for some fixed $\Delta > 0$. In such a case, we expect that all the significant components of $\breve{\gamma}(x)$ in $|x| \leqslant \Delta$ be recovered precisely if $M\varrho\alpha(\varrho) \approx \Delta$ (see, e.g., Fig. 3.3), where one can essentially ignore (i.e., zero pad) the coefficients in $|x| > \Delta$ since they are negligibly small. Note that the condition $M\varrho\alpha(\varrho) \approx \Delta$ can be fulfilled even with a moderately large number of antennas. Moreover, for any fixed Δ, one can select ϱ very close to 1 (no spatial oversampling, thus, maximum spatial resolution) as $M \to \infty$, thus, reconstructing the DL covariance matrix without paying any penalty in spatial DoF.
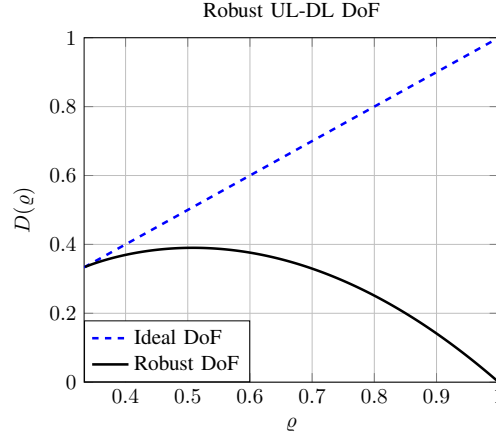
Robust UL-DL DoF

Figure 3.4: Comparison between the ideal and achievable (robust) UL-DL DoF (due to frequency variation) as a function of spatial oversampling factor ϱ. It is seen that for $\varrho \leqslant \frac{1}{3}$, both DoFs are equal. Moreover, the maximum robust DoF is achieved for $\varrho \approx 0.5$.

Remark 3.3. Throughout this chapter we have assumed, as in almost all massive MIMO deployments, $f_{\mathrm{ul}} < f_{\mathrm{dl}}$, thus, $\nu = \frac{f_{\mathrm{ul}}}{f_{\mathrm{dl}}} < 1$. In such a case, the desired interpolation window $[0, \frac{M\varrho}{\nu}]$ needed for UL goes much beyond the observation window $[0, M\varrho]$ in the DL (see, e.g., Fig. 3.1 and Fig. 3.3), which makes a reliable UL-DL covariance extrapolation quite challenging. In contrast, when $f_{\mathrm{ul}} > f_{\mathrm{dl}}$ and $\nu = \frac{f_{\mathrm{ul}}}{f_{\mathrm{dl}}} > 1$, the interpolation window $[0, \frac{M\varrho}{\nu}]$ for the DL is a smaller subset of observation window $[0, M\varrho]$ in the UL. As a result, the UL-DL covariance extrapolation is indeed much easier, even with a moderate number of antennas M, and can be done with much less spatial oversampling, i.e., with a larger ϱ.

The Chebyshev Differential Equation

In this section, we develop suitable approximation techniques and derive upper bounds on the approximation error using Chebyshev functions. We will use these results in Section 3.4.4 to obtain suitable minimax upper bounds and to finally prove Theorem 3.1. We first consider the Chebyshev second order *ordinary differential equation* (ODE) given by [143]

$$(1 - t^2)y'' - ty' + \varsigma^2 y = 0, \tag{3.17}$$

where $\varsigma \in \mathbb{R}_+$ is a fixed parameter. Since the coefficients of the ODE (3.17) are differentiable infinitely many times in a neighborhood of $t = 0$, the ODE has an analytic solution as a power series $y(t) = \sum_{n=0}^{\infty} a_n t^n$ around $t = 0$. Moreover, since the ODE (3.17) has singular points at $t = \pm 1$, this series solution has a convergence radius of at most 1 around $t = 0$ [14, 29]. As we will see in Section 3.4.4, we will need the solutions of this ODE in the interval $[-\eta, \eta]$ for $\eta = \sin(\frac{\pi\varrho}{2})$, which will be included in the region $(-1, 1)$ as $\varrho \in (0, 1)$.

Replacing the power series in (3.17), we obtain the following recursion for the coefficients:

$$a_{n+2} = \frac{n^2 - \varsigma^2}{(n+1)(n+2)} a_n, \; n = 0, 1, \ldots. \tag{3.18}$$

It is seen that the resulting recursion has order 2, thus, it yields two linearly independent solution $y_e(t)$ and $y_o(t)$ for the initial values ($a_0 = 1, a_1 = 0$) and ($a_0 = 0, a_1 = 1$). In fact, y_e and y_o are even and odd functions of t respectively, thus, they are linearly independent and span the two-dimensional space (as the ODE is second order) of the solutions of the ODE (3.17), namely, any arbitrary solution $y(t)$ can be written as a linear combination of y_e and y_o. It is also worthwhile to mention that when $\varsigma = 2k_0$ is an even integer, $y_e(t)$ is an even polynomial of degree $2k_0$, whereas $y_o(t)$ has infinitely many terms in its power series. Similarly, when $\varsigma = 2k_0 + 1$ is an odd integer, $y_o(t)$ is an odd polynomial of order $2k_0 + 1$, whereas $y_e(t)$ has infinitely many terms in its power series. These polynomial solutions correspond to Chebyshev polynomials of even and odd order.

Explicit Formulas for the Solutions and Error Bounds
A direct calculation shows that $\cos(\varsigma \sin^{-1}(t))$ and $\sin(\varsigma \sin^{-1}(t))$ satisfy the Chebyshev ODE with parameter ς. Since $\cos(\varsigma \sin^{-1}(t))$ is an even function over $t \in (-1, 1)$, entails a power series expansion around $t = 0$, and satisfies $\cos(\varsigma \sin^{-1}(t)) \vert_{t=0} = 1$ and

$$\frac{d}{dt} \cos(\varsigma \sin^{-1}(t)) \vert_{t=0} = 0 \;,$$

from $y_e(0) = 1, y_e'(0) = 0$ and the uniqueness of the solutions of ODE (3.17), it should correspond to $y_e(t)$. Similarly, we can check that $y_o(t) = \frac{\sin(\varsigma \sin^{-1}(t))}{\varsigma}$. We will mainly focus on $y_e(t)$ with a parameter $\varsigma = 2s$ for $s \in [0, M]$. We have

$$y_e(t) = 1 + \sum_{k=1}^{\infty} a_{2k} t^{2k}, \tag{3.19}$$

where a_{2k} from (3.18) is given by

$$a_{2k}(s) = \prod_{n=0}^{k-1} \frac{(2n)^2 - (2s)^2}{(2n+1)(2n+2)}, \tag{3.20}$$

where we also represented explicitly the dependence of a_{2k} on s. We will keep s fixed in this section and will drop the explicit dependence on s for notation simplicity.

We first write $y_e(t)$ as follows

$$y_e(t) = 1 + \sum_{k=1}^{M-1} a_{2k} t^{2k} + \sum_{k=M}^{\infty} a_{2k} t^{2k} =: y_{e,M}(t) + E_M(t),$$

where $E_M(t)$ denotes the truncation error consisting of the terms with exponents larger than or equal to $2M$. We first prove the following key result.

Proposition 3.1. *Let $E_M(t)$ be the truncation error of order M as define before. Then, we have the following:*

1. *$E_M(t) = (-1)^{\lfloor s \rfloor} \sum_{k=M}^{\infty} |a_{2k}| t^{2k}$ for all $t \in (-1, 1)$, where $\lfloor s \rfloor$ denotes the largest integer smaller than s. In particular, $E_M(t)$ has the same sign for all $t \in (-1, 1)$.*

2. *For any fixed $\eta \in (0, 1)$, the error $E_M(t)$ converges to 0 uniformly for $t \in [-\eta, \eta]$ with the maximum error occurring at the boundary $t = \pm \eta$, i.e., $\max_{t \in [-\eta, \eta]} |E_M(t)| = |E_M(\pm \eta)|$.*

3. *Over the interval $[-\eta, \eta]$, the truncation error $E_M(t)$ is upper bounded by $a_{2M} \frac{\eta^{2M}}{1 - \eta^2}$.*

4. *For any fixed $\eta \in (0, 1)$, the derivative of the truncation error $E_M(t)$ also converges uniformly to 0 in the interval $t \in [-\eta, \eta]$ and satisfies $\max_{t \in [-\eta, \eta]} |E_M'(t)| \leqslant 2|a_{2M}| \frac{\eta^{2M-1}(M - (M-2)\eta)}{(1 - \eta^2)^2}$.*

□

Proof. To prove part 1, first note from (3.20) that the coefficients $a_{2k}(s)$ have alternating signs for $k \leqslant s$ (due to multiplication by the negative factor $(2k)^2 - (2s)^2$), whereas all the coefficients with $k > s$ have the same sign (since $(2k)^2 - (2s)^2$ is positive). As $s \in [0, M]$, this implies that all the coefficients a_{2k} for $k \geqslant M$ have the same sign which can be checked to be $(-1)^{\lfloor s \rfloor}$. As a result, we can write $E_M(t) = (-1)^{\lfloor s \rfloor} \sum_{M}^{\infty} |a_{2k}| t^{2k}$.

To prove part 2, note that from part 1 and the fact that t^{2k} are increasing functions of t, it immediately results that the maximum of $E_M(t)$ over $t \in [-\eta, \eta]$ is achieved at the boundary point $t = \pm \eta$. Since $\eta \in (-1, 1)$, from the convergence of the series at η, it results that $|E_M(\pm \eta)|$ converges to 0 as M tends to infinity. This implies the convergence of $E_M(t)$ to 0 for all $t \in [-\eta, \eta]$ as $M \to \infty$.

To prove part 3, note that from the recursion equation for the coefficients a_{2k} in (3.20), we have $\frac{a_{2k+2}}{a_{2k}} = \frac{(2k)^2 - (2s)^2}{(2k+1)(2k+2)}$. Since $s \in [0, M]$, it is seen that $|\frac{a_{2k+2}}{a_{2k}}| \leqslant 1$ for $k \geqslant M$, thus, $|a_{2k}|$ is a decreasing sequence of k for $k \geqslant M$. As a result, over the interval $t \in [-\eta, \eta]$, we have that

$$\max_{t \in [-\eta, \eta]} |E_M(t)| = |E_M(\pm \eta)| = \sum_{k=M}^{\infty} |a_{2k}| \eta^{2k} \tag{3.21}$$

$$\leqslant |a_{2M}| \sum_{k=M}^{\infty} \eta^{2k} = |a_{2M}| \frac{\eta^{2M}}{1-\eta^2}. \tag{3.22}$$

Finally to prove the last part, note that we have $E_M'(t) = (-1)^{\lfloor s \rfloor} \sum_{k=M}^{\infty} |a_{2k}|(2k)t^{2k-1}$, thus, $|E_M'(t)| = \sum_{k=M}^{\infty} |a_{2k}|(2k)t^{2k-1}$. This implies that

$$|E_M'(t)| \leqslant \sum_{k=M}^{\infty} |a_{2k}|(2k)\eta^{2k-1} = |E_M'(\eta)|, \tag{3.23}$$

for $t \in [-\eta, \eta]$, thus, the maximum of $|E_M'(t)|$ is achieved at the boundary point $t = \pm\eta$. Moreover, by applying the ratio test [104], we can see that multiplication of the coefficients a_{2k} by $2k$ does not change the radius of the convergence of the series, thus, $|E_M'(\pm\eta)|$ converges to zero as $M \to \infty$, which implies the uniform convergence of $E_M'(t)$ in the interval $[-\eta, \eta]$. Also, to obtain the final expression, note that

$$|E_M'(t)| \leqslant |E_M'(\pm\eta)| = \sum_{k=M}^{\infty} |a_{2k}|(2k)\eta^{2k} \tag{3.24}$$

$$\overset{(a)}{\leqslant} |a_{2M}| \sum_{k=M}^{\infty} (2k)\eta^{2k-1} \tag{3.25}$$

$$= |a_{2M}| \frac{d}{d\eta} \sum_{k=M}^{\infty} \eta^{2k} \tag{3.26}$$

$$= |a_{2M}| \frac{d}{d\eta} \frac{\eta^{2M}}{1-\eta^2} \tag{3.27}$$

$$= 2|a_{2M}| \frac{\eta^{2M-1}(M-(M-2)\eta)}{(1-\eta^2)^2} \tag{3.28}$$

where in (a) we used the fact that $|a_{2k}|$ is a decreasing sequence of k for $k \geqslant M$. This completes the proof. $\qquad\square$

3.4.3 More Refined Error Analysis

In this part, we will focus on the scaling law of the coefficient $|a_{2M}(s)|$ as a function of s. Our goal is to find a scaling law of an exponential form $|a_{2M}(s)| \leqslant g(\frac{s}{M})^{2M}$ for some continuous function $g : [0,1] \to \mathbb{R}_+$. We will combine this with the error bound derived in Proposition 3.1 to prove that the truncation error will vanish for all $s \in [0, M]$ inside the probing window for which $g(\frac{s}{M})\eta < 1$. From (3.20), we have

$$|a_{2M}(s)| = \prod_{n=0}^{M-1} \frac{|(2n)^2 - (2s)^2|}{(2n+1)(2n+2)} \tag{3.29}$$

$$= \frac{(2M)^{2M}}{(2M)!} \prod_{n=0}^{M-1} |(\frac{n}{M})^2 - (\frac{s}{M})^2|. \tag{3.30}$$

Applying the Stirling's approximation for a positive integer l

$$\sqrt{2\pi l}(\frac{l}{e})^l \leqslant l! \leqslant e\sqrt{l}(\frac{l}{e})^l, \tag{3.31}$$

we can upper/lower bound $h(s) := \frac{1}{2M} \log |a_M(s)|$ as follows

$$f(s) - \frac{\log(2e^2 M)}{2M} \leqslant h(s) \leqslant f(s) - \frac{\log(4\pi M)}{2M}, \tag{3.32}$$

where we defined $f(s)$ as

$$f(s) = 1 + \frac{1}{2M} \sum_{n=0}^{M-1} \log |(\frac{n}{M})^2 - (\frac{s}{M})^2|. \tag{3.33}$$

We focus on a scaling regime where s and M grow proportionally such that $\frac{s}{M} \to \alpha \in [0, 1]$, where we can approximate $f(s)$ by the following integral

$$f(\alpha) := 1 + \frac{1}{2} \int_0^1 \log(|t^2 - \alpha^2|)dt, \tag{3.34}$$

where for simplicity we used the same notation f for the function with the normalized argument. With this approximation, we see that for $\frac{s}{M} \to \alpha \in [0, 1]$, we have

$$\lim_{M\to\infty} \frac{1}{2M} \log |a_M(s)| = \lim_{M\to\infty} h(s) = f(\alpha), \tag{3.35}$$

We can also evaluate the logarithmic integral in (3.34) as

$$\int_0^1 \log(|t^2 - \alpha^2|) = \int_0^1 \log(t + \alpha)dt + \int_0^\alpha \log(\alpha - t)dt$$

$$+ \int_\alpha^1 \log(t - \alpha)dt \tag{3.36}$$

$$= \int_\alpha^{1+\alpha} \log(t)dt + \int_0^\alpha \log(t)dt + \int_0^{1-\alpha} \log(t)dt \tag{3.37}$$

$$= \int_0^{1+\alpha} \log(t)dt + \int_0^{1-\alpha} \log(t)dt \tag{3.38}$$

$$= (1 + \alpha) \log(1 + \alpha) + (1 - \alpha) \log(1 - \alpha) - 2. \tag{3.39}$$

Comparison of the finite-M and the asymptotic bound
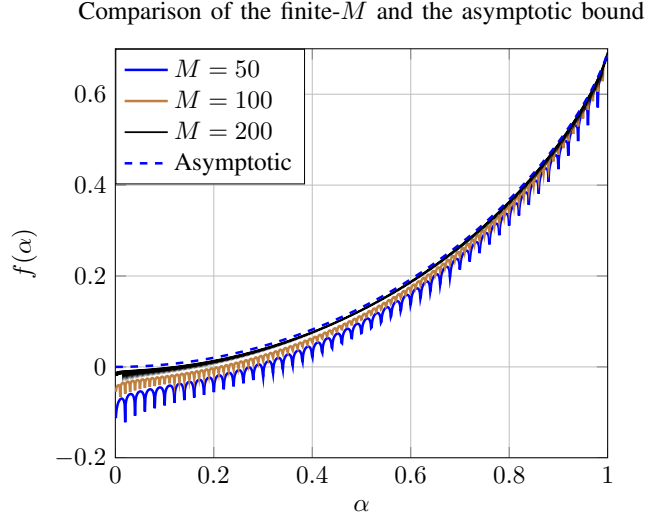
Figure 3.5: Comparison between the asymptotic bound for $M \to \infty$ and the finite-M bound for $M \in \{50, 100, 200\}$.

This implies that

$$f(\alpha) = \frac{1}{2}\Big((1+\alpha)\log(1+\alpha) + (1-\alpha)\log(1-\alpha)\Big) \tag{3.40}$$

$$= (1 - h_2(\frac{1+\alpha}{2}))\log(2), \tag{3.41}$$

where $h_2(x) = -x\log_2(x) - (1-x)\log_2(1-x)$ is the binary entropy function [31] for $x \in [0,1]$. Fig. 3.5 illustrates $f(\alpha)$ for $\alpha \in [0,1]$ and its comparison with $\frac{1}{2M}|a_{2M}(s)|$ for $\frac{s}{M} \in [0,1]$ in the finite-M regime for $M \in \{50, 100, 200\}$. It is seen that the approximation is quite tight even for $M = 100$, where it is also seen that $f(\alpha)$ is an upper bound on $\frac{1}{2M}|a_{2M}(s)|$ for $\frac{s}{M} \in [0,1]$. It is also worthwhile to mention that $|a_{2M}(s)| = 0$ for all the integers $s \in \mathbb{Z}_M$ since for an integer s the even solution of Chebyshev ODE is a polynomial of order $2s$ with zero coefficients for terms with order higher than $2s$, thus, $\log(|a_{2M}(s)|) \to -\infty$ for all $\frac{s}{M} \in \{0, \frac{1}{M}, \dots, 1\}$, but $f(\alpha)$ is well-defined and continuous for all $\alpha \in [0,1]$. The following proposition summarizes some of the properties of $f(\alpha)$.

Proposition 3.2. *Let $f(\alpha)$ for $\alpha \in [0,1]$ be as before. Then, $f(\alpha)$ is a positive, convex, and increasing function of α for $\alpha \in [0,1]$. Moreover, $f(\alpha) \in [0, \log(2)] \approx [0, 0.6931]$.* □

Proof. The proof simply follows from the properties of the binary entropy function h_2. The positivity follows from the fact that $h_2(\frac{1+\alpha}{2}) \in [0,1]$ for $\alpha \in [0,1]$. The increasing property follows from the fact that $x \mapsto h_2(x)$ is a decreasing function of x for $x \in [\frac{1}{2}, 1]$, thus, $h_2(\frac{1+\alpha}{2})$ is a decreasing and $f(\alpha)$ is an increasing function of α for $\alpha \in [0,1]$. The convexity also follows from the fact that $x \mapsto h_2(x)$ is a concave function of x and $\alpha \mapsto \frac{1+\alpha}{2}$

is an affine function of α. The last part also follows immediately from the monotonicity of f and the fact that $h_2(\frac{1}{2}) = 1$ and $h_2(1) = 0$. This completes the proof. $\qquad\square$

3.4.4 Bound on the Width of the Graph

In this section, we will use the results obtained in Section 3.4.2 to derive upper bounds on the width of the graph $\mathcal{G}$ associated to a fixed normalized measure γ as introduced in Section 3.4.1. We will use this to prove Theorem 3.1. Recall that the definition of the graph $\mathcal{G}$ in (3.11) and its width in a fixed probing point $s \in [0, M\varrho]$ in (3.12).

A Simple Minimax Bound

Let $s \in [0, M\varrho]$ be a fixed probing point and define $\phi(\xi, s) = e^{j\pi s\xi}$ for $\xi \in [-1, 1]$. In this section, we consider the following problem. Suppose that $\xi \in [-1, 1]$ is fixed but unknown. The goal is to estimate $\phi(\xi, s)$ at a specific s from the samples $\mathcal{E}_\xi = \{\phi(\xi, m\varrho) : m \in \mathbb{Z}_M\}$ (note the explicit dependence of $\mathcal{E}_\xi$ on ξ). Recall that ϱ is the spatial oversampling factor as introduced before. Let $\widehat{\phi} : \mathcal{E}_\xi \to \mathbb{C}$ be an estimator for $\phi(\xi, s)$ from the available samples $\mathcal{E}_\xi$. We define the worst-case error of $\widehat{\phi}$ by

$$e_{\widehat{\phi}} = \sup_{\xi \in [-1,1]} |\phi(\xi, s) - \widehat{\phi}(\mathcal{E}_\xi)|. \tag{3.42}$$

Here, we will mainly focus on linear estimators, where $\widehat{\phi}$ is a linear function of the observations $\mathcal{E}_\xi$, where this linear function (linear estimator) can depend on the probing point s. We denote the set of all such linear estimators for a given s by $\mathcal{L}_s$. We define the minimax error over the class of linear estimators in $\mathcal{L}_s$ by

$$e_M(s) = \inf_{\widehat{\phi} \in \mathcal{L}_s} \sup_{\xi \in [-1,1]} |\phi(\xi, s) - \widehat{\phi}(\mathcal{E}_\xi)|. \tag{3.43}$$

We prove the following result.

Proposition 3.3. *Let $s \in [0, M\varrho]$ be a fixed probing point. Let γ be a fixed positive normalized measure and let $\mathcal{G}$ and $w_{\mathcal{G}}(s)$ be the graph corresponding to γ and its width at s. Let also $e_M(s)$ be as in (3.43). Then, we have:*

1. *$e_M(s) \leqslant 1$ for all s.*

2. *$w_{\mathcal{G}}(s) \leqslant 2e_M(s)$.* $\qquad\square$

Proof. The first part simply follows by setting $\widehat{\phi} = 0$ to be the zero estimator (which is linear and belongs to $\mathcal{L}_s$) and the fact that $|\phi(\xi, s)| = 1$.

To prove the second part, note that by definition of $e_M(s)$, for any $\epsilon > 0$, there is a linear estimator $\widehat{\phi}$ such that

$$|\phi(\xi, s) - \widehat{\phi}(\mathcal{E}_\xi)| \leqslant e_M(s) + \epsilon, \tag{3.44}$$

over the whole set $\xi \in [-1, 1]$. Consider a positive normalized measure $\mu \in \Gamma_\gamma$ (see the definition of Γ_γ in (3.10)). Note that since the estimator $\widehat{\phi}$ is a linear function of $\mathcal{E}_\xi$, we have that

$$
\begin{aligned}
\int \mu(\xi) \widehat{\phi}(\mathcal{E}_\xi) d\xi &= \widehat{\phi}\Big(\int \mu(\xi) \mathcal{E}_\xi d\xi \Big) \\
&= \widehat{\phi}\Big(\breve{\mu}(0), \dots, \breve{\mu}\big((M-1)\varrho\big) \Big) \\
&\overset{(a)}{=} \widehat{\phi}\Big(\breve{\gamma}(0), \dots, \breve{\gamma}\big((M-1)\varrho\big) \Big) =: c_\gamma(s),
\end{aligned} \tag{3.45}
$$

where (a) follows from the fact that $\mu \in \Gamma_\gamma$, thus, it has the same Fourier transform as $\breve{\gamma}$ at sampling points $\{m\varrho : m \in \mathbb{Z}_M\}$. Also, note that $c_\gamma(s)$ in the last expression depends only on γ, s but not on a specific $\mu \in \Gamma_\gamma$. From (3.44) and (3.45), we obtain that

$$
\begin{aligned}
|\breve{\mu}(s) - c_\gamma(s)| &= \Big| \int \mu(\xi) \phi(\xi, s) d\xi - \int \mu(\xi) \widehat{\phi}(\mathcal{E}_\xi) d\xi \Big| \\
&= \Big| \int \mu(\xi) \big(\phi(\xi, s) d\xi - \widehat{\phi}(\mathcal{E}_\xi) \big) \Big| \\
&\leqslant \int \mu(\xi) \big| \phi(\xi, s) - \widehat{\phi}(\mathcal{E}_\xi) \big| d\xi \tag{3.46} \\
&\leqslant e_M(s) + \epsilon. \tag{3.47}
\end{aligned}
$$

Since this is true for any $\epsilon > 0$, we have that $|\breve{\mu}(s) - c_\gamma(s)| \leqslant e_M(s)$. Note that this results is valid for any $\mu \in \Gamma_\gamma$, thus, by applying the triangle inequality, we have that $|\breve{\mu}(s) - \breve{\varepsilon}(s)| \leqslant 2 e_M(s)$ for any arbitrary measure $\mu, \varepsilon \in \Gamma_\gamma$. From the definition of the width of $\mathcal{G}$ in (3.12), this implies that $w_{\mathcal{G}}(s) \leqslant 2 e_M(s)$. This completes the proof. $\qquad \square$

We will use Proposition 3.3 in the following to find an upper bound on $w_{\mathcal{G}}(s)$ at any probing point s by finding suitable upper bound for $e_M(s)$.

Minimax Error of the Real Part

Let us first derive an upper bound on the minimax error of estimating the real part of $\phi(\xi, s)$ from the samples $\mathcal{E}_\xi = \{\phi(\xi, k\varrho) : k \in \mathbb{Z}_M\}$. This boils down in the linear minimax

estimation of $\cos(\pi s\xi)$ from the real and the imaginary parts of $\mathcal{E}_\xi$ given by

$$\mathbb{R}(\mathcal{E}_\xi) := \Big\{1, \cos(\pi\varrho\xi), \ldots, \cos\big(\pi\varrho(M-1)\xi\big)\Big\}, \tag{3.48}$$

$$\mathfrak{I}(\mathcal{E}_\xi) := \Big\{\sin(\pi\varrho\xi), \ldots, \sin\big(\pi\varrho(M-1)\xi\big)\Big\}. \tag{3.49}$$

Since we are looking for an upper bound on the error, it is sufficient to consider only a subset of linear estimators that use only the real part $\mathbb{R}(\mathcal{E}_\xi)$. Denoting by $\mathbf{c} = (c_0, \ldots, c_{M-1})^\mathsf{T} \in \mathbb{R}^M$ the coefficients of such a linear estimator, we can upper bound the minimax estimation error of the real part by

$$e_M^{(r)}(s) = \inf_{\mathbf{c}\in\mathbb{R}^M} \max_{\xi\in[-1,1]} |\cos(\pi s\xi) - \sum_{k=0}^{M-1} c_m \cos(m\pi\varrho\xi)|$$

$$= \inf_{\mathbf{c}\in\mathbb{R}^M} \max_{\xi^\circ\in[-\varrho,\varrho]} |\cos(\frac{\pi s}{\varrho}\xi^\circ) - \sum_{m=0}^{M-1} c_m \cos(m\pi\xi^\circ)|$$

$$= \inf_{\mathbf{c}} \max_{t\in[-\eta,\eta]} \Big|\cos(\frac{2s}{\varrho}\sin^{-1}(t)) - \sum_{m=0}^{M-1} c_m \cos(2m\sin^{-1}(t))\Big|,$$

where the superscript "r" refers to the real part, and where we made the change of variable $\xi^\circ = \varrho\xi$ and $t = \sin(\frac{\pi\xi^\circ}{2})$, and defined $\eta = \sin(\frac{\pi\varrho}{2})$, where $\eta \in (0,1)$ since $\varrho \in (0,1)$. Note that we have

$$\sum_{m=0}^{M-1} c_m \cos(2m\sin^{-1}(t)) \overset{(a)}{=} \sum_{m=0}^{M-1} c_m \cos(m\pi + 2m\cos^{-1}(t))$$

$$= \sum_{m=0}^{M-1} c_m(-1)^m \cos(2m\cos^{-1}(t))$$

$$\overset{(b)}{=} \sum_{m=0}^{M-1} c_m(-1)^m T_{2m}(t) \tag{3.50}$$

where in (a) we used the identity $\sin^{-1}(t) + \cos^{-1}(t) = \frac{\pi}{2}$ for $t \in (-1,1)$, and where in (b) we used the fact that for an integer m, $\cos(2m\cos^{-1}(t))$ coincides with the Chebyshev polynomial $T_{2m}(t)$ of order $2m$. Note that $\{T_{2m}(t) : m \in \mathbb{Z}_M\}$ forms a basis for the M-dim linear space of all even polynomials in $t \in (-1,1)$ of order at most $2(M-1)$. However, this space is also spanned by the monomials $\{t^{2m} : m \in \mathbb{Z}_M\}$. As a result, using (3.50), we can write the desired minimax error bound $e_M^{(r)}(s)$ more directly as

$$e_M^{(r)}(s) = \inf_{\mathbf{b}\in\mathbb{R}^M} \max_{t\in[-\eta,\eta]} \Big|\cos(\frac{2s}{\varrho}\sin^{-1}(t)) - \sum_{m=0}^{M-1} b_{2m}t^{2m}\Big|,$$

where we defined $\mathbf{b} = (b_0, b_2, \ldots, b_{2(M-1)})^{\mathsf{T}} \in \mathbb{R}^M$. Replacing the coefficients $\mathbf{b}$ with the coefficients of Taylor's expansion of $\cos(\frac{2s}{\varrho} \sin^{-1}(t))$, i.e., $\{a_{2m}(\frac{s}{\varrho}) : m \in \mathbb{Z}_M\}$ with the notation introduced in Section 3.4.2 (note that $\frac{s}{\varrho} \in [0, M]$), and using Proposition 3.1, we obtain the following upper bound

$$e_M^{(r)}(s) \leqslant |a_{2M}(\frac{s}{\varrho})| \frac{\eta^{2M}}{1 - \eta^2} \approx \left(\sin(\frac{\pi \varrho}{2}) g(\frac{s}{M\varrho}) \right)^{2M} \tag{3.51}$$

where in the last expression we replaced $\eta = \sin(\frac{\pi \varrho}{2})$ and used (3.35) and replaced $g = e^f$.

Minimax Error of the Imaginary Part

We repeat similar steps to derive an upper bound on the minimax estimation of the imaginary part of $\phi(\xi, s)$, where this time we estimate the imaginary part of $\phi(\xi, s)$ from $\Im(\mathcal{E}_\xi)$. More specifically, we consider the following minimax error

$$e_M^{(i)}(s) = \inf_{\mathbf{s}} \max_{\xi \in [-1,1]} \left| \sin(\pi s \xi) - \sum_{m=0}^{M-1} s_m \sin(m \pi \varrho \xi) \right|$$

$$= \inf_{\mathbf{s}} \max_{\xi^\circ \in [-\varrho, \varrho]} \left| \sin(\frac{\pi s}{\varrho} \xi^\circ) - \sum_{m=0}^{M-1} s_m \sin(m \pi \xi^\circ) \right|$$

$$= \inf_{\mathbf{s}} \max_{t \in [-\eta, \eta]} \left| \sin(\frac{2s}{\varrho} \sin^{-1}(t)) - \sum_{m=0}^{M-1} s_m \sin(2m \sin^{-1}(t)) \right|,$$

where the superscript (i) refers to the imaginary part, where we defined $\mathbf{s} = (s_1, \ldots, s_{M-1})^{\mathsf{T}} \in \mathbb{R}^{M-1}$, and where we made the change of variable $\xi^\circ = \varrho \xi$ and $t = \sin(\frac{\pi \xi^\circ}{2})$ as before. We obtain an upper bound on $e_M^{(i)}(s)$ via the estimation of the real part done before. We first define

$$e_M^{(r)}(s, t) := \cos(\frac{2s}{\varrho} \sin^{-1}(t)) - \sum_{m=0}^{M-1} a_{2m} \cos(2m \sin^{-1}(t)),$$

as the truncation error in the estimation of the real part, where $t \in [-\eta, \eta]$ as before, and where have used the same coefficients $a_{2m} = a_{2m}(\frac{s}{\varrho})$ as in the real case. Taking the derivative of $e_M^{(r)}(s, t)$ with respect to t and some simplification yields

$$\max_{t \in [-\eta, \eta]} \left| \frac{\varrho \sqrt{1 - t^2}}{2s} \frac{d}{dt} e_M^{(r)}(s, t) \right|$$

$$:= \left\| \sin(\frac{2s}{\varrho} \sin^{-1}(t)) - \sum_{m=0}^{M-1} \frac{m}{s} a_{2m} \sin(2m \sin^{-1}(t)) \right\|_\infty \tag{3.52}$$

$$\overset{(a)}{\geqslant} |e_M^{(i)}(s)| \tag{3.53}$$

where in (a) we used the fact that (3.52) can be interpreted as the estimation error of an estimator with coefficients $s_m = \frac{m}{s} a_{2m}$, which can be lower bounded by $|e_M^{(i)}(s)|$ by definition. Applying Proposition 3.1, we can bound the derivative of the truncation error $e_M^{(r)}(s,t)$ in $t \in [-\eta, \eta]$ by

$$\max_{t \in [-\eta, \eta]} |\frac{d}{dt} e_M^{(r)}(s,t)| \leqslant 2|a_{2M}(\frac{s}{\varrho})| \frac{\eta^{2M-1}(M-(M-2)\eta)}{(1-\eta^2)^2}.$$

From (3.53), this yields

$$e_M^{(i)}(s) \leqslant \max_{t \in [-\eta, \eta]} |\frac{\varrho\sqrt{1-t^2}}{2s}| \max_{t \in [-\eta, \eta]} |\frac{d}{dt} e_M^{(r)}(s,t)| \tag{3.54}$$

$$\leqslant |a_{2M}(\frac{s}{\varrho})| \frac{\eta^{2M-1}(M-(M-2)\eta)}{s(1-\eta^2)^2} \tag{3.55}$$

$$\overset{M \to \infty}{\approx} \left(\sin(\frac{\pi\varrho}{2}) g(\frac{s}{M\varrho}) \right)^{2M} \tag{3.56}$$

where we used $\eta = \sin(\frac{\pi\varrho}{2})$ as before.

3.4.5 Proof of Theorem 3.1

Combining the upper bound on the minimax error of the real part $e_M^{(r)}(s)$ and that of the imaginary part $e_M^{(i)}(s)$, we obtain an upper bound on the minimax error (3.43) as follows

$$e_M(s) \leqslant e_M^{(r)}(s) + e_M^{(i)}(s) \leqslant C' \left(\sin(\frac{\pi\varrho}{2}) g(\frac{s}{M\varrho}) \right)^{2M} \tag{3.57}$$

where C' is a universal constant independent of M and s. From Proposition 3.3, this yields the following upper bound on the width of the graph $\mathcal{G}$ associated to a given measure γ

$$w_{\mathcal{G}}(s) \leqslant \min\{2e_M(s), 2\} \leqslant \min\left\{ C\left(\sin(\frac{\pi\varrho}{2}) g(\frac{s}{M\varrho}) \right)^{2M}, 2 \right\},$$

for some universal constant C independent of M, $s \in [0, M\varrho]$, and the measure γ. This proves Theorem 3.1.

3.5 Algorithms for UL-DL Covariance Transformation

Now we turn to the second question posed at the beginning of this chapter, regarding algorithms for UL-DL covariance transformation. Perhaps the most general way to perform

the transformation (resampling) task is to first find a *feasible* ASF, that is a non-negative measure $\widehat{\gamma} : [-1, 1] \rightarrow \mathbb{R}_+$, that is consistent with the observed UL covariance entries in the sense that

$$\int_{-1}^{1} \widehat{\gamma}(\xi) e^{j\pi\varrho m \xi} d\xi = [\boldsymbol{\sigma}_{\mathsf{ul}}]_m, \ m \in \mathbb{Z}_M, \tag{3.58}$$

where $[\boldsymbol{\sigma}_{\mathsf{ul}}]_m = \int_{-1}^{1} \gamma(\xi) e^{j\pi\varrho m \xi} d\xi$ is generated by the ground-truth ASF γ. Recall that the feasibility set, consisting of all such measures was introduced in (3.10) as the closed convex set Γ_γ. The set Γ_γ can be seen as the intersection of two convex sets:

$$\Gamma_\gamma = \mathcal{D}_+ \bigcap \mathcal{D}_\gamma, \tag{3.59}$$

where the constraint sets, encoding non-negativity and data consistency, are respectively given as

$$\mathcal{D}_+ = \{\mu \, : \, \mu(\xi) \geqslant 0 \, \text{for all} \, \xi \in [-1, 1]\}, \tag{3.60a}$$

$$\mathcal{D}_\gamma = \{\mu \, : \, \int_{-1}^{1} \mu(\xi) e^{j\pi\varrho m \xi} d\xi = [\boldsymbol{\sigma}_{\mathsf{ul}}]_m = \, \text{for all} \, m \in \mathbb{Z}_M\}. \tag{3.60b}$$

Note that $\gamma \in \mathcal{D}_+ \bigcap \mathcal{D}_\gamma$ and therefore the intersection Γ_γ is always non-empty. Given an estimate of the ASF as $\widehat{\gamma}$, one can estimate the DL covariance simply by computing

$$[\widehat{\boldsymbol{\sigma}}_{\mathsf{dl}}]_m = \int_{-1}^{1} \widehat{\gamma}(\xi) e^{j\pi\varrho \frac{m}{\nu} \xi} d\xi, \ m \in \mathbb{Z}_M. \tag{3.61}$$

In the absence of further structure, the first step above can be formulated as

$$\begin{aligned} &\text{find} &&\mu \\ &\text{subject to} &&\mu \in \Gamma_\gamma = \mathcal{D}_+ \bigcap \mathcal{D}_\gamma. \end{aligned} \tag{$\mathcal{P}_{\mathsf{feas}}$}$$

Such problems are generally known as *convex feasibility problems* [6]. Obviously, the solution to such a problem may not be unique in general, hence it can be seen as an ill-posed inverse problem. In the absence of further structure, such as sparsity of the ASF in a known basis, it seems that any measure in Γ_γ is a valid solution.

Remark 3.4. The theory developed earlier in this chapter on the minimax interpolation error bounds can be seen as a "positive" result: for any measure belonging to Γ_γ, the transformation error for most covariance entries is fairly small. Therefore, all algorithms that return solutions within Γ_γ must have similar estimation errors, especially for those covariance entries that correspond to lattice points in the S-region of the interpolation window (see Fig. 3.3). Therefore, the transformation algorithms that will be discussed below are only decent examples of methods that do not rely on specific assumptions on

the channel and treat the problem in its generality. One can in principle use any method that solves $\mathcal{P}_{\mathsf{feas}}$, and yet be assured of a robust covariance transformation as a result of Theorem 3.1.

3.5.1 Transformation via Alternating Projection

A popular approach to solving convex feasibility problems, is to use an alternating projection method. Starting from an initial point $\gamma^{(0)}$, this method constructs a sequence of functions $\{\gamma^{(t)}, t = 0, 1, \ldots\}$ by projecting the solution of the previous step first to one of the involved convex constraint sets $(\mathcal{D}_+)$ and then to the other one $(\mathcal{D}_\gamma)$ [6]. Excellent results with such an approach were reported in [87] for arbitrary array geometries. In particular, the authors of this work employ an *extrapolated alternating projection method* (EAPM) [7], briefly explained as follows. Assume γ belongs to the Hilbert space of functions with bounded ℓ_2-norm over $[-1, 1]$, equipped with the standard Euclidean inner product. Starting from $\gamma^{(0)}$, this projection method is updated as

$$\gamma^{(t+1)} = (1 - \tau K_t)\gamma^{(t)} + \tau K_t \mathcal{P}_{\mathcal{D}_\gamma}\left(\mathcal{P}_{\mathcal{D}_+}\left(\gamma^{(t)}\right)\right), \ t = 0, 1, \ldots, \tag{3.62}$$

where K_t is known as an *extrapolation parameter* that can be computed as a function of the current estimate $\gamma^{(t)}$. The operator $\mathcal{P}_{\mathcal{D}_+}$ maps a measure μ to its orthogonal projection on $\mathcal{D}_+$ and is given as $\mathcal{P}_{\mathcal{D}_+}(\mu) = \max\{\mu, 0\}$. Also, $\mathcal{P}_{\mathcal{D}_\gamma}$ maps μ to its orthogonal projection onto the (infinite-dimensional) hyperplane $\mathcal{D}_\gamma$ through a closed-form expression (see [87]). One can show that, under mild conditions, the sequence $\{\gamma^{(t)}\}_{t \in \mathbb{Z}_+}$ generated via (3.62) converges weakly to a point in $\Gamma_\gamma = \mathcal{D}_+ \bigcap \mathcal{D}_\gamma$ [7]. Denoting the convergence point by $\gamma^{(\infty)}$, one can estimate the entries of the DL channel covariance as

$$[\hat{\boldsymbol{\sigma}}_{\mathsf{dl}}]_m = \int_{-1}^{1} \gamma^{(\infty)}(\xi) e^{j\pi\varrho\frac{m}{\nu}\xi} d\xi, \ m \in \mathbb{Z}_M. \tag{3.63}$$

The projection method is one of the methods that we will repeatedly refer to in the simulation results section. The strong aspect of this method is that, in principle, it returns an infinite-dimensional estimate of the ASF and does not rely on a parametric representation.[3] On the flip side, to the best of our knowledge, the projection method returns a function in $L_2([-1, 1])$, i.e. the space of square-integrable functions over $[-1, 1]$. However, we know that an important class of power spectra include discrete measures such as Dirac impulses, representing LoS and specular scattering components. Obviously, these types of spectra do not belong to $L_2([-1, 1])$ and it is not clear from [87] (and the

---

[3]Despite this fact, any practical implementation of the projection method by a software relies on a discretization of the angle domain, so that one can compute the update per iteration in closed form.

companion papers [22] and [86]) how one can extend these results to cover such cases. This is however not to claim that the projection method fails in the task of UL-DL covariance transformation, since in this case we are interested not in the ASF per se, but in its inner product with a set of complex exponentials associated with the DL spatial sampling lattice points (see (3.63)).

A simpler version of the alternating projection method can be considered, by removing the non-negativity constraint from $\mathcal{P}_{\mathsf{feas}}$. This significantly reduces the complexity of the algorithm, since the solution can be obtained in closed form [87]. The drawback is a higher transformation error, as we will see in the simulation results of this chapter.

3.5.2 Transformation via MUSIC-Assisted Non-Negative Least-Squares

Apart from infinite-dimensional solutions to the ASF estimation problem, one can instead consider finite- but high-dimensional approximations of the ASF, which relaxes the problem to that of estimating a vector of coefficients at the cost of a (small) approximation error. We discussed this method in Chapter 2, where the idea is to approximate the ASF as a superposition of the elements of a family of densities

$$\Psi = \Psi_d \bigcup \Psi_c, \ |\Psi| = G_c + \widehat{K} = G, \tag{3.64}$$

where $\Psi_c := \{\psi_i, \ i \in \mathbb{Z}_{G_c}\}$ is a family of $G_c \gg 1$ compact-support, non-negative functions, whose support partitions the angular domain, i.e.

$$\bigcup_i \mathrm{supp}(\psi_i) = [-1, 1], \ \mathrm{supp}(\psi_i) \bigcap \mathrm{supp}(\psi_j) = \varnothing, \ \forall i \neq j. \tag{3.65}$$

Standard examples include the family of rectangular densities

$$\Psi_{c,\mathsf{rect}} = \{ \frac{G_c}{2}\mathsf{rect}_{[-1+\frac{2i}{G_c}, -1+\frac{2(i+1)}{G_c})} : \ i \in \mathbb{Z}_G\}.$$

In addition, $\Psi_d := \{\delta(\cdot - \widehat{\xi}_k), \ k = 0 \ldots, \widehat{K}\}$ is a family of Dirac impulses located at $\{\widehat{\xi}_k\}_{k=0}^{\widehat{K}-1}$. In Chapter 2, we explained how one can estimate line spectral components from the covariance matrix (or its empirical estimate) and provided evidence, showing that under mild conditions on the amplitude of the line spectral components, one can obtain a consistent estimate of the set $\{\widehat{\xi}_k\}_{k=0}^{\widehat{K}-1}$ when the covariance dimension grows large ($M \to \infty$; see Chapter 2, Section 2.4.1). In practice, we observe that these estimates are indeed close to their true values $\{\xi_k\}_{k=0}^{K-1}$, even for moderately large dimensions. Therefore, we employ

the *hybrid* density family $\Psi = \Psi_d \bigcup \Psi_c$ and approximate the ASF as

$$\gamma \approx \sum_{i=0}^{G_c-1} [\boldsymbol{\gamma}]_i \psi_i + \sum_{i=G_c}^{G_c+\widehat{K}-1} [\boldsymbol{\gamma}]_i \delta(\xi - \widehat{\xi}_{i-G_c}) =: \sum_{i=0}^{G-1} [\boldsymbol{\gamma}]_i \phi_i, \quad (3.66)$$

where $\boldsymbol{\gamma}$ is the $(G_c + \widehat{K})$-dimensional vector of approximation coefficients, and where we have defined $\phi_i = \psi_i$, for $i = 0, \ldots, G_c - 1$ and $\phi_i = \delta(\cdot - \widehat{\xi}_{i-G_c})$, for $i = G_c, \ldots, G - 1$. With this approximation we have

$$[\boldsymbol{\sigma}_{\mathsf{ul}}]_m = \int_{-1}^{1} \gamma(\xi) e^{j\pi\varrho m\xi} c\xi \approx \sum_{i \in \mathbb{Z}_G} [\boldsymbol{\gamma}]_i \int_{-1}^{1} \phi_i(\xi) e^{j\pi\varrho m\xi} d\xi = [\mathbf{A}]_{m,\cdot} \boldsymbol{\gamma}, \, m \in \mathbb{Z}_M \quad (3.67)$$

where $[\mathbf{A}]_{m,\cdot}$ is the m-th row of a $M \times G$ "sensing matrix" with entries

$$[\mathbf{A}]_{m,n} = \begin{cases} \int_{-1}^{1} \psi_n(\xi) e^{j\pi\varrho m\xi} d\xi, & n = 0, \ldots, G_c - 1, \\ e^{j\pi\varrho m\widehat{\xi}_{n-G_c}}, & n = G_c, \ldots, G - 1, \end{cases}, \, m \in \mathbb{Z}_M. \quad (3.68)$$

The goal then is to estimate $\boldsymbol{\gamma}$ such that it is non-negative and consistent with the given UL covariance entries. However, since we are considering a finite-dimensional approximation, we have to *loosen* the data consistency constraint. This suggests a natural way of estimating the vector $\boldsymbol{\gamma}$, such that the approximation error is minimized. Combined with the non-negativity constraint on $\boldsymbol{\gamma}$, this results in the following non-negative least-squares (NNLS) convex program:

$$\begin{aligned} \text{minimize}_{\boldsymbol{\gamma}} \quad & C(\boldsymbol{\mu}) := \|\mathbf{M}(\mathbf{A}\boldsymbol{\mu} - \boldsymbol{\sigma}_{\mathsf{ul}})\|^2 \\ \text{subject to} \quad & \boldsymbol{\mu} \in \mathcal{D}_+^d, \end{aligned} \qquad (\mathcal{P}_{\mathsf{nnls}}^d)$$

where $\mathbf{M} = \text{diag}\left(\left[\sqrt{M}, \sqrt{2(M-1)}, \ldots, \sqrt{2}\right]^{\mathsf{T}}\right)$, and $\mathcal{D}_+^d = \mathbb{R}_+^G$ is the G-dimensional non-negative orthant and a discretized version of $\mathcal{D}_+$. Recall that the matrix $\mathbf{M}$ is used so that the cost in $\mathcal{P}_{\mathsf{nnls}}^d$ is equal to $\|\mathcal{T}(\mathbf{A}\boldsymbol{\mu}) - \boldsymbol{\Sigma}_{\mathsf{ul}}\|_{\mathsf{F}}^2$, where $\mathcal{T}$ is the Toeplitzification operator. Given the solution of $\mathcal{P}_{\mathsf{nnls}}^d$ as $\boldsymbol{\gamma}_{\mathsf{nnls}}^\star$, we estimate the DL covariance as

$$\widehat{\boldsymbol{\Sigma}}_{\mathsf{dl}} = \mathcal{T}(\widehat{\boldsymbol{\sigma}}_{\mathsf{dl}}), \ \widehat{\boldsymbol{\sigma}}_{\mathsf{dl}} = \mathbf{A}\boldsymbol{\gamma}_{\mathsf{nnls}}^\star. \quad (3.69)$$

The following proposition shows an equivalence between the two problems $\mathcal{P}_{\mathsf{feas}}$ and $\mathcal{P}_{\mathsf{nnls}}^d$ in a special case.

Proposition 3.4. *Assume that γ belongs to the span of a set of densities Ψ whose elements have non-overlapping support, unless perhaps on a finite number of points. Then $\mathcal{P}_{\mathsf{feas}}$ and $\mathcal{P}^d_{\mathsf{nnls}}$ have identical solution sets.*

Proof. Since $\gamma \in \mathrm{span}(\Psi)$, there exists a vector of coefficients $\boldsymbol{\gamma} \in \mathbb{R}^G$ such that $\gamma = \sum_{i \in \mathbb{Z}_G}[\boldsymbol{\gamma}]_i\phi_i$ (with strict equality). Then the cost of $\mathcal{P}^d_{\mathsf{nnls}}$ achieves an optimal (minimum) value of zero, since $\boldsymbol{\sigma}_{\mathsf{ul}} = \mathbf{A}\boldsymbol{\gamma}$ and

$$C(\boldsymbol{\gamma}) = \|\mathbf{M}\left(\mathbf{A}\boldsymbol{\gamma} - \boldsymbol{\sigma}_{\mathsf{ul}}\right)\|^2 = 0.$$

Therefore the set of minimizers of $\mathcal{P}^d_{\mathsf{nnls}}$ is given as $\mathcal{S}_{\mathsf{nnls}} = \{\boldsymbol{\mu} : C(\boldsymbol{\mu}) = 0, \ \boldsymbol{\mu} \in \mathcal{D}^d_+\}$. Since $\mathbf{M}$ is diagonal with non-zero diagonal elements, we conclude that in this case

$$\mathcal{S}_{\mathsf{nnls}} = \{\boldsymbol{\mu} : \mathbf{A}\boldsymbol{\mu} = \boldsymbol{\sigma}_{\mathsf{ul}}, \ \boldsymbol{\mu} \in \mathcal{D}^d_+\}.$$

Besides, we can write $\mathcal{P}_{\mathsf{feas}}$ as the problem of finding a vector $\boldsymbol{\mu}$ such that

$$\sum_{i=0}^{G-1}[\boldsymbol{\mu}]_i\phi_i(\xi) \geqslant 0, \ \xi \in [-1, 1], \tag{3.70a}$$

$$\int_{-1}^{1}\sum_{i=0}^{G-1}[\boldsymbol{\mu}]_i\phi_i(\xi)e^{jm\pi\varrho\xi}d\xi = [\boldsymbol{\sigma}_{\mathsf{ul}}]_m. \tag{3.70b}$$

Note that the non-negativity constraint in 3.70a has a subtle meaning: it is equivalent to the Dirac impulses ϕ_i, $i = G_c, \ldots, G$ to have non-negative coefficients ($[\boldsymbol{\mu}]_i \geqslant 0$ for $i = G_c, \ldots, G$) *and* $\sum_{i=0}^{G_c-1}[\boldsymbol{\mu}]_i\phi_i \geqslant 0$. Since the densities ϕ_i, $i = 0, \ldots, G_c - 1$ have disjoint support, the combination of these conditions is equivalent to $[\boldsymbol{\mu}]_i \geqslant 0$ for all $i = 0, \ldots, G-1$, i.e. $\gamma \in \mathcal{D}^d_+$. In addition, the constraint (3.70b) is equivalent to $\mathbf{A}\boldsymbol{\mu} = \boldsymbol{\sigma}_{\mathsf{ul}}$. This gives the solution set of $\mathcal{P}_{\mathsf{feas}}$ as $\mathcal{S}_{\mathsf{feas}} = \{\boldsymbol{\mu} : \mathbf{A}\boldsymbol{\mu} = \boldsymbol{\sigma}_{\mathsf{ul}}, \ \boldsymbol{\mu} \in \mathcal{D}^d_+\}$ and therefore we have $\mathcal{S}_{\mathsf{feas}} = \mathcal{S}_{\mathsf{nnls}}$. $\qquad\square$

3.6 Extension to Arbitrary Array Geometries

So far we discussed the UL-DL covariance transformation problem specifically for ULAs. Here, we briefly discuss extensions to arbitrary (full-dimensional) array geometries. For arbitrary array geometries, one can generally parameterize the AoAs by points over the unit sphere $\mathbb{S}^2 := \{\boldsymbol{\xi} \in \mathbb{R}^3 : \|\boldsymbol{\xi}\| = 1\}$ in $\mathbb{R}^3$ and represent the angular ASF of each user as a non-negative measure γ over $\mathbb{S}^2$. Denoting by $\mathcal{R} := \{\mathbf{r}_m \in \mathbb{R}^3 : m \in \mathbb{Z}_M\}$ the position of antenna elements in the BS array, we can define UL-DL array responses as $\mathbf{a}_{\mathsf{ul}}(\boldsymbol{\xi})$ and

$\mathbf{a}_{\mathsf{dl}}(\boldsymbol{\xi})$, where

$$[\mathbf{a}_{\mathsf{ul}}(\boldsymbol{\xi})]_k = e^{jm\frac{2\pi}{\lambda_{\mathsf{ul}}}\langle\boldsymbol{\xi},\mathbf{r}_m\rangle}, [\mathbf{a}_{\mathsf{dl}}(\boldsymbol{\xi})]_m = e^{jm\frac{2\pi}{\lambda_{\mathsf{dl}}}\langle\boldsymbol{\xi},\mathbf{r}_m\rangle}, \tag{3.71}$$

for $m \in \mathbb{Z}_M$. Similarly, we can introduce the Fourier transform of the angular ASF γ as

$$\breve{\gamma}(\mathbf{r}) = \int_{\mathbb{S}} e^{j\pi\langle\boldsymbol{\xi},\mathbf{r}\rangle}\gamma(\boldsymbol{\xi})d\boldsymbol{\xi}, \tag{3.72}$$

where the integral is taken over the set of all AoAs parameterized with $\boldsymbol{\xi} \in \mathbb{S}$. The channel covariance matrices in UL and DL are given, respectively, as

$$[\boldsymbol{\Sigma}_{\mathsf{ul}}]_{m,l} = \int_{\mathbb{S}} e^{j\pi\langle\boldsymbol{\xi},\frac{\mathbf{r}_m-\mathbf{r}_l}{\frac{\lambda_{\mathsf{ul}}}{2}}\rangle}\gamma(\boldsymbol{\xi})d\boldsymbol{\xi} = \breve{\gamma}(\frac{\mathbf{r}_m - \mathbf{r}_l}{\frac{\lambda_{\mathsf{ul}}}{2}}), \tag{3.73}$$

$$[\boldsymbol{\Sigma}_{\mathsf{dl}}]_{k,l} = \int_{\mathbb{S}} e^{j\pi\langle\boldsymbol{\xi},\frac{\mathbf{r}_m-\mathbf{r}_l}{\frac{\lambda_{\mathsf{dl}}}{2}}\rangle}\gamma(\boldsymbol{\xi})d\boldsymbol{\xi} = \breve{\gamma}(\frac{\mathbf{r}_m - \mathbf{r}_l}{\frac{\lambda_{\mathsf{dl}}}{2}}). \tag{3.74}$$

Denoting by $\mathcal{D} := \mathcal{R} - \mathcal{R} = \{\mathbf{r}_m - \mathbf{r}_l : m, l \in \mathbb{Z}_M\}$ the Minkowski difference [120] of the antenna geometry $\mathcal{R} := \{\mathbf{r}_m : m \in \mathbb{Z}_M\}$, we can see that the UL-DL covariance transformation in this setup can be posed, similar to that stated in Problem 1, as the problem of resampling the Fourier transform $\breve{\gamma}$ of γ at the DL sampling set $\mathcal{D}_{\mathsf{dl}} := \frac{\mathcal{D}}{\frac{\lambda_{\mathsf{dl}}}{2}}$ from its value at UL sampling set $\mathcal{D}_{\mathsf{ul}} := \frac{\mathcal{D}}{\frac{\lambda_{\mathsf{ul}}}{2}}$, where for an $\alpha \in \mathbb{R}$ we denote by $\alpha\mathcal{D} := \{\alpha\boldsymbol{\phi} : \boldsymbol{\phi} \in \mathcal{D}\}$ the component-wise scaling of the elements of $\mathcal{D}$ by the factor α. Interestingly, it is seen that, as in the one-dimensional case of the ULA, the set of DL sampling positions is simply given by $\mathcal{D}_{\mathsf{dl}} = \frac{1}{\nu} \times \mathcal{D}_{\mathsf{ul}}$, which is a scaled version of the set of UL sampling positions $\mathcal{D}_{\mathsf{ul}}$ by a factor $\frac{1}{\nu}$ (larger than 1 for $\nu < 1$), where $\nu = \frac{\lambda_{\mathsf{dl}}}{\lambda_{\mathsf{ul}}} = \frac{f_{\mathsf{ul}}}{f_{\mathsf{dl}}}$, denotes the ratio between the UL and DL carrier frequencies.

This is illustrated for a circular array geometry in Fig. 3.6, where it is seen that $\mathcal{D}_{\mathsf{ul}}$ and $\mathcal{D}_{\mathsf{dl}}$ each consist of $O(M^2)$ points. For the ULA, the underlying geometry is one-dimensional and consists of $\mathcal{R} = \{md\boldsymbol{\xi}_0 : m \in \mathbb{Z}_M\}$ for some unit vector $\boldsymbol{\xi}_0 \in \mathbb{S}^2$ and some antenna spacing d. Thus, the difference set $\mathcal{D} = \mathcal{R} - \mathcal{R} = \{md\boldsymbol{\xi}_0 : m = -(M-1), \ldots, (M-1)\}$ is also one-dimensional, lies along $\boldsymbol{\xi}_0$, and consists of $2M - 1$ points (rather than $O(M^2)$ points as in a circular array). Also, for the ULA, the resulting UL-DL covariance matrices are Toeplitz matrices that depend on a one-dimensional normalized measure that is obtained by projecting $\gamma(\boldsymbol{\xi})$ onto the collection of 2-dim planes that are orthogonal to $\boldsymbol{\xi}_0$, and can be represented by $\gamma(d\xi)$ in terms of the parameter $\xi := \langle\boldsymbol{\xi}_0, \boldsymbol{\xi}\rangle$ that takes on values in $[-1, 1]$ as $\boldsymbol{\xi}$ varies over the sphere $\mathbb{S}^2$. Thus, the UL-DL covariance transformation problem boils down to that stated in Problem 1. We expect that when

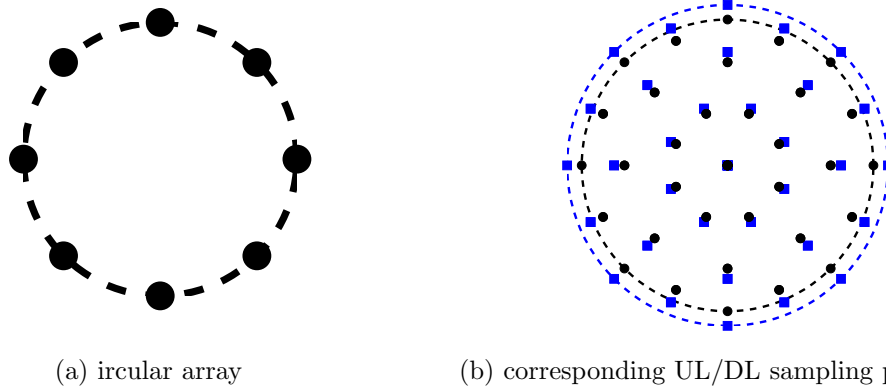(a) ircular array (b) corresponding UL/DL sampling positions

Figure 3.6: Illustration of a circular array geometry $\mathcal{R}$ and the corresponding UL (solid circle) and DL (solid square) sampling positions $\mathcal{D}_{\mathrm{ul}} = \frac{\mathcal{R} - \mathcal{R}}{\frac{\lambda_{\mathrm{ul}}}{2}}$ and $\mathcal{D}_{\mathrm{dl}} = \frac{1}{\nu} \times \mathcal{D}_{\mathrm{ul}}$.

$\nu < 1$, thus, the DL sampling set $\mathcal{D}_{\mathrm{dl}}$ is an expanded version of UL one $\mathcal{D}_{\mathrm{ul}}$, as in the 1-dim case of ULA, one needs to impose some spatial oversampling by a factor $\varrho < 1$ in our notation to guarantee a stable UL-DL covariance extrapolation[4], namely, the array elements should be closely spaced (measured in terms of $\frac{\lambda_{\mathrm{ul}}}{2}$).

Also note that, for any array geometry $\mathcal{R}$, the Minkowski difference $\mathcal{D} = \mathcal{R} - \mathcal{R}$ is a symmetric set, i.e., $\mathcal{D} = -\mathcal{D}$, centered at the origin $\mathbf{0}$, and $\mathcal{D}_{\mathrm{ul}}$ and $\mathcal{D}_{\mathrm{dl}}$ correspond to the directional scaling of this set with respect to the origin by a factor $\frac{2}{\lambda_{\mathrm{ul}}}$ and $\frac{2}{\lambda_{\mathrm{dl}}}$ respectively. As a result, one can always identify a well-defined boundary between $\mathcal{D}_{\mathrm{ul}}$ and $\mathcal{D}_{\mathrm{dl}}$. For example, for a circular array illustrated in Fig. 3.6, this boundary corresponds to all DL sampling points (solid squares) that lie outside the boundary circle corresponding to the UL sampling points (solid circles). Earlier in this chapter, we used the properties of the series solutions of Chebyshev ordinary differential equation (since we had a one-dimensional variable ξ) to derive bounds on the required spatial oversampling factor ϱ and to specify the subset of reliable samples as in (3.15) in a minimax setup. It would be interesting to derive similar bounds using perhaps tools from partial differential equations (since $\boldsymbol{\xi} \in \mathbb{S}^1$ is two-dimensional). We leave this as an interesting problem to be further investigated beyond this dissertation.

3.7 Simulation Results

In this section, we study different aspects of the UL-DL covariance transformation problem via numerical simulations. We consider the IMT FDD band as in the LTE standard [111]

---

[4]Note that for general array configurations the UL/DL sampling sets can be highly irregular (non-uniform), and one needs to make sure that the more general Landau sampling theorem [74, 125] is fulfilled (rather than the uniform sampling in the sampling theorem), which puts restrictions on the minimum density of the sampling sets (see, e.g., Fig. 3.6).

Figure 3.7: UL-DL interpolation error on the DL sampling positions for different number of antennas $M \in \{50, 100\}$ when $\varrho = 1.05 > 1$ and does not fulfill the sampling theorem.

with a UL band $\mathcal{F}_{\mathsf{ul}} = [1920, 1980]\,\text{MHz}$ and a DL band $\mathcal{F}_{\mathsf{ul}} = [2110, 2170]\,\text{MHz}$. We set the carrier frequencies as $f_{\mathsf{ul}} = 1950$ and $f_{\mathsf{dl}} = 2140$ and $\nu = \frac{f_{\mathsf{ul}}}{f_{\mathsf{dl}}} \approx 0.9 < 1$. We assume that the ULA at the BS scans the angular range $\Theta = [-\theta_{\max}, \theta_{\max}]$ with a $\theta_{\max} = 60$ degrees. We assume that the ASF γ is the piece-wise constant density

$$\gamma = \mathsf{rect}_{[0.6, 0.8]} + 4\,\mathsf{rect}_{[0.8, 1]}, \tag{3.75}$$

where for $\mathcal{X} \subset [-1, 1]$, we denote by $\mathsf{rect}_{\mathcal{X}}$ a rectangular pulse of amplitude 1 in $\mathcal{X}$ and 0 elsewhere. Note that γ is a normalized measure and $\breve{\gamma}(0) = \gamma([-1, 1]) = 1$. For the simulations, we apply the NNLS covariance transformation method solved as in $\mathcal{P}^d_{\mathsf{nnls}}$ introduced in Section 3.5 over a grid of size $G = 4M$, where M denotes the number of antennas.

3.7.1 Aliasing Effect (Grating Lobes) for $\varrho > 1$

We first consider the following simulation to illustrate the importance of the spatial over-sampling factor ϱ. We assume that the antenna spacing d violates the conditions of sampling theorem, that is, $\varrho > 1$ and $d = \frac{\varrho \lambda_{\mathsf{ul}}}{2 \sin(\theta_{\max})}$ is larger than $\frac{\lambda_{\mathsf{ul}}}{2 \sin(\theta_{\max})}$. In this case, even if $M \to \infty$, one might not be able to recover γ uniquely due to the aliasing, which is also known as the grating lobe effect in array processing literature [67, 127]. Fig. 3.7 illustrates the UL-DL transformation error at the DL sampling set. It is seen that even for

Interpolation Error for $\varrho = 0.9$
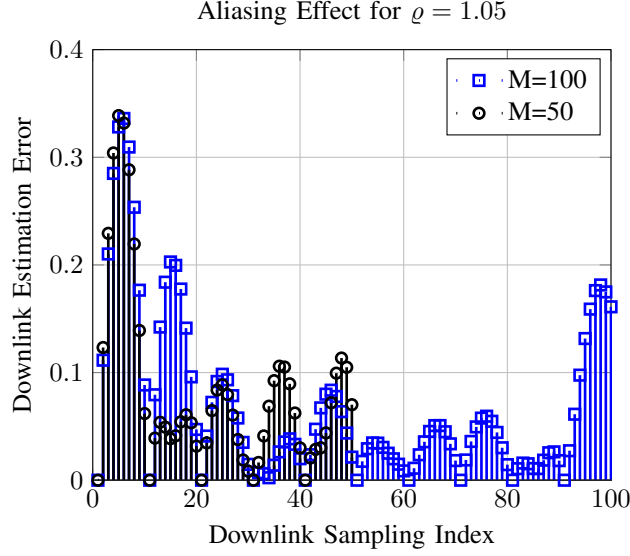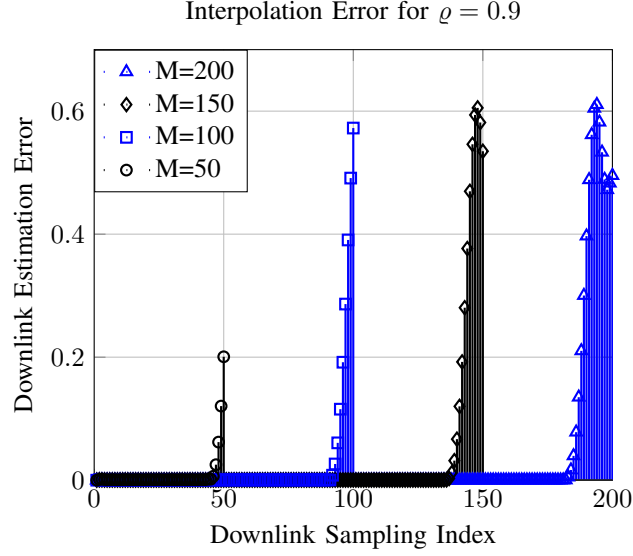

Figure 3.8: UL-DL transformation error on the DL sampling positions for different number of antennas M when $\varrho = 0.9 < 1$ and fulfills the sampling theorem.

$\varrho = 1.05$, which is only slightly larger than 1, the UL-DL transformation completely fails, thus, illustrating the importance of $\varrho < 1$. Also, note that, in this case, the interpolation error does not vanish by increasing the number of antennas M.

We repeat the simulations, where this time we assume that $\varrho = 0.9 < 1$ and fulfills the sampling theorem. Fig. 3.8 illustrates the simulation results, where we again assume that the ASF γ is as in (3.75). It is seen that the proposed transformation algorithm estimates $\breve{\gamma}$ very well at all DL sampling locations, where the transformation error grows quite fast on the boundary locations, which consists of around 10% of the samples. It is also seen that the error at the boundary points grows with increasing the number of antennas, which is compatible with Fig. 3.5, where the error exponent grows by increasing the number of antennas M and approaches the positive error exponent $f(\alpha)$ as in (3.34) asymptotically as $M \to \infty$.

The simulation results also illustrate that our proposed UL-DL interpolation algorithm seems to perform better than what predicted from Theorem 3.1. This is partly due to the fact that the result of Theorem 3.1 has a minimax nature and considers the worst-case γ. It would be interesting to sharpen the result of Theorem 3.1 for a more structured γ via a more refined analysis, which requires further research.

3.7.2 The Effect of Increasing the Array Size

A natural question that arises from our analysis, is how increasing the number of antennas M effects the ergodic rate in a massive MIMO system. As described in section 3.4.2 there

exists an underlying trade-off between the covariance transformation quality and the array degrees of freedom, scaling with $O(M\varrho)$. Increasing M while ϱ is constant, improves the degrees of freedom, but increases the interpolation error as well. In this section, we examine the effect of increasing M, while keeping a constant degrees of freedom, i.e. $M\varrho = \text{const}$, by adjusting the spatial oversampling factor. Our implied message here is that, one can achieve the advantages of FDD massive MIMO by paying in terms of the number of antennas implemented at the BS.

The simulation settings for this experiment are as follows. We assume that there are $K = 5$ groups of users, occupying disjoint spatial intervals $[-\nu + \frac{2\nu}{5}k, -\nu + \frac{2\nu}{5}(k+1)]$, $k = 0, \ldots, K-1$. In each group there exists one user, whose ASF consists of 30 randomly located equal-amplitude delta functions in the respective interval. We assume that the UL covariance for each user is available at the BS.[5] To estimate the DL covariances, the BS incorporates the NNLS method by solving $\mathcal{P}^d_{\mathsf{nnls}}$ method. In order to transmit data in the DL, the BS incorporates a statistical beamforming method. There exist several statistical beamforming schemes, among which we select the one proposed in [140]. Note that, the choice of the statistical beamforming method is not critical here and we adopt the mentioned method for simplicity. Let $\boldsymbol{\Sigma}_k$, $k = 0, \ldots, K-1$ denote the DL covariance matrix of the users. Then the beamforming vector for user k is given by

$$\mathbf{u}_\ell = \mathbf{u}_{\max}\{(\mathbf{I}_M + \sum_{k \neq \ell} \boldsymbol{\Sigma}_k)^{-1}\boldsymbol{\Sigma}_\ell\}, \tag{3.76}$$

where $\mathbf{u}_{\max}(\mathbf{X})$ denotes the eigenvector corresponding to the maximum eigenvalue of $\mathbf{X}$. The precoding matrix is given as $\mathbf{V} = \left(\tilde{\mathbf{U}}^{\mathsf{H}}\right)^\dagger \mathbf{D}^{1/2}$, where $\tilde{\mathbf{U}} = [\mathbf{u}_1, \ldots, \mathbf{u}_K]$ and $\mathbf{D}$ is a diagonal matrix that makes the columns of $\mathbf{V}$ to have unit norm. Then, the signal received at user ℓ can be written as

$$y_\ell = \mathbf{h}_\ell^{\mathsf{H}}\mathbf{V}\mathbf{s} + n_\ell = b_{\ell,\ell}[\mathbf{s}]_\ell + \sum_{k \neq \ell} b_{\ell,k}[\mathbf{s}]_k + n_\ell,$$

where $\mathbf{s} \in \mathbb{C}^K$ is the symbols vector, $n_\ell \sim \mathcal{CN}(0, N_0)$ is the AWGN with $N_0 = \frac{P}{K\mathrm{SNR}}$, where P is the total transmission power at the BS. Also the coefficients $(b_{\ell,1}, \ldots, b_{\ell,K})$ are given by the elements of the $1 \times K$ row vector $\mathbf{h}_\ell^{\mathsf{H}}\mathbf{V}$. The ergodic sum-rate with perfect CSI at the receiver is given by

$$R_{\mathrm{sum}} = \sum_{\ell=1}^{K} \mathbb{E}\left[\log\left(1 + \frac{|b_{\ell,\ell}|^2}{1 + \sum_{k \neq \ell} |b_{\ell,k}|^2}\right)\right], \tag{3.77}$$

---

[5]Otherwise the UL covariance can be estimated via UL pilot observations, using a standard Toeplitz PSD covariance estimator.
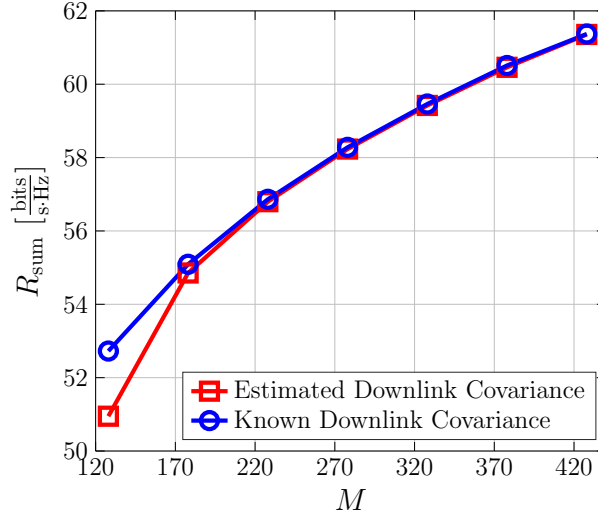
Figure 3.9: Sum-rate comparison as a function of M with SNR = 20 dB. Here the degrees of freedom is set to $M\varrho = 128$.

which we calculate using Monte Carlo simulations.

Fig. 3.9 compares the ergodic sum-rate as a function of the number of antennas M for two scenarios, where in one the DL covariance is known to the BS and in the other, it is estimated from the UL covariance. As we can see, the sum-rate in the estimated DL covariance case increases by increasing M and approaches that of the case with perfect DL covariance knowledge. This means that, if we increase the number of antennas at the BS and simultaneously decrease the spatial oversampling factor inversely proportional to the number of antennas, DL covariance estimation as proposed throughout this chapter imposes no sum-rate degradation.

3.7.3 MUSIC+NNLS vs Alternating Projection

To compare the performance of transformation based on NNLS and transformation via the alternating projection method proposed in [87], we compare their accuracy in terms of Downlink covariance estimation error. Note that the algorithm in [87] returns an ASF in the feasibility set $\Gamma_\gamma = \mathcal{D}_+ \bigcap \mathcal{D}_\gamma$ (see $\mathcal{P}_{\text{feas}}$). Therefore, the estimate depends on the solution method and two different algorithms may yield different ASF estimates. Therefore, a comparison of the proposed NNLS method and the projection method is reasonable.

Error Metrics
We use the following two error metrics to compare the estimated DL covariance with the true DL covariance.

1. **Normalized Euclidean Distance.** Let $\boldsymbol{\Sigma} \in \mathbb{S}_+^M$ be the true channel covariance and let $\widehat{\boldsymbol{\Sigma}}$ denote its estimate. A standard error metric is given by the normalized Euclidean norm of the difference, given as

$$E_{\mathsf{euc}} = \frac{\|\boldsymbol{\Sigma} - \widehat{\boldsymbol{\Sigma}}\|_{\mathsf{F}}}{\|\boldsymbol{\Sigma}\|_{\mathsf{F}}}. \qquad (3.78)$$

2. **Relative Efficiency.** In wireless communication it is sometimes crucial to know, how similar their associated "dominant" subspaces. Let $p \in \{1, \ldots, M\}$ denote a subspace dimension parameter, $\mathbf{U}_p = [\mathbf{U}_{\cdot,1}, \ldots, \mathbf{U}_{\cdot,p}] \in \mathbb{C}^{M \times p}$ the p eigenvectors of $\boldsymbol{\Sigma}$ corresponding to its largest p eigenvalues. Similarly, define $\widehat{\mathbf{U}}_p$ as the matrix of the p eigenvectors of $\widehat{\boldsymbol{\Sigma}}$ corresponding to its p largest eigenvalues. We define the relative efficiency parameter as [51]

$$\varepsilon(p) = 1 - \frac{\langle \boldsymbol{\Sigma}, \widehat{\mathbf{U}}_p \widehat{\mathbf{U}}_p^{\mathsf{H}} \rangle}{\langle \boldsymbol{\Sigma}, \mathbf{U}_p \mathbf{U}_p^{\mathsf{H}} \rangle}. \qquad (3.79)$$

It is easy to see that $\varepsilon(p) \in [0, 1]$, for all $p = 1, \ldots, M$. Disregarding the estimated eigenvalues, this metric compares the power captured by the estimated p-dominant subspace to that captured by the true p-dominant subspace. The closer $\varepsilon(p)$ is to zero, the better the p-dominant subspace of $\boldsymbol{\Sigma}$ is estimated.

Averages of these metrics $\mathbb{E}[E_{\mathsf{euc}}]$ and $\mathbb{E}[\varepsilon(p)]$ are taken over all sources of randomness and compared for different methods.

In order to generate random ASF's, we consider a similar setting to the one used in Section 2.6 of Chapter 2, where we for the sake of simulation, the parametric ASF model is used to generate random spectra, so that the error metrics are averaged over a variety of possible ASF's. According to the ASF parametric model

$$\gamma(\xi) = \gamma_d(\xi) + \gamma_c(\xi) = \sum_{k \in \mathbb{Z}_K} c_k \delta(\xi - \xi_k) + \sum_{k \in \mathbb{Z}_{K'}} \widetilde{c}_k g_k(\xi),$$

where K is the number of discrete paths, $c_k \in \mathbb{R}_+$ is the coefficient of the k-th discrete path, ξ_k is its AoA, where K' is the number of continuous scattering elements, $\widetilde{c}_k$ is a scaling factor associated with the k-th diffuse scatterer, and g_k is a continuous, non-negative measure over an open interval $\mathcal{X}_k \subseteq (-1, 1)$. Therefore, we can generate an ASF by specifying $2K + 3\widetilde{K}$ real-valued parameters, namely $\{\xi_k, c_k\}_{k=0}^{K-1}$ and $\{a_k, b_k, \widetilde{c}_k\}_{k=0}^{\widetilde{K}-1}$. In addition, to have control over the relative power of the discrete and continuous ASF

Parameter	Distribution
K	$\mathrm{Unif}(\{1, 2, 3, 4\})$
$\widetilde{K}$	$\mathrm{Unif}(\{1, 2, 3, 4\})$
α	$\in \{0, \frac{1}{2}, 1\} \sim$ deterministic defined within the specific simulation result
ξ_k	$\sim \mathrm{Unif}([-1, 1]), \ k \in \mathbb{Z}_K$
c_k	$\sim 1 + \|\mathcal{CN}(0, 1)\|, \ k \in \mathbb{Z}_K$
a_k	$\sim \mathrm{Unif}([-1, 1]), \ k \in \mathbb{Z}_{\widetilde{K}}$
b_k	$\sim \mathrm{Unif}([0.1, 0.1 + \mathsf{w}]), \ k \in \mathbb{Z}_{\widetilde{K}}$, where w represents the scatterer's width will be set to different values depending on the specific simulation result
$\widetilde{c}_k$	$\sim 1 + \|\mathcal{CN}(0, 1)\|, \ k \in \mathbb{Z}_{\widetilde{K}}$

Table 3.1: Simulation Parameters

components, we generate the ASF as

$$\gamma = \alpha \gamma_d^{\mathsf{nor}} + (1 - \alpha) \gamma_c^{\mathsf{nor}}, \tag{3.80}$$

where $\gamma_d^{\mathsf{nor}} = \gamma_d / \int_{-1}^{1} \gamma_d(\xi) d\xi$, and $\gamma_c^{\mathsf{nor}} = \gamma_c / \int_{-1}^{1} \gamma_c(\xi) d\xi$ are normalized discrete and continuous ASF components, respectively. α controls the contribution of each part to the overall ASF γ, i.e. if $\alpha = 0$, we have a purely continuous ASF, while if $\alpha = 1$, we have a purely discrete one. Then, since α, K, and $\widetilde{K}$ are themselves to be specified, in total we need to determine $2K + 3\widetilde{K} + 3$ parameters. We select these parameters independently, each according to a distribution listed as in Table 3.1. We average the results over 1000 random ASF realizations.

Continuous vs Discrete ASF Components As discussed throughout this dissertation, in many propagation environments, the channel power is received at the BS through a combination of LoS and specular as well as diffuse components. We modeled the former as discrete measures (Dirac impulses) and the latter as continuous measures in the spectrum. It is interesting to see how the UL-DL transformation algorithms perform in various discrete-continuous *power-splitting* scenarios. This split of power can be modeled via the ASF expression in (3.80) by changing the scalar $\alpha \in [0, 1]$. When $\alpha \to 1$, the ASF is more and more dominated by its discrete component such that at $\alpha = 1$, the ASF is purely discrete. This is in fact the case that is previously studied in plenty of works in the literature dedicated to estimating line-spectral angular power spectra. In contrast, when $\alpha \to 0$, the ASF is dominated by its continuous component such that at $\alpha = 0$ the ASF is purely continuous. We apply the MUSIC-assisted NNLS and alternating projections methods for three different values of $\alpha = 0, 0.5, 1$. The shape of the ASF is generated semi-randomly according to Table 3.1. In all cases we choose the dictionary used

Figure 3.10: covariance estimation error curves for the average normalized Euclidean distance and relative efficiency metrics against the channel dimension M for three values of the continuous-discrete power-splitting ratio α.

for approximating the continuous ASF component to be the rectangular density family $\Psi_{c,\mathrm{rect}} = \{\mathrm{rect}_{[-1+\frac{2i}{G_c}, -1+\frac{2(i+1)}{G_c})}, i = 0, \ldots, G_c - 1\}$ of size $G_c = 3M$.

Fig. 3.10 illustrates the error curves for four different algorithms, namely, the alternating projections (AP) method with and without the ASF non-negativity constraint (AP Alg. 2 and AP Alg. 1, respectively) and the NNLS method with and without MUSIC.

First, comparison between alternating projections Alg. 1 and Alg. 2 illustrates the important point that adding the non-negativity constraint reduces the transformation error. Although, this comes at the cost of more algorithmic complexity, but it eventually improves channel estimation MSE. Second, the comparison between the NNLS method with and without the MUSIC step shows the importance of the MUSIC step in estimating the support of the discrete ASF component. Using the support information, the algorithm adds appropriate columns to dictionary $\mathbf{A}$ in (3.67), which can represent the contribution of the discrete part to the covariance entries. Without adding these columns, the contribution of the rank-1 matrices $\{\mathbf{a}(\xi_k)\mathbf{a}(\xi_k)^{\mathsf{H}}\}_{k=0}^{K-1}$ can not be well-approximated by the rest of the dictionary columns.

Finally, these results show the superiority of the MUSIC-assisted NNLS method to the other methods in all cases except for the case in which no discrete component exists ($\alpha =$

0). The reason is that, this method is able to accurately compensate for the contribution of the discrete ASF component by estimating the support of this component through a preliminary MUSIC step. In contrast, alternating projection methods do not differentiate between discrete and continuous ASF components and assume the ASF to belong to the space $L_2([-1,1])$. When the contribution of the discrete ASF is high ($\alpha \to 1$), this causes a performance degradation in the transformation.

The Effect of Sparsity

Channel sparsity as a structural assumption has been exploited in previous works as a tool for improved covariance estimation/transformation by restricting the class of admissible ASF to those that are sparse. In a finite-dimensional setup, angular sparsity of the channel is equivalent to the ASF having a number of non-zero components s that is much smaller than channel dimension ($s \ll M$), or is otherwise *well-approximated* by a sparse vector. In an infinite dimensional case, one can translate the concept of sparsity to the ASF having a support size that is much smaller than the volume of the angular domain. If the discrete ASF component has small support, then one can approximate the number of angular directions having significant power as

$$s = K + \frac{\left| \bigcup_{k=0}^{\widetilde{K}-1} \operatorname{supp}(g_k) \right|}{2} M, \tag{3.81}$$

where K is the number of discrete AoAs and g_k, $k = 0, \ldots, \widetilde{K} - 1$ are the continuous functions representing the diffuse scattering components. Now, consider a ULA consisting of $M = 100$ antennas, with a $d = \frac{\lambda_{\mathsf{ul}}}{2}$ uniform spacing. In order to study the effect of sparsity on UL-DL covariance transformation, we compare the error metrics as a function of channel sparsity order. In this case we use the following generative model:

$$\gamma_d(\xi) = \sum_{k=0}^{K_i-1} c_k \delta(\xi - \xi_k), \tag{3.82a}$$

$$\gamma_c(\xi) = \sum_{k=0}^{1} c_k g_{k,i}(\xi) \tag{3.82b}$$

$$\gamma = \alpha \gamma_d^{\mathsf{nor}} + (1 - \alpha) \gamma_c^{\mathsf{nor}}, \tag{3.82c}$$

where $\gamma_d^{\mathsf{nor}} = \gamma_d / \int_{-1}^{1} \gamma_d(\xi) d\xi$, and $\gamma_c^{\mathsf{nor}} = \gamma_c / \int_{-1}^{1} \gamma_c(\xi) d\xi$ are normalized discrete and continuous ASF components, respectively, K_i is the number of discrete impulses associated with the i-th sparsity order. The sparsity order is chosen as $s_i = 10i + 10$, and the number of impulses as $K_i = 1 + i$ for $i = 0, \ldots, 6$. The discrete AoA $\{\xi_k\}_{k=0}^{K_i-1}$ are generated uniformly and independently at random over $[-1, 1]$ with coefficients, generated indepen-
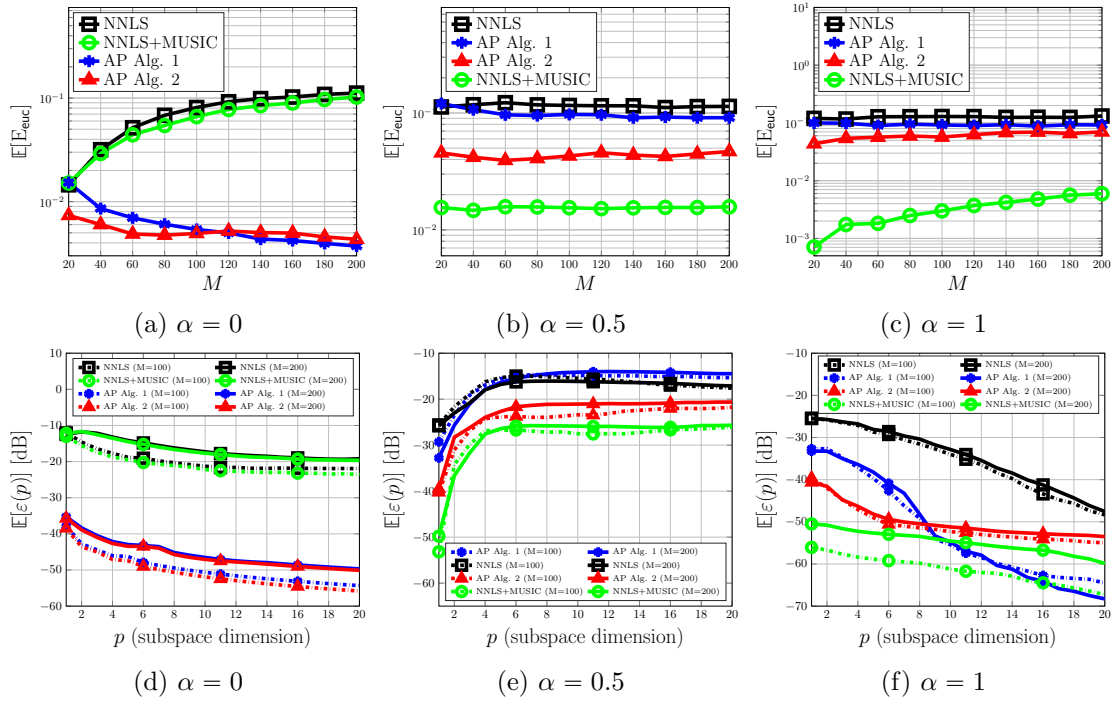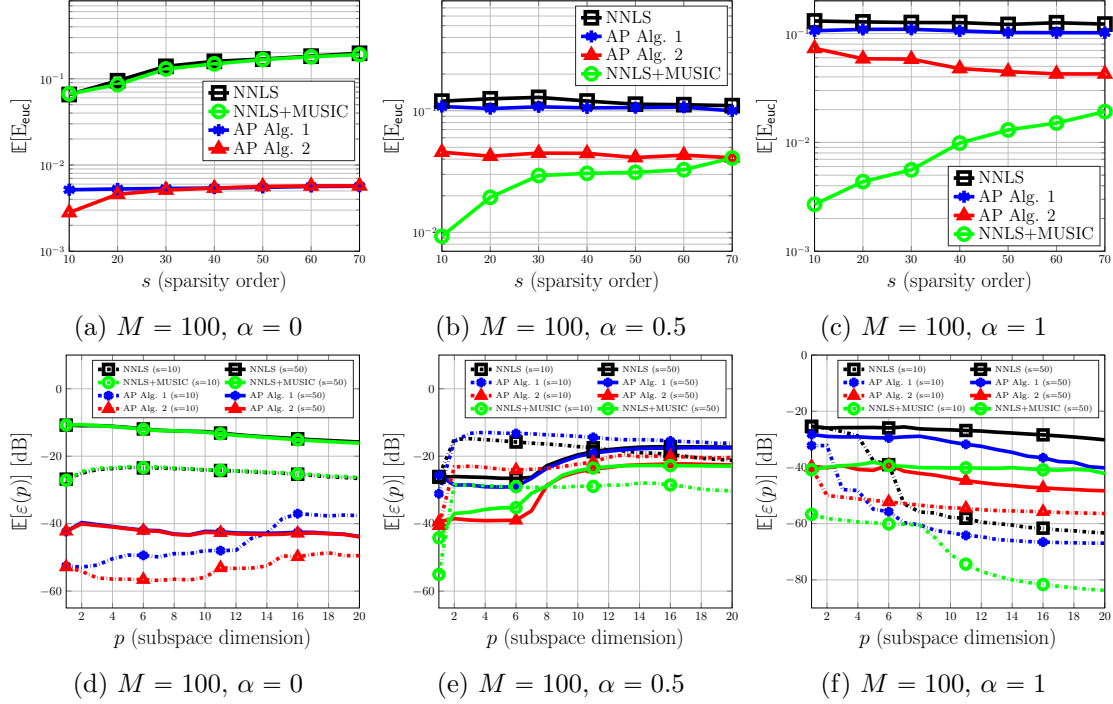
Figure 3.11: covariance estimation error curves for the normalized Euclidean distance and relative efficiency metrics against the channel sparsity order s for three values of the continuous-discrete power-splitting ratio α.

dently according to the expression provided in Table 3.1, and $g_{k,i}$ are generated such that $\mathrm{supp}(g_{0,i}) \subset [-1,0)$ and $\mathrm{supp}(g_{1,i}) \subset [0,1)$ for all i, so that the two continuous components have non-overlapping support. Also, we generate the scatterer width parameter b_k deterministically as

$$b_k = \frac{s_i - K_i}{M}, \ k = 0, 1,$$

for $i = 0, \ldots, 6$. The mean parameter is then generated according to $a_0 \sim \mathrm{Unif}([-1 + \frac{b_0}{2}, -\frac{b_0}{2}])$ and $a_1 \sim \mathrm{Unif}([\frac{b_1}{2}, 1 - \frac{b_1}{2}])$.

Fig. 3.11 illustrates the error curves of this experiment for $\alpha = 0, 0.5, 1$. The results show that, when the ASF contains a discrete component ($\alpha \neq 0$), the MUSIC-assisted NNLS method has the best performance for all values of the sparsity order (see Figs. 3.11b and 3.11c). In contrast, when $\alpha = 0$ the methods based on alternating projections have a superior performance (see Fig. 3.11a). Also, for larger sparsity orders, the performance of NNLS (with or without MUSIC) deteriorates. The reason for this phenomenon is that the NNLS method implicitly favors sparse solutions and when the channel is not sparse, this causes a not-so-large increase in the transformation error.

4 Downlink Precoding

4.1 Problem Statement

Consider a base station (BS) equipped with a uniform linear array (ULA) with M antennas, operating in frequency division duplex (FDD) mode: Uplink (UL) transmission from a user to the BS takes place over a frequency band $\mathcal{F}_{\mathsf{ul}} = [f_{\mathsf{ul}} - \frac{\mathrm{BW}_{\mathsf{ul}}}{2}, f_{\mathsf{ul}} + \frac{\mathrm{BW}_{\mathsf{ul}}}{2}]$ with carrier frequency f_{ul} and a bandwidth $\mathrm{BW}_{\mathsf{ul}}$, and Downlink (DL) transmission from a BS to the user takes place over a band $\mathcal{F}_{\mathsf{dl}} = [f_{\mathsf{dl}} - \frac{\mathrm{BW}_{\mathsf{dl}}}{2}, f_{\mathsf{dl}} + \frac{\mathrm{BW}_{\mathsf{dl}}}{2}]$ with carrier frequency f_{dl} and a bandwidth $\mathrm{BW}_{\mathsf{dl}}$. For example, one mode of operation in the 3GPP standard uses the $\mathcal{F}_{\mathsf{ul}} = [1920, 1980]$ MHz band for UL transmission and the $\mathcal{F}_{\mathsf{dl}} = [2110, 2170]$ MHz band for DL, so that $f_{\mathsf{ul}} = 1950$ MHz, $\mathrm{BW}_{\mathsf{ul}} = 60$ MHz, $f_{\mathsf{dl}} = 2140$ MHz, and $\mathrm{BW}_{\mathsf{dl}} = 60$ MHz [111]. This is in contrast to the time division duplexing (TDD) operation mode which uses the same band for UL and DL transmission, and uses time duplexing for UL/DL data reception/transmission. In TDD the BS can directly use the obtained channel state information (CSI) during UL for data transmission in the DL. This is possible based on the reasonable assumption that the UL and DL channels are identical, since the time difference between the two reception and transmission is so short that the change in the instantaneous channel is negligible. This phenomena is known as *channel reciprocity*, which does not hold in FDD mode, since the instantaneous UL and DL channels correspond to two distinct frequency bands (the gap between these bands is equal to190 MHz for the example above).

Channel reciprocity simplifies the task of DL data transmission in a TDD system: the BS uses the estimated CSI during UL to design a precoder and transmits multiplexed data to a set of users. In FDD, however, the BS has to first estimate the DL channels. This is done by broadcasting a number of T_{dl} pilot symbols to the multi-user system and receiving the corresponding pilot measurements from each user in UL, which the BS uses to estimate DL channels and design the precoder. In a massive MIMO system this proves to be a challenging task, due to the high channel dimension ($M \gg 1$). In order to train M antennas, a conventional DL training scheme requires a minimum pilot dimension of $T_{\mathsf{dl}} = M$. Hence, with such a scheme, the number of BS antennas M cannot be made arbitrarily large. For example, consider a typical case taken from the LTE system [111], where groups

of users are scheduled over resource blocks spanning 14 OFDM symbols $\times$ 12 subcarriers, for a total dimension of $T = 168$ symbols in the time-frequency plane. Consider a typical massive MIMO configuration serving $K \sim 20$ users with $M \geqslant 200$ antennas (see e.g. [79]). In this case, since $M \geqslant T$, the entire resource block dimension would be consumed by the DL pilot, leaving no room for data communication. Furthermore, feeding back the M-dimensional measurements (or estimated/quantized channel vectors) represents also a significant feedback overhead for the UL [19, 63, 69, 77, 137].

Designing a massive MIMO system that operates in FDD mode requires developing methods that overcome these dimensionality issues. In particular, we must address the problem of designing a "suitable" channel precoder that enables simultaneous data transmission to a set of users. This chapter proposes a novel solution to this problem.

4.1.1 Related Work

Several works have proposed to reduce both the DL training and UL feedback overheads by exploiting the sparse structure of the massive MIMO channel. In particular, these works assume that propagation between the BS array and the user antenna occurs through a limited number of scattering clusters, with limited support in the Angle-of-Arrival/Angle-of-Departure (AoA-AoD) domain.[1] Hence, by decomposing the angle domain into discrete "virtual beam" directions, the M-dimensional user channel vectors admit a sparse representation in the beam-space domain [5, 108]. Building on this idea, a large number of works (see, e.g., [32, 40, 44, 46, 71, 98, 112, 135]) have proposed to use "compressed pilots", i.e., a reduced DL pilot dimension $T_{\mathsf{dl}} < M$, in order to estimate the channel vectors using *Compressed Sensing* (CS) techniques [20, 41]. For example, in [5] the sparse representation of channel multipath components in angle, delay and Doppler domains was exploited to propose CS methods for channel estimation using far fewer measurements than required by conventional least-squares (LS) methods. In [46], the authors note that the angles of the multipath channel components are common among all the subcarriers in the OFDM signaling. Then they propose to exploit the common sparsity pattern of the channel coefficients to further reduce the number of required pilot measurements. This gives rise to a so-called Multiple Measurement Vector (MMV) setup, which is typically applied when multiple snapshots of a random vector with common sparse support can be acquired and jointly processed [24, 43]. This was adapted to FDD in the massive MIMO regime, where the frequent idea is to probe the channel using compressed DL pilots, receiving the measurements at the BS via feedback and performing channel estimation there. A recent work

---

[1]From the BS perspective, AoD for the DL and AoA for the UL indicate the same domain. Hence, we shall simply refer to this as the "angle domain", while the meaning of departure (DL) or arrival (UL) is clear from the context.

based on this approach was presented in [98], starting with the observation that, as shown in many experimental studies [45,58,65,73], the propagation between the BS antenna array and the users occurs along scattering clusters that may be common to multiple users, since they all belong to the same scattering environment. In turn, this yields that the channel sparse representations (in the angle/beam domain) share a common part of their support. Hence, [98] considers a scheme where the users feed back their noisy DL pilot measurements to the BS and the latter runs a *joint recovery* algorithm, coined as *Joint Orthogonal Matching Pursuit* (J-OMP), able to take advantage of the common sparsity. It follows that in the presence of common sparsity, J-OMP improves upon the basic CS schemes that estimate each user channel separately.

More recent CS-based methods, in addition, make use of the *angular reciprocity* between the UL and the DL channels in FDD systems to improve channel estimation. Namely, this refers to the fact that the directions (angles) of propagation for the UL and DL channel are invariant over the frequency range spanning the UL and DL bands, which is generally very small with respect to the carrier frequency (e.g., UL/DL separation of the order of 100 MHz, for carrier frequencies ranging between 2 and 6 GHz) [2,59,134]. In [135] the sparse set of AoAs is estimated from a preamble transmission phase in the UL, and this information is used for user grouping and channel estimation in the DL according to the well-known *joint spatial division and multiplexing* (JSDM) paradigm [1,90]. In [40] the authors proposed a dictionary learning-based approach for training DL channels. First, in a preliminary learning phase the BS "learns" a pair of UL-DL dictionaries that are able to sparsely represent the channel. Then, these dictionaries are used for a joint sparse estimation of instantaneous UL-DL channels. An issue with this method is that the dictionary learning phase requires off-line training and must be re-run if the propagation environment around the BS changes (e.g., due to large moving objects such as truck and buses, or a new building). In addition, the computation involved in the instantaneous channel estimation is prohibitively demanding for real-time operations with a large number of antennas ($M > 100$). In [32] the authors propose estimating the DL channel using a Sparse Bayesian Learning framework, aiming at joint-maximum a posteriori (MAP) estimation of the off-grid AoAs and multipath power coefficients by observing instantaneous UL channel measurements. This method has the drawback that it fundamentally assumes discrete and separable (in the AoA domain) multipath components and that the number of signal paths (number of channel AoAs) is known a priori. Hence, the method simply cannot be applied in scenarios with diffuse (continuous) scattering, where the scattering power is distributed over an interval of the angular domain with non-negligible width. Such scattering are observed and modeled for various types of communication channels, and they do not necessarily admit a sparse angular representation [93,101,103,121].
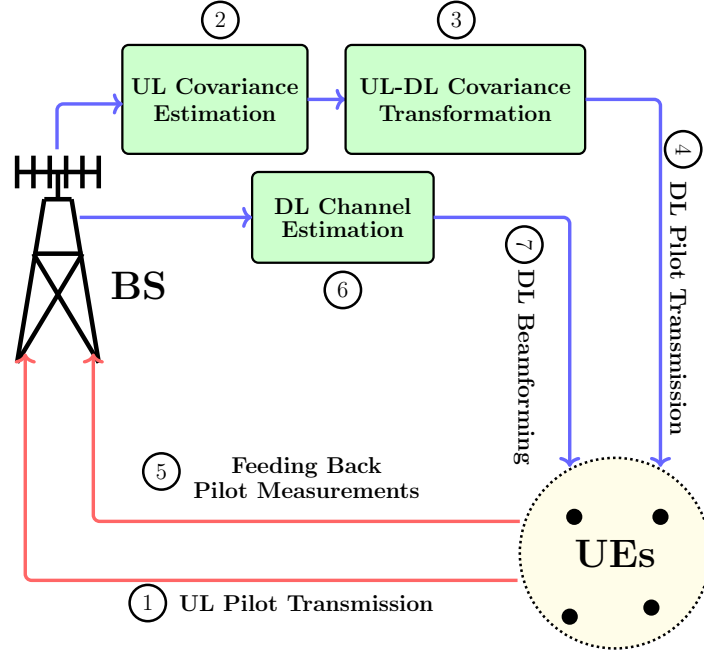
Figure 4.1: Overall diagram of our scheme for an FDD massive MIMO system.

Throughout this chapter we develop the idea of *active channel sparsification* (ACS) in which, for a given DL pilot dimension, the BS selects a set of angular directions to transmit data to the users, such that the number of DL data streams that the system can support is maximized. We show that this method does not assume a sparse structure on the channel and yields excellent performance with any limited pilot dimension. This chapter also completes our proposed scheme for realizing a multi-user FDD massive MIMO system. To recapitulate, this scheme consists of the following modules: (1) UL channel covariance estimation from orthogonal UL pilots received at the BS, (2) UL-DL covariance transformation, and (3) DL instantaneous channel training and multi-user precoding. The diagram presented in Fig. 4.1 illustrates the overall proposed scheme. In what follows we first discuss the concept of ACS for a multi-user system with a ULA geometry. Then we extend the idea to arbitrary array geometries.

4.2 System Setup: Multi-User MIMO with a massive ULA

In this section, for the sake of clarity, we develop the ACS method for the simple case of a ULA with M antennas and a multi-user scenario in which K users are to be served. The DL channel vector of each user $\mathbf{h}_k \in \mathbb{C}^M$, $k \in \mathbb{Z}_K$ is assumed to be complex Gaussian distributed, with zero mean and a covariance $\boldsymbol{\Sigma}_k$ (to simplify notation, we have dropped the "dl" subscript, since we only consider DL transmission and all mentioned covariances

are DL covariances). In chapter 2 we studied the problem of covariance estimation and in chapter 3 we studied the problem of UL-DL covariance transformation. The DL channel covariance of a user can be estimated, for example, by first estimating the UL covariance (using the tools developed in chapter 2) and then transforming it to the DL covariance (using the tools developed in chapter 3). Nevertheless, in order to address the problem of DL multi-user precoding and study its properties apart and isolated from other components of a FDD massive MIMO scheme, in this chapter we assume that the *true* channel covariances $\{\boldsymbol{\Sigma}_k\}_k$ are known to the BS.

4.2.1 Common Eigenbasis of ULA Covariances

The channel covariance of a ULA channel is Toeplitz Hermitian (see Section 2.3 in Chapter 2). In the massive MIMO regime where $M \gg 1$, finding a set of (approximate) common eigenvectors for all the covariances is possible by considering the circulant approximation of Toeplitz matrices in large dimensions ($M \gg 1$) that follows as an application of the well-known Szegö's Theorem (see details in [1] and references therein).

Consider the channel covariance of user k with eigendecomposition $\boldsymbol{\Sigma}_k = \mathbf{U}_k \boldsymbol{\Lambda}_k \mathbf{U}_k^{\mathsf{H}}$. Define the diagonal matrices $\mathring{\boldsymbol{\Lambda}}_k$, $k \in \mathbb{Z}_K$ for which $[\mathring{\boldsymbol{\Lambda}}_k]_{m,m} = [\mathbf{F}^{\mathsf{H}} \boldsymbol{\Sigma}_k \mathbf{F}]_{m,m}$, where $\mathbf{F}$ is the $M \times M$ DFT matrix, whose (m,n)-th entry is given by $[\mathbf{F}]_{m,n} = \frac{1}{\sqrt{M}} e^{-j2\pi \frac{mn}{M}}$, $m, n \in \mathbb{Z}_M$. There are several ways to define a circulant approximation [142], among which we choose the following:

$$\mathring{\boldsymbol{\Sigma}}_k = \mathbf{F} \mathring{\boldsymbol{\Lambda}}_k \mathbf{F}^{\mathsf{H}}. \tag{4.1}$$

According to Szegö's theorem, for large M, $\mathring{\boldsymbol{\Lambda}}_k$ converges to the diagonal eigenvalue matrix $\boldsymbol{\Lambda}_k$ of $\boldsymbol{\Sigma}_k$, i.e. $\mathring{\boldsymbol{\Lambda}}_k \to \boldsymbol{\Lambda}_k$ as $M \to \infty$ (see [48] for a more rigorous statement). Hence, within a small error for large M, the set of common eigenvectors for all the users is provided by the columns of the $M \times M$ DFT matrix. As a consequence, the DL channel covariance of user k is characterized simply via a vector of eigenvalues $\boldsymbol{\lambda}_k \in \mathbb{R}^M$, with m-th element $[\boldsymbol{\lambda}_k]_m = [\mathring{\boldsymbol{\Lambda}}_k]_{m,m}$. In addition, the DFT matrix forms a unitary basis for (approximately) expressing any user channel vector via a Karhunen-Loeve expansion. In particular, let $\mathbf{f}_m := [\mathbf{F}]_{\cdot,m}$ denote the m-th column of $\mathbf{F}$. We can express the DL channel vector of user k as

$$\mathbf{h}_k \approx \sum_{m=0}^{M-1} g_{k,m} \sqrt{[\boldsymbol{\lambda}_k]_m} \, \mathbf{f}_m, \tag{4.2}$$

where $\{g_{k,m} \sim \mathcal{CN}(0,1)\}_{k,m}$ are i.i.d. standard complex Gaussian random variables. The columns of $\mathbf{F}$ are very similar to array response vectors and in fact, each column with index $m \in [M]$ of the DFT matrix can be seen as the array response to an angular direction $\theta = \sin^{-1}(\frac{m}{M} \sin \theta_{\max})$ where $\theta_{\max}$ denotes the array maximum angular aperture

and $[\boldsymbol{\lambda}_k]_m$ can be seen as the power of the channel vector associated with user k along that direction. Due to the limited number of local scatterers as seen at the BS and the large number of antennas of the array, one can hypothesize that only a few entries of $\boldsymbol{\lambda}_k$ are significantly large, implying that the DL channel vector $\mathbf{h}_k$ is sparse in the Fourier basis. This sparsity in the beam-space domain is precisely what has been exploited in the CS-based works discussed in Section 4.1.1, in order to reduce the DL pilot dimension T_{dl}. It is also evident that this channel representation potentially yields the common sparsity across users, as exploited by J-OMP in [98]. As seen in the next section, our proposed approach does *not* rely on any intrinsic channel sparsity assumption, but adopts a novel *artificial sparsification* technique that smartly reduces the effective channel dimension to enable channel estimation regardless of its sparsity.

4.3 Active Channel Sparsification and DL Channel Probing

In this section we consider the estimation of the *instantaneous realization* of the DL user channel vectors, as a first step to design a DL precoder. As in [1], we consider the concatenation of the physical channel with a fixed precoder, i.e., a linear transformation that may depend on the user channel statistics (notably, on their covariance matrices), but is independent of the instantaneous channel realizations, which in fact must be estimated via the closed-loop DL probing and channel state feedback mechanism.

The BS transmits a training space-time matrix $\boldsymbol{\Psi}$ of dimension $T_{\mathsf{dl}} \times M'$, such that each row $\boldsymbol{\Psi}_{i,\cdot}$ is transmitted simultaneously from the $M' \leqslant M$ inputs of a precoding matrix $\mathbf{B}$ of dimension $M' \times M$, and where M' is a suitable intermediate dimension that will be determined later. The precoded DL training length (in time-frequency symbols) spans therefore T_{dl} dimensions, and the DL training phase is repeated at each resource block of dimension T. Stacking the T_{dl} DL training symbols in a column vector, the corresponding observation at the user k receiver is given by

$$\mathbf{y}_k = \boldsymbol{\Psi}\mathbf{B}\mathbf{h}_k + \mathbf{z}_k = \boldsymbol{\Psi}\mathbf{h}_k^{\mathsf{eff}} + \mathbf{z}_k, \tag{4.3}$$

where $\mathbf{B}$ is the precoding matrix, $\mathbf{h}_k$ is the channel vector of user k, and we have defined $\mathbf{h}_k^{\mathsf{eff}} := \mathbf{B}\mathbf{h}_k$ as the *effective channel* vector, formed by the concatenation of the actual DL channel (antenna-to-antenna) with the precoder $\mathbf{B}$. We consider additive white Gaussian noise (AWGN) with distribution $\mathbf{z}_k \sim \mathcal{CN}(\mathbf{0}, N_0\mathbf{I}_{T_{\mathsf{dl}}})$. The training and precoding matrices are normalized such that

$$\mathsf{tr}(\boldsymbol{\Psi}\mathbf{B}\mathbf{B}^{\mathsf{H}}\boldsymbol{\Psi}^{\mathsf{H}}) = T_{\mathsf{dl}}P_{\mathsf{dl}}, \tag{4.4}$$

where P_{dl} denotes the total BS transmit power and we define the DL signal-to-noise-ratio as SNR= P_{dl}/N_0. Notice that most works on channel estimation focus on the estimation of the actual channels $\{\mathbf{h}_k\}_k$. This is recovered in our setting by letting $\mathbf{B} = \mathbf{I}_M$. However, our goal here is to design a *sparsifying precoder* $\mathbf{B}$ such that each user's effective channel has low dimension (in the beam-space representation) and yet the collection of effective channels for $k \in \mathbb{Z}_k$ form a high-rank matrix. In this way, each user's channel can be estimated using a small pilot overhead T_{dl}, but the BS is still able to serve many data streams using spatial multiplexing in the DL (in fact, as many as the rank of the effective matrix).

4.3.1 Necessity and Implication of Stable Channel Estimation

Suppose that the channel representation (4.2) holds exactly and that the eigenvalue vectors $\boldsymbol{\lambda}_k$ have support $\mathcal{S}_k = \{m : [\boldsymbol{\lambda}_k]_m \neq 0\}$ with sparsity level $s_k = |\mathcal{S}_k|$. We hasten to point out that the above are convenient *design assumptions*, made in order to obtain a tractable problem, and that the precoder designed according to our simplifying assumptions are applied to the actual physical channels. Under these assumptions, the following lemma yields necessary and sufficient conditions of stable estimation of the channel vectors $\mathbf{h}_k$.

Lemma 4.1. *Consider the sparse Gaussian vector* $\mathbf{h}_k$ *with support set* $\mathcal{S}_k$ *given by the RHS of (4.2). Let* $\widehat{\mathbf{h}}_k$ *denote any estimator for* $\mathbf{h}_k$ *based on the observation[2]* $\mathbf{y}_k = \boldsymbol{\Psi}\mathbf{h}_k + \mathbf{z}_k$, *and let* $\mathbf{R}_e = \mathbb{E}[(\mathbf{h}_k - \widehat{\mathbf{h}}_k)(\mathbf{h}_k - \widehat{\mathbf{h}}_k)^{\mathsf{H}}]$ *denote the corresponding estimation error covariance matrix. If* $T_{\mathsf{dl}} \geqslant s_k$ *there exist pilot matrices* $\boldsymbol{\Psi} \in \mathbb{C}^{T_{\mathsf{dl}} \times M}$ *for which* $\lim_{N_0 \downarrow 0} \mathsf{tr}(\mathbf{R}_e) = 0$ *for all support sets* $\mathcal{S}_k : |\mathcal{S}_k| = s_k$. *Conversely, for any support set* $\mathcal{S}_k : |\mathcal{S}_k| = s_k$ *any pilot matrix* $\boldsymbol{\Psi} \in \mathbb{C}^{T_{\mathsf{dl}} \times M}$ *with* $T_{\mathsf{dl}} < s_k$ *yields* $\lim_{N_0 \downarrow 0} \mathsf{tr}(\mathbf{R}_e) > 0$. $\quad\quad\square$

Proof. See appendix 4.6.1. $\quad\quad\square$

As a direct consequence of Lemma 4.1, we have that any scheme relying on intrinsic channel sparsity cannot yield stable estimation if $T_{\mathsf{dl}} < s_k$ for some users $k \in \mathbb{Z}_K$. Furthermore, we need to impose that the effective channel sparsity (after the introduction of the sparsifying precoder $\mathbf{B}$) is less than or equal to the desired DL pilot dimension T_{dl}. It is important to note that the requirement of estimation stability is *essential* in order to achieve high spectral efficiency in high SNR conditions, irrespective of the DL precoding scheme. In fact, if the estimation mean squared error (MSE) of the user channels does not vanish as $N_0 \downarrow 0$, the system self-interference due to the imperfect channel knowledge grows proportionally to the signal power, yielding a Signal-to-Interference plus Noise Ratio (SINR) that saturates to a constant when SNR becomes large. Hence, for sufficiently high

---

[2]Note that this coincides with (4.3) with $\mathbf{B} = \mathbf{I}_M$, i.e., without the sparsifying precoder.

SNR, the best strategy would consist of transmitting just a single data stream, since any form of multiuser precoding would inevitably lead to an interference limited regime, where the sum-rate remains bounded while SNR $\rightarrow \infty$ [33]. Conversely, it is also well-known that when the channel estimation error vanishes as $O(N_0)$ for $N_0 \downarrow 0$, the high-SNR sum rate behaves as if the channel was perfectly known and can be achieved by very simple linear precoding [19]. A possible solution to this problem consists of serving only the users whose channel support s_k is not larger than T_{dl}. This is assumed *implicitly* in all CS-based schemes, and represents a major intrinsic limitation of the CS-based approaches. In contrast, by artificially sparsifying the user channels, we manage to serve all users given a fixed DL pilot dimension T_{dl}.

4.3.2 Sparsifying Precoder Design

Before proceeding to discuss the design of the sparsifying precoder, let us introduce some graph-theoretic terms [38]. A bipartite graph is a graph whose vertices (nodes) can be divided into two sets $\mathcal{V}_1$ and $\mathcal{V}_2$, such that every edge in the set of graph edges $\mathcal{E}$ connects a vertex in $\mathcal{V}_1$ to one in $\mathcal{V}_2$. We denote such a graph by $\mathcal{L} = (\mathcal{V}_1, \mathcal{V}_2, \mathcal{E})$. A subgraph of $\mathcal{L}$ is a graph $\mathcal{L}' = (\mathcal{V}_1', \mathcal{V}_2', \mathcal{E}')$ such that $\mathcal{V}_1' \subseteq \mathcal{V}_1$, $\mathcal{V}_2' \subseteq \mathcal{V}_2$ and $\mathcal{E}' \subseteq \mathcal{E}$. With regards to $\mathcal{L}$, the following terms shall be defined and later used.

- **Degree of a vertex:** for a vertex $x \in \mathcal{V}_1 \cup \mathcal{V}_2$, the degree of x refers to the number of edges in $\mathcal{E}$ incident to x and is denoted by $\deg_{\mathcal{L}}(x)$.

- **Neighbors of a vertex:** the neighbors of a vertex $x \in \mathcal{V}_1 \cup \mathcal{V}_2$ are the set of vertices $y \in \mathcal{V}_1 \cup \mathcal{V}_2$ connected to x. This set is denoted by $\mathcal{N}_{\mathcal{L}}(x)$.

- **Matching:** a matching in $\mathcal{L}$ is a subset of edges in $\mathcal{E}$ without common vertices.

- **Maximal matching:** a maximal matching $\mathcal{M}$ of $\mathcal{L}$ is a matching with the property that if any edge outside $\mathcal{M}$ and in $\mathcal{E}$ is added to it, it is no longer a matching.

- **Perfect matching:** a perfect matching in $\mathcal{L}$ is a matching that covers all vertices of $\mathcal{L}$.

We propose to design the sparsifying precoder using a graphical model, where a bipartite graph is formed by a set of vertices representing users on one side and another set of vertices representing beams (the DFT columns) on the other side. An edge of the bipartite graph between a beam and a user represents the presence of that beam in the user's channel angular profile, with its weight denoting the user channel power along that beam. Now, we wish to design the precoder $\mathbf{B}$ such that the support of every effective user channel vector $\mathbf{h}_k^{\mathsf{eff}} = \mathbf{B}\mathbf{h}_k$ is not larger than T_{dl}, such that all users have a chance of being served.

Let $\check{\mathbf{H}} = \mathbf{L} \odot \mathbf{G} \in \mathbb{C}^{M \times K}$ denote the matrix of DL channel coefficients expressed in the DFT basis (4.2), in which each column of $\check{\mathbf{H}} = [\check{\mathbf{h}}_0, \ldots, \check{\mathbf{h}}_{K-1}]$ represents the coefficients vector of a user ($\mathbf{h}_k = \mathbf{F}\check{\mathbf{h}}_k$, $k \in \mathbb{Z}_K$), $\mathbf{L}$ is a $M \times K$ matrix with elements $[\mathbf{L}]_{m,k} = \sqrt{[\boldsymbol{\lambda}_k]_m}$, $\mathbf{G} \in \mathbb{C}^{M \times K}$ has i.i.d. elements $[\mathbf{G}]_{m,k} = g_{k,m} \sim \mathcal{CN}(0,1)$, and where $\odot$ denotes the Hadamard (elementwise) product. Let $\mathbf{A}$ denote a one-bit thresholded version of $\mathbf{L}$, such that $[\mathbf{A}]_{m,k} = 1$ if $[\boldsymbol{\lambda}_k]_m \geqslant \mathsf{th}$ and $[\mathbf{A}]_{m,k} = 0$, if $[\boldsymbol{\lambda}_k]_m < \mathsf{th}$, for some suitable small threshold $\mathsf{th} > 0$, used to identify the significant angular coefficients, and consider the $M \times K$ bipartite graph $\mathcal{L} = (\mathcal{A}, \mathcal{K}, \mathcal{E})$ with adjacency matrix $\mathbf{A}$ and weights $w_{m,k} = [\boldsymbol{\lambda}_k]_m$ on the edges $(m, k) \in \mathcal{E}$.

Given a pilot dimension T_{dl}, our goal consists in selecting a subgraph $\mathcal{L}' = (\mathcal{A}', \mathcal{K}', \mathcal{E}')$ of $\mathcal{L}$ in which each node on either side of the graph consists of a degree at least 1 and such that:

1. For all $k \in \mathcal{K}'$ we have $\deg_{\mathcal{L}'}(k) \leqslant T_{\mathsf{dl}}$, where $\deg_{\mathcal{L}'}$ denotes the degree of a node in the selected subgraph.

2. The sum of weights of the edges incident to any node $k \in \mathcal{K}'$ in the subgraph $\mathcal{L}'$ is greater than a threshold, i.e. $\sum_{m \in \mathcal{N}_{\mathcal{L}'}(k)} w_{m,k} \geqslant P_{\mathsf{th}}$, $\forall k \in \mathcal{K}'$.

3. The channel matrix $\check{\mathbf{H}}_{\mathcal{A}',\mathcal{K}'}$ obtained from $\check{\mathbf{H}}$ by selecting $a \in \mathcal{A}'$ ("selected beam directions") and $k \in \mathcal{K}'$ ("selected users") has large rank.

The first criterion enables stable estimation of the effective channel of any selected user with only T_{dl} common pilot dimensions and T_{dl} complex symbols of feedback per selected user. The second criterion makes sure that the effective channel strength of any selected user is greater than a desired threshold, since we do not want to spend resources on probing and serving users with weak effective channels (where "weak" is quantitatively determined by the value of P_{th}). Therefore P_{th} is a parameter that serves to obtain a trade-off between the rank of the effective matrix (which ultimately determines the number of spatially multiplexed DL data streams) and the beamforming gain (i.e., the power effectively conveyed along each selected user effective channel). The third criterion is motivated by the fact that the DL pre-log factor is given by $\mathrm{rank}(\check{\mathbf{H}}_{\mathcal{A}',\mathcal{K}'}) \times \max\{0, 1 - T_{\mathsf{dl}}/T\}$, and it is obtained by serving a number of users equal to the rank of the effective channel matrix. The following lemmas relate the rank of the effective channel matrix to a graph-theoretic quantity, namely, the size of the maximal matching.

Lemma 4.2. *[Skeleton or "$\mathbf{CUR}$" decomposition] Consider $\check{\mathbf{H}} \in \mathbb{C}^{M \times K}$, of rank r. Let $\mathbf{Q}$ be an $r \times r$ non-singular intersection submatrix obtained by selecting r rows and r columns of $\check{\mathbf{H}}$. Then, we have $\check{\mathbf{H}} = \mathbf{CUR}$, where $\mathbf{C} \in \mathbb{C}^{M \times r}$ and $\mathbf{R} \in \mathbb{C}^{r \times K}$ are the matrices of the selected columns and rows forming the intersection $\mathbf{Q}$ and $\mathbf{U} = \mathbf{Q}^{-1}$.* □
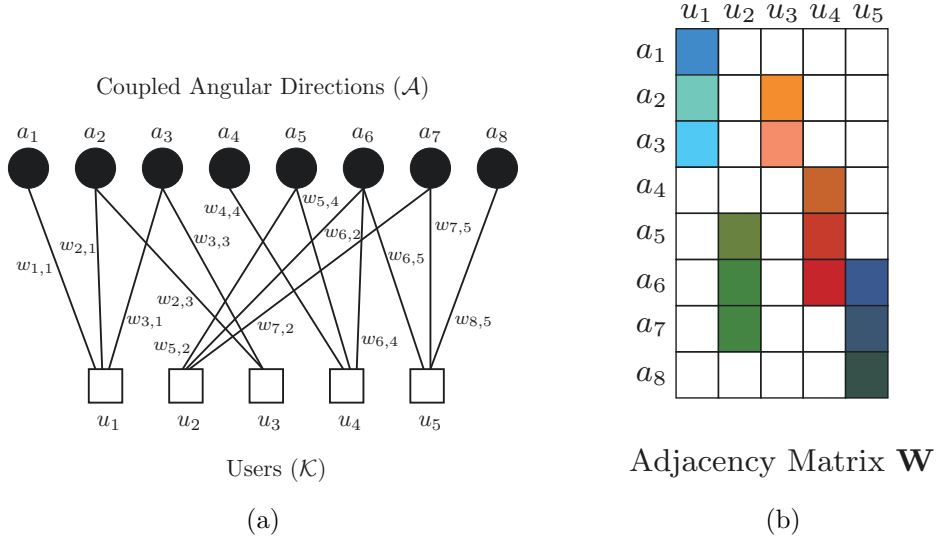
Figure 4.2: (a) An example bipartite graph $\mathcal{L}$. (b) The corresponding weighted adjacency matrix $\mathbf{W}$.

Proof. See [47]. □

Lemma 4.3. *[Rank and perfect matchings] Let $\mathbf{Q}$ denote an $r \times r$ matrix with some elements identically zero, and the non-identically zero elements independently drawn from a continuous distribution. Consider the associated bipartite graph with adjacency matrix $\mathbf{A}$ such that $\mathbf{A}_{i,j} = 1$ if $\mathbf{Q}_{i,j}$ is not identically zero, and $\mathbf{A}_{i,j} = 0$ otherwise. Then, $\mathbf{Q}$ has rank r with probability 1 iff the associated bipartite graph contains a perfect matching.* ▫

Proof. A similar theorem can be found in [124], but we provide a direct proof in Appendix 4.6.2 for the sake of completeness. □

Lemmas 4.2 and 4.3 result in the following corollary.

Corollary 4.1. *The rank r of a random matrix $\check{\mathbf{H}} \in \mathbb{C}^{M \times K}$ with either identically zero elements or elements independently drawn from a continuous distribution is given, with probability 1, by the size of the largest intersection submatrix whose associated bipartite graph (defined as in Lemma 4.3) contains a perfect matching.* ▫

This corollary holds in our case where the non-zero elements of $\check{\mathbf{H}}$ are drawn from the complex Gaussian distribution. Using Corollary 4.1 this problem can be formulated as:

Problem 2. *Let T_{dl} denote the available DL pilot dimension and let $\mathcal{M}(\mathcal{A}', \mathcal{K}')$ denote a matching of the subgraph $\mathcal{L}'(\mathcal{A}', \mathcal{K}', \mathcal{E}')$ of the bipartite graph $\mathcal{L}(\mathcal{A}, \mathcal{K}, \mathcal{E})$. Find the solution of the following optimization problem:*

$$\underset{\mathcal{A}' \subseteq \mathcal{A}, \mathcal{K}' \subseteq \mathcal{K}}{maximize} \quad \left| \mathcal{M}\left(\mathcal{A}', \mathcal{K}'\right) \right| \tag{4.5a}$$

$$\text{subject to} \qquad deg_{\mathcal{L}'}(k) \leqslant T_{\mathsf{dl}} \; \forall k \in \mathcal{K}', \tag{4.5b}$$

$$\sum_{a \in \mathcal{N}_{\mathcal{L}'}(k)} w_{a,k} \geqslant P_{\mathsf{th}}, \; \forall k \in \mathcal{K}'. \tag{4.5c}$$

$$\diamondsuit$$

The following theorem shows that Problem 2 can be solved in a tractable way.

Theorem 4.1. *The optimization problem in (4.5) is equivalent to the mixed integer linear program (MILP) below:*

$$\underset{x_m, y_k, z_{m,k}}{maximize} \qquad \sum_{m \in \mathcal{A}} \sum_{k \in \mathcal{K}} z_{m,k} \tag{4.6a}$$

$$\text{subject to} \qquad z_{m,k} \leqslant [\mathbf{A}]_{m,k} \; \forall m \in \mathcal{A}, k \in \mathcal{K}, \tag{4.6b}$$

$$\sum_{k \in \mathcal{K}} z_{m,k} \leqslant x_m \; \forall m \in \mathcal{A}, \tag{4.6c}$$

$$\sum_{m \in \mathcal{A}} z_{m,k} \leqslant y_k \; \forall k \in \mathcal{K}, \tag{4.6d}$$

$$\sum_{m \in \mathcal{A}} [\mathbf{A}]_{m,k} x_m \leqslant T_{\mathsf{dl}} y_k + M(1 - y_k) \; \forall k \in \mathcal{K} \tag{4.6e}$$

$$P_{\mathsf{th}} \, y_k \leqslant \sum_{m \in \mathcal{A}} [\mathbf{W}]_{m,k} x_m \; \forall k \in \mathcal{K}, \tag{4.6f}$$

$$x_m \leqslant \sum_{k \in \mathcal{K}} [\mathbf{A}]_{m,k} y_k \; \forall m \in \mathcal{A}, \tag{4.6g}$$

$$x_m, y_k \in \{0, 1\} \; \forall a \in \mathcal{A}, k \in \mathcal{K}, \tag{4.6h}$$

$$z_{m,k} \in [0, 1] \; \forall m \in \mathcal{A}, k \in \mathcal{K}, \tag{4.6i}$$

where $\mathbf{W}$ *is the* $|\mathcal{A}| \times |\mathcal{K}|$ *weighted adjacency matrix in which* $[\mathbf{W}]_{m,k} = w_{m,k}$. *The solution sub-graph is given by the set of nodes* $\mathcal{A}' = \{m : x_m^* = 1\}$ *and* $\mathcal{K}' = \{k : y_k^\star = 1\}$, *with* $\{x_m^\star\}_{m=0}^{M-1}$ *and* $\{y_k^\star\}_{k=0}^{K-1}$ *being a solution of* (4.6). □

Proof. See Appendix 4.6.3. ☐

The solution to this optimization, however, is not necessarily unique, i.e. there may exist several sub-graphs with the same (maximum) matching size. In order to limit the solution set we introduce a regularization term to the objective of (4.6) to favor solutions containing more active beams. The regularized form of (4.6) is given as

$$\underset{x_m, y_k, z_{m,k}}{maximize} \qquad \sum_{m \in \mathcal{A}} \sum_{k \in \mathcal{K}} z_{m,k} + \epsilon \sum_{m \in \mathcal{A}} x_m$$
$$\text{subject to} \quad \{x_m, y_k, z_{m,k}\}_{m \in \mathcal{A}, k \in \mathcal{K}} \in \mathcal{S}_{\text{feasible}}, \tag{$\mathcal{P}_{\text{MILP}}$}$$

where the feasibility set $\mathcal{S}_{\text{feasible}}$ encodes the constraints (4.6a)-(4.6i). Here the regularization factor ϵ is chosen to be a small positive value such that it does not effect the matching size of the solution sub-graph. In fact choosing $\epsilon < \frac{1}{M}$ ensures this, since then $\epsilon \sum_{m \in \mathcal{A}} x_m < 1$ and a solution to $\mathcal{P}_{\text{MILP}}$ must have the same matching size as a solution to 4.6, otherwise the objective of $\mathcal{P}_{\text{MILP}}$ can be improved by choosing a solution with a larger matching size. The introduced MILP can be efficiently solved using an off-the-shelf optimization toolbox. In the simulation results of this chapter, we have used the MATLAB intlinprog, which adopts a Branch and Bound method to find the solution to an MILP [85].

4.3.3 Channel Estimation and Multiuser Precoding

For a given set of user DL covariance matrices, let $\{x_m^\star\}_m$ and $\{y_k^\star\}_k$ denote the MILP solution and denote by $\mathcal{B} = \{m : x_m^* = 1\} = \{m_1, m_2, \ldots, m_{M'}\}$ the set of selected beam directions of cardinality $|\mathcal{B}| = M'$ and by $\mathcal{K} = \{k : y_k^\star = 1\}$ the set of selected users of cardinality $|\mathcal{K}| = K'$. The resulting sparsifying precoding matrix $\mathbf{B}$ in (4.3) is simply obtained as

$$\mathbf{B} = \mathbf{F}_{\mathcal{B}}^{\mathsf{H}}, \tag{4.7}$$

where $\mathbf{F}_{\mathcal{B}} = [\mathbf{f}_{m_1}, \ldots, \mathbf{f}_{m_{M'}}]$ and $\mathbf{f}_m$ denotes the m-th column of the $M \times M$ unitary DFT matrix $\mathbf{F}$. Given a DFT column $\mathbf{f}_m$, we have

$$\mathbf{B}\mathbf{f}_m = \begin{cases} \mathbf{0} & \text{if } m \notin \mathcal{B} \\ \mathbf{u}_i & \text{if } m = m_i \in \mathcal{B} \end{cases}$$

where $\mathbf{u}_i$ denotes a $M' \times 1$ vector with all zero components but a single "1" in the i-th position. Using the above property and (4.2), the effective DL channel vectors take on the form

$$\mathbf{h}_k^{\text{eff}} = \mathbf{B} \sum_{m \in \mathcal{S}_k} g_{k,m} \sqrt{[\boldsymbol{\lambda}_k]_m} \mathbf{f}_m = \sum_{i : m_i \in \mathcal{B} \bigcap \mathcal{S}_k} \sqrt{[\boldsymbol{\lambda}_k]_{m_i}} g_{k,m} \mathbf{u}_i. \tag{4.8}$$

In words, the effective channel of user k is a vector with non-identically zero elements only at the positions corresponding to the intersection of the beam directions in $\mathcal{S}_k$, along which the physical channel of user k carries positive energy, and in $\mathcal{B}$, selected by the sparsifying precoder. The non-identically zero elements are independent Gaussian coefficients $\sim \mathcal{CN}(0, [\boldsymbol{\lambda}_k]_{m_i})$. Notice also that, by construction, the number of non-identically zero coefficients are $|\mathcal{B} \bigcap \mathcal{S}_k| \leq T_{\text{dl}}$ and their positions (encoded in the vectors $\mathbf{u}_i$ in (4.8)), plus an estimate of their variances $[\boldsymbol{\lambda}_k]_{m_i}$ are known to the BS. Hence, the effective channel vectors can be estimated from the T_{dl}-dimensional DL pilot observation (4.3) with an estimation MSE that vanishes as 1/SNR. The pilot observation in the form (4.3) is obtained at the user k receiver. In this work, we assume that each user sends its pilot observa-

tions using T_{dl} channel uses in the UL, using analog unquantized feedback, as analyzed for example in [19, 69]. At the BS receiver, after estimating the UL channel from the UL pilots, the BS can apply linear MMSE estimation and recovers the channel state feedback which takes on the same form as (4.3) with some additional noise due to the noisy UL transmission.

Remark 4.1. As an alternative, one can consider quantized feedback using T_{dl} channel uses in the UL (see [19, 69] and references therein). Digital quantized feedback yields generally a better end-to-end estimation MSE in the absence of feedback errors. However, the effect of decoding errors on the channel state feedback is difficult to characterize in a simple manner since it depends on the specific joint source-channel coding scheme employed. Hence, in this work we restrict to the simple analog feedback.

With the above precoding, we have $\mathbf{B}\mathbf{B}^{\mathsf{H}} = \mathbf{I}_{M'}$. Also, we can choose the DL pilot matrix $\boldsymbol{\Psi}$ to be proportional to a random unitary matrix of dimension $T_{\mathsf{dl}} \times M'$, such that $\boldsymbol{\Psi}\boldsymbol{\Psi}^{\mathsf{H}} = P_{\mathsf{dl}}\mathbf{I}_{T_{\mathsf{dl}}}$. In this way, the DL pilot phase power constraint (4.4) is automatically satisfied. The estimation of $\mathbf{h}_k^{\mathsf{eff}}$ from the DL pilot observations (4.3) (with suitably increased AWGN variance due to the noisy UL feedback) is completely straightforward and shall not be treated here in details.

For the sake of completeness, we conclude this section with the DL precoded data phase and the corresponding sum rate performance metric that we shall later use for numerical analysis and comparison with other schemes. Let $\widehat{\mathbf{H}}_{\mathrm{eff}} = [\widehat{\mathbf{h}}_{\mathrm{eff}}^{(1)}, \dots, \widehat{\mathbf{h}}_{\mathrm{eff}}^{(K')}]$ be the matrix of the estimated effective DL channels for the selected users. We consider the ZF beamforming matrix $\mathbf{V}$ given by the column-normalized version of the Moore-Penrose pseudoinverse of the estimated channel matrix, i.e.

$$\mathbf{V} = \left(\widehat{\mathbf{H}}_{\mathrm{eff}}\right)^{\dagger} \mathbf{J}^{1/2},$$

where $\left(\widehat{\mathbf{H}}_{\mathrm{eff}}\right)^{\dagger} = \widehat{\mathbf{H}}_{\mathrm{eff}} \left(\widehat{\mathbf{H}}_{\mathrm{eff}}^{\mathsf{H}}\widehat{\mathbf{H}}_{\mathrm{eff}}\right)^{-1}$ and $\mathbf{J}$ is a diagonal matrix that makes the columns of $\mathbf{V}$ to have unit norm. A channel use of the DL precoded data transmission phase at the k-th user receiver takes on the form

$$r_k = (\mathbf{h}_k)^{\mathsf{H}} \mathbf{B}^{\mathsf{H}} \mathbf{V} \mathbf{P}^{1/2} \mathbf{d} + n_k, \tag{4.9}$$

where $\mathbf{d} \in \mathbb{C}^{K' \times 1}$ is a vector of unit-energy user data symbols and $\mathbf{P}$ is a diagonal matrix defining the power allocation to the DL data streams. The transmit power constraint is given by $\mathsf{tr}(\mathbf{B}^{\mathsf{H}}\mathbf{V}\mathbf{P}\mathbf{V}^{\mathsf{H}}\mathbf{B}) = \mathsf{tr}(\mathbf{V}^{\mathsf{H}}\mathbf{V}\mathbf{P}) = \mathsf{tr}(\mathbf{P}) = P_{\mathsf{dl}}$, where we used $\mathbf{B}\mathbf{B}^{\mathsf{H}} = \mathbf{I}_{M'}$ and the fact that $\mathbf{V}^{\mathsf{H}}\mathbf{V}$ has unit diagonal elements by construction. In particular, in the simulation results section we use the simple uniform power allocation $P_k = P_{\mathsf{dl}}/K'$ to each k-th user

data stream. In the case of perfect ZF beamforming, i.e., for $\widehat{\mathbf{H}}_{\text{eff}} = \mathbf{H}_{\text{eff}}$, we have that (4.9) reduces to $r_k = \sqrt{J_k P_k} d_k + n_k$, where J_k is the k-th diagonal element of the norm normalizing matrix $\mathbf{J}$, P_k is the k-th diagonal element of the power allocation matrix $\mathbf{P}$, and d_k is the k-th user data symbol. Since in general $\widehat{\mathbf{H}}_{\text{eff}} \neq \mathbf{H}_{\text{eff}}$, due to non-zero estimation error, the received symbol at user k receiver is given by $r_k = b_{k,k} d_k + \sum_{k' \neq k} b_{k,k'} d_{k'} + n_k$, where the coefficients $(b_{k,1}, \ldots, b_{k,K'})$ are given by the elements of the $1 \times K'$ row vector $(\mathbf{h}_k)^{\mathsf{H}} \mathbf{B}^{\mathsf{H}} \mathbf{V} \mathbf{P}^{1/2}$ in (4.9). Of course, in the presence of an accurate channel estimation we expect that $b_{k,k} \approx \sqrt{J_k P_k}$ and $b_{k,k'} \approx 0$ for $k' \neq k$. For simplicity, in this chapter we compare the performance of the proposed scheme with that of the state-of-the-art CS-based scheme in terms of ergodic sum rate, assuming that all coefficients $(b_{k,1}, \ldots, b_{k,K'})$ are known to the corresponding receiver k. Including the DL training overhead, this yields the rate expression (see [18])

$$R_{\text{sum}} = \left(1 - \frac{T_{\text{dl}}}{T}\right) \sum_{k \in \mathcal{K}} \mathbb{E}\left[\log\left(1 + \frac{|b_{k,k}|^2}{1 + \sum_{k' \neq k} |b_{k,k'}|^2}\right)\right]. \tag{4.10}$$

4.4 Extension to Arbitrary Array Geometries

A necessary step before performing sparsification is that all of the estimated DL covariance matrices share the same (approximate) eigenbasis, namely a set of common eigenvectors that represent the beam space. As discussed earlier in this chapter, for a massive ULA, this common eigenbasis is given by the DFT basis due to an application of Szegö's theorem for large Hermitian Toeplitz matrices (see Section 4.2.1). For an arbitrary array geometry, the covariance is not necessarily Toeplitz. We are not aware of any work suggesting the existence of an (approximate) common eigenbasis for MIMO covariances of generic array geometries. In fact, one has reasons to believe the contrary, that is, in general a common eigenbasis does not exist, even when the channel dimension grows infinitely large ($M \to \infty$). Then, what is a suitable strategy to obtain an *approximate* common eigenbasis? In order to extend the ACS method to arrays with arbitrary design, here we propose a method for obtaining an approximate common eigenbasis among the channel covariances of a multi-user system with an arbitrary array geometry.

4.4.1 Beam-Space Design for Arbitrary Array Geometries

Consider a BS equipped with an array of arbitrary geometry consisting of M antennas, that communicating with K users. Suppose the BS collects N noisy pilot samples per user as

$$\mathbf{y}_k[i] = \mathbf{h}_k[i] + \mathbf{z}_k[i], \ i \in \mathbb{Z}_N, \tag{4.11}$$

where $\mathbf{h}_k[i]$ is the i-th realization of user k's channel vector, and $\mathbf{z}_k[i] \sim \mathcal{CN}(\mathbf{0}, N_0\mathbf{I}_M)$ is a random realization of the AWGN. The channels are all assumed zero-mean, complex Gaussian with covariance $\mathbf{\Sigma}_k = \mathbb{E}\left[\mathbf{h}_k\mathbf{h}_k^{\mathsf{H}}\right]$, $k \in \mathbb{Z}_K$. Define the eigendecomposition of $\mathbf{\Sigma}_k$ as $\mathbf{\Sigma}_k = \mathbf{U}_k\mathbf{\Lambda}_k\mathbf{U}_k^{\mathsf{H}}$, where $\mathbf{U}_k$ is the unitary matrix of eigenvectors ($\mathbf{U}_k^{\mathsf{H}}\mathbf{U}_k = \mathbf{I}_M$) and $\mathbf{\Lambda}_k$ is the diagonal matrix of non-negative eigenvalues. We note that the eigenbasis of distinct covariances are generally different. This makes the joint processing of the channels and the precoding design highly difficult. Hence, we are interested in obtaining an approximate common eigenbasis among all covariances $\{\mathbf{\Sigma}_k\}_k$, given the noisy pilot samples $\{\mathbf{y}_k[i]\}_{k,i}$. If the covariances are jointly diagonalizable, i.e. if there exists a unitary matrix $\mathbf{U}^c$ such that $\mathbf{U}_1 = \mathbf{U}_2 = \ldots = \mathbf{U}_K = \mathbf{U}^c$, then it is desirable to obtain $\mathbf{U}^c$ as the common eigenbasis. If the covariances are *not* jointly diagonalizable, then we want to obtain a unitary matrix $\mathbf{U}^{\star}$ that "best" diagonalizes the covariances.

We introduce a parametric decomposition of the channel covariances that imposes the joint diagonalizability criterion on the estimation model, in which each covariance is decomposed as

$$\mathbf{\Sigma}_k = \mathbf{U}\mathbf{\Lambda}_k\mathbf{U}^{\mathsf{H}}, \tag{4.12}$$

where $\mathbf{U} = [\mathbf{u}_0, \ldots, \mathbf{u}_{M-1}] \in \mathbb{C}^{M \times M}$ is a unitary matrix ($\mathbf{U}^{\mathsf{H}}\mathbf{U} = \mathbf{I}_M$) representing the to-be-estimated common eigenbasis and $\mathbf{\Lambda}_k = \mathrm{diag}(\boldsymbol{\lambda}_k)$ is a diagonal matrix with non-negative diagonal elements given as the vector $\boldsymbol{\lambda}_k$ for $k \in \mathbb{Z}_K$. Both $\mathbf{U}$ and $\{\boldsymbol{\lambda}_k\}_k$ are to be estimated. We propose a method that, given noisy pilot samples, estimates these parameters by maximizing the associated likelihood function.

The Likelihood Function

Define the matrices $\mathbf{Y}_k := [\mathbf{y}_k[0], \ldots, \mathbf{y}_k[N-1]] \in \mathbb{C}^{M \times N}$, where $\mathbf{Y}_k$ includes pilot samples of user k as its columns. Given $\{\mathbf{y}_k[i]\}_{k,i}$, one can write the likelihood function as

$$
\begin{aligned}
&p\left(\{\mathbf{y}_k[i]\}_{k,i}|\{\mathbf{\Sigma}_k\}_k, N_0\right) = \\
&\prod_{k=0}^{K-1} \frac{1}{(2\pi)^{\frac{M}{2}} \det(\mathbf{\Sigma}_k + N_0\mathbf{I})^{\frac{N}{2}}} \exp\left(-\frac{1}{2N}\mathrm{trace}(\mathbf{Y}_k^{\mathsf{H}}(\mathbf{\Sigma}_k + N_0\mathbf{I})^{-1}\mathbf{Y}_k)\right) = \\
&\prod_{k=0}^{K-1} \frac{1}{(2\pi)^{\frac{M}{2}} \det(\mathbf{\Sigma}_k + N_0\mathbf{I})^{\frac{N}{2}}} \exp\left(-\frac{1}{2}\mathrm{trace}\left((\mathbf{\Sigma}_k + N_0\mathbf{I})^{-1}\hat{\mathbf{\Sigma}}_{\mathbf{y},k}\right)\right),
\end{aligned}
\tag{4.13}
$$

where we have defined $\hat{\mathbf{\Sigma}}_{\mathbf{y},k} := \frac{1}{N}\mathbf{Y}_k\mathbf{Y}_k^{\mathsf{H}}$, as the sample covariance of the k-th pilot sample set. Taking the $-\log(\cdot)$ of the likelihood function, scaling it, omitting constant terms, and replacing $\mathbf{\Sigma}_k$ with $\mathbf{U}\mathbf{\Lambda}_k\mathbf{U}^{\mathsf{H}}$ from (4.12), one can show that maximizing the likelihood function is equivalent to minimizing the following ML cost as a function of the parameters

U and $\{\boldsymbol{\lambda}_k\}_k$:

$$\mathcal{C}\left(\mathbf{U}, \{\boldsymbol{\lambda}_k\}_k\right) := \sum_{k=1}^{K} \log \det \left(\mathbf{U}\mathrm{diag}(\boldsymbol{\lambda}_k)\mathbf{U}^{\mathsf{H}} + N_0\mathbf{I}\right) + \mathrm{trace}\left(\left(\mathbf{U}\mathrm{diag}(\boldsymbol{\lambda}_k)\mathbf{U}^{\mathsf{H}} + N_0\mathbf{I}\right)^{-1}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\right).$$
(4.14)

Since

$$\det\left(\mathbf{U}\mathrm{diag}(\boldsymbol{\lambda}_k)\mathbf{U}^{\mathsf{H}} + N_0\mathbf{I}\right) = \det\left(\mathrm{diag}(\boldsymbol{\lambda}_k + N_0\mathbf{1})\right),$$

and

$$\mathrm{tr}\left(\left(\mathbf{U}\mathrm{diag}(\boldsymbol{\lambda}_k)\mathbf{U}^{\mathsf{H}} + N_0\mathbf{I}\right)^{-1}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\right) = \mathrm{tr}\left(\mathrm{diag}\left(\boldsymbol{\lambda}_k + N_0\mathbf{1}\right)^{-1}\mathbf{U}^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{U}\right),$$

we have

$$\mathcal{C}\left(\mathbf{U}, \{\boldsymbol{\lambda}_k\}_k\right) = \sum_{m,k} \log(\lambda_{k,m} + N_0) + \frac{\mathbf{u}_m^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m}{\lambda_{k,m} + N_0},$$
(4.15)

where $\lambda_{k,m} > 0$ is the m-th element of $\boldsymbol{\lambda}_k$. The ML optimization problem then can be formulated as

$$\begin{aligned}
\underset{\{\mathbf{u}_m\}_m, \{\boldsymbol{\lambda}_k\}_k}{\text{minimize}} \quad & \sum_{m,k} \log(\lambda_{k,m} + N_0) + \frac{\mathbf{u}_m^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_k\mathbf{u}_m}{\lambda_{k,m} + N_0} \\
\text{subject to} \quad & \mathbf{u}_m^{\mathsf{H}}\mathbf{u}_n = \delta_{m,n}, \ m, n \in \mathbb{Z}_M.
\end{aligned}$$
(4.16)

It is easy to show that, for given $\{\mathbf{u}_m\}_m$, the function $g(x) = \log(x + N_0) + \frac{\mathbf{u}_m^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m}{x + N_0}$, $x \geqslant 0$ achieves its minimum at $x = \mathbf{u}_m^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m - N_0$. Therefore, we take the minimization in (4.16) first with respect to $\{\lambda_{k,m}\}_{k,m}$, and simplify (4.16) to

$$\begin{aligned}
\underset{\{\mathbf{u}_m\}_m}{\text{minimize}} \quad & f(\mathbf{U}) = \sum_{m,k} \log\left(\mathbf{u}_m^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m\right) \\
\text{subject to} \quad & \mathbf{u}_m^{\mathsf{H}}\mathbf{u}_n = \delta_{m,n}, \ m, n \in \mathbb{Z}_M.
\end{aligned}$$
$(\mathcal{P}_1)$

This presents an optimization problem over the manifold of unitary matrices $\mathcal{U} = \{\mathbf{U} \in \mathbb{C}^{M \times M} : \mathbf{U}^{\mathsf{H}}\mathbf{U} = \mathbf{I}_M\}$. To solve $\mathcal{P}_1$, we propose a gradient projection method and show that it converges to a stationary point of the cost function f.

Intermezzo: ML and Jointly Diagonalizable Covariances

Assume that the observation sample covariance has converged to its expectation, i.e. the true covariance: $\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k} \to \boldsymbol{\Sigma}$. One can show that, if the true channel covariances are jointly diagonalizable, then the global optimum of $\mathcal{P}_1$ is given by the common eigenbasis. To see this, first note that the channel covariance of user k can be decomposed as

$$\boldsymbol{\Sigma}_k = \mathbf{U}^c\boldsymbol{\Lambda}_k\mathbf{U}^{c\,\mathsf{H}},$$
(4.17)

for $k = 0, \ldots, K - 1$, where $\mathbf{U}^c \in \mathbb{C}^{M \times M}$, $\mathbf{U}^{c\,\mathsf{H}} \mathbf{U}^c = \mathbf{I}_M$ denotes the common eigenbasis.

Definition 4.1 (Majorization)**.** *For* $\mathbf{x} \in \mathbb{R}^M$*, define* $\mathbf{x}^{\downarrow}$ *as the vector containing the elements of* $\mathbf{x}$ *in descending order. Let* $\mathbf{y} \in \mathbb{R}^M$ *be another vector such that* $\sum_{i=0}^{M-1} \mathbf{x}_i = \sum_{i=0}^{M-1} \mathbf{y}_i$*. We say* $\mathbf{x}$ *majorizes* $\mathbf{y}$ *(*$\mathbf{x} > \mathbf{y}$*) iff*

$$\sum_{i=0}^{m} \mathbf{x}_i^{\downarrow} \geqslant \sum_{i=0}^{m} \mathbf{y}_i^{\downarrow},$$

for all $m \in \mathbb{Z}_M$*.*

We have the following theorem on global optimality of $\mathbf{U}^c$ for the ML optimization problem $(\mathcal{P}_1)$.

Theorem 4.2. *Let* $\mathbf{\Sigma}_k$*,* $k = 0, \ldots, K - 1$ *be a set of jointly diagonalizable covariance matrices as in* (4.17)*. Then* $\mathbf{U}^{\star} = \mathbf{U}^c$ *is a global optimum of the optimization problem*

$$\underset{\mathbf{U}}{minimize}\, f(\mathbf{U}) = \sum_{m,k} \log \left(\mathbf{u}_m^{\mathsf{H}} \left(\mathbf{\Sigma}_k + N_0 \mathbf{I}_M \right) \mathbf{u}_m \right) \ \ s.t. \ \mathbf{u}_m^{\mathsf{H}} \mathbf{u}_n = \delta_{m,n}. \tag{4.18}$$

Proof. For any unitary $\mathbf{U}$, define the vector $\boldsymbol{\sigma}_k(\mathbf{U}) \in \mathbb{R}^M$ where

$$[\boldsymbol{\sigma}_k(\mathbf{U})]_m = \mathbf{u}_m^{\mathsf{H}} \left(\mathbf{\Sigma}_k + N_0 \mathbf{I} \right) \mathbf{u}_m.$$

In particular $\boldsymbol{\sigma}_k(\mathbf{U}^c)$ is the vector of eigenvalues of $\mathbf{\Sigma}_k + N_0 \mathbf{I}$. Using the properties of eigenvalue decomposition one can show $\boldsymbol{\sigma}_k(\mathbf{U}^c) > \boldsymbol{\sigma}_k(\mathbf{U})$ for all $\mathbf{U} \in \mathcal{U}$ and all $k = 0, \ldots, K-1$. In addition, the function $h(\mathbf{x}) = \sum_i \log(\mathbf{x}_i)$ is Schur-concave [94] and therefore $\sum_m \log([\boldsymbol{\sigma}_k(\mathbf{U}^c)]_m) \leqslant \sum_m \log([\boldsymbol{\sigma}_k(\mathbf{U})]_m)$ for all k. Hence, $f(\mathbf{U}^c) \leqslant f(\mathbf{U})$ for all $\mathbf{U} \in \mathcal{U}$, proving $\mathbf{U}^c$ to be the global minimizer of f over $\mathcal{U}$. $\qquad \square$

This serves as a sanity check for the fact that the ML problem returns reasonable solutions. At least when the channel covariances in fact share a common eigenbasis, the ML optimizer coincides with it.

ML via Projected Gradient Descent

Now we turn to the solving the ML problem $\mathcal{P}_1$. We use a projected gradient descent (PGD) method to minimize the objective cost function f. The PGD is a well-known iterative optimization algorithm [9]. Starting from an initial point $\mathbf{U}^{(0)}$, this method consists of the two following steps per iteration:

$$\tilde{\mathbf{U}}^{(t)} = \mathbf{U}^{(t)} - \alpha_t \nabla f(\mathbf{U}^{(t)}) \qquad\qquad \text{(Gradient Step)}$$

$$\mathbf{U}^{(t+1)} = \mathcal{P}_{\mathcal{U}}(\tilde{\mathbf{U}}^{(t)}) \qquad\qquad \text{(Projection Step)}$$

where $\alpha_t > 0$ is a step size, $\nabla f(\mathbf{U}^{(t)}) \in \mathbb{C}^{M \times M}$ is the gradient of f at $\mathbf{U}^{(t)}$ and $\mathcal{P}_{\mathcal{U}} :$ $\mathbb{C}^{M \times M} \to \mathcal{U}$ denotes the orthogonal projection operator onto the set of unitary matrices. The explicit expression of this operator is given in [80]; Nevertheless, we derive it here through the following lemma for the sake of completeness.

Lemma 4.4. *Let* $\mathbf{V} \in \mathbb{C}^{M \times M}$ *be a matrix with singular value decomposition* $\mathbf{V} = \mathbf{SDT}^{\mathsf{H}}$ *where* $\mathbf{S}$ *and* $\mathbf{T}$ *are unitary matrices of left and right eigenvectors and* $\mathbf{D} = diag(\mathbf{d})$ *is non-negative diagonal. Then, the orthogonal projection of* $\mathbf{V}$ *onto the set of unitary matrices is given by* $\mathcal{P}_{\mathcal{U}}(\mathbf{V}) = \mathbf{ST}^{\mathsf{H}}$.

Proof. The orthogonal projection of $\mathbf{V}$ is given by the minimizer of $g(\mathbf{U}) = \|\mathbf{V} - \mathbf{U}\|_F^2$ over the set of unitary matrices, where $\mathbf{U}^{\mathsf{H}}\mathbf{U} = \mathbf{I}_M$. We can write

$$\begin{aligned} g(\mathbf{U}) = \|\mathbf{V} - \mathbf{U}\|_F^2 &= \|\mathbf{U}\|_{\mathsf{F}}^2 + \|\mathbf{V}\|_{\mathsf{F}}^2 - 2\mathbb{R}\{\langle \mathbf{V}, \mathbf{U} \rangle\} \\ &= M + \|\mathbf{V}\|_{\mathsf{F}}^2 - 2\mathbb{R}\{\langle \mathbf{V}, \mathbf{U} \rangle\}, \end{aligned} \qquad (4.19)$$

where the inner product is defined as $\langle \mathbf{V}, \mathbf{U} \rangle = \text{trace}(\mathbf{U}^{\mathsf{H}}\mathbf{V})$ and we used the fact that $\text{trace}(\mathbf{UU}^{\mathsf{H}}) = \text{trace}(\mathbf{I}_M) = M$. According to Von Neumann's trace inequality we have $|\langle \mathbf{V}, \mathbf{U} \rangle| = |\text{trace}(\mathbf{U}^{\mathsf{H}}\mathbf{V})| \leqslant \langle \mathbf{s_U}, \mathbf{d} \rangle$, [88], where $\mathbf{s}$ denotes the singular values vector of $\mathbf{U}$. In the special case where $\mathbf{U}$ is unitary, we have $\mathbf{s_U} = [1, \ldots, 1]^{\mathsf{T}}$ and $|\langle \mathbf{V}, \mathbf{U} \rangle| \leqslant \langle \mathbf{s_U}, \mathbf{d} \rangle = \sum_i \mathbf{d}_i$. Now, using (4.19) we have $g(\mathbf{U}) \geqslant M + \|\mathbf{V}\|_{\mathsf{F}}^2 - 2\sum_i \mathbf{d}_i$, where the RHS of the inequality is independent of $\mathbf{U}$. We show that the lower bound on $g(\mathbf{U})$ is achieved by $\mathbf{U}^{\star} = \mathbf{ST}^{\mathsf{H}}$. This is seen by the fact that

$$\begin{aligned} g(\mathbf{ST}^{\mathsf{H}}) &= M + \|\mathbf{SDT}^{\mathsf{H}}\|_{\mathsf{F}}^2 - 2\mathbb{R}\{\langle \mathbf{SDT}^{\mathsf{H}}, \mathbf{ST}^{\mathsf{H}} \rangle\} \\ &= M + \|\mathbf{SDT}^{\mathsf{H}}\|_{\mathsf{F}}^2 - 2\text{trace}(\mathbf{D}) \\ &= M + \|\mathbf{V}\|_{\mathsf{F}}^2 - 2\sum_i \mathbf{d}_i, . \end{aligned} \qquad (4.20)$$

This completes the proof. □

This lemma provides an explicit formula for the orthogonal projection of a matrix into the unitary set as

$$\mathcal{P}_{\mathcal{U}} : \mathbb{C}^{M \times M} \to \mathcal{U}, \ \mathbf{V} = \mathbf{SDT}^{\mathsf{H}} \to \mathbf{ST}^{\mathsf{H}}. \qquad (4.21)$$

The following theorem presents the main result of this section.

Theorem 4.3. *Let* $\mathbf{U}^{(0)} \in \mathcal{U}$ *be an initial point and consider the gradient projection update rule*

$$\mathbf{U}^{(t+1)} = \mathcal{P}_{\mathcal{U}}\left(\mathbf{U}^{(t)} - \alpha_t \nabla f(\mathbf{U}^{(t)})\right), \ t = 0, 1, \ldots, \qquad (4.22)$$

with $\alpha_t \in (0, \frac{1}{L})$ *for all* t, *where* L *is the Lipschitz constant of* $\nabla f(\mathbf{U})$. *Then the sequence* $\{\mathbf{U}^{(t)}, t = 0, 1, \ldots\}$ *converges to a stationary point of* $f(\mathbf{U})$.

In order to prove Theorem 4.3, we need to first prove some useful properties of the ML optimization problem.

Lipschitz Continuity of the Cost Gradient

As a first step, we prove that the cost gradient $\nabla f(\mathbf{U})$ is Lipschitz continuous over $\mathcal{U}$. Note that the manifold $\mathcal{U}$ is a subset of the closed convex ball $\mathcal{B}$ ($\mathcal{U} \subset \mathcal{B}$) where $\mathcal{B} = \{\mathbf{U} : \|\mathbf{U}\|_{\mathsf{F}} \leq \sqrt{M}\}$. One can show that $f(\mathbf{U})$ has Lipschitz continuous gradient over $\mathcal{B}$, i.e. there exists a constant L, such that $\|\nabla f(\mathbf{U}) - \nabla f(\mathbf{U}')\|_{\mathsf{F}} \leq L\|\mathbf{U} - \mathbf{U}'\|_{\mathsf{F}}$, for all $\mathbf{U}, \mathbf{U}' \in \mathcal{B}$. One way to prove this is by showing that the Hessian of $f(\mathbf{U})$ has bounded operator norm over $\mathcal{B}$. Define the complex Hessian as the $M^2 \times M^2$ square matrix $\nabla^2 f(\mathbf{U})$ whose elements are given as [50]

$$[\nabla^2 f(\mathbf{U})]_{m,n} = \frac{\partial^2 f(\mathbf{U})}{\partial[\text{vec}(\mathbf{U})]_m \partial[\text{vec}(\mathbf{U})]_n^*}, \tag{4.23}$$

for $m, n \in \mathbb{Z}_{M^2}$, where $\text{vec}(\mathbf{U}) = [\mathbf{u}_1^{\mathsf{T}}, \ldots, \mathbf{u}_M^{\mathsf{T}}]^{\mathsf{T}}$ is the vectorized version of $\mathbf{U}$. Some calculation shows that the Hessian is a block-diagonal matrix with its m-th diagonal block given as

$$\mathbf{D}_f^{(m)} = \sum_{k=0}^{K-1} \frac{\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}^{\mathsf{T}}}{\mathbf{u}_m^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m} - \sum_{k=0}^{K-1} \frac{\left(\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m\mathbf{u}_m^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\right)^{\mathsf{T}}}{\left(\mathbf{u}_m^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m\right)^2}, \tag{4.24}$$

so that we have $\nabla^2 f(\mathbf{U}) = \text{blkdiag}\left(\mathbf{D}_f^{(1)}, \ldots, \mathbf{D}_f^{(M)}\right)$. Note that both terms on the RHS of (4.24) are PSD and therefore $\mathbf{D}_f^{(m)}$ is the difference of two PSD matrices.

Lemma 4.5. *The Hessian matrix* $\nabla^2 f(\mathbf{U})$ *is bounded in operator norm.*

Proof. Define the operator norm of a matrix $\mathbf{A} \in \mathbb{C}^{M^2 \times M^2}$ as

$$\|\mathbf{A}\|_{op} = \sup_{\|\mathbf{x}\|=1} \frac{\|\mathbf{A}\mathbf{x}\|}{\|\mathbf{x}\|}, \tag{4.25}$$

where $\|\cdot\|$ is the ℓ_2 norm. For a block-diagonal matrix such as $\nabla^2 f(\mathbf{U})$, the operator norm is equal to the maximum of the operator norms of each individual block, i.e. $\|\nabla^2 f(\mathbf{U})\|_{op} = \max_m \|\mathbf{D}_f^{(m)}\|_{op}$. Using (4.24), the operator norm of block m is bounded as

$$\|\mathbf{D}_f^{(m)}\|_{op} \leq \max\left\{\sum_k \frac{\|\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\|_{op}}{\mathbf{u}_m^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m}, \sum_k \frac{\|\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m\|^2}{(\mathbf{u}_m^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m)^2}\right\}, \tag{4.26}$$

where we used the fact that $\mathbf{D}_f^{(m)}$ is the difference of two PSD matrices and therefore its operator norm is bounded by the maximum of the operator norms of the two. Also, since the matrix $\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m\mathbf{u}_m^{\mathsf{H}}\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}$ is of rank one, its operator norm is equal to $\|\widehat{\boldsymbol{\Sigma}}_{\mathbf{y},k}\mathbf{u}_m\|^2$. Finally, since sample covariances are assumed to be non-singular, both arguments in $\max\{\cdot\}$ are finite. Taking the maximum over all M bounds also results in a finite value and the proof is complete. □

Next we show that the Lipschitz constant of $\nabla f(\mathbf{U})$ is related to the operator norm of $\nabla^2 f(\mathbf{U})$.

Lemma 4.6. *For a twice differentiable function f with Hessian bounded in operator norm as $\|\nabla^2 f\|_{op} \leqslant L$ for all $\mathbf{U}$, the gradient ∇f is Lipschitz continuous with Lipschitz constant L.*

Proof. We show that the Lipschitz continuity condition $\|\nabla f(\mathbf{U}) - \nabla f(\mathbf{U}')\|_{\mathsf{F}} \leqslant L\|\mathbf{U} - \mathbf{U}'\|_{\mathsf{F}}$ holds for any $\mathbf{U}, \mathbf{U}'$ via the following sequence of inequalities:

$$
\begin{aligned}
\|\nabla f(\mathbf{U}) - \nabla f(\mathbf{U}')\|_{\mathsf{F}} &\overset{(a)}{\leqslant} \sup_{\|\mathbf{B}\|_{\mathsf{F}}=1} \left|\left\langle \mathbf{B}, \nabla f(\mathbf{U}) - \nabla f(\mathbf{U}')\right\rangle\right| \\
&= \sup_{\|\mathbf{B}\|_{\mathsf{F}}=1} \left|\left\langle \mathrm{vec}(\mathbf{B}), \int_0^1 \nabla^2 f\left(\mathbf{U}' + t(\mathbf{U} - \mathbf{U}')\right) \mathrm{vec}(\mathbf{U} - \mathbf{U}')\, dt\right\rangle\right| \\
&\leqslant \sup_{\|\mathbf{B}\|_{\mathsf{F}}=1} \int_0^1 \left|\left\langle \mathrm{vec}(\mathbf{B}), \nabla^2 f\left(\mathbf{U}' + t(\mathbf{U} - \mathbf{U}')\right) \mathrm{vec}(\mathbf{U} - \mathbf{U}')\right\rangle\right| dt \\
&\overset{(b)}{\leqslant} \sup_{\|\mathbf{B}\|_{\mathsf{F}}=1} \sup_{t\in[0,1]} \|\nabla^2 f\left(\mathbf{U}' + t(\mathbf{U} - \mathbf{U}')\right)\|_{op}\|\mathrm{vec}(\mathbf{U} - \mathbf{U}')\|\|\mathrm{vec}(\mathbf{B})\| \\
&= \sup_{t\in[0,1]} \|\nabla^2 f\left(\mathbf{U}' + t(\mathbf{U} - \mathbf{U}')\right)\|_{op}\|\mathrm{vec}(\mathbf{U} - \mathbf{U}')\| \\
&\overset{(c)}{\leqslant} L\|\mathbf{U} - \mathbf{U}'\|_{\mathsf{F}}.
\end{aligned}
$$
(4.27)

Inequality (a) holds by taking into account the fact that for the particular value of $\mathbf{B}$ as $\mathbf{B}_0 = \frac{\nabla f(\mathbf{U}) - \nabla f(\mathbf{U}')}{\|\nabla f(\mathbf{U}) - \nabla f(\mathbf{U}')\|_{\mathsf{F}}}$ we have $\langle \mathbf{B}_0, \nabla f(\mathbf{U}) - \nabla f(\mathbf{U}')\rangle = \|\nabla f(\mathbf{U}) - \nabla f(\mathbf{U}')\|_{\mathsf{F}}$. Inequality (b) comes from an application of the Cauchy-Schwarz inequality and the definition of the operator norm. Finally, inequality (c) holds due to the assumption on the boundedness of the Hessian operator norm, i.e. $\|\nabla^2 f(\mathbf{U})\|_{op} \leqslant L$ for all $\mathbf{U}$, and the proof is complete. □

The next lemma emerges as a consequence of the discussion above.

Lemma 4.7. *For any pair of matrices $\mathbf{U}, \mathbf{U}' \in \mathcal{B}$ we have*

$$
f(\mathbf{U}) \leqslant f(\mathbf{U}') + \langle \nabla f(\mathbf{U}'), \mathbf{U} - \mathbf{U}'\rangle + \frac{L}{2}\|\mathbf{U} - \mathbf{U}'\|_{\mathsf{F}}^2,
$$
(4.28)

where L is the gradient Lipschitz constant.

Proof. See [9], proposition 6.1.2. □

This lemma is used as a tool to prove the convergence of PGD to a stationary point, as outlined by Theorem 4.3.

Proof of Theorem 4.3

We start by replacing $\mathbf{U}'$ with $\mathbf{U}^{(t)}$ in (4.28) and defining the RHS of (4.28) as the proxy function

$$f_{\text{proxy}}^{(t)}(\mathbf{U}) = f(\mathbf{U}^{(t)}) + \langle \nabla f(\mathbf{U}^{(t)}), \mathbf{U} - \mathbf{U}^{(t)} \rangle + \frac{L}{2} \| \mathbf{U} - \mathbf{U}^{(t)} \|_{\mathsf{F}}^2, \tag{4.29}$$

at point $\mathbf{U}^{(t)}$, such that we have

$$f(\mathbf{U}) \leqslant f_{\text{proxy}}^{(t)}(\mathbf{U}) \tag{4.30}$$

for all $\mathbf{U} \in \mathcal{U}$ and $f(\mathbf{U}^{(t)}) = f_{\text{proxy}}^{(t)}(\mathbf{U}^{(t)})$. Now let us show that the point $\mathbf{U}^{(t+1)} = \mathcal{P}_{\mathcal{U}}\left(\mathbf{U}^{(t)} - \alpha_t \nabla f(\mathbf{U}^{(t)})\right)$ is indeed a minimizer of $f_{\text{proxy}}^{(t)}(\mathbf{U})$ over $\mathcal{U}$ with $\alpha_t = \frac{1}{L}$. To see this, note that we can expand $f_{\text{proxy}}^{(t)}(\mathbf{U})$ as

$$f_{\text{proxy}}^{(t)}(\mathbf{U}) = \langle \nabla f(\mathbf{U}^{(t)}), \mathbf{U} \rangle - L \langle \mathbf{U}^{(t)}, \mathbf{U} \rangle + \text{const.},$$

for all unitary $\mathbf{U}$. Then, minimizing $f_{\text{proxy}}^{(t)}(\mathbf{U})$ is equivalent to the maximization problem:

$$\underset{\mathbf{U}^{\mathsf{H}}\mathbf{U}=\mathbf{I}_M}{\text{maximize}} \, \langle \mathbf{U}^{(t)} - \frac{1}{L} \nabla f(\mathbf{U}^{(t)}), \mathbf{U} \rangle.$$

But the maximum of this objective is achieved at the point $\mathbf{U}^{\star} = \mathbf{S}_t \mathbf{T}_t^{\mathsf{H}}$, where $\mathbf{S}_t$ and $\mathbf{T}_t$ are matrices of left and right eigenvectors in the SVD form $\mathbf{U}^{(t)} - \frac{1}{L} \nabla f(\mathbf{U}^{(t)}) = \mathbf{S}_t \mathbf{D}_t \mathbf{T}_t^{\mathsf{H}}$. This implies that $\mathbf{U}^{\star} = \mathbf{S}_t \mathbf{T}_t^{\mathsf{H}} = \mathcal{P}_{\mathcal{U}}\left(\mathbf{U}^{(t)} - \frac{1}{L} \nabla f(\mathbf{U}^{(t)})\right) = \mathbf{U}^{(t+1)}$ and $\mathbf{U}^{(t+1)}$ is a minimizer of $f_{\text{proxy}}^{(t)}(\mathbf{U})$. The chain of inequalities below immediately follows:

$$f(\mathbf{U}^{(t+1)}) \overset{(a)}{\leqslant} f_{\text{proxy}}^{(t)}(\mathbf{U}^{(t+1)}) \overset{(b)}{\leqslant} f_{\text{proxy}}^{(t)}(\mathbf{U}^{(t)}) \overset{(c)}{=} f(\mathbf{U}^{(t)}), \tag{4.31}$$

where (a) follows from (4.30), (b) follows from the fact that $\mathbf{U}^{(t+1)}$ is a minimizer of $f_{\text{proxy}}^{(t)}(\mathbf{U})$, and (c) is a result of $f(\mathbf{U}^{(t)}) = f_{\text{proxy}}^{(t)}(\mathbf{U}^{(t)})$. Therefore we have $f(\mathbf{U}^{(t+1)}) \leqslant f(\mathbf{U}^{(t)})$, for $t = 0, 1, \ldots$ and since $f(\mathbf{U})$ is bounded from below, the gradient projection sequence $\{\mathbf{U}^{(t)}, t = 0, 1, \ldots\}$ converges to a stationary point of $f(\mathbf{U})$. □

Theorem 4.3 guarantees that the sequence generated by PGD converges to a stationary point of the likelihood function. This gives a suitable common eigenbasis that, in a sense,

approximately diagonalizes all the user covariance matrices. This basis can serve as the beam-space representation of the channel. Then, as we will show later in this chapter, it can be used to implement ACS for arbitrary array geometries.

Extension of ACS to Arbitrary Array Geometries

We can directly extend the ACS technique for FDD massive MIMO channels with non-ULA geometries. In Section 4.4.1 we proposed a method of designing a common eigenbasis. Given user channel covariances $\{\mathbf{\Sigma}_k\}_k$, or their sample covariance, this method yields a common eigenbasis $\mathbf{U}^\star$, and the user-dependent eigenvalue matrices $\mathbf{\Lambda}_k^\star = \mathrm{diag}(\boldsymbol{\lambda}_k^\star)$, $k \in \mathbb{Z}_K$, and approximates covariance k as

$$\mathbf{\Sigma}_k \approx \mathbf{U}^\star \mathbf{\Lambda}_k^\star \mathbf{U}^{\star\,\mathsf{H}}. \tag{4.32}$$

The eigenbasis $\mathbf{U}^\star$ consists of the array virtual beams. Since this beam-space is shared among all users, we can define the bipartite user-beam graph introduced in Section 4.3.2. In this case, the edge weight between a user k and a beam m is given by $w_{m,k} = [\mathbf{\Lambda}_k^\star]_m$. Then we can solve the same matching-size maximization problem in (4.5) through the MILP. Let $\{x_m^\star\}_m$ denote the MILP solution for the binary variables representing beam nodes and $\{y_k^\star\}_k$ its solution for binary variables representing user nodes. Also let $\mathcal{B} = \{m : x_m^\star = 1\}$ define the set of active beams and $\mathcal{K} = \{k : y_k^\star = 1\}$ the set of active users. The sparsifying precoder in this case is given as

$$\mathbf{B} = \mathbf{U}_{\mathcal{B}}^{\star\,\mathsf{H}}. \tag{4.33}$$

The rest of the channel training and precoding procedure is performed just like the ULA case.

Remark 4.2. Note that since $\mathbf{U}^\star$ is only an approximate eigenbasis of $\mathbf{\Sigma}_k$, we can not guarantee the coefficients of the linear expansion of a random channel vector $\mathbf{h}_k$ in terms of the columns of $\mathbf{U}^\star$ to be independent random variables with a continuous distribution. Hence, we can not prove that maximizing the matching size in the beam-user bipartite graph is equivalent to maximizing the channel matrix rank. The reason is that the conditions of Lemma 4.3 are violated, since we can not assume a distribution on the coefficients. Nevertheless, we hope that the error of approximating the covariances as in (4.32) is not large, such that $\mathbf{U}^\star$ is close to $\mathbf{U}_k$ for all k. This would lead the coefficients of the expansion in terms of the columns of $\mathbf{U}^\star$ to be close to the Gaussian coefficients of the Karhunen-Loeve expansion. Then maximizing the matching size will maximize the channel matrix rank.

4.5 Simulation Results

With empirical simulations, we compare the performance of the proposed ACS method to two of the most recent CS-based methods proposed in [98] and [40] in terms of channel estimation error and sum-rate. In [98], the authors proposed a method based on common probing of the DL channel with random Gaussian pilots. The DL pilot measurements $\mathbf{y}_k$ at users $k = 1, \ldots, K$ (similar to (4.4), but with a different pilot matrix) are fed back and collected by the BS, which recovers the channel vectors using a joint orthogonal matching pursuit (J-OMP) technique able to exploit the possible common sparsity between the user channels. In [40], a method based on dictionary learning for sparse channel estimation was proposed. In this scheme, the BS jointly *learns* sparsifying dictionaries for the UL and DL channels by collecting channel measurements at different cell locations (e.g., via an off-line learning phase). The actual user channel estimation is posed as a norm-minimization convex program using the trained dictionaries and with the constraint that UL and DL channels share the same support over their corresponding dictionaries. Following the terminology used in [40], we refer to this method as JDLCM.

4.5.1 Channel Estimation Error and Sum-Rate vs Pilot Dimension

For this comparison, we considered $M = 128$ antennas at the BS, $K = 13$ users, and resource blocks of size $T = 128$ symbols. For our proposed method, the BS computes the users' sample UL covariance matrices by taking $N_{\mathsf{ul}} = 1000$ UL pilot observations and then applies the NNLS-based UL-DL covariance transformation scheme explained in Section 3.5 of Chapter 3. Given the obtained DL channel covariance matrix estimates, we first perform the circulant approximation and extract the vector of approximate eigenvalues as in (4.1). Then, we compute the sparsifying precoder $\mathbf{B}$ via the MILP solution as given in Section 4.3.2. In the results presented here, we set the parameter P_{th} in the MILP to a small value in order to favor a high rank of the resulting effective channel matrix over the beamforming gain.[3] After probing the effective channel of the selected users along these active beam directions via a random unitary pilot matrix $\boldsymbol{\Psi}$, we calculate their MMSE estimate using the estimated DL covariance matrices.

Eventually, for all the three methods, we compute the ZF beamforming matrix based on the obtained channel estimates. In addition, instead of considering all selected users, in both cases we apply the Greedy ZF user selection approach of [39], that yields a significant benefit when the number of users is close to the rank of the effective channel matrix. As said before, the DL SNR is given by SNR $= P_{\mathsf{dl}}/N_0$ and during the simulations we

---

[3]This approach is appropriate in the medium to high-SNR regime. For low SNR, it is often convenient to increase P_{th} in order to serve less users with a larger beamforming energy transfer per user.
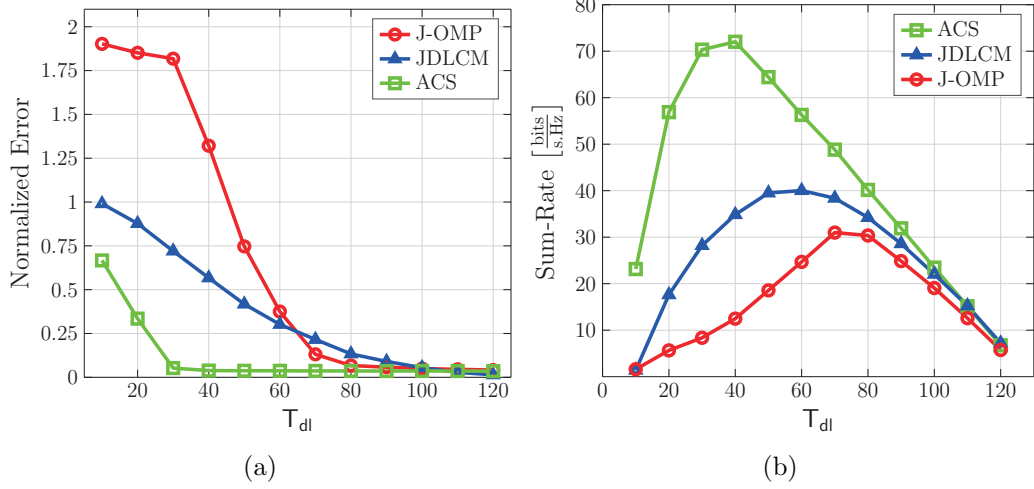
Figure 4.3: (a) Normalized channel estimation error, and (b) achievable sum-rate as a function of DL pilot dimension with $SNR = 20$ dB, $M = 128$ and $K = 13$.

consider ideal noiseless feedback for simplicity, i.e., we assume that the BS receives the measurements in (4.3) without extra feedback noise to the system.[4]. The sparsity order of each channel vector is given as an input to the J-OMP method, but not to the other two methods. This represents a genie-aided advantage for J-OMP, that we introduce here for simplicity. As the simulation geometry, we consider three MPC clusters with random locations within the angular range (parametrized by ξ rather than θ) $[-1, 1)$. We denote by Ξ the i-th interval and set each interval size to be $|\Xi_i| = 0.2$, $i = 1, 2, 3$. The ASF for each user is obtained by selecting at random two out of three such clusters, such that the overlap of the angular components among users is large. The ASF is non-zero over the angular intervals corresponding to the chosen MPCs and zero elsewhere, i.e., $\gamma_k(d\xi) = \beta \mathbf{1}_{\Xi_{i_1} \cup \Xi_{i_2}}$, where $\beta = 1/\int_{-1}^{1} \gamma_k(d\xi)$ and $i_1, i_2 \in \{1, 2, 3\}$.

The described arrangement results in each generated channel vector being roughly $s_k = 0.2 \times M \approx 26$-sparse. To measure channel estimation error we use the normalized Euclidean distance as follows. Let $\mathbf{H} \in \mathbb{C}^{M \times K'}$ define the matrix whose columns correspond to the channel vectors of the K' served users and let $\widehat{\mathbf{H}}$ denote the estimation of $\mathbf{H}$. Then the normalized Euclidean error is defined as

$$E_{\mathsf{euc}} = \mathbb{E}\left[\frac{\|\mathbf{H} - \widehat{\mathbf{H}}\|_{\mathsf{F}}^2}{\|\mathbf{H}\|_{\mathsf{F}}^2}\right].$$

---

[4]Notice that by introducing noisy feedback the relative gain w.r.t. J-OMP is even larger, since CS schemes are known to be more noise-sensitive than plain MMSE estimation using estimated DL covariance matrices
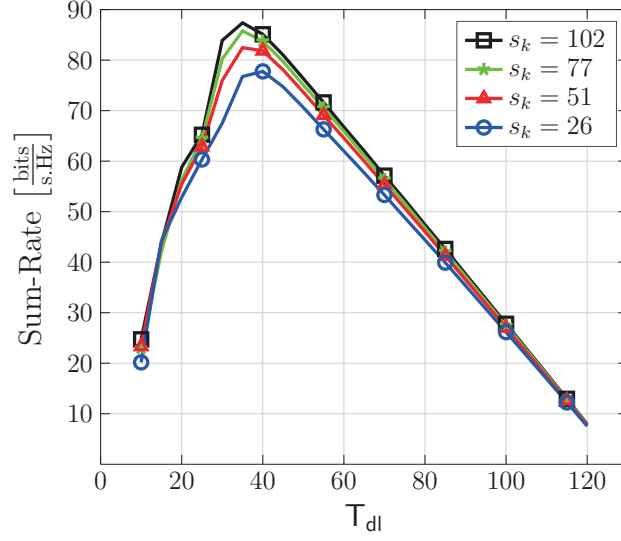
Figure 4.4: Sum-rate vs T_{dl} for various channel sparsity orders. Here SNR = 20 dB, $M = 128$ and $K = 13$.

Fig. 4.3a shows the normalized channel estimation error for the J-OMP, JDLCM and our proposed Active Channel Sparsification (ACS) method as a function of the DL pilot dimension T_{dl} with $SNR = 20$ dB. Our ACS method outperforms the other two by a large margin, especially for low DL pilot dimensions. When the pilot dimension is below channel sparsity order, CS-based methods perform very poorly, since the number of channel measurements is less than the inherent channel dimension. Fig. 4.3b compares the achievable sum-rate for the three methods. Again our ACS method shows a much better performance compared to J-OMP and JDLCM. This figure also shows that there is an optimal DL pilot dimension that maximizes the sum-rate. This optimal value is $T_{dl} \approx 40$ for our proposed method, $T_{dl} \approx 60$ for JDLCM and $T_{dl} \approx 70$ for the J-OMP method.

4.5.2 The Effect of Channel Sparsity

Depending on the geometry and user location, channels may show different levels of sparsity in the angular domain. In contrast to CS-based methods, our proposed method is highly flexible with regards to various channel sparsity orders, thanks to the active sparsification method. In this section, we investigate how sparsity order effects channel estimation error as well as sum-rate within the framework of our proposed method. We suppose that the user ASF's consist of two clusters chosen at random among the three. But now we vary the size of the angular interval each of the clusters occupies ($|\Xi_i| = 0.2, 0.4, 0.6, 0.8$) and see how it effects the error and sum-rate metrics. The sparsification, channel probing and transmission are performed as described before. Since each ASF consists of two clus-
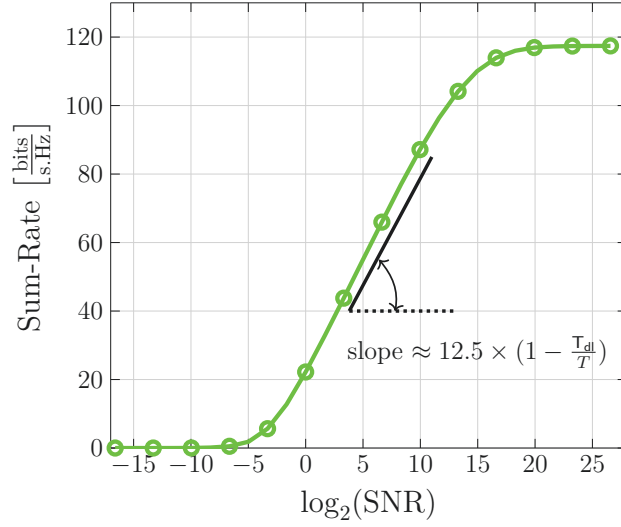
Figure 4.5: Sum-rate as a function of $\log_2 (\text{SNR})$ with $M = 128$ and $K = 13$.

ters and $M = 128$ channel sparsity order (roughly) takes on the values $s_k = 26, 51, 77, 102$ for all users $k \in \mathbb{Z}_{K'}$. For each value of the pilot dimension we perform a Monte Carlo simulation to empirically calculate the sum-rate. Fig. 4.4 illustrates the results. Notice that in these results we fix the channel coefficient power along each scattering component, such as richer (less sparse) channels convey more signal energy. This corresponds to the physical fact that the more scattered signal energy is collected at the receiving antennas the higher the received signal energy is. As we can see in Fig. 4.4, for a fixed T_{dl}, when the number of non-zero channel coefficients increases (i.e., the channel is less sparse), we generally have a larger sum-rate. The main reason is that, with less sparse channels, the beamforming gain is larger due to the fact that more scattering components contribute to the channel. Therefore, we can generally say that with our method, for a fixed pilot dimension, less sparse channels are *better*. Of course, this is not the case for CS-based techniques, or techniques based on the "sparsity assumption" of a small number of discrete angular components, which tend to collapse and yield very bad results when such sparsity assumptions are not satisfied.

4.5.3 The Multiplexing Gain

An interesting final observation is to examine the system sum-rate vs. SNR with our proposed method, and in particular show that there is indeed a regime of intermediate SNR for which the slope of the sum-rate curve yields quite faithfully the number of spatially multiplexed data streams. We performed a simulation with $M = 128$ antennas and $K = 13$ users and a pilot dimension of $T_{\text{dl}} = 60$. The pre-log factor determines the slope of the sum-

rate vs $\log_2(\text{SNR})$ curve, in an intermediate regime where the sum-rate is not saturated, and yet the spectral efficiency is large.[5] As illustrated in Fig. 4.5, this slope is equal to $12.5 \times (1 - \frac{T_{dl}}{T})$. Notice that the Greedy ZF scheme decides to serve a number of users that may be less than K in an opportunistic fashion, such that the expected number of served users (DL data streams) in this SNR regime is indeed slightly less than the maximum possible $K = 13$. Hence, the agreement between the sum-rate slope in this regime and the number of served DL data streams is exactly what can be expected, thus showing the relevance of maximizing the rank of the effective matrix in the proposed optimization of the sparsifying precoder.

4.5.4 Performance of the ML-Based Beam-Space Design

In Section 4.4.1 of this chapter we developed a method of MIMO beam-space design for arbitrary array geometries that finds, via maximizing a likelihood function, the common unitary matrix that jointly diagonalizes a set of user channel covariances. Here, we provide simulation results to study the performance of this technique.

Jointly Diagonalizable Covariances

The case of jointly diagonalizable covariances is especially interesting, as we know by Theorem 4.2 that the global optimum of the ML problem is given by the shared CES $\mathbf{U}^c$ (see (4.17)). In order to assess the performance of our method, we compare the ML cost as a function of the number of random samples per process N, to the cost at the global minimum. Consider a signal dimension (number of antennas) $M = 16$ and a number of processes (number of users) $K = 8$. We generate a random unitary matrix as the CES $\mathbf{U}^c$ by calculating the eigenvectors of a random matrix of size $M \times M$ with i.i.d complex Gaussian elements. Also, for each process $k \in [K]$, we generate a random vector of eigenvalues $\boldsymbol{\lambda}_k$ with i.i.d, positive elements given as $\boldsymbol{\lambda}_{k,m} = |\rho_m|$ where $\rho_m \sim \mathcal{N}(0, 1)$. Then we form the covariance matrix of user k as $\widetilde{\boldsymbol{\Sigma}} = \mathbf{U}^c \text{diag}(\boldsymbol{\lambda}_k) \mathbf{U}^{c\mathsf{H}}$. We also normalize the covariances to have trace equal to one. This way we have randomly generated covariances with a shared CES. Now, having the covariances, we can generate random realizations for each process for different sample sizes N. We run a Monte Carlo simulation with 1000 iterations, at each iteration generating covariances as stated above, then for each sample size N we run our proposed PGD method which converges to a point $\mathbf{U}^\star(N)$ (explicitly noting the dependence on N). Then, we average the cost function $\mathbf{U}^\star(N)$ in $(\mathcal{P}_1)$ over the Monte Carlo iterations for each value of N and compare it to the average cost at the global optimum $f(\mathbf{U}^c)$. Note that the latter of course is not a function of N. The PGD method is initialized with a random unitary matrix. Theorem 4.3 guarantees convergence

---

[5]This saturation is due to the non-vanishing covariance estimation error and happens at around SNR = 60 dB.
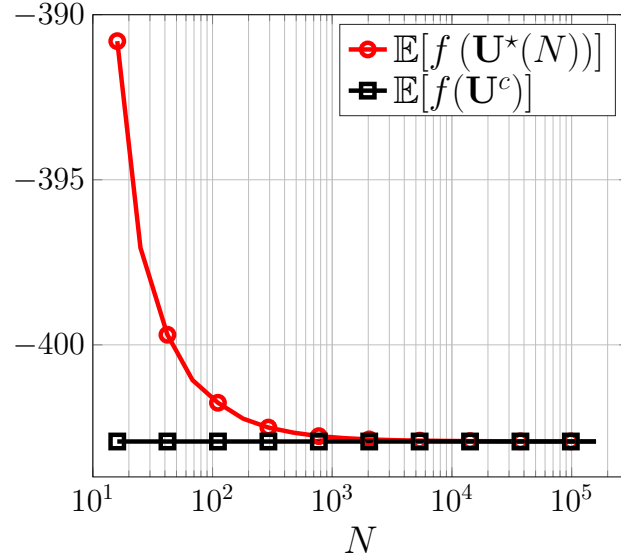
Figure 4.6: The average cost f as a function of the number of samples, for the solution of PGD $\mathbf{U}^\star(N)$ and the global optimum $\mathbf{U}^c$. We have set $M = 16$ and $K = 8$.

to a stationary point when the step size is chosen as $\alpha_t \in (0, 1/L)$. However, practically we can be more ambitious by choosing larger step sizes as $\alpha_t = \frac{\alpha_0}{t}$, $t = 1, 2, \ldots$ with $\alpha_0 = 2$ and our simulation results show that even with this choice, PGD converges.

Fig. 4.6 illustrates the result. We denote the solution of our method with $\mathbf{U}^\star(N)$, to explicitly highlight its dependence on the number of samples N. The interesting fact about this result is that, as the number of samples gets larger, the PGD with random initialization always converges to the global optimum, as its cost value is the same as that in the optimal point $\mathbf{U}^c$. This is an empirical evidence for the convergence of our proposed method to the global solution for jointly diagonalizable covariances.

Non-Jointly Diagonalizable Covariances

As a different scenario, we consider covariances that are not jointly diagonalizable. This is done by generating a different random unitary eigenvector matrix for each process separately as $\widetilde{\boldsymbol{\Sigma}} = \mathbf{U}_k \mathrm{diag}(\boldsymbol{\lambda}_k)\mathbf{U}_k^{\mathsf{H}}$. The eigenvalue vectors $\boldsymbol{\lambda}_k$ are generated as before and we normalize the covariances to have unit trace. So, in this case we do not have a CES and the PGD method yields a unitary matrix that approximately jointly diagonalizes the covariances. Since we do not have the global optimum in this case, we compare our method to the JADE algorithm, which is a classic Jacobian-based method for joint covariance diagonalization (we do not explain the details of this method here due to space limitations and refer the reader to [21] for a full account).
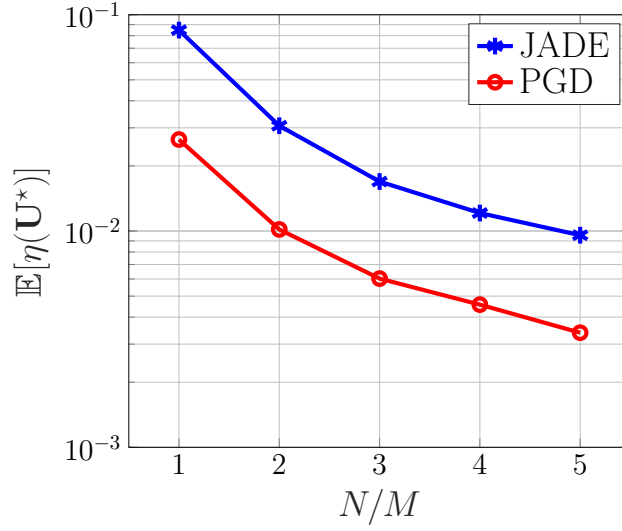
Figure 4.7: The average diagonalization metric as function of the sampling ratio N/M, for the solution of our proposed PGD method vs the JADE method. We set $M = 16$ and $K = 8$.

One way to measure the performance of joint diagonalization methods is by defining the metric $\eta : \mathbb{C}^{M \times M} \to [0, 1]$:

$$\eta(\mathbf{U}) = 1 - \frac{1}{K} \sum_{k=1}^{K} \frac{\|\mathrm{diag}(\mathbf{U}^{\mathsf{H}} \boldsymbol{\Sigma}_k \mathbf{U})\|}{\|\boldsymbol{\Sigma}_k\|_{\mathsf{F}}}, \tag{4.34}$$

where $\mathrm{diag}(\cdot)$ with a matrix argument as in (4.34) denotes the M-dim vector of the diagonal elements of its argument. The smaller the value of $\eta(\mathbf{U})$ is, the better $\mathbf{U}$ jointly diagonalizes the covariances. In the extreme case, If $\mathbf{U}$ diagonalizes all covariance matrices, we have $\eta(\mathbf{U}) = 0$.

The joint diagonalization metric is empirically averaged over 1000 Monte Carlo simulations and for different sample sizes for the solutions of our proposed PGD method and the JADE method. Fig. 4.7 illustrates the results. It clearly shows that for the ranges of sample sizes considered here, the proposed PGD method outperforms the classic JADE method, yielding smaller values of the diagonalization metric on average, and hence achieving a better joint diagonalization of the covariances.

4.5.5 CES for ULA: the PGD Solution vs the Fourier Basis

For a ULA it is usually taken for granted that the CES is given by the Fourier basis vectors. While this is true in an asymptotic sense thanks to the Szegö theorem, it does not hold
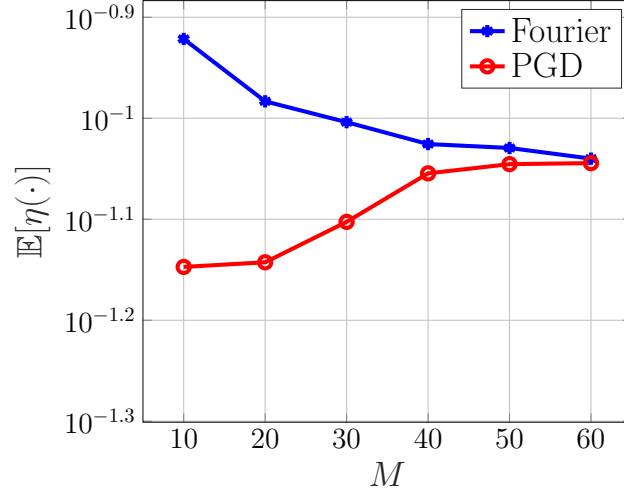
Figure 4.8: The average diagonalization metric as function of the number of antennas M, for the solution of our proposed PGD method vs the Fourier basis. Here we have set $K = 5$.

for small to moderate array sizes. We conclude our simulations by showing that, in fact the unitary basis produced by our proposed method better diagonalizes ULA covariances, compared to the DFT basis.

We consider $K = 5$ M-dimensional ULA covariances, generated randomly according to $\boldsymbol{\Sigma} = \int_{-1}^{1} \gamma(d\xi)\mathbf{a}(\xi)\mathbf{a}(\xi)^{\mathsf{H}}$, by generating random ASF's according to the model that was described in Section 2.6 of Chapter 2. Since each covariance is associated with a user that occupies a limited angular range as seen from the BS, we generate the covariances such that each of them has an "effective" rank of $r_k = \mathrm{effrank}(\widetilde{\boldsymbol{\Sigma}}_k) = \lceil \frac{M}{2} \rceil$. The effective rank is equivalent to the number of significant covariance eigenvalues as well as channel angular sparsity.

We plot the expected joint diagonalization metric $\mathbb{E}[\eta(\cdot)]$ as a function of the number of antennas for the unitary matrix yielded by our method as well as for the Fourier basis $\mathbf{F}$ where $[\mathbf{F}]_{m,n} = \frac{1}{\sqrt{M}} e^{j2\pi \frac{(m-1)(n-1)}{M}}$, $m, n \in [M]$. We assume that the sample covariance has converged, i.e. $\widehat{\boldsymbol{\Sigma}}_k = \widehat{\boldsymbol{\Sigma}}_k$ for all k. The expectation $\mathbb{E}[\eta(\cdot)]$ is taken over random covariance realizations and is calculated empirically over 100 Monte-Carlo loops. Fig. 4.8 illustrates the result. As we can see, the basis given by the PGD method achieves better diagonalization (smaller η values) than the Fourier basis, which shows that using our method is in fact preferable even for the diagonalization of ULA covariances. Also, as M increases, the two bases have closer diagonalization metrics since we are approaching the asymptotic regime in which the Fourier basis approximately diagonalizes the Toeplitz ULA covariances.

4.6 Appendices

4.6.1 Proof of Lemma 4.1

The proof follows by using the representation $\mathbf{h}_k = \sum_{m \in \mathcal{S}_k} g_{k,m}\sqrt{[\boldsymbol{\lambda}_k]_m}\mathbf{f}_m$ (see (4.2)), which holds exactly by assumption. Estimating $\mathbf{h}_k$ is equivalent to estimating the vector of KL Gaussian i.i.d. coefficients $\mathbf{g}_k = (g_{k,m} : m \in \mathcal{S}_k) \in \mathbb{C}^{s_k \times 1}$. Define the $M \times s_k$ DFT submatrix $\mathbf{F}_{\mathcal{S}_k} = (\mathbf{f}_m : m \in \mathcal{S}_k)$, and the corresponding diagonal $s_k \times s_k$ matrix of the non-zero eigenvalues $[\boldsymbol{\Lambda}_k]_{\mathcal{S}_k,\mathcal{S}_k}$. After some simple standard algebra, the MMSE estimation error covariance of $\mathbf{g}_k$ from $\mathbf{y}_k$ in (4.3) with $\mathbf{B} = \mathbf{I}_M$ can be written in the form

$$
\begin{aligned}
\widetilde{\mathbf{R}}_e = \mathbf{I}_{s_k} &- ([\boldsymbol{\Lambda}_k]_{\mathcal{S}_k,\mathcal{S}_k})^{1/2}\,\mathbf{F}_{\mathcal{S}_k}^{\mathsf{H}}\,\boldsymbol{\Psi}^{\mathsf{H}} \\
&\times \left(\boldsymbol{\Psi}\mathbf{F}_{\mathcal{S}_k}[\boldsymbol{\Lambda}_k]_{\mathcal{S}_k,\mathcal{S}_k}\mathbf{F}_{\mathcal{S}_k}^{\mathsf{H}}\boldsymbol{\Psi}^{\mathsf{H}} + N_0\mathbf{I}_{T_{\mathsf{dl}}} \right)^{-1} \boldsymbol{\Psi}\mathbf{F}_{\mathcal{S}_k}\,([\boldsymbol{\Lambda}_k]_{\mathcal{S}_k,\mathcal{S}_k})^{1/2}\,.
\end{aligned}
\tag{4.35}
$$

Using the fact that $\mathbf{R}_e = \mathbf{F}_{\mathcal{S}_k}([\boldsymbol{\Lambda}_k]_{\mathcal{S}_k,\mathcal{S}_k})^{1/2}\widetilde{\mathbf{R}}_e([\boldsymbol{\Lambda}_k]_{\mathcal{S}_k,\mathcal{S}_k})^{1/2}\mathbf{F}_{\mathcal{S}_k}^{\mathsf{H}}$, such that $\mathsf{tr}(\mathbf{R}_e) = \mathsf{tr}([\boldsymbol{\Lambda}_k]_{\mathcal{S}_k,\mathcal{S}_k}\widetilde{\mathbf{R}}_e)$, we have that $\mathsf{tr}(\mathbf{R}_e)$ and $\mathsf{tr}(\widetilde{\mathbf{R}}_e)$ have the same vanishing order with respect to N_0. In particular, it is sufficient to consider the behavior of $\mathsf{tr}(\widetilde{\mathbf{R}}_e)$ as a function of N_0. Now, using the Sherman-Morrison-Woodbury matrix inversion lemma [57], after some algebra omitted for the sake of brevity we arrive at

$$
\mathsf{tr}(\widetilde{\mathbf{R}}_e) = s_k - \sum_{i=1}^{s_k} \frac{\mu_i}{N_0 + \mu_i},
\tag{4.36}
$$

where μ_i is the i-th eigenvalue of the $s_k \times s_k$ matrix

$$
\mathbf{A} = ([\boldsymbol{\Lambda}_k]_{\mathcal{S}_k,\mathcal{S}_k})^{1/2}\mathbf{F}_{\mathcal{S}_k}^{\mathsf{H}}\boldsymbol{\Psi}^{\mathsf{H}}\boldsymbol{\Psi}\mathbf{F}_{\mathcal{S}_k}([\boldsymbol{\Lambda}_k]_{\mathcal{S}_k,\mathcal{S}_k})^{1/2}
$$

Next, notice that

$$
\begin{aligned}
\mathsf{rank}(\mathbf{A}) &= \mathsf{rank}(\mathbf{F}_{\mathcal{S}_k}^{\mathsf{H}}\boldsymbol{\Psi}^{\mathsf{H}}\boldsymbol{\Psi}\mathbf{F}_{\mathcal{S}_k}) \\
&= \mathsf{rank}(\mathbf{F}_{\mathcal{S}_k}\mathbf{F}_{\mathcal{S}_k}^{\mathsf{H}}\boldsymbol{\Psi}^{\mathsf{H}}) \leqslant \min\{s_k, T_{\mathsf{dl}}\}.
\end{aligned}
\tag{4.37}
$$

In fact, $[\boldsymbol{\Lambda}_k]_{\mathcal{S}_k,\mathcal{S}_k}$ is diagonal with strictly positive diagonal elements, such that left and right multiplication by $([\boldsymbol{\Lambda}_k]_{\mathcal{S}_k,\mathcal{S}_k})^{1/2}$ yields rank-preserving row and column scalings, the matrix $\mathbf{F}_{\mathcal{S}_k}\mathbf{F}_{\mathcal{S}_k}^{\mathsf{H}}$ is the orthogonal projector onto the s_k-dimensional column-space of $\mathbf{F}_{\mathcal{S}_k}$ and has rank s_k, while the matrix $\boldsymbol{\Psi}^{\mathsf{H}} \in \mathbb{C}^{M \times T_{\mathsf{dl}}}$ has the same rank of $\boldsymbol{\Psi}^{\mathsf{H}}\boldsymbol{\Psi}$, that is at most T_{dl}.

For $T_{\mathsf{dl}} \geqslant s_k$ the existence of matrices $\boldsymbol{\Psi}$ such that the rank upper bound (4.37) holds with equality (i.e., for which $\mathsf{rank}(A) = s_k$ for any support set $\mathcal{S}_k$ of size s_k) is shown as follows. Generate a random $\boldsymbol{\Psi}$ with i.i.d. elements $\sim \mathcal{CN}(0,1)$. Then, the columns

of $\mathbf{F}_{\mathcal{S}_k}^{\mathsf{H}} \mathbf{\Psi}^{\mathsf{H}}$ form a collection of $T_{\mathsf{dl}} \geqslant s_k$ mutually independent s_k-dimensional Gaussian vectors with i.i.d. $\sim \mathcal{CN}(0,1)$ components. The event that these vectors span a space of dimension less than s_k is a null event (zero probability). Hence, such randomly generated matrix satisfies the rank equality in (4.37) with probability 1. As a consequence, for $T_{\mathsf{dl}} \geqslant s_k$ we have that $\mu_i > 0$ for all $i \in [s_k]$ and (4.36) vanishes as $O(N_0)$ as $N_0 \downarrow 0$. In contrast, if $T_{\mathsf{dl}} < s_k$, by (4.37) for any matrix $\mathbf{\Psi}$ at most T_{dl} eigenvalues μ_i in (4.36) are non-zero and $\lim_{N_0 \downarrow 0} s_k - \sum_{i=1}^{s_k} \frac{\mu_i}{N_0 + \mu_i} \geqslant s_k - T_{\mathsf{dl}} > 0$. $\qquad \square$

4.6.2 Proof of Lemma 4.3

The determinant of $\mathbf{Q}$ is given by the expansion $\det(\mathbf{Q}) = \sum_{\iota \in \boldsymbol{\pi}_r} \mathrm{sgn}(\iota) \prod_i [\mathbf{Q}]_{i,\iota(i)}$, where ι is a permutation of the set $\{1, 2, \ldots, r\}$, where $\boldsymbol{\pi}_r$ is the set of all such permutations and where $\mathrm{sgn}(\iota)$ is either 1 or -1. The product $\prod_i [\mathbf{Q}]_{i,\iota(i)}$ is non-zero only for the perfect matchings in the bipartite graph. Hence, if the bipartite graph contains a perfect matching, then $\det(\mathbf{Q}) \neq 0$ with probability 1 (and $\mathrm{rank}(\mathbf{Q}) = r$), since the non-identically zero entries of $\mathbf{W}$ are drawn from a continuous distribution. If it does not contain a perfect matching, then $\det(\mathbf{Q}) = 0$ and therefore $\mathrm{rank}(\mathbf{Q}) < r$. $\qquad \square$

4.6.3 Proof of Theorem 4.1

First, without loss of generality let us assume that $\mathcal{L}$ contains no isolated nodes (since these would be discarded anyway). As before, the $|\mathcal{A}| \times |\mathcal{K}|$ weighted adjacency matrix is denoted by $\mathbf{W}$ where $[\mathbf{W}]_{m,k} = w_{m,k}$. An example of the bipartite graph $\mathcal{L}$ and its corresponding weighted adjacency matrix $\mathbf{W}$ is illustrated in Figs. 4.2a and 4.2b. Given the bipartite graph $\mathcal{L}(\mathcal{A}, \mathcal{K}, \mathcal{E})$, we select the subgraph $\mathcal{L}'(\mathcal{A}', \mathcal{K}', \mathcal{E}')$, so that the constraint (4.5b) is satisfied. We introduce the binary variables $\{x_m, m \in \mathcal{A}\}$ and $\{y_k, k \in \mathcal{K}\}$ to indicate if beam m and user k are selected, respectively. As such, the constraint (4.5b) is equivalent to the set of constraints:

$$x_m \leqslant \sum_{k \in \mathcal{K}} [\mathbf{A}]_{m,k} y_k \ \forall m \in \mathcal{A} \tag{4.38a}$$

$$y_k \leqslant \sum_{m \in \mathcal{A}} [\mathbf{A}]_{m,k} x_m \ \forall k \in \mathcal{K} \tag{4.38b}$$

$$\sum_{m \in \mathcal{A}} [\mathbf{A}]_{m,k} x_m \leqslant T_{\mathsf{dl}} y_k + M(1 - y_k) \ \forall k \in \mathcal{K} \tag{4.38c}$$

In particular, (4.38a) ensures that if the beam m is selected (i.e., $x_m = 1$), there must be some $k \in \mathcal{K}$ such that $(m, k) \in \mathcal{E}$ is selected as well, whereas if beam m is not selected, then this constraint is redundant. Similarly, in (4.38b) if user k is selected (i.e., $y_k = 1$),

there must be some $m \in \mathcal{A}$ such that $(m, k) \in \mathcal{E}$ is selected as well. Furthermore, (4.38c) guarantees that if user k is chosen (i.e., $y_k = 1$), the number of chosen beams with $x_m = 1$ is no more than T_{dl}, and otherwise this constraint is redundant. Meanwhile, the constraint (4.5c) is written as:

$$P_{\mathsf{th}}\, y_k \leqslant \sum_{m \in \mathcal{A}} [\mathbf{W}]_{m,k} x_m \ \forall k \in \mathcal{K} \tag{4.39}$$

which ensures that if user k is chosen (i.e., $y_k = 1$) then the sum weights of the selected beams (i.e., $m \in \mathcal{N}_{\mathcal{L}'}(k)$ if $x_m = 1$) is no less than P_{th}, while if user k is not chosen (i.e., $y_k = 0$) then this constraint is not required and redundant. A closer look reveals that the constraint (4.39) renders the one (4.38b) redundant, because when $y_k = 1$ in (4.39) there must exist at least one $m \in \mathcal{A}$ with $x_m = 1$. Second, given the selected subgraph $\mathcal{L}'(\mathcal{A}', \mathcal{K}', \mathcal{E}')$, we find a matching $\mathcal{M}(\mathcal{A}', \mathcal{K}')$ with maximum cardinality. To this end, we introduce another set of binary variables $\{z_{mk}, m \in \mathcal{A}, k \in \mathcal{K}\}$ to indicate if an edge $(a, k) \in \mathcal{E}$ is chosen to form the maximum matching in $\mathcal{L}'(\mathcal{A}', \mathcal{K}', \mathcal{E}')$. Following the canonical linear program formulation of the maximum cardinality matching for bipartite graphs, we translate the objective in (4.5) into the following optimization:

$$\underset{z_{m,k} \in \{0,1\}}{\text{maximize}} \quad \sum_{m \in \mathcal{A}'} \sum_{k \in \mathcal{K}'} [\mathbf{A}]_{m,k} z_{m,k} \tag{4.40a}$$

$$\text{subject to} \quad \sum_{k \in \mathcal{K}'} [\mathbf{A}]_{m,k} z_{m,k} \leqslant 1 \ \forall m \in \mathcal{A}', \tag{4.40b}$$

$$\sum_{m \in \mathcal{A}'} [\mathbf{A}]_{m,k} z_{m,k} \leqslant 1 \ \forall k \in \mathcal{K}', \tag{4.40c}$$

Now, to transport the optimization problem on $\mathcal{L}'$ to the original setting on $\mathcal{L}$, we need to guarantee that $\mathcal{M}(\mathcal{A}', \mathcal{K}') \subseteq \mathcal{E}'$, i.e., $z_{mk} = 1$ only if $m \in \mathcal{A}'$ ($x_m = 1$), and $k \in \mathcal{K}'$ ($y_k = 1$). This is obtained for a given configuration of the variables $\{x_m\}$ and $\{y_k\}$ which define $\mathcal{L}'$, by adding constraints to (4.40) and yields

$$\underset{z_{m,k} \in \{0,1\}}{\text{maximize}} \quad \sum_{m \in \mathcal{A}} \sum_{k \in \mathcal{K}} [\mathbf{A}]_{m,k} z_{m,k} \tag{4.41a}$$

$$\text{subject to} \quad \sum_{k \in \mathcal{K}} [\mathbf{A}]_{m,k} z_{m,k} \leqslant 1 \ \forall m \in \mathcal{A}, \tag{4.41b}$$

$$\sum_{m \in \mathcal{A}} [\mathbf{A}]_{m,k} z_{m,k} \leqslant 1 \ \forall k \in \mathcal{K}, \tag{4.41c}$$

$$[\mathbf{A}]_{m,k} z_{m,k} \leqslant x_m \ \forall k \in \mathcal{K}, m \in \mathcal{A}, \tag{4.41d}$$

$$[\mathbf{A}]_{m,k} z_{m,k} \leqslant y_k \ \forall k \in \mathcal{K}, m \in \mathcal{A}, \tag{4.41e}$$

where (4.41d)-(4.41e) impose that the edge set $\{(m,k) : z_{m,k} = 1\}$ should be a subset of $\mathcal{E}'$. A further inspection on these constraints yields the following equivalent simplified form:

$$\underset{z_{m,k}\in\{0,1\}}{\text{maximize}} \sum_{m\in\mathcal{A}}\sum_{k\in\mathcal{K}} z_{m,k} \tag{4.42a}$$

$$\text{subject to } z_{m,k} \leqslant [\mathbf{A}]_{m,k}, \ \forall m \in \mathcal{A}, k \in \mathcal{K}, \tag{4.42b}$$

$$\sum_{k\in\mathcal{K}} z_{m,k} \leqslant x_m, \ \forall m \in \mathcal{A}, \tag{4.42c}$$

$$\sum_{m\in\mathcal{A}} z_{m,k} \leqslant y_k, \ \forall k \in \mathcal{K}, \tag{4.42d}$$

where the additional constraint (4.42b) turns all the terms of the type $[\mathbf{A}]_{m,k}z_{m,k}$ in (4.41) to $z_{m,k}$ in (4.42), the constraint (4.42c) results from the combination of the constraints (4.41b) and (4.41d), and (4.42d) results from the combination of (4.41c) with (4.41e). The formulation in (4.42) can be seen as a modified maximum cardinality bipartite matching with selective vertices, in which the vertices with $x_m = 1$ and $y_k = 1$ are selected to participate in the maximum cardinality matching. The eventual mixed integer linear program is given as in (4.6). Notice that we have relaxed the binary constraint on $\{z_{m,k}, m \in \mathcal{A}, k \in \mathcal{K}\}$ to the linear constraint (4.6i) based on the following lemma.

Lemma 4.8. *The problem* $\mathcal{P}_{\mathrm{MILP}}$ *as stated in* (4.6) *always has binary-valued solutions for* $\{z_{m,k}, m \in \mathcal{A}, k \in \mathcal{K}\}$. □

Proof. It suffices to show that $z_{m,k}$ are binary, given that x_m and y_k are binary. First, if either x_m, $m \in \mathcal{A}$ or y_k, $k \in \mathcal{K}$ are 0, then $z_{a,k} = 0$. So, we only need to focus on the case where $x_m = y_k = 1$, $m \in \mathcal{A}, k \in \mathcal{K}$. In that case, the constraints of $\mathcal{P}_{\mathrm{MILP}}$ with respect to $z_{m,k}$, $m \in \mathcal{A}, k \in \mathcal{K}$ form a convex polytope. This polytope is called the bipartite matching polytope, which is integral, i.e. all of its extreme points have integer (and in this case binary) values (see [110, Corollary 18.1b. and Theorem 18.2.]). Therefore, given $x_m, y_k \in \{0, 1\}$, $\forall m \in \mathcal{A}, k \in \mathcal{K}$, $\mathcal{P}_{\mathrm{MILP}}$ reduces to a linear program with respect to the variables $z_{m,k}$ and the optimal solutions are the integral extreme points of the corresponding polyhedra and the proof is complete. □

Conclusion

This dissertation has addressed a few theoretical and practical problems in taking advantage of the potentials of multi-user massive antenna systems that operate in FDD mode. We motivated the fact that FDD massive MIMO is competitive with (and sometimes preferable to) its TDD counterpart in various communication scenario. We presented algorithms for efficient channel covariance estimation and transformation. Universal minimax error bounds were driven for the problem of covariance transformation, suggesting its theoretical feasibility and a rigorous framework to design transformers. We also proposed a novel solution to the DL channel precoding in FDD massive MIMO. By maximizing the channel matrix rank with restricted training pilot dimension and feedback overhead, our proposed active channel sparsification method achieves excellent results in multiplexing data to an arbitrary set of users with any given (fixed by the standard) channel training resource dimension. We backed up our theoretical proposals with extensive simulation results, showing that they are mostly superior to their alternatives in the literature and provide strong evidence for the claim that FDD is compatible with massive MIMO and inherits its long list of benefits.

Bibliography

[1] Ansuman Adhikary, Junyoung Nam, Jae-Young Ahn, and Giuseppe Caire. Joint spatial division and multiplexing: the large-scale array regime. *IEEE Trans. on Inform. Theory*, 59(10):6441–6463, 2013.

[2] Anum Ali, Nuria González-Prelcic, and Robert W Heath. Millimeter wave beam-selection using out-of-band spatial information. *IEEE Transactions on Wireless Communications*, 17(2):1038–1052, 2017.

[3] Donia Ben Amor, Michael Joham, and Wolfgang Utschick. Bilinear precoding for FDD massive MIMO system with imperfect covariance matrices. In *WSA 2020; 24th International ITG Workshop on Smart Antennas*, pages 1–6. VDE, 2020.

[4] Jinho Baik and Jack W Silverstein. Eigenvalues of large sample covariance matrices of spiked population models. *Journal of multivariate analysis*, 97(6):1382–1408, 2006.

[5] Waheed U Bajwa, Jarvis Haupt, Akbar M Sayeed, and Robert Nowak. Compressed channel sensing: A new approach to estimating sparse multipath channels. *Proceedings of the IEEE*, 98(6):1058–1076, 2010.

[6] Heinz H Bauschke and Jonathan M Borwein. On projection algorithms for solving convex feasibility problems. *SIAM review*, 38(3):367–426, 1996.

[7] Heinz H Bauschke, Patrick L Combettes, and Serge G Kruk. Extrapolation algorithm for affine-convex feasibility problems. *Numerical Algorithms*, 41(3):239–274, 2006.

[8] Alain Berlinet and Christine Thomas-Agnan. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.

[9] Dimitri P Bertsekas and Athena Scientific. *Convex optimization algorithms*. Athena Scientific Belmont, 2015.

[10] Peter J Bickel, Elizaveta Levina, et al. Covariance regularization by thresholding. *The Annals of Statistics*, 36(6):2577–2604, 2008.

[11] Emil Björnson, Jakob Hoydis, Marios Kountouris, and Merouane Debbah. Massive MIMO systems with non-ideal hardware: Energy efficiency, estimation, and capacity limits. *IEEE Transactions on Information Theory*, 60(11):7112–7139, 2014.

[12] Emil Björnson, Luca Sanguinetti, and Merouane Debbah. Massive MIMO with imperfect channel covariance information. In *2016 50th Asilomar Conference on Signals, Systems and Computers*, pages 974–978. IEEE, 2016.

[13] Federico Boccardi, Robert W Heath, Angel Lozano, Thomas L Marzetta, and Petar Popovski. Five disruptive technology directions for 5G. *IEEE Communications Magazine*, 52(2):74–80, 2014.

[14] William E Boyce, Richard C DiPrima, and Charles W Haines. *Elementary differential equations and boundary value problems*, volume 9. Wiley New York, 1969.

[15] James A Cadzow. Signal enhancement-a composite property mapping algorithm. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 36(1):49–62, 1988.

[16] T Tony Cai, Cun-Hui Zhang, Harrison H Zhou, et al. Optimal rates of convergence for covariance matrix estimation. *The Annals of Statistics*, 38(4):2118–2144, 2010.

[17] Tony Cai, Weidong Liu, and Xi Luo. A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association*, 106(494):594–607, 2011.

[18] Giuseppe Caire. On the ergodic rate lower bounds with applications to massive MIMO. *IEEE Transactions on Wireless Communications*, 17(5):3258–3268, 2018.

[19] Giuseppe Caire, Nihar Jindal, Mari Kobayashi, and Niranjay Ravindran. Multiuser MIMO achievable rates with downlink training and channel state feedback. *IEEE Transactions on Information Theory*, 56(6):2845–2866, 2010.

[20] Emmanuel J Candès and Michael B Wakin. An introduction to compressive sampling. *IEEE signal processing magazine*, 25(2):21–30, 2008.

[21] Jean-François Cardoso and Antoine Souloumiac. Blind beamforming for non-Gaussian signals. In *IEE proceedings F (radar and signal processing)*, volume 140, pages 362–370. IET, 1993.

[22] Renato LG Cavalcante, Lorenzo Miretti, and S Stańczak. Error bounds for FDD massive MIMO channel covariance conversion with set-theoretic methods. In *2018 IEEE Global Communications Conference (GLOBECOM)*, pages 1–7. IEEE, 2018.

[23] Peter WC Chan, Ernest S Lo, Ray R Wang, Edward KS Au, Vincent KN Lau, Roger S Cheng, Wai Ho Mow, Ross D Murch, and Khaled Ben Letaief. The evolution path of 4G networks: FDD or TDD? *IEEE Communications Magazine*, 44(12):42–50, 2006.

[24] Jie Chen and Xiaoming Huo. Theoretical results on sparse representations of multiple-measurement vectors. *IEEE Transactions on Signal Processing*, 54(12):4634–4643, 2006.

[25] Yuejie Chi, Yonina C Eldar, and Robert Calderbank. PETRELS: Parallel subspace estimation and tracking by recursive least squares from partial observations. *IEEE Transactions on Signal Processing*, 61(23):5947–5959, 2013.

[26] Junil Choi, David J Love, and Patrick Bidigare. Downlink training techniques for FDD massive MIMO systems: Open-loop and closed-loop training with memory. *IEEE Journal of Selected Topics in Signal Processing*, 8(5):802–814, 2014.

[27] Sheng-Fu Chuang, Wen-Rong Wu, and Yen-Ting Liu. High-resolution AoA estimation for hybrid antenna arrays. *IEEE Transactions on Antennas and Propagation*, 63(7):2955–2968, 2015.

[28] Kai Lai Chung, Ruth J Williams, and RJ Williams. *Introduction to stochastic integration*, volume 2. Springer, 1990.

[29] Earl A Coddington and Norman Levinson. *Theory of ordinary differential equations*. Tata McGraw-Hill Education, 1955.

[30] Shane F Cotter, Bhaskar D Rao, Kjersti Engan, and Kenneth Kreutz-Delgado. Sparse solutions to linear inverse problems with multiple measurement vectors. *IEEE Transactions on Signal Processing*, 53(7):2477–2488, 2005.

[31] Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.

[32] Jisheng Dai, An Liu, and Vincent KN Lau. FDD massive MIMO channel estimation with arbitrary 2d-array geometry. *IEEE Transactions on Signal Processing*, 66(10):2584–2599, 2018.

[33] Arash Gholami Davoodi and Syed Ali Jafar. Aligned image sets under channel uncertainty: Settling conjectures on the collapse of degrees of freedom under finite precision CSIT. *IEEE Transactions on Information Theory*, 62(10):5603–5618, 2016.

[34] Alexis Decurninge, Maxime Guillaud, and Dirk Slock. Riemannian coding for co-variance interpolation in massive MIMO frequency division duplex systems. In *2016 IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM)*, pages 1–5. IEEE, 2016.

[35] Alexis Decurninge, Maxime Guillaud, and Dirk TM Slock. Channel covariance estimation in massive MIMO frequency division duplex systems. In *Globecom Workshops (GC Wkshps), 2015 IEEE*, pages 1–6. IEEE, 2015.

[36] Vittorio Degli-Esposti, Veli-Matti Kolmonen, Enrico M Vitucci, and Pertti Vainikainen. Analysis and modeling on co-and cross-polarized urban radio propagation for dual-polarized MIMO wireless systems. *IEEE transactions on antennas and propagation*, 59(11):4247–4256, 2011.

[37] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977.

[38] Reinhard Diestel. *Graph theory (Graduate texts in mathematics)*, volume 173. Springer Heidelberg, 2005.

[39] Goran Dimic and Nicholas D Sidiropoulos. On downlink beamforming with greedy user selection: performance analysis and a simple new algorithm. *IEEE Transactions on Signal processing*, 53(10):3857–3868, 2005.

[40] Yacong Ding and Bhaskar D Rao. Dictionary learning-based sparse channel representation and estimation for FDD massive MIMO systems. *IEEE Transactions on Wireless Communications*, 17(8):5437–5451, 2018.

[41] David L Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.

[42] Noureddine El Karoui et al. Operator norm consistent estimation of large-dimensional sparse covariance matrices. *The Annals of Statistics*, 36(6):2717–2756, 2008.

[43] Yonina C Eldar and Holger Rauhut. Average case analysis of multichannel sparse recovery using convex relaxation. *IEEE Transactions on Information Theory*, 56(1):505–519, 2010.

[44] Jun Fang, Xingjian Li, Hongbin Li, and Feifei Gao. Low-rank covariance-assisted downlink training and channel estimation for FDD massive MIMO systems. *IEEE Transactions on Wireless Communications*, 16(3):1935–1947, 2017.

[45] Xiang Gao, Ove Edfors, Fredrik Rusek, and Fredrik Tufvesson. Linear pre-coding performance in measured very-large MIMO channels. In *Vehicular Technology Conference (VTC Fall), 2011 IEEE*, pages 1–5. IEEE, 2011.

[46] Zhen Gao, Linglong Dai, Zhaocheng Wang, and Sheng Chen. Spatially common sparsity based adaptive channel estimation and feedback for FDD massive MIMO. *IEEE Transactions on Signal Processing*, 63(23):6169–6183, 2015.

[47] Sergei A Goreinov, Eugene E Tyrtyshnikov, and Nickolai L Zamarashkin. A theory of pseudoskeleton approximations. *Linear algebra and its applications*, 261(1-3):1–21, 1997.

[48] Robert M Gray. *Toeplitz and circulant matrices: A review.* now publishers inc, 2006.

[49] Jiann-Ching Guey and L Daniel Larsson. Modeling and evaluation of MIMO systems exploiting channel reciprocity in TDD mode. In *IEEE 60th Vehicular Technology Conference, 2004. VTC2004-Fall. 2004*, volume 6, pages 4265–4269. IEEE, 2004.

[50] Robert Clifford Gunning and Hugo Rossi. *Analytic functions of several complex variables*, volume 368. American Mathematical Soc., 2009.

[51] Saeid Haghighatshoar and Giuseppe Caire. Massive MIMO channel subspace estimation from low-dimensional projections. *IEEE Transactions on Signal Processing*, 65(2):303–318, 2016.

[52] Saeid Haghighatshoar, Mahdi Barzegar Khalilsarai, and Giuseppe Caire. Multi-band covariance interpolation with applications in massive MIMO. In *arXiv preprint arXiv:1801.03714*, 2018.

[53] Saeid Haghighatshoar, Mahdi Barzegar Khalilsarai, and Giuseppe Caire. Multi-band covariance interpolation with applications in massive MIMO. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 386–390. IEEE, 2018.

[54] Yantao Han, Jiqing Ni, and Gaoke Du. The potential approaches to achieve channel reciprocity in FDD system with frequency correction algorithms. In *Communications and Networking in China (CHINACOM), 2010 5th International ICST Conference on*, pages 1–5. IEEE, 2010.

[55] Edwin Hewitt and Karl Stromberg. *Real and abstract analysis: a modern treatment of the theory of functions of a real variable.* Springer-Verlag, 2013.

[56] Bertrand M Hochwald and TL Maretta. Adapting a downlink array from uplink measurements. *IEEE Transactions on Signal Processing*, 49(3):642–653, 2001.

[57] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 1990.

[58] Jakob Hoydis, Cornelis Hoek, Thorsten Wild, and Stephan ten Brink. Channel measurements for large antenna arrays. In *Wireless Communication Systems (ISWCS), 2012 International Symposium on*, pages 811–815. IEEE, 2012.

[59] Klaus Hugl, Kimmo Kalliola, and Juha Laurila. Spatial reciprocity of uplink and downlink radio channels in FDD systems. *Proc. COST 273 Technical Document TD (02)*, 66:7, 2002.

[60] Klaus Hugl, Juha Laurila, and Ernst Bonek. Downlink beamforming for frequency division duplex systems. In *Seamless Interconnection for Universal Services. Global Telecommunications Conference. GLOBECOM'99.(Cat. No. 99CH37042)*, volume 4, pages 2097–2101. IEEE, 1999.

[61] David R Hunter and Kenneth Lange. A tutorial on MM algorithms. *The American Statistician*, 58(1):30–37, 2004.

[62] Zhiyuan Jiang, Andreas F Molisch, Giuseppe Caire, and Zhisheng Niu. Achievable rates of FDD massive MIMO systems with spatial channel correlation. *IEEE Transactions on Wireless Communications*, 14(5):2868–2882, 2015.

[63] Nihar Jindal. MIMO broadcast channels with finite-rate feedback. *IEEE Transactions on information theory*, 52(11):5045–5060, 2006.

[64] Michel Journée, Yurii Nesterov, Peter Richtárik, and Rodolphe Sepulchre. Generalized power method for sparse principal component analysis. *Journal of Machine Learning Research*, 11(2), 2010.

[65] Florian Kaltenberger, David Gesbert, Raymond Knopp, and Marios Kountouris. Correlation and capacity of measured multi-user MIMO channels. In *Personal, Indoor and Mobile Radio Communications, 2008. PIMRC 2008. IEEE 19th International Symposium on*, pages 1–5. IEEE, 2008.

[66] Sato Ken-Iti. *Lévy processes and infinitely divisible distributions*. Cambridge university press, 1999.

[67] H Kim and Mats Viberg. Two decades of array signal processing research. *IEEE signal magazine*, 13(4):67–94, 1996.

[68] Fima C Klebaner. *Introduction to stochastic calculus with applications*. World Scientific Publishing Company, 2005.

[69] Mari Kobayashi, Nihar Jindal, and Giuseppe Caire. Training and feedback optimization for multiuser MIMO downlink. *IEEE Transactions on Communications*, 59(8):2228–2240, 2011.

[70] Hamid Krim and Mats Viberg. Two decades of array signal processing research: the parametric approach. *IEEE signal processing magazine*, 13(4):67–94, 1996.

[71] Ping-Heng Kuo, HT Kung, and Pang-An Ting. Compressive sensing based channel feedback protocols for spatially-correlated massive antenna arrays. In *Wireless Communications and Networking Conference (WCNC), 2012 IEEE*, pages 492–497. IEEE, 2012.

[72] Andreas E Kyprianou. The lévy–itô decomposition and path structure. In *Fluctuations of Lévy Processes with Applications*, pages 35–69. Springer, 2014.

[73] Persefoni Kyritsi, Donald C Cox, Reinaldo A Valenzuela, and Peter W Wolniansky. Correlation analysis based on MIMO channel measurements in an indoor environment. *IEEE Journal on Selected areas in communications*, 21(5):713–720, 2003.

[74] HJ Landau. Necessary density conditions for sampling and interpolation of certain entire functions. *Acta Mathematica*, 117(1):37–52, 1967.

[75] Erik G Larsson, Ove Edfors, Fredrik Tufvesson, and Thomas L Marzetta. Massive MIMO for next generation wireless systems. *IEEE Communications Magazine*, 52(2):186–195, 2014.

[76] Junho Lee, Gye-Tae Gil, and Yong H Lee. Channel estimation via orthogonal matching pursuit for hybrid MIMO systems in millimeter wave communications. *IEEE Transactions on Communications*, 64(6):2370–2386, 2016.

[77] David J Love, Robert W Heath, and Thomas Strohmer. Grassmannian beamforming for multiple-input multiple-output wireless systems. *IEEE transactions on information theory*, 49(10):2735–2747, 2003.

[78] Angel Lozano, Robert W Heath, and Jeffrey G Andrews. Fundamental limits of cooperation. *IEEE Transactions on Information Theory*, 59(9):5213–5226, 2013.

[79] Steffen Malkowsky, Joao Vieira, Liang Liu, Paul Harris, Karl Nieman, Nikhil Kundargi, Ian C Wong, Fredrik Tufvesson, Viktor Öwall, and Ove Edfors. The world's first real-time testbed for massive MIMO: Design, implementation, and validation. *IEEE Access*, 5:9073–9088, 2017.

[80] Jonathan H Manton. Optimization algorithms exploiting unitary constraints. *IEEE Transactions on Signal Processing*, 50(3):635–650, 2002.

[81] Vladimir Alexandrovich Marchenko and Leonid Andreevich Pastur. Distribution of eigenvalues for some sets of random matrices. *Matematicheskii Sbornik*, 114(4):507–536, 1967.

[82] Thomas L Marzetta. How much training is required for multiuser MIMO? In *Fortieth Asilomar Conference on Signals, Systems and Computers, 2006. ACSSC'06.*, pages 359–363. IEEE, 2006.

[83] Thomas L. Marzetta. Noncooperative cellular wireless with unlimited numbers of base station antennas. *IEEE Trans. on Wireless Commun.*, 9(11):3590–3600, Nov. 2010.

[84] Thomas L Marzetta, Erik G Larsson, Hong Yang, and Hien Quoc Ngo. *Fundamentals of Massive MIMO*. Cambridge University Press, 2016.

[85] MATLAB. *version 9.9.0 (R2010b)*. The MathWorks Inc., Natick, Massachusetts, 2020.

[86] Lorenzo Miretti, Renato LG Cavalcante, and Slawomir Stańczak. Downlink channel spatial covariance estimation in realistic FDD massive MIMO systems. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 161–165. IEEE, 2018.

[87] Lorenzo Miretti, Renato Luis Garrido Cavalcante, and Slawomir Stanczak. FDD massive MIMO channel spatial covariance conversion using projection methods. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3609–3613. IEEE, 2018.

[88] Leon Mirsky. A trace inequality of John von Neumann. *Monatshefte für mathematik*, 79(4):303–306, 1975.

[89] O Najim, Pascal Vallet, Guillaume Ferré, and Xavier Mestre. On the statistical performance of MUSIC for distributed sources. In *2016 IEEE Statistical Signal Processing Workshop (SSP)*, pages 1–5. IEEE, 2016.

[90] Junyoung Nam, Ansuman Adhikary, Jae-Young Ahn, and Giuseppe Caire. Joint spatial division and multiplexing: Opportunistic beamforming, user grouping and simplified downlink scheduling. *IEEE J. of Sel. Topics in Sig. Proc. (JSTSP)*, 8(5):876–890, 2014.

[91] Christos H Papadimitriou and Kenneth Steiglitz. *Combinatorial optimization: algorithms and complexity.* Courier Corporation, 1998.

[92] Sungwoo Park and Robert W Heath. Spatial channel covariance estimation for the hybrid MIMO architecture: A compressive sensing-based approach. *IEEE Transactions on Wireless Communications*, 17(12):8047–8062, 2018.

[93] Juan Pascual-García, José-María Molina-García-Pardo, Maria-Teresa Martinez-Ingles, Jose-Victor Rodriguez, and Noelia Saurin-Serrano. On the importance of diffuse scattering model parameterization in indoor wireless channels at mm-wave frequencies. *Ieee Access*, 4:688–701, 2016.

[94] Josip E Peajcariaac and Yung Liang Tong. *Convex functions, partial orderings, and statistical applications.* Academic Press, 1992.

[95] Yuexing Peng, Yonghui Li, and Peng Wang. An enhanced channel estimation method for millimeter wave systems with massive antenna arrays. *IEEE Communications Letters*, 19(9):1592–1595, 2015.

[96] Antonios Pitarokoilis, Saif Khan Mohammed, and Erik G Larsson. Uplink performance of time-reversal MRC in massive MIMO systems subject to phase noise. *IEEE Transactions on Wireless Communications*, 14(2):711–723, 2014.

[97] Ada SY Poon, Robert W Brodersen, and David NC Tse. Degrees of freedom in multiple-antenna channels: A signal space approach. *IEEE Transactions on Information Theory*, 51(2):523–536, 2005.

[98] Xiongbin Rao and Vincent KN Lau. Distributed compressive CSIT estimation and feedback for FDD multi-user massive MIMO systems. *IEEE Transactions on Signal Processing*, 62(12):3261–3271, 2014.

[99] Pradeep Ravikumar, Martin J Wainwright, Garvesh Raskutti, Bin Yu, et al. High-dimensional covariance estimation by minimizing ℓ_1-penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.

[100] Cássio B Ribeiro, Esa Ollila, and Visa Koivunen. Propagation parameter estimation in MIMO systems using mixture of angular distributions model. In *Proceedings.(ICASSP'05). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005.*, volume 4, pages iv–885. IEEE, 2005.

[101] Andreas Richter. Estimation of radio channel parameters: Models and algorithms. ISLE, 2005.

[102] Andreas Richter, Jussi Salmi, and Visa Koivunen. Distributed scattering in radio channels and its contribution to MIMO channel capacity. In *2006 First European Conference on Antennas and Propagation*, pages 1–7. IEEE, 2006.

[103] Andreas Richter and Reiner S Thomä. Parametric modeling and estimation of distributed diffuse scattering components of radio channels. 2003.

[104] Walter Rudin et al. *Principles of mathematical analysis*, volume 3. McGraw-hill New York, 1964.

[105] M Salehi and J Proakis. Digital communications. *McGraw-Hill Education*, 31:32, 2007.

[106] J Salmi, J Poutanen, K Haneda, A Richter, V-M Kolmonen, P Vainikainen, and AF Molisch. Incorporating diffuse scattering in geometry-based stochastic MIMO channel models. In *Proceedings of the Fourth European Conference on Antennas and Propagation*, pages 1–5. IEEE, 2010.

[107] Matilde Sánchez-Fernández, Vahid Jamali, Jaime Llorca, and Antonia Tulino. Gridless multidimensional angle of arrival estimation for arbitrary 3D antenna arrays. *arXiv preprint arXiv:2008.12323*, 2020.

[108] Akbar M Sayeed. Deconstructing multiantenna fading channels. *IEEE Transactions on Signal Processing*, 50(10):2563–2579, 2002.

[109] Ralph Schmidt. Multiple emitter location and signal parameter estimation. *IEEE transactions on antennas and propagation*, 34(3):276–280, 1986.

[110] Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer Science & Business Media, 2003.

[111] Stefania Sesia, Issam Toufik, and Matthew Baker. *LTE-the UMTS long term evolution: from theory to practice.* John Wiley & Sons, 2011.

[112] Min Soo Sim, Jeonghun Park, Chan-Byoung Chae, and Robert W Heath. Compressed channel feedback for correlated massive MIMO systems. *Journal of Communications and Networks*, 18(1):95–104, 2016.

[113] Martin Slawski, Matthias Hein, et al. Non-negative least squares for high-dimensional linear models: Consistency and sparse recovery without regularization. *Electronic Journal of Statistics*, 7:3004–3056, 2013.

[114] Suraj Srivastava, Amrita Mishra, Anupama Rajoriya, Aditya K Jagannatham, and Gerd Ascheid. Quasi-static and time-selective channel estimation for block-sparse millimeter wave hybrid MIMO systems: Sparse bayesian learning (SBL) based approaches. *IEEE Transactions on Signal Processing*, 67(5):1251–1266, 2018.

[115] Gerhard Steinböck, Jean-Marc Conrat, Troels Pedersen, and Bernard H Fleury. On initialization and search procedures for iterative highresolution channel parameter estimators. In *COST2100 9th Management Committee Meeting*, number 09, page 956, 2009.

[116] Petre Stoica, Prabhu Babu, and Jian Li. SPICE: A sparse covariance-based estimation method for array processing. *IEEE Transactions on Signal Processing*, 59(2):629–638, 2010.

[117] Petre Stoica and Arye Nehorai. MUSIC, maximum likelihood, and Cramer-Rao bound. *IEEE Transactions on Acoustics, speech, and signal processing*, 37(5):720–741, 1989.

[118] Ying Sun, Prabhu Babu, and Daniel P Palomar. Majorization-minimization algorithms in signal processing, communications, and machine learning. *IEEE Transactions on Signal Processing*, 65(3):794–816, 2016.

[119] Till Tantau. *The TikZ and PGF Packages.*

[120] Terence Tao and Van H Vu. *Additive combinatorics*, volume 105. Cambridge University Press, 2006.

[121] RS Thomä, M Landmann, A Richter, and U Trautwein. Multidimensional high-resolution channel sounding. In *Smart Antennas in Europe–State-of-the-Art*, volume 3. Hindawi Publishing Corporation, 2005.

[122] Michael E Tipping. Sparse bayesian learning and the relevance vector machine. *Journal of machine learning research*, 1(Jun):211–244, 2001.

[123] David Tse and Pramod Viswanath. *Fundamentals of wireless communication.* Cambridge university press, 2005.

[124] William T Tutte. The factorization of linear graphs. *Journal of the London Mathematical Society*, 1(2):107–111, 1947.

[125] Michael Unser. Sampling-50 years after Shannon. *Proceedings of the IEEE*, 88(4):569–587, 2000.

[126] C Van Den Berg, Jens Peter Reus Christensen, and Paul Ressel. *Harmonic Analysis on Semigroups: Theory of Positive Definite and Related Functions*, volume 100. Springer Science & Business Media, 2012.

[127] Harry L Van Trees. *Optimum array processing: Part IV of detection, estimation and modulation theory*, volume 1. Wiley Online Library, 2002.

[128] Harry L Van Trees. *Optimum array processing: Part IV of detection, estimation, and modulation theory.* John Wiley & Sons, 2004.

[129] Deepak Vasisht, Swarun Kumar, Hariharan Rahul, and Dina Katabi. Eliminating channel feedback in next-generation cellular networks. In *Proceedings of the 2016 conference on ACM SIGCOMM 2016 Conference*, pages 398–411. ACM, 2016.

[130] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

[131] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

[132] Hubert Stanley Wall. *Analytic theory of continued fractions.* Courier Dover Publications, 2018.

[133] Eugene P Wigner. On the distribution of the roots of certain symmetric matrices. *Annals of Mathematics*, pages 325–327, 1958.

[134] Hongxiang Xie, Feifei Gao, Shi Jin, Jun Fang, and Ying-Chang Liang. Channel estimation for TDD/FDD massive MIMO systems with channel covariance computing. *arXiv preprint arXiv:1710.00704*, 2017.

[135] Hongxiang Xie, Feifei Gao, Shun Zhang, and Shi Jin. A unified transmission strategy for TDD/FDD massive MIMO systems with spatial basis expansion model. *IEEE Transactions on Vehicular Technology*, 66(4):3170–3184, 2017.

[136] Wei Xu, Xian Wu, Xiaodai Dong, Hua Zhang, and Xiaohu You. Dual-polarized massive MIMO systems under multi-cell pilot contamination. *IEEE Access*, 4:5998–6013, 2016.

[137] Haifan Yin, David Gesbert, Miltiades Filippou, and Yingzhuang Liu. A coordinated approach to channel estimation in large-scale multiple-antenna systems. *IEEE Journal on selected areas in communications*, 31(2):264–273, 2013.

[138] Alan L Yuille and Anand Rangarajan. The concave-convex procedure (cccp). In *Advances in neural information processing systems*, pages 1033–1040, 2002.

[139] Benno Zerlin, Michael Joham, Wolfgang Utschick, and Josef A Nossek. Covariance-based linear precoding. *IEEE Journal on Selected Areas in Communications*, 24(1):190–199, 2005.

[140] Cheng Zhang, Zhaohua Lu, Yongming Huang, Jing Zhang, and Luxi Yang. Statistical beamforming for FDD massive mimo downlink systems. In *Communications in China (ICCC), 2015 IEEE/CIC International Conference on*, pages 1–6. IEEE, 2015.

[141] Lizhong Zheng and David N. C. Tse. Communication on the Grassmann manifold: A geometric approach to the noncoherent multiple-antenna channel. *IEEE Transactions on Information Theory*, 48(2):359–383, 2002.

[142] Zhihui Zhu and Michael B Wakin. On the asymptotic equivalence of circulant and Toeplitz matrices. *IEEE Transactions on Information Theory*, 63(5):2975–2992, 2017.

[143] Daniel Zwillinger. *Handbook of differential equations*, volume 1. Gulf Professional Publishing, 1998.