DISSERTATION

LINEAR SYSTEM DESIGN FOR COMPRESSION AND FUSION

Submitted by

Yuan Wang

Department of Statistics

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Fall 2013

Doctoral Committee:

    Advisor: Haonan Wang
    Co-Advisor: Louis L. Scharf

    F. Jay Breidt
    Rockey J. Luo

ABSTRACT

LINEAR SYSTEM DESIGN FOR COMPRESSION AND FUSION

This is a study of measurement compression and fusion design. The idea common to both problems is that measurements can often be linearly compressed into lower-dimensional spaces without introducing too much excess mean-squared error or excess volume in a concentration ellipse. The question is how to design these compressions to minimize the excesses at any given dimension.

The first part of this work is motivated by sensing and wireless communication, where data compression or dimension reduction may be used to reduce the required communication bandwidth. The high-dimensional measurements are converted into low-dimensional representations through linear compression. Our aim is to compress a noisy measurement, allowing for the fact that the compressed measurement will be transmitted over a noisy channel. We review optimal compression with no transmission noise and show its connection with canonical coordinates. When the compressed measurement is transmitted with noise, we give the closed-form expression for the optimal compression matrix with respect to the trace and determinant of the error covariance matrix. We show that the solutions are canonical coordinate solutions, scaled by coefficients which account for canonical correlations and transmission noise variance, followed by a coordinate transformation into the sub-dominant invariant subspace of the channel noise.

The second part of this work is a problem of integrating multiple sources of measurements. We consider two multiple-input-multiple-output (MIMO) channels, a primary channel and a secondary channel, with dependent input signals. The primary channel carries the signal of interest, and the secondary channel carries a signal that shares a

joint distribution with the primary signal. The problem of particular interest is designing the secondary channel, with a fixed primary channel. We formulate the problem as an optimization problem, in which the optimal secondary channel maximizes an information-based criterion. An analytical solution is provided in a special case. Two fast-to-compute algorithms, one extrinsic and the other intrinsic, are proposed to approximate the optimal solutions in general cases. In particular, the intrinsic algorithm exploits the geometry of the unit sphere, a manifold embedded in Euclidean space. The performances of the proposed algorithms are examined through a simulation study. A discussion of the choice of dimension for the secondary channel is given, leading to rules for dimension reduction.

ACKNOWLEDGEMENTS

DEDICATION

To my motherland.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

1.1 Communication Systems and MIMO Channels

Communication devices such as radios, wire or cordless telephones, wifi, and remote controllers, have emerged in everyone's daily life. Regardless of the functions of these devices, the principle is essentially the transfer of information among different objects over time and space. Figure 1.1 describes a general communication system: The input signal, which could be a human voice, a television picture, or an electronic waveform, is modified by the transmitter for efficient transmission. The channel is a medium through which the transmitted signal is sent to the receiver. The receiver then reprocesses the received signal by undoing the signal modifications made at the transmitter and in the channel. Modern communication systems are classified into two categories: wireline communication systems and wireless communication system. The wireline channel generally requires a physical transmission medium such as cable or wire, while the wireless system sends signals without any electrical conductor.
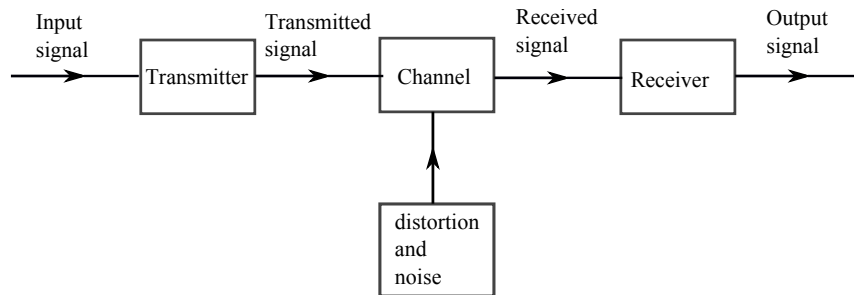
Figure 1.1: A communication system.

One of the most widely used channels is the multiple-input-multiple-output (MIMO) channel. MIMO channels arise in many different systems, for example, the wireless channel with multiple antennas at both transmitter and receiver (see Foschini and Gans

1

[1998], Tarokh et al. [1998], and Telatar [1999]), the wireline digital subscriber line channel with multiple twisted-pairs of telephone subscriber lines. Readers can refer to Ginis and Cioffi [2002], Honig et al. [1990], Lee and Petersen [1976], Salz [1985] for introductions to various communication channels. Generally speaking, a MIMO channel has multiple dimensions at both the transmitter and receiver. A generic MIMO channel, as shown in Figure 1.2, can be expressed as

$$\boldsymbol{x} = \boldsymbol{H}\boldsymbol{\theta} + \boldsymbol{u}, \qquad (1.1.1)$$

where $\boldsymbol{\theta} \in \mathbb{R}^p$ is the input signal, $\boldsymbol{H} \in \mathbb{R}^{n \times p}$ is the channel matrix, $\boldsymbol{x} \in \mathbb{R}^n$ is the output signal, and $\boldsymbol{u} \in \mathbb{R}^n$ is the channel noise. Note that the block diagram in Figure 1.2 is used to describe the model throughout this dissertation. The linear functions are represented by blocks and the circles stand for additive noise. The line with a single arrowhead depicts functional flow from left to right, with the dimension of the signal in the path shown on top of the slash.
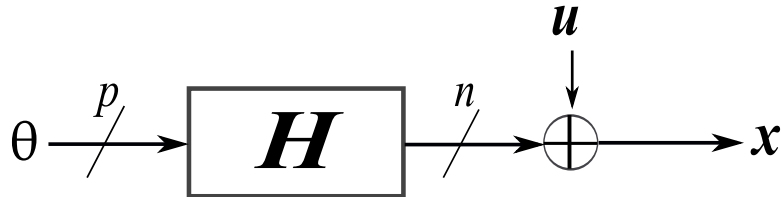
Figure 1.2: A generic MIMO channel.

Next, we will introduce several design problems common in communication with MIMO channel.

1.2 MIMO Channel Design Review

For the MIMO channel (1.1.1), the channel state information refers to the channel matrix $\boldsymbol{H}$ and the statistical properties of the noise $\boldsymbol{u}$. This information describes how a signal is propagated from the transmitter to the receiver. Throughout this dissertation, we will assume perfect channel state information is available. In this case, the

transmission and reception can be adapted to each channel realization using signal processing techniques. The MIMO channel design, which controls the phase and amplitude of the signal, is important for efficient transmission. The first study dates back to the 1970s in Lee and Petersen [1976] and Salz [1985]. In this section, we will introduce the design of the transmitter and the receiver, respectively. These techniques have applications in areas such as wireless communication, seismology, acoustics, radar, sonar, and biomedicine. See Monzingo and Miller [1980], Van Veen and Buckley [1988], Johnson and Dudgeon [1993], Krim and Viberg [1996], Van Trees [2002], and references therein.

1.2.1 Rank-Reduced Filtering

In signal processing, filtering is used to produce an estimate of a desired random process by filtering an observed noisy process. As shown in Figure. 1.3, the linear filter estimates the input signal by linearly combining the elements of the received signal via a matrix $\boldsymbol{B}$ as

$$\widehat{\boldsymbol{\theta}} = \boldsymbol{B}\boldsymbol{x}. \tag{1.2.1}$$

It is well known that the optimal linear filtering that minimizes the mean squared error of $\boldsymbol{\theta}$ is the Wiener filter or LMMSE filter. Recently, rank-reduced filtering has emerged in signal processing problems where data or model reduction, or high computational efficiency is required. The reduced-rank LMMSE (or Wiener filter) is first brought up in Scharf [1990] by minimizing mean squared error, and further developed in Scharf and Thomas [1998], Scharf and Mullis [1998], and Schreier and Scharf [2006]. Hua et al. [2001] later give a unified review for the class of optimal reduced rank estimators with respect to three commonly used measures of loss: trace, determinant, and weighted trace of the error covariance. The optimal rank-reduced filter matrix is given by the singular value decomposition of the coherence matrix between the channel input $\boldsymbol{\theta}$ and the output $\boldsymbol{x}$. In fact, the optimal filtering matrix $\boldsymbol{B}_m$ with rank m returns the first m canonical

coordinates of $\boldsymbol{x}$ corresponding to the first m largest canonical correlations. These may be half canonical coordinates (see Scharf [1990]) or (full) canonical coordinates (Hua et al. [2001]).
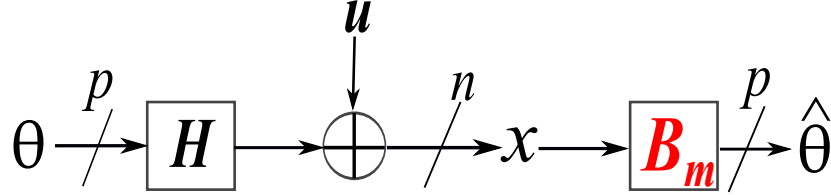
Figure 1.3: Reduced rank filtering.

1.2.2   Precoding Design

Consider the generic MIMO channel (1.1.1). The linear transmitter controls the phase and amplitude of the input signal $\boldsymbol{\theta}$ via the matrix $\boldsymbol{P}$ as

$$s = P\theta, \qquad (1.2.2)$$

The signal $\boldsymbol{s}$ is sent over a MIMO channel, and the resulted output signal is

$$x = HP\theta + u, \qquad (1.2.3)$$

as shown in Figure. 1.4.
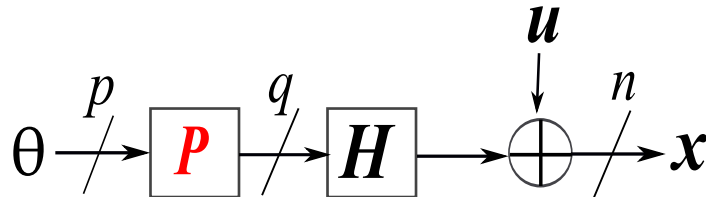
Figure 1.4: Precoding.

The transmit processing $\boldsymbol{P}$ is called precoding. In general, the precoding is a signal processing technique that operates on the signal before transmission. Mathematically, the function of the precoding matrix is the same as the transmitter. One of the first results on precoder design is introduced in Vojcic and Jang [1998] for CDMA system,

4

with a fixed linear MIMO channel. Linear precoding has been studied extensively for its simplicity and optimal performance from an information viewpoint (see Skoglund and Jöngren [2003]). Different authors have considered various criteria: channel capacity, mean squared error, signal-to-noise-ratio, mutual information, or bit-error-rate. (See Visotsky and Madhow [2001], Sampath et al. [2001], Scaglione et al. [2002], Palomar et al. [2003], Ding et al. [2003], Cover and Thomas [2005], Wiesel et al. [2006], Vu and Paulraj [2007] and references therein.)

One of the differences between transmitter design and receiver design is that, to conserve the total transmit power, the transmitter or precoder must satisfy the power constraint

$$\text{tr}(\boldsymbol{P}\boldsymbol{P}^T) \leq c \qquad (1.2.4)$$

for some positive constant c. In other words, the sum of transmitted power over all subchannels is bounded by c. In many scenarios, the optimal precoder assigns power to subchannels in a waterfilling manner and gives more power to strong subchannels and less or no power to weak subchannels.

1.3   Compression and Fusion Design

In this section, we will briefly introduce the two design problems of interest in this dissertation. The idea common to both problems is that measurements can often be linearly compressed into lower-dimensional measurement spaces without introducing too much excess mean-squared error or excess volume in a concentration ellipse.

1.3.1   Compression Design

In a communication system, when the communication bandwidth is limited, one can precompress measurements to a lower-dimensional space before transmission. Such data compression or dimension reduction greatly reduces the communication burden. In this
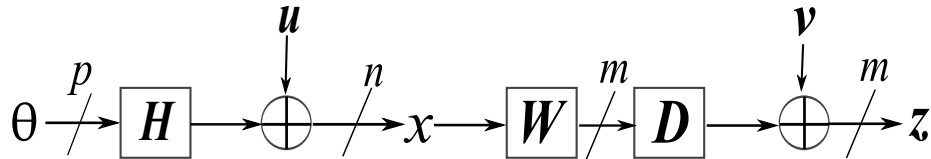
Figure 1.5: The linear compression with transmission noise.

work, we are interested in designing the linear compression matrix when measurements are noisy, and the compressed measurement is to be transmitted over a noisy channel. The system is given in Figure 1.5, where $\boldsymbol{\theta}$ is the signal of interest, $\boldsymbol{x}$ is a noisy measurement of $\boldsymbol{\theta}$. The measurement $\boldsymbol{x}$ is compressed through the matrix $\boldsymbol{W}$ and then transmitted over a MIMO channel with channel matrix $\boldsymbol{D}$ and additive noise $\boldsymbol{v}$. The key term of interest is the compression matrix $\boldsymbol{W}$.

The special case with no measurement noise or transmission noise is the standard framework for compressed sensing, (see Candès et al. [2006] and Donoho [2006]), so stability studies for compressed sensing establish performance guarantees when measurements are noise-free. When the compressed measurement is to be transmitted noise-free, the compression-estimation task is essentially a canonical correlation analysis problem. We further notice that dimension-reduction design is then equivalent to reduced-rank filtering and estimation introduced in Section 1.2.1. Reduced-rank estimators transform the measurement to a lower-dimensional measurement space, with dimension equaling rank. When there is no measurement noise but with transmission noise, Carson et al. [2012] derived the optimal projection of a high-dimensional, noise-free signal to maximize mutual information between the signal and the compressed measurement. This problem is a special case of the precoding and equalization design problem in MIMO communication systems with identity channel matrix, which has been discussed in Section 1.2.2.

The compression design here is for compression of a noisy measurement, followed by transmission of the compressed measurement over a noisy channel. This is the model

6

considered in Schizas et al. [2007] and here in Chapter 2. Schizas et al. [2007] provide the optimal compression and filtering operators to minimize the mean squared error. We minimize determinant of error covariance, which in the multivariate normal case maximizes mutual information. Moreover, we show that the mean-squared error and mutual information solutions are both scaled and rotated canonical coordinate solutions. They are different by their design of the scaling matrix and the choice of the coordinate system.

1.3.2 Fusion Design

For the design problems introduced in the proceeding sections, there is a single source of measurement. In the fusion problem, we allow multiple sources of measurements. More specifically, we consider a two-channel system as shown in Figure 1.6. The top panel is a linear MIMO channel with input signal $\boldsymbol{\theta}$ and a noisy measurement $\boldsymbol{x}$, and the bottom panel is another linear channel with the input signal $\boldsymbol{\phi}$ and a noisy measurement $\boldsymbol{y}$. We assume that the input signals $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are correlated random quantities and the signal $\boldsymbol{\theta}$ is the term of interest. We call $\boldsymbol{x}$ the primary channel and $\boldsymbol{y}$ the secondary channel since $\boldsymbol{\phi}$ plays the role of a nuisance parameter. Such a system is quite common in practice when there exist multiple sources of measurements. For example, the elements of the primary signal $\boldsymbol{\theta}$ may be the complex scattering coefficients of several radar-scattering targets and the elements of the secondary signal $\boldsymbol{\phi}$ may be intensities in an optical map of these same optical-scattering targets. The measurement $\boldsymbol{x}$ is then a range-doppler map and the measurement $\boldsymbol{y}$ is an optical image.

The objective in this study is to design the secondary channel, with the primary channel fixed, such that combining the measurements $\boldsymbol{x}$ and $\boldsymbol{y}$ brings the largest improvement in differential information rate. We will design the secondary channel matrix $\boldsymbol{G}$, or equivalently the precoding matrix for a channel with identity channel matrix. An information-based criterion will be used to quantify the gain, subject to a total power

7
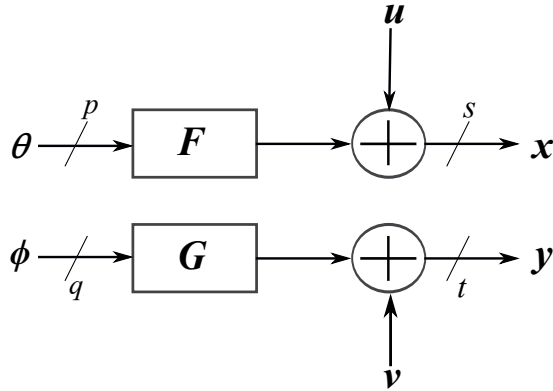
Figure 1.6: A two-channel system with two linear channels.

constraint. More details will be given in Chapter 3. We must point out that the main difference between the fusion design and the compression design is the existence of the primary channel. In fact, without the primary channel, this is simply the precoding design problem introduced in Section 1.2.2 and a special case of compression design. The optimal secondary channel depends on the primary channel, the joint distribution of $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, and the relation between $\boldsymbol{y}$ and $\boldsymbol{\phi}$. The optimization problem is more complicated than the compression design. We obtain the analytical solution, in Section 3.3, when the conditional covariance of $\boldsymbol{\phi}$ given $\boldsymbol{\theta}$ is proportional to the identity matrix. For general cases, we approximate the optimal channel by numerical algorithms in Section 3.4.

1.4 Notation

The set of length m real vectors is denoted by $\mathbb{R}^m$ and the set of $m \times n$ real matrices is denoted $\mathbb{R}^{m \times n}$. Bold upper case letters denote matrices, boldface lower case letters denote column vectors, and italics denote scalars. The scalar x_i denotes the ith element of vector $\boldsymbol{x}$, and $\boldsymbol{X}_{i,j}$ denotes the element of $\boldsymbol{X}$ at row i and column j. The diagonal matrix with diagonal elements $\boldsymbol{x}$ is denoted as $\mathrm{diag}(\boldsymbol{x})$. The $n \times n$ identity matrix is denoted by $\boldsymbol{I}_n$. The transpose, inverse, pseudo inverse, trace and determinant of a matrix are denoted by $(\cdot)^T$, $(\cdot)^{-1}$, $(\cdot)^+$, $\mathrm{tr}(\cdot)$ and $\det(\cdot)$, respectively.

A covariance matrix is denoted by bold upper case $\boldsymbol{Q}$ with specified subscripts: $\boldsymbol{Q}_{zz}$ denotes the covariance matrix of a random vector $\boldsymbol{z}$; $\boldsymbol{Q}_{z_1 z_2}$ is the cross-covariance matrix between $\boldsymbol{z}_1$ and $\boldsymbol{z}_2$; $\boldsymbol{Q}_{z_1 z_1 | z_2}$ is the conditional covariance matrix of $\boldsymbol{z}_1$ given $\boldsymbol{z}_2$.

OPTIMUM COMPRESSION WITH TRANSMISSION OVER A NOISY CHANNEL[1]

2.1 Introduction

In a distributed sensor network, one can precompress observations to lower dimensional measurements before transmitting them around the network or to a fusion center. Such data compression or dimension reduction reduces the communication burden, but increases mean-squared error and reduces information rate. In this chapter, we are interested in designing the linear compression matrix that minimizes the mean-squared error or maximizes the information rate at the optimal compression ratio, under a power constraint.
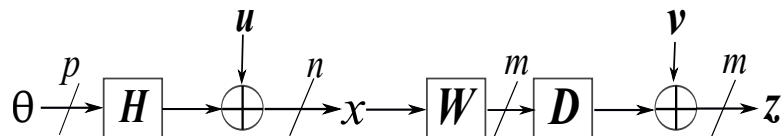
Figure 2.1: Linear compression of a noisy measurement $\boldsymbol{x}$ with channel noise $\boldsymbol{v}$.

The diagram of Figure 2.1 frames the problem of interest in this chapter. In this figure, $\boldsymbol{\theta} \in \mathbb{R}^p$ is a signal of interest. The signal $\boldsymbol{\theta}$ is carried through a sensor by a linear transformation $\boldsymbol{H} \in \mathbb{R}^{n \times p}$ and then observed as the noisy and transformed measurement $\boldsymbol{x} = \boldsymbol{H}\boldsymbol{\theta} + \boldsymbol{u} \in \mathbb{R}^n$. This noisy measurement $\boldsymbol{x}$ is to be compressed with the linear transformation $\boldsymbol{W} \in \mathbb{R}^{m \times n}(m < n)$ and then transmitted through a noisy channel. The channel transforms the measurement by a channel matrix $\boldsymbol{D} \in \mathbb{R}^{m \times m}$ and adds noise to produce a measurement $\boldsymbol{z} \in \mathbb{R}^m$. Our goal is to design the compressor $\boldsymbol{W}$ so that the noisy and compressed measurement $\boldsymbol{z}$ may be processed for an estimator

---

of the signal $\boldsymbol{\theta}$ whose error covariance has minimum trace or minimum determinant. The minimum trace solution minimizes mean-squared error of the estimate and the minimum determinant solution minimizes volume of the error concentration ellipsoid. In the Gaussian case, it maximizes differential information rate.

We are not the first to consider this problem and its variants. In fact as we will show, this chapter is an extension of the original work of Schizas, Giannakis, and Luo Schizas et al. [2007], which in turn generalizes the work of Scharf [1990], Scharf and Thomas [1998], Scharf and Mullis [1998], Hua et al. [2001], Schreier and Scharf [2006], Scaglione et al. [2002], and Pérez-Cruz et al. [2010]. The innovation of our work is this. First we replace the *mean-squared error criterion* of Schizas et al. [2007] with a maximum information rate criterion, and second we show that our designs and theirs may be cast as *scaled and rotated canonical coordinate designs*. This finding is important, for it generalizes the theory of canonical coordinates to a much more general class of problems than the class for which they were originally designed in Hotelling [1936]. The maximum information rate designs of this work require a different proof technique than the proof technique of Schizas et al. [2007].

Let us place our work in the context of prior art, by again making reference to Figure 2.1. The problem addressed by Schizas et al. [2007] is to design the compression matrix $\boldsymbol{W}$ so that the measurement $\boldsymbol{z}$ may be filtered to produce a minimum mean-squared error estimate of the signal $\boldsymbol{\theta}$. We generalize this problem to the maximization of information rate and show that canonical coordinates are central to both criteria. The literature on reduced-rank filtering assumes that the channel matrix $\boldsymbol{D}$ is identity and the channel noise $\boldsymbol{v}$ is zero. The result of Carson et al. [2012], which assumes the sensor matrix $\boldsymbol{H}$ is identity, the sensor noise $\boldsymbol{u}$ is zero, and the channel matrix $\boldsymbol{D}$ is identity, is a special case of precoding and equalizing.

So we may summarize by saying that the theory of *canonical coordinates* treats the problem of compression when there is *noise at the input to the compressor* and

the theory of *scaled and rotated canonical coordinates* developed in this work treats the problem of compression when there is *noise at the input and the output of the compressor*. Noise at the output brings an important element of design to the compression problem, for it forces a constraint on the power out of the compressor $\boldsymbol{W}$, a constraint that leads to rather complicated reasoning about Lagrangians and the KKT conditions for optimality, as for example in the prior work of Schizas et al. [2007], Scaglione et al. [2002], and Pérez-Cruz et al. [2010].

The rest of this chapter is organized as follows. In Section 2.2, we briefly introduce the problem of interest. In Section 2.3, in the channel-noise-free case, the compression matrix returns half canonical coordinates for trace minimization, and full canonical coordinates for determinant minimization. In Section 2.4, when the compressed measurement is transmitted over a noisy channel, the compression matrix for trace or determinant minimization returns a scaled and rotated canonical coordinate design. Moreover, the scaling matrix, which accounts for canonical correlations and channel noise variance, has a mercury/waterfilling interpretation. In Section 2.5, we extend the trace and determinant criteria to differentiable functions of the error covariance and establish a unified factorization for the optimal compression matrix. Section 2.6 concludes the chapter.

2.2 Problem Statement

Suppose that $\boldsymbol{\theta} \in \mathbb{R}^p$ is a random signal of interest. Consider the linear model, as depicted in Figure 2.1,

$$\begin{aligned} \boldsymbol{x} &= \boldsymbol{H}\boldsymbol{\theta} + \boldsymbol{u} \\ \boldsymbol{z} &= \boldsymbol{D}\boldsymbol{W}\boldsymbol{x} + \boldsymbol{v}. \end{aligned} \tag{2.2.1}$$

Here $\boldsymbol{x} \in \mathbb{R}^n$ is a noisy measurement of $\boldsymbol{\theta} \in \mathbb{R}^p$, $\boldsymbol{W} \in \mathbb{R}^{m \times n}$ $(m \leq n)$ is the compression matrix, and $\boldsymbol{W}\boldsymbol{x}$ is the signal to be transmitted over a noisy channel with a full-rank channel matrix $\boldsymbol{D} \in \mathbb{R}^{m \times m}$ and random noise $\boldsymbol{v} \in \mathbb{R}^m$. Note that the dimension of the signal $\boldsymbol{W}\boldsymbol{x}$ is smaller than that of the original signal $\boldsymbol{x}$. It is assumed that the channel

noise $\boldsymbol{v}$ has mean $\boldsymbol{0}$, and is independent of $\boldsymbol{\theta}$, $\boldsymbol{u}$ and $\boldsymbol{x}$. Our objective is to design the compression matrix $\boldsymbol{W}$ such that the compressed measurement $\boldsymbol{z}$ is optimal according to a pre-specified performance metric.

We use linear estimation which is optimal in the multivariate normal case. In particular, given a measurement $\boldsymbol{z}$, the best linear unbiased estimator (BLUE) of $\boldsymbol{\theta}$ is

$$\widehat{\boldsymbol{\theta}}_{\boldsymbol{z}} = \boldsymbol{\mu}_{\boldsymbol{\theta}} + \boldsymbol{Q}_{\boldsymbol{\theta}\boldsymbol{z}}\boldsymbol{Q}_{\boldsymbol{z}\boldsymbol{z}}^{+}(\boldsymbol{z} - \boldsymbol{\mu}_{\boldsymbol{z}}),$$

where $\boldsymbol{\mu}_{\boldsymbol{\theta}}, \boldsymbol{\mu}_{\boldsymbol{z}}$ are the means of $\boldsymbol{\theta}$ and $\boldsymbol{z}$ respectively, and $\boldsymbol{Q}_{\boldsymbol{z}\boldsymbol{z}}^{+}$ is the pseudo inverse of $\boldsymbol{Q}_{\boldsymbol{z}\boldsymbol{z}}$. The error covariance matrix of $\widehat{\boldsymbol{\theta}}_{\boldsymbol{z}}$, denoted by $\boldsymbol{Q}_{ee}$, is

$$\boldsymbol{Q}_{ee} = \mathbb{E}[(\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}_{\boldsymbol{z}})(\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}_{\boldsymbol{z}})^{T}] = \boldsymbol{Q}_{\boldsymbol{\theta}\boldsymbol{\theta}} - \boldsymbol{Q}_{\boldsymbol{\theta}\boldsymbol{z}}\boldsymbol{Q}_{\boldsymbol{z}\boldsymbol{z}}^{+}\boldsymbol{Q}_{\boldsymbol{z}\boldsymbol{\theta}}.$$

Under model (2.2.1), $\boldsymbol{Q}_{ee}$ can be written as a function of $\boldsymbol{W}$; that is,

$$\boldsymbol{Q}_{ee} = \boldsymbol{Q}_{\boldsymbol{\theta}\boldsymbol{\theta}} - \boldsymbol{Q}_{\boldsymbol{\theta}\boldsymbol{x}}\boldsymbol{W}^{T}\boldsymbol{D}^{T}(\boldsymbol{D}\boldsymbol{W}\boldsymbol{Q}_{\boldsymbol{x}\boldsymbol{x}}\boldsymbol{W}^{T}\boldsymbol{D}^{T} + \boldsymbol{Q}_{\boldsymbol{v}\boldsymbol{v}})^{-1}\boldsymbol{D}\boldsymbol{W}\boldsymbol{Q}_{\boldsymbol{x}\boldsymbol{\theta}} \qquad (2.2.2)$$

We assume that the covariance matrices $\boldsymbol{Q}_{\boldsymbol{\theta}\boldsymbol{\theta}}$, $\boldsymbol{Q}_{\boldsymbol{x}\boldsymbol{x}}$, $\boldsymbol{Q}_{\boldsymbol{v}\boldsymbol{v}}$ and the cross covariance matrix $\boldsymbol{Q}_{\boldsymbol{\theta}\boldsymbol{x}}$ are known. In practice, the covariance matrices are determined from the physics of a problem or estimated from a two-channel experiment that generates realizations of $(\boldsymbol{\theta}, \boldsymbol{x})$. Only the second order moments are required, not the exact distribution of the random signals.

The performance of the compression matrix is determined by evaluating functions of the resulting error covariance $\boldsymbol{Q}_{ee}$. In the literature, the most prominent functions are the determinant criterion, $\det(\boldsymbol{Q}_{ee})$, the average-variance criterion, $(\mathrm{tr}(\boldsymbol{Q}_{ee}^{-1}))^{-1}$, the smallest-eigenvalue criterion, $\lambda_{\min}(\boldsymbol{Q}_{ee})$, and the trace criterion, $\mathrm{tr}(\boldsymbol{Q}_{ee})$. See Pukelsheim [1993] for more detailed review and discussion. All these criteria provide a reasonable measure of "largeness" of the error covariance matrix $\boldsymbol{Q}_{ee}$. Consequently, the optimal compression matrix $\boldsymbol{W}$ can be obtained by solving an optimization problem using one of the aforementioned criterion. In Sections 2.3 and 2.4, we will focus on two classical

criteria: $\text{tr}(\boldsymbol{Q}_{ee})$ and $\det(\boldsymbol{Q}_{ee})$. The first measure $\text{tr}(\boldsymbol{Q}_{ee})$ is the mean squared error of $\widehat{\boldsymbol{\theta}}_{\boldsymbol{z}}$. The second measure $\det(\boldsymbol{Q}_{ee})$ is the volume of the error concentration ellipsoid. When $\boldsymbol{z}$ and $\boldsymbol{\theta}$ are jointly Gaussian distributed, minimizing $\det(\boldsymbol{Q}_{ee})$ is equivalent to maximizing the mutual information between $\boldsymbol{z}$ and $\boldsymbol{\theta}$, or the differential information rate at which measurement $\boldsymbol{z}$ brings information about $\boldsymbol{\theta}$ (see Cover and Thomas [2005]). For simplicity, let us refer to the problems where we try to minimize $\text{tr}(\boldsymbol{Q}_{ee})$ and $\det(\boldsymbol{Q}_{ee})$ as the *min-trace* and *min-det* problems, respectively. In Section 2.5, we will explore more general criteria, a class of differentiable functions of the error covariance matrix.

2.3 Channel-Noise-Free Compression Design

In this section, we study a special case of (2.2.1) in which the compressed measurement can be transmitted perfectly, i.e., $\boldsymbol{z} = \boldsymbol{W}\boldsymbol{x}$. In particular, the error covariance matrix is

$$\boldsymbol{Q}_{ee} = \boldsymbol{Q}_{\boldsymbol{\theta}\boldsymbol{\theta}} - \boldsymbol{Q}_{\boldsymbol{\theta}\boldsymbol{x}}\boldsymbol{W}^T(\boldsymbol{W}\boldsymbol{Q}_{\boldsymbol{x}\boldsymbol{x}}\boldsymbol{W}^T)^-\boldsymbol{W}\boldsymbol{Q}_{\boldsymbol{x}\boldsymbol{\theta}}, \qquad (2.3.1)$$

where $(\cdot)^-$ is the pseudo-inverse. The solutions of the min-trace and min-det problems can be obtained by directly applying the results on optimal reduced-rank filtering in Scharf [1990], Scharf and Thomas [1998], Scharf and Mullis [1998], Hua et al. [2001], and Schreier and Scharf [2006].

First, we will discuss a notion of *canonical coordinates*. The basic idea is to transfer $(\boldsymbol{\theta}, \boldsymbol{x})$ to canonical coordinates $(\tilde{\boldsymbol{\theta}}, \tilde{\boldsymbol{x}})$ which have a diagonal cross-covariance matrix. For the min-trace problem, we consider the singular value decomposition (SVD) of the half coherence matrix

$$\boldsymbol{Q}_{\boldsymbol{\theta}\boldsymbol{x}}\boldsymbol{Q}_{\boldsymbol{x}\boldsymbol{x}}^{-T/2} = \boldsymbol{F}\boldsymbol{K}\boldsymbol{G}^T, \qquad (2.3.2)$$

where $\boldsymbol{K} \in \mathbb{R}^{p \times n}$ is a diagonal matrix with diagonal elements $k_1 \geq \ldots \geq k_{\min\{n,p\}} \geq 0$, and $\boldsymbol{F} \in \mathbb{R}^{p \times p}$ and $\boldsymbol{G} \in \mathbb{R}^{n \times n}$ are orthogonal matrices. The vectors $\tilde{\boldsymbol{\theta}} = \boldsymbol{F}^T\boldsymbol{\theta}$ and

14

$\tilde{\boldsymbol{x}} = \boldsymbol{G}^T \boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2} \boldsymbol{x}$ are the half canonical coordinates for $\boldsymbol{\theta}$ and $\boldsymbol{x}$, respectively. Note that the cross-covariance matrix between $\tilde{\boldsymbol{\theta}}$ and $\tilde{\boldsymbol{x}}$ is the diagonal matrix $\boldsymbol{K}$ given in (2.3.2).

For the min-det problem, the choice of the canonical coordinates is different. In this case, we consider an SVD of the coherence matrix

$$\boldsymbol{Q}_{\boldsymbol{\theta\theta}}^{-1/2} \boldsymbol{Q}_{\boldsymbol{\theta x}} \boldsymbol{Q}_{\boldsymbol{xx}}^{-T/2} = \boldsymbol{F K G}^T, \qquad (2.3.3)$$

where $\boldsymbol{K} \in \mathbb{R}^{p \times n}$ is a diagonal matrix with diagonal elements $k_1 \geq \ldots \geq k_{\min\{n,p\}} \geq 0$, and $\boldsymbol{F} \in \mathbb{R}^{p \times p}$ and $\boldsymbol{G} \in \mathbb{R}^{n \times n}$ are orthogonal matrices. Now, the vectors $\tilde{\boldsymbol{\theta}} = \boldsymbol{F}^T \boldsymbol{Q}_{\boldsymbol{\theta\theta}}^{-1/2} \boldsymbol{\theta}$ and $\tilde{\boldsymbol{x}} = \boldsymbol{G}^T \boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2} \boldsymbol{x}$ are the full canonical coordinates of $\boldsymbol{\theta}$ and $\boldsymbol{x}$, respectively. Note that for the simplicity of our notation, we choose to re-use the variables $\boldsymbol{F}, \boldsymbol{K}$, and $\boldsymbol{G}$ for both SVDs.

The optimal compression matrix is given in Proposition 2.1, which is a re-statement of the results of Scharf [1990] and Hua et al. [2001].

Proposition 2.1. *For the min-trace and min-det problems, the optimal compression matrix $\boldsymbol{W}_0^* \in \mathbb{R}^{m \times n}$ can be written as*

$$\boldsymbol{W}_0^* = \boldsymbol{G}_m^T \boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2} \qquad (2.3.4)$$

where $\boldsymbol{G}_m$ consists of the first m columns of $\boldsymbol{G}$. The matrix $\boldsymbol{G}$ is defined in (2.3.2) for the min-trace problem and in (2.3.3) for the min-det problem. Moreover, for any $m \times m$ nonsingular matrix $\boldsymbol{T}$, $\boldsymbol{T W}_0^*$ is also an optimal compression matrix.

Proposition 2.1 figures prominently in our derivation of scaled and rotated canonical coordinates for optimum compression with channel noise. It is also worth mentioning that $\boldsymbol{W}_0^* \boldsymbol{x}$ returns the first m canonical coordinates in $\tilde{\boldsymbol{x}}$. Let $\boldsymbol{W}_{tr,0}^*$ and $\boldsymbol{W}_{det,0}^*$ denote the optimal compression matrices given in Proposition 2.1. Straightforward calculation yields that, using the compression matrix $\boldsymbol{W}_{tr,0}^*$, the minimum MSE of $\widehat{\boldsymbol{\theta}}_{\boldsymbol{z}}$ is, as given in Scharf [1990],

$$\text{tr}(\boldsymbol{Q}_{ee}(\boldsymbol{W}_{tr,0}^*)) = \text{tr}(\boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}}) + \sum_{i=m+1}^{\min\{n,p\}} k_i^2,$$

15

where $\boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}}$ is the error covariance for the BLUE of $\boldsymbol{\theta}$ given $\boldsymbol{x}$, and $\sum_{i=m+1}^{\min\{n,p\}} k_i^2$ is the minimum increase of the MSE. In addition, using $\boldsymbol{W}_{det,0}^*$, the resulting minimum volume of the error concentration ellipsoid (see Scharf and Thomas [1998], Scharf and Mullis [1998], and Schreier and Scharf [2006]) is

$$\det(\boldsymbol{Q}_{ee}(\boldsymbol{W}_{det,0}^*)) = \det \boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}} \prod_{i=m+1}^{\min\{n,p\}} (1 - k_i^2)^{-1}. \tag{2.3.5}$$

Note that, in the min-det problem, the diagonal elements of $\boldsymbol{K}$, i.e., $k_1, \ldots, k_{\min\{n,p\}}$, are the full canonical correlations that measure cosines of principle angles between $\boldsymbol{\theta}$ and $\boldsymbol{x}$ (see Schreier and Scharf [2006]). In general, the k_i's take values between 0 and 1, but in (2.3.5), we assume the k_i's are strictly less than 1. It is easy to see that $\det(\boldsymbol{Q}_{ee}(\boldsymbol{W}_{det,0}^*)) \geq \det(\boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}})$, which shows that compression indeed discards some information about $\boldsymbol{\theta}$ by compressing $\boldsymbol{x}$ to a lower-dimensional measurement.

2.4 Compression Design with Sensor Noise and Channel Noise

Now we extend the results in Section 2.3 by considering the linear compression of the noisy measurement to be transmitted over a noisy channel. We assume the channel noise $\boldsymbol{v}$ has mean zero and covariance matrix $\boldsymbol{Q}_{\boldsymbol{vv}}$, and $\boldsymbol{v}$ is independent of $\boldsymbol{\theta}$ and $\boldsymbol{x}$.

A significant feature of the design for noisy transmission is the need for a *power constraint* on the compression matrix, for otherwise the design problem is not well-defined. In this chapter, we restrict the compression matrix $\boldsymbol{W}$ subject to $\text{tr}(\boldsymbol{W}\boldsymbol{Q}_{\boldsymbol{xx}}\boldsymbol{W}^T) \leq c$ for some pre-specified constant c.

Define $\boldsymbol{Q}_{\boldsymbol{\omega\omega}} = \boldsymbol{D}^{-1}\boldsymbol{Q}_{\boldsymbol{vv}}(\boldsymbol{D}^{-1})^T$ with the eigendecomposition $\boldsymbol{Q}_{\boldsymbol{\omega\omega}} = \boldsymbol{U}_{\boldsymbol{\omega}}\boldsymbol{\Sigma}_{\boldsymbol{\omega}}\boldsymbol{U}_{\boldsymbol{\omega}}^T$, where $\boldsymbol{U}_{\boldsymbol{\omega}}$ is an $m \times m$ orthogonal matrix and $\boldsymbol{\Sigma}_{\boldsymbol{\omega}} \in \mathbb{R}^{m \times m}$ is a diagonal matrix with diagonal elements $\sigma_{\boldsymbol{\omega},1}^2, \ldots, \sigma_{\boldsymbol{\omega},m}^2$ with $0 < \sigma_{\boldsymbol{\omega},1}^2 \leq \ldots \leq \sigma_{\boldsymbol{\omega},m}^2$.

16

2.4.1 Min-Trace Compression with Channel Noise

Under the power constraint, Schizas et al. [2007] have derived the optimal compression matrix to minimize $\mathrm{tr}(\boldsymbol{Q}_{ee})$. In Theorem 2.1, we re-state their result as a *scaled and rotated canonical coordinate design*.

Theorem 2.1. *An optimal compression matrix $\boldsymbol{W}_{tr}^*$ minimizing $\mathrm{tr}(\boldsymbol{Q}_{ee})$ is given by Schizas et al. [2007]*

$$\boldsymbol{W}_{tr}^* = \boldsymbol{U}_{\boldsymbol{\omega}} \boldsymbol{\Sigma}_{tr}^* \boldsymbol{G}^T \boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2} \tag{2.4.1}$$

Here the matrices $\boldsymbol{G}$ and $\boldsymbol{K}$ are given in (2.3.2), $\boldsymbol{\Sigma}_{tr}^*$ is an $m \times n$ diagonal matrix with diagonal elements

$$\sigma_{ii} = \begin{cases} \sqrt{k_i \sigma_{\boldsymbol{\omega},i}/\sqrt{\mu} - \sigma_{\boldsymbol{\omega},i}^2} & i = 1, \ldots, \kappa \\ 0 & i = \kappa+1, \ldots, m, \end{cases} \tag{2.4.2}$$

with κ the maximum integer between 1 and $\mathrm{rank}(\boldsymbol{K})$ such that $\sigma_{ii}^2 > 0$ for $i = 1, \ldots, \kappa$, and

$$\mu = \left((c + \sum_{i=1}^{\kappa} \sigma_{\boldsymbol{\omega},i}^2)^{-1} \sum_{i=1}^{\kappa} \sigma_{\boldsymbol{\omega},i} k_i \right)^2.$$

It can be seen that $\boldsymbol{W}_{tr}^*$ factors into whitening $\boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}$, canonical coordinate transformation $\boldsymbol{G}^T$, scaling $\boldsymbol{\Sigma}_{tr}^*$ and rotation $\boldsymbol{U}_{\boldsymbol{\omega}}$ into the sub-dominant invariant subspace of $\boldsymbol{Q}_{\boldsymbol{\omega}\boldsymbol{\omega}}$.

2.4.2 Min-Det Compression with Channel Noise

The optimal compression matrix $\boldsymbol{W}$ to minimize $\det(\boldsymbol{Q}_{ee})$ under a power constraint solves the optimization problem,

$$\boldsymbol{W}_{det}^* = \underset{\boldsymbol{W} \in \mathbb{R}^{m \times n}}{\arg\min} \det(\boldsymbol{Q}_{ee}) \quad \text{subject to} \quad \mathrm{tr}(\boldsymbol{W}\boldsymbol{Q}_{\boldsymbol{xx}}\boldsymbol{W}^T) \leq c. \tag{2.4.3}$$

The matrix $\boldsymbol{W}$ has mn degrees of freedom. But let's restrict $\boldsymbol{W}$ to a subset of $\mathbb{R}^{m \times n}$, over which the local minimizer of $\det(\boldsymbol{Q}_{ee})$ can be expressed explicitly. For a given $n \times n$ orthogonal matrix $\boldsymbol{V}$, define

$$\Omega_{\boldsymbol{V}} = \{\boldsymbol{U}_{\boldsymbol{\omega}} \boldsymbol{\Pi}_m \boldsymbol{\Sigma} \boldsymbol{\Pi}_n^T \boldsymbol{V}^T \boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}, \text{where } \boldsymbol{\Pi}_m \in \mathbb{R}^{m \times m}, \boldsymbol{\Pi}_n \in \mathbb{R}^{n \times n} \text{ are permutation}$$
$$\text{matrices, and } \boldsymbol{\Sigma} \in \mathbb{R}^{m \times n} \text{ is diagonal with } \sum_{i=1}^{m} \sigma_{ii}^2 \le c\}. \tag{2.4.4}$$

For fixed matrices $\boldsymbol{U}_{\boldsymbol{\omega}}$ and $\boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}$, the set $\Omega_{\boldsymbol{V}}$ is a subset of the constrained space of problem (2.4.3). Therefore, the local minimizer of $\det(\boldsymbol{Q}_{ee})$ over $\boldsymbol{W} \in \Omega_{\boldsymbol{V}}$ generally gives a suboptimal solution for problem (2.4.3). However, in Lemma 2.1, we show that, for a suitable choice of $\boldsymbol{V}$, the suboptimal solution on $\Omega_{\boldsymbol{V}}$ is a global optimal solution for problem (2.4.3).

Lemma 2.1. *Suppose that $\boldsymbol{G}$ is the orthogonal matrix given in (2.3.3). Then,*

$$\min_{\boldsymbol{W} \in \Omega_{\boldsymbol{G}}} \det(\boldsymbol{Q}_{ee}) = \det(\boldsymbol{Q}_{ee}(\boldsymbol{W}_{det}^*)).$$

The proof is given in Section 2.7.2. Following a similar proof, we can show that Lemma 2.1 holds for the min-trace problem as well, with $\boldsymbol{G}$ given in (2.3.2).

From Lemma 2.1, it can be seen that the local minimizer over $\Omega_{\boldsymbol{G}}$ is also a global minimizer of (2.4.3). For any $\boldsymbol{W} \in \Omega_{\boldsymbol{G}}$, we have

$$\det(\boldsymbol{Q}_{ee}(\boldsymbol{W})) = \det(\boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}}) \det\left(\boldsymbol{I}_n + \boldsymbol{\Pi}_n^T \boldsymbol{\Gamma} \boldsymbol{\Pi}_n (\boldsymbol{I}_n + \boldsymbol{\Sigma}^T \boldsymbol{\Pi}_m^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Pi}_m \boldsymbol{\Sigma})^{-1}\right). \tag{2.4.5}$$

Here $\boldsymbol{\Gamma} = \boldsymbol{K}^T (\boldsymbol{I}_p - \boldsymbol{K}\boldsymbol{K}^T)^{-1} \boldsymbol{K}$, with $\boldsymbol{K}$ given in (2.3.3), is an $n \times n$ positive semi-definite diagonal matrix with diagonal elements $\gamma_i^2 = k_i^2/(1 - k_i^2)$ for $i = 1, \ldots, \min\{n, p\}$ and 0 otherwise. We require $0 \le k_i < 1$ for all i, and consequently, $\gamma_1^2, \ldots, \gamma_n^2$ is a finite decreasing sequence. The permutation matrices $\boldsymbol{\Pi}_n$ and $\boldsymbol{\Pi}_m$ reorder the diagonal elements of $\boldsymbol{\Gamma}$ and $\boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1}$, respectively. In fact, for any $\boldsymbol{W} \in \Omega_{\boldsymbol{G}}$, $\boldsymbol{\Pi}_n$ reorders the canonical coordinates $\boldsymbol{G}^T \boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2} \boldsymbol{x}$ and determines which m coordinates will be transmitted, and

18

the permutation matrix $\boldsymbol{\Pi}_m$ reorders the selected coordinates and determines which sub-channel the coordinates will be transmitted over. The optimal compression matrix can be obtained by minimizing $\det(\boldsymbol{Q}_{ee}(\boldsymbol{W}))$ with respect to the permutation matrices $\boldsymbol{\Pi}_m$, $\boldsymbol{\Pi}_n$ and the diagonal matrix $\boldsymbol{\Sigma}$. The computational complexity of this optimization has been greatly reduced since there are just $2m + n$ degrees of freedom in the permutation matrices $\boldsymbol{\Pi}_m$, $\boldsymbol{\Pi}_n$ and the diagonal matrix $\boldsymbol{\Sigma}$. We give in Theorem 2.2 the closed-form expression for the optimal compression matrix $\boldsymbol{W}_{det}^*$.

Theorem 2.2. *Suppose the matrix $\boldsymbol{Q}_{\boldsymbol{\omega\omega}}$ has distinct eigenvalues, i.e., $0 < \sigma_{\boldsymbol{\omega},1}^2 < \ldots < \sigma_{\boldsymbol{\omega},m}^2$. Then, the optimal compression matrix $\boldsymbol{W}_{det}^*$ solving problem (2.4.3) is*

$$\boldsymbol{W}_{det}^* = \boldsymbol{U}_{\boldsymbol{\omega}} \boldsymbol{\Sigma}_{det}^* \boldsymbol{G}^T \boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2} \tag{2.4.6}$$

Here $\boldsymbol{G}$ is given in (2.3.3) where the matrix $\boldsymbol{K}$ contains singular values $0 \leq k_i < 1$ for all i; $\boldsymbol{\Sigma}_{det}^* \in \mathbb{R}^{m \times n}$ is a diagonal matrix with diagonal elements $\sigma_{11}^*, \ldots, \sigma_{mm}^*$ such that

$$\sigma_{ii}^{*2} = \begin{cases} \frac{1}{2}\sigma_{\boldsymbol{\omega},i}^2 \left(-2 - \gamma_i^2 + \sqrt{\gamma_i^4 + 4\frac{\gamma_i^2}{\mu\sigma_{\boldsymbol{\omega},i}^2}}\right) & i = 1, \ldots, \kappa \\ 0 & i = \kappa + 1, \ldots, m \end{cases} \tag{2.4.7}$$

where κ is the maximum integer between 1 and m such that $\sigma_{ii}^{*2} > 0$ or equivalently $\sigma_{\boldsymbol{\omega},i}^2/k_i^2 < 1/\mu$ for $i = 1, \ldots, \kappa$. The value of μ is nonnegative and uniquely solves $\sum_{i=1}^{\kappa} \sigma_{ii}^{*2} = c$. Moreover, the diagonal element of $\boldsymbol{\Sigma}_{det}^*$ are decreasingly ordered, i.e., $\sigma_{11}^* \geq \ldots \geq \sigma_{mm}^* \geq 0$.

The proof of Theorem 2.2 is given in Section 2.7.3.

Given that $\gamma_1^2 \geq \ldots \geq \gamma_n^2$ and $\sigma_{\boldsymbol{\omega},1}^2 \leq \ldots \leq \sigma_{\boldsymbol{\omega},m}^2$, the optimal permutation matrices $\boldsymbol{\Pi}_m$ and $\boldsymbol{\Pi}_n$ are both identity matrices. This indicates that the canonical coordinates of the measurement with higher canonical correlation between the canonical coordinates of $\boldsymbol{\theta}$ are transmitted over the subchannels with less noise. The decreasingly ordered sequence of scalings $\sigma_{11}^* \geq \ldots \geq \sigma_{mm}^* \geq 0$ shows that the subchannels with higher canonical correlation and lower noise are assigned higher power.

19

Example 2.1. *A Circulant Model. Consider the model $\boldsymbol{x} = \boldsymbol{\theta} + \boldsymbol{u}$ where $\boldsymbol{\theta}$ and $\boldsymbol{u}$ have circulant covariances. Suppose that the channel matrix $\boldsymbol{D} = \boldsymbol{I}_m$ and the channel noise $\boldsymbol{v}$ has a circulant covariance. Then both $\boldsymbol{Q}_{\boldsymbol{\theta\theta}}^{-1/2}\boldsymbol{Q}_{\boldsymbol{\theta x}}\boldsymbol{Q}_{\boldsymbol{xx}}^{-T/2}$ and $\boldsymbol{Q}_{\boldsymbol{\omega\omega}}$ are circulant matrices with Discrete Fourier Transform (DFT) representations*

$$\boldsymbol{Q}_{\boldsymbol{\theta\theta}}^{-1/2}\boldsymbol{Q}_{\boldsymbol{\theta x}}\boldsymbol{Q}_{\boldsymbol{xx}}^{-T/2} = \boldsymbol{U}_n\boldsymbol{S}_{\boldsymbol{\theta x}}\boldsymbol{U}_n^T; \boldsymbol{Q}_{\boldsymbol{\omega\omega}} = \boldsymbol{V}_m\boldsymbol{S}_{\boldsymbol{\omega\omega}}\boldsymbol{V}_m^T$$

where $\boldsymbol{U}_n \in \mathbb{R}^{n \times n}$ and $\boldsymbol{V}_m \in \mathbb{R}^{m \times m}$ are the DFT matrices and $\boldsymbol{S}_{\boldsymbol{\theta x}}$ and $\boldsymbol{S}_{\boldsymbol{\omega\omega}}$ are diagonal matrices. Let $\boldsymbol{u}_{(1)}, \ldots, \boldsymbol{u}_{(\kappa)}$ be the columns of the DFT matrix $\boldsymbol{U}_n$ sorted according to the κ largest canonical correlations, and $\boldsymbol{v}_{(1)}, \ldots, \boldsymbol{v}_{(\kappa)}$ be the columns of the DFT matrix $\boldsymbol{V}_m$ sorted according to the smallest eigenvalues of $\boldsymbol{Q}_{\boldsymbol{\omega\omega}}$. It can be seen that the optimal compression matrix selects the first κ sorted DFT coordinates $\boldsymbol{u}_{(1)}^T\boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}\boldsymbol{x}, \ldots, \boldsymbol{u}_{(\kappa)}^T\boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}\boldsymbol{x}$ and sends them over the κ sorted DFT modes $\boldsymbol{v}_{(1)}, \ldots, \boldsymbol{v}_{(\kappa)}$ of the channel.

Simple calculation shows that, with compression, the minimum determinant of the error covariance is

$$\det(\boldsymbol{Q}_{ee}(\boldsymbol{W}_{det}^*)) = \det \boldsymbol{Q}_{\boldsymbol{\theta\theta|x}} \prod_{i=\kappa+1}^{\min\{n,p\}} \frac{1}{1-k_i^2} \prod_{i=1}^{\kappa} \left(1 + \frac{2}{\sqrt{1+4(\gamma_i^2\sigma_{\boldsymbol{\omega},i}^2\mu)^{-1}}-1}\right). \quad (2.4.8)$$

The first term on the right hand side is the minimum volume of the error concentration ellipsoid with no dimension reduction; the second term scales this volume according to canonical correlations of discarded canonical coordinates; the third term scales the volume by a term that depends on the channel noise variance, the power c, and the full canonical correlations. The integer κ is the number of subchannels assigned with positive power, and κ/n is the optimal compression ratio for a given power c.

In Theorem 2.2, it is assumed that all eigenvalues of $\boldsymbol{Q}_{\boldsymbol{\omega\omega}}$ are distinct. Notice that, if some eigenvalues have multiplicity greater than 1, one can perturb $\boldsymbol{Q}_{\boldsymbol{\omega\omega}}$ by $\delta\boldsymbol{Q}$ such that the matrix $\widetilde{\boldsymbol{Q}}_{\boldsymbol{\omega\omega}} = \boldsymbol{Q}_{\boldsymbol{\omega\omega}} + \delta\boldsymbol{Q}$ has distinct eigenvalues. Moreover, we can restrict $\delta\boldsymbol{Q}$ such that the eigenspace of $\widetilde{\boldsymbol{Q}}_{\boldsymbol{\omega\omega}}$ is fixed. Because the optimal entries in (2.4.7) are

continuous functions of $\sigma_{\boldsymbol{\omega},1}^2, \ldots, \sigma_{\boldsymbol{\omega},m}^2$, we can obtain the optimal compression matrix by letting $\delta \boldsymbol{Q}$ go to zero.

It is worth mentioning that, for a sufficiently large c, we have $1/\mu > \sigma_{\boldsymbol{\omega},i}^2/k_i^2$ (or equivalently $\sigma_{ii}^2 > 0$) for all $i = 1, \ldots, m$. Consequently, the solution given in Theorem 2.2 is also an optimal compression for the channel-noise-free case. We simply let the nonsingular matrix $\boldsymbol{T}$ in Proposition 2.1 be $\boldsymbol{T} = \boldsymbol{U}_{\boldsymbol{\omega}} \operatorname{diag}(\sigma_{11}^*, \ldots, \sigma_{mm}^*)$. When c goes to infinity, the third part in (2.4.8) goes to 1, and the minimum determinant of $\boldsymbol{Q}_{ee}$ converges to the channel-noise-free case. On the other hand, the diagonal elements of $\boldsymbol{\Sigma}_{det}^*$ go to infinity. Therefore, we can see that the optimization problem is ill-posed without a (finite) power constraint.

Finally, we comment on the canonical correlations. Under our current framework, all canonical correlations, k_i, are less than 1. In fact, the factorization in (2.4.6) still holds when $k_i = 1$ with a different scaling matrix $\boldsymbol{\Sigma}_{det}^*$. In the sensor-noise-free case, suppose that $\boldsymbol{x} = \boldsymbol{H}\boldsymbol{\theta}$ and the matrix $\boldsymbol{H} \in \mathbb{R}^{n \times p}$ has rank p. The full canonical correlations between $\boldsymbol{\theta}$ and $\boldsymbol{x}$ are all 1. In this case, the compressor $\boldsymbol{W}$ operates on $\boldsymbol{H}\boldsymbol{\theta}$ directly and the design of $\boldsymbol{W}$ becomes a precoder design problem Scaglione et al. [2002], Pérez-Cruz et al. [2010], Carson et al. [2012]. The optimal scaling matrix $\boldsymbol{\Sigma}_{det}^*$ has diagonal elements

$$\sigma_{ii}^* = \begin{cases} \sqrt{1/\mu - \sigma_{\boldsymbol{\omega},i}^2} & \sigma_{\boldsymbol{\omega},i}^2 < 1/\mu \\ 0 & \sigma_{\boldsymbol{\omega},i}^2 \geq 1/\mu \end{cases} \tag{2.4.9}$$

where the value of μ is determined by the power constraint $\sum_{i=1}^m \sigma_{ii}^2 = c$.

2.4.3 A Mercury/Waterfilling Interpretation

First, consider a sensor-noise-free case, $\boldsymbol{x} = \boldsymbol{H}\boldsymbol{\theta}$. The optimal compressor has been discussed in Section 2.4.2, with the factorization in (2.4.6) and $\boldsymbol{\Sigma}_{det}^*$ given in (2.4.9). The scaling matrix $\boldsymbol{\Sigma}_{det}^*$ distributes the power among all the m subchannels according to a *waterfilling* policy, Cover and Thomas [2005], with a graphical display given in Figure 2.2. There are m vessels, each of which represents a subchannel. The goal is to

pour water of total volume c into these vessels. Here, each vessel has its own solid base with height $\sigma^2_{\boldsymbol{w},i}$. Recall that $\sigma^2_{\boldsymbol{\omega},1}, \ldots, \sigma^2_{\boldsymbol{\omega},m}$ are the eigenvalues of $\boldsymbol{Q_{\omega\omega}}$. They can be viewed as the variances of the channel noise in the channel coordinates since the optimal compression rotates the scaled canonical coordinates by $\boldsymbol{U_\omega}$, the eigenvectors of $\boldsymbol{Q_{\omega\omega}}$. The desired water level $1/\mu$ is determined by the power constraint, or equivalently, the total volume of water equals c. The optimal compressor *pours water into each vessel until the water level reaches* $1/\mu$. As a result, the water height in each vessel gives the power assigned to the corresponding subchannel. Note that less power will be allocated to noisier subchannels, and no power will be assigned to subchannels with noise variance larger than $1/\mu$.
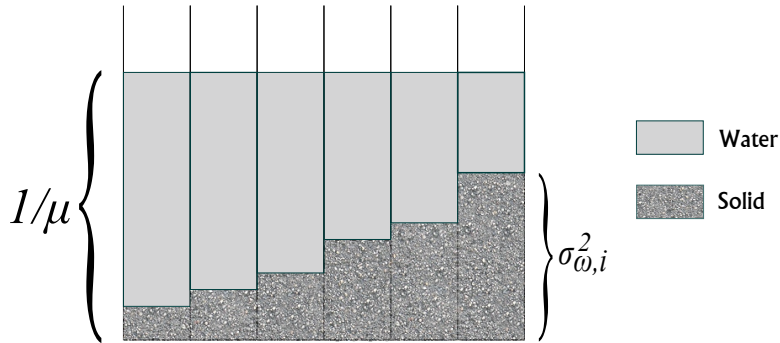
Figure 2.2: Waterfilling without sensor noise. The total volume of water is c, and the water height over the solid base on the ith vessel gives the power for the ith subchannel.

In general, $\boldsymbol{x}$ is a noisy measurement of $\boldsymbol{\theta}$ and the full canonical correlations are strictly less than 1. Therefore, the optimal power allocation policy needs to be adjusted according to the canonical correlations. As a consequence of Theorem 2.2, the solution can be interpreted as a *mercury/waterfilling* policy, which is a three-step procedure that has been introduced in Lozano et al. [2006]:

1. For the ith vessel, fill in the solid base with height $\sigma^2_{\boldsymbol{\omega},i}/k_i^2$.

2. Compute μ from the power constraint. For the vessels with base height less than

$1/\mu$, fill in mercury in the vessel until the height reaches

$$\max\left\{\frac{1}{\mu} - \frac{1}{2}\sigma_{\boldsymbol{\omega},i}^2\left(-2 - \gamma_i^2 + \sqrt{\gamma_i^4 + 4\frac{\gamma_i^2}{\sigma_{\boldsymbol{\omega},i}^2\mu}}\right), \frac{\sigma_{\boldsymbol{\omega},i}^2}{k_i^2}\right\}.$$

3. Pour water into all vessels until the height of water in each vessel reaches $1/\mu$.

In this mercury/waterfilling policy, $1/\mu$ is the parameter in the formula for water volume σ_{ii}^2 that minimizes $\det(\boldsymbol{Q}_{ee})$ under the constraint that the total volume of water is c. Given the value of μ, the determinant of the error covariance is minimized when the value of σ_{ii}^2 equals the height of water in the corresponding vessel.

The height of the solid base, $\sigma_{\boldsymbol{\omega},i}^2/k_i^2$, is the variance of the channel noise in the ith vessel divided by the ith squared canonical correlation. A higher solid base means a less informative channel with high channel noise and weak correlation with $\boldsymbol{\theta}$. For any vessel with base height exceeding $1/\mu$, neither mercury nor water will be added, or equivalently, no power will be assigned to the corresponding subchannel.

While the base height determines whether water will be added, the mercury stage regulates the water level for each vessel. Without adding mercury, the optimal power allocation will have variable solid-plus-water levels among different vessels. The mercury is added to balance the sensor noise contained in $\boldsymbol{x}$ and the channel noise added in transmission. Recall that no mercury is added in the special case when $\boldsymbol{x} = \boldsymbol{\theta}$. The water height in each vessel is the optimal power assigned to the corresponding subchannel. As demonstrated in Theorem 2.2, the water height for each vessel is decreasingly ordered.

2.4.4 Scaled and Rotated Canonical Coordinate Design

Theorems 2.1 and 2.2 suggest a common architecture for compression, which specializes to all previous designs for reduced-rank filtering and for reduced rank precoding and equalizing. The optimal compressor can be factored into four component matrices. As shown in Figure 2.4, the first matrix $\boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}$ whitens the noisy measurement $\boldsymbol{x}$. The
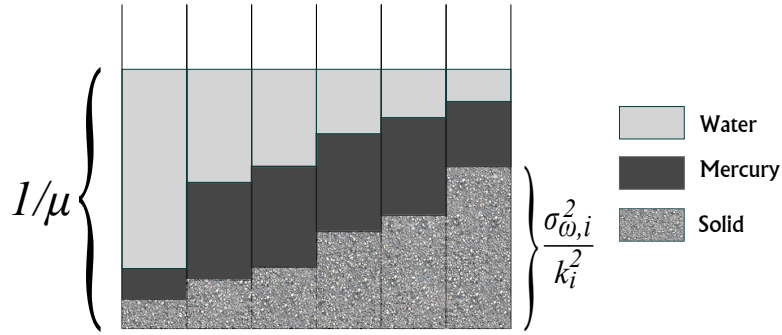
Figure 2.3: A mercury/waterfilling policy. The total volume of water is c, and the water height over mercury on the ith vessel gives σ_{ii}^2.
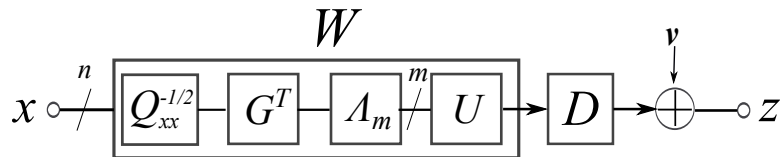
Figure 2.4: Scaled canonical coordinate transformation for compressing a noisy measurement with transmission over a noisy channel.

second matrix $\boldsymbol{G}^T$ transforms the whitened measurement into a canonical coordinate system. For the min-det problem, the full canonical coordinates, $\boldsymbol{G}^T \boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2} \boldsymbol{x}$, are uncorrelated and have unit variance. The third matrix $\boldsymbol{\Sigma}^* \in \mathbb{R}^{m \times n}$ is diagonal. The role of $\boldsymbol{\Sigma}^*$ is to extract the first m full canonical coordinates and distribute power across the canonical channels. The ith canonical coordinate is scaled to have power σ_{ii}^2. For the min-det problem, when $\gamma_i^2 = 0$ (i.e., $k_i = 0$), the corresponding scaling is $\sigma_{ii} = 0$, which means those canonical coordinates uncorrelated or weakly correlated with $\boldsymbol{\theta}$ will be eliminated. In general, the diagonal elements of $\boldsymbol{\Sigma}^*$ have a mercury/waterfilling interpretation. The matrix $\boldsymbol{U}_{\boldsymbol{\omega}}$ rotates the compressed canonical coordinates into the sub-dominant invariant subspace of the matrix $\boldsymbol{Q}_{\boldsymbol{\omega\omega}}$.

The difference between the trace and determinant designs is in the canonical coordinates and in the values of scaling constants in the diagonal scaling matrix.

24

2.5 A Unified Framework for Optimal Compression

In the previous sections, our interest has centered on optimal compression under two commonly used criteria: trace and determinant. Next, we consider the problem of designing a compression matrix to minimize a general criterion:

$$\boldsymbol{W}^* = \underset{\boldsymbol{W} \in \mathbb{R}^{m \times n}}{\arg\min} \varphi(\boldsymbol{Q}_{ee}) \text{ subject to } \text{tr}(\boldsymbol{W}\boldsymbol{Q}_{xx}\boldsymbol{W}^T) \le c. \tag{2.5.1}$$

Here φ is a differentiable function on the space of $p \times p$ positive definite matrices.

Denote the first derivative of φ by φ'. Then φ' is a mapping from $\mathbb{R}^{p \times p}$ to $\mathbb{R}^{p \times p}$, with

$$(\varphi'(\boldsymbol{Q}_{ee}))_{ij} = \lim_{t \to 0} \frac{\varphi(\boldsymbol{Q}_{ee} + t\boldsymbol{J}_{ij}) - \varphi(\boldsymbol{Q}_{ee})}{t},$$

where $\boldsymbol{J}_{ij}$ is the $p \times p$ single-entry matrix with 1 at (i, j) and 0 elsewhere.

We first establish a unified factorization of $\boldsymbol{W}^*$ in the following theorem.

Theorem 2.3. *Suppose that the diagonal matrix $\boldsymbol{\Sigma}_{\boldsymbol{\omega}}$ has distinct diagonal elements. Then for any minimizer $\boldsymbol{W}^*$ of (2.5.1), there exists an $m \times m$ permutation matrix $\boldsymbol{\Pi}_m^*$, an $m \times n$ diagonal matrix $\boldsymbol{\Sigma}^*$, and an $n \times n$ orthogonal matrix $\boldsymbol{V}^*$ such that*

$$\boldsymbol{W}^* = \boldsymbol{U}_{\boldsymbol{\omega}}\boldsymbol{\Pi}_m^*\boldsymbol{\Sigma}^*\boldsymbol{V}^{*T}\boldsymbol{Q}_{xx}^{-1/2}. \tag{2.5.2}$$

Proof of Theorem 2.3 is given in Section 2.7.4. Theorem 2.3 suggests that the optimal compression matrix can be expressed as a sequence of operations, including whitening $(\boldsymbol{Q}_{xx}^{-1/2})$, coordinate system transformation $(\boldsymbol{V}^T)$, scaling $(\boldsymbol{\Sigma})$, re-ordering $(\boldsymbol{\Pi}_m)$ and rotation to the invariant subspace of the channel noise $(\boldsymbol{U}_{\boldsymbol{\omega}})$.

In general, searching for the global minimizer of the optimization problem (2.5.1) is rather challenging. Enlightened by the factorization in (2.5.2), we first consider searching for the optimal compressor for a fixed orthogonal matrix $\boldsymbol{V}$. This more restricted optimization problem can be carried out using the KKT conditions, which is not computationally costly, as $\boldsymbol{\Pi}_m$ is a permutation matrix, and $\boldsymbol{\Sigma}$ is a diagonal matrix.

The choice of $\boldsymbol{V}$ is the key to the compression design, and there is no general solution for it. The following lemma provides a necessary condition for the optimal coordinate system $\boldsymbol{V}^*$.

Lemma 2.2. *The optimal orthogonal matrix $\boldsymbol{V}^*$ defined in (2.5.2) satisfies the condition*

$$[\boldsymbol{V}^{*T}\boldsymbol{L}^T\varphi'(\boldsymbol{Q}_{ee}^*)\boldsymbol{L}\boldsymbol{V}^*, (\boldsymbol{I}_n + \boldsymbol{\Delta}^*)^{-1}] = \boldsymbol{0}_{n\times n}, \tag{2.5.3}$$

where $[\boldsymbol{A}, \boldsymbol{B}] = \boldsymbol{A}\boldsymbol{B} - \boldsymbol{B}\boldsymbol{A}$, $\boldsymbol{L} = \boldsymbol{Q}_{\boldsymbol{\theta x}}\boldsymbol{Q}_{\boldsymbol{xx}}^{-T/2} \in \mathbb{R}^{p\times n}$ is the half coherence matrix between $\boldsymbol{\theta}$ and $\boldsymbol{x}$, $\boldsymbol{Q}_{ee}^*$ is the error covariance for $\boldsymbol{\theta}$ corresponding to the compression $\boldsymbol{W}^*$, and $\boldsymbol{\Delta}^* = \boldsymbol{\Sigma}^{*T}\boldsymbol{\Pi}_m^{*T}\boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1}\boldsymbol{\Pi}_m^*\boldsymbol{\Sigma}^* \in \mathbb{R}^{n\times n}$ is a diagonal matrix.

The proof is given in Section 2.7.5. An equivalent statement of Lemma 2.2 is that $(\boldsymbol{I}_n + \boldsymbol{\Delta}^*)^{-1}$ and $\boldsymbol{V}^{*T}\boldsymbol{L}^T\varphi'(\boldsymbol{Q}_{ee}^*)\boldsymbol{L}\boldsymbol{V}^*$ commute. In general, the solution for (2.5.3) is intractable, mainly because the terms $\boldsymbol{Q}_{ee}^*$ and $\boldsymbol{\Delta}^*$ contain unknown $\boldsymbol{\Pi}_m^*, \boldsymbol{\Sigma}^*, \boldsymbol{\Pi}_n^*$. Nevertheless, since $(\boldsymbol{I}_n + \boldsymbol{\Delta}^*)^{-1}$ is a diagonal matrix, we may choose an orthogonal matrix $\boldsymbol{V}$ such that $\boldsymbol{V}^T\boldsymbol{L}^T\varphi'(\boldsymbol{Q}_{ee}^*)\boldsymbol{L}\boldsymbol{V}$ is diagonal for any $\boldsymbol{\Pi}_m^*, \boldsymbol{\Sigma}^*, \boldsymbol{\Pi}_n^*$. Next we give two specific examples.

Example 2.2. *Consider a linear criterion φ, i.e.,*

$$\varphi(\boldsymbol{A} + \boldsymbol{B}) = \varphi(\boldsymbol{A}) + \varphi(\boldsymbol{B});$$

$$\varphi(\alpha\boldsymbol{A}) = \alpha\varphi(\boldsymbol{A}),$$

for any $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{p\times p}$ and $\alpha \in \mathbb{R}^1$. Then, the derivative $\varphi'(\boldsymbol{Q}_{ee})$, denoted by $\boldsymbol{M}$, is a constant known matrix. One choice of $\boldsymbol{V}$ to satisfy (2.5.3) is that the columns of $\boldsymbol{V}$ are the eigenvectors of $\boldsymbol{L}^T\boldsymbol{M}\boldsymbol{L}$. As a special case, when $\varphi(\boldsymbol{Q}_{ee}) = \mathrm{tr}(\boldsymbol{Q}_{ee})$, we have $\varphi'(\boldsymbol{Q}_{ee}) = \boldsymbol{I}_p$ and $\boldsymbol{V} = \boldsymbol{G}$ in (2.3.2) is a feasible choice.

Example 2.3. *Consider $\varphi(\boldsymbol{Q}_{ee}) = \det(\boldsymbol{Q}_{ee})$ with $\varphi'(\boldsymbol{Q}_{ee}) = \det(\boldsymbol{Q}_{ee})\boldsymbol{Q}_{ee}^{-1}$. Given the full canonical coordinate system in (2.3.3), one can show that*

$$\boldsymbol{V}^T\boldsymbol{L}^T\varphi'(\boldsymbol{Q}_{ee}^*)\boldsymbol{L}\boldsymbol{V} = \boldsymbol{V}^T\boldsymbol{G}\boldsymbol{K}^T(\boldsymbol{I} - \boldsymbol{K}\boldsymbol{G}^T\boldsymbol{V}^T(\boldsymbol{I}_n + \boldsymbol{\Delta}^*)^{-1}\boldsymbol{V}\boldsymbol{G}\boldsymbol{K}^T)^{-1}\boldsymbol{K}\boldsymbol{G}^T\boldsymbol{V}. \tag{2.5.4}$$

It can be seen that when $\boldsymbol{V} = \boldsymbol{G}$ with $\boldsymbol{G}$ given in (2.3.3), the matrix in (2.5.4) is diagonal and (2.5.3) holds for any $\boldsymbol{\Delta}^*$.

In the proceeding examples, a candidate coordinate system $\boldsymbol{V}$ is provided for each problem by solving condition (2.5.3). For the min-trace problem, the half canonical coordinate system indeed is optimal, according to the results in Schizas et al. [2007]. For the mid-det problem in Section 2.4.2, the full canonical coordinate system is optimal. Moreover, for the weighted min-trace problem where $\varphi(\boldsymbol{Q}_{ee}) = \text{tr}(\boldsymbol{A}\boldsymbol{Q}_{ee}\boldsymbol{A}^T)$, one can check that the weighted half canonical coordinate system, given by the SVD of the weighted half coherence matrix $\boldsymbol{A}\boldsymbol{Q}_{\boldsymbol{\theta x}}\boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}$, is optimal.

2.6 Summary

In this chapter we have considered the problem of compressing a noisy measurement for transmission over a noisy channel, introduced in Schizas et al. [2007]. This problem generalizes the problem of reduced rank filtering (Scharf [1990], Scharf and Thomas [1998], Scharf and Mullis [1998], Hua et al. [2001], and Schreier and Scharf [2006]) and the problem of reduced rank precoder and equalizer design (Scaglione et al. [2002] and Pérez-Cruz et al. [2010]), producing those designs as special cases. We have shown that designs for minimizing trace or determinant of an error covariance matrix share a common architecture. In this architecture, a noisy sensor measurement is first transformed into a system of canonical coordinates. These coordinates are then scaled and rotated into the sub-dominant subspace of the channel noise. The difference between the two designs resides in the definition of canonical coordinates and in the determination of the scaling constants. A generalization to differentiable functions of error covariance leads to a factorization theorem that supports practical design for general criteria.

2.7 Proofs

2.7.1 Proof for Proposition 2.1

We first prove the min-trace result of Proposition 2.1. Let $\widetilde{\boldsymbol{W}} = \boldsymbol{W}\boldsymbol{Q}_{xx}^{1/2}$. Then,

$$
\begin{aligned}
\mathrm{tr}(\boldsymbol{Q}_{ee}) &= \mathrm{tr}(\boldsymbol{Q}_{\theta\theta} - \boldsymbol{Q}_{\theta x}\boldsymbol{Q}_{xx}^{T/2}\widetilde{\boldsymbol{W}}^T(\widetilde{\boldsymbol{W}}\widetilde{\boldsymbol{W}}^T)^+\widetilde{\boldsymbol{W}}\boldsymbol{Q}_{xx}^{1/2}\boldsymbol{Q}_{x\theta}) \\
&= \mathrm{tr}(\boldsymbol{Q}_{\theta\theta}) - \mathrm{tr}(\widetilde{\boldsymbol{W}}^T(\widetilde{\boldsymbol{W}}\widetilde{\boldsymbol{W}}^T)^+\widetilde{\boldsymbol{W}}\boldsymbol{Q}_{xx}^{1/2}\boldsymbol{Q}_{x\theta}\boldsymbol{Q}_{\theta x}\boldsymbol{Q}_{xx}^{T/2}) \\
&= \mathrm{tr}(\boldsymbol{Q}_{\theta\theta}) - \mathrm{tr}(\widetilde{\boldsymbol{W}}^T(\widetilde{\boldsymbol{W}}\widetilde{\boldsymbol{W}}^T)^+\widetilde{\boldsymbol{W}}\boldsymbol{G}\boldsymbol{K}^T\boldsymbol{K}\boldsymbol{G}^T)
\end{aligned}
$$

where the matrices $\boldsymbol{K}$ and $\boldsymbol{G}$ are given by the SVD

$$
\boldsymbol{Q}_{\theta x}\boldsymbol{Q}_{xx}^{T/2} = \boldsymbol{F}\boldsymbol{K}\boldsymbol{G}^T.
$$

Suppose that $\widetilde{\boldsymbol{W}}$ has the SVD $\widetilde{\boldsymbol{W}} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T$. Then,

$$
\mathrm{tr}(\boldsymbol{Q}_{ee}) = \mathrm{tr}(\boldsymbol{Q}_{\theta\theta}) - \mathrm{tr}\left(\begin{bmatrix} \boldsymbol{I}_m & 0 \\ 0 & 0 \end{bmatrix}\boldsymbol{V}^T\boldsymbol{G}\boldsymbol{K}^T\boldsymbol{K}\boldsymbol{G}^T\boldsymbol{V}\right)
$$

Note that the matrix $\boldsymbol{G}^T\boldsymbol{V}$ is an orthogonal matrix. Therefore $\boldsymbol{V}^T\boldsymbol{G}\boldsymbol{K}^T\boldsymbol{K}\boldsymbol{G}^T\boldsymbol{V}$ has the same eigenvalues as the diagonal matrix $\boldsymbol{K}^T\boldsymbol{K}$. Therefore, by the bounds on eigenvalues of matrix products given in Komaroff [1990], it is easy to see that for any $\boldsymbol{W}$,

$$
\mathrm{tr}(\boldsymbol{Q}_{ee}) \geq \mathrm{tr}(\boldsymbol{Q}_{\theta\theta}) - \sum_{i=1}^{m} k_i^2,
$$

and the minimum is attained at $\widetilde{\boldsymbol{W}} = \boldsymbol{G}_m^T$ which yields the optimal compression matrix minimizing $\mathrm{tr}(\boldsymbol{Q}_{ee})$ is $\boldsymbol{W}_0^* = \boldsymbol{G}_m^T\boldsymbol{Q}_{xx}^{-1/2}$. Note that the minimum equals $\mathrm{tr}(\boldsymbol{Q}_{\theta\theta|x}) + \sum_{i=m+1}^{\min\{n,p\}} k_i^2$ since $\mathrm{tr}(\boldsymbol{Q}_{\theta\theta|x}) = \mathrm{tr}(\boldsymbol{Q}_{\theta\theta}) - \sum_{i=1}^{\min\{n,p\}} k_i^2$.

For the min-det result of Proposition 2.1, the proof follows the same track. Let $\widetilde{\boldsymbol{W}} = \boldsymbol{W}\boldsymbol{Q}_{xx}^{1/2}$. Then we have

$$
\begin{aligned}
\det(\boldsymbol{Q}_{ee}) &= \det(\boldsymbol{Q}_{\theta\theta} - \boldsymbol{Q}_{\theta x}\boldsymbol{Q}_{xx}^{T/2}\widetilde{\boldsymbol{W}}^T(\widetilde{\boldsymbol{W}}\widetilde{\boldsymbol{W}}^T)^+\widetilde{\boldsymbol{W}}\boldsymbol{Q}_{xx}^{1/2}\boldsymbol{Q}_{x\theta}) \\
&= \det(\boldsymbol{Q}_{\theta\theta}) \times \det(\boldsymbol{I} - \widetilde{\boldsymbol{W}}^T(\widetilde{\boldsymbol{W}}\widetilde{\boldsymbol{W}}^T)^+\widetilde{\boldsymbol{W}}\boldsymbol{Q}_{xx}^{1/2}\boldsymbol{Q}_{\theta\theta}\boldsymbol{Q}_{\theta\theta}^{-1}\boldsymbol{Q}_{\theta\theta}\boldsymbol{Q}_{xx}^{T/2}) \\
&= \det(\boldsymbol{Q}_{\theta\theta}) \times \det(\boldsymbol{I} - \widetilde{\boldsymbol{W}}^T(\widetilde{\boldsymbol{W}}\widetilde{\boldsymbol{W}}^T)^+\widetilde{\boldsymbol{W}}\boldsymbol{G}\boldsymbol{K}^T\boldsymbol{K}\boldsymbol{G}^T)
\end{aligned}
$$

where the matrices $\boldsymbol{K}$ and $\boldsymbol{G}$ are given by the SVD

$$\boldsymbol{Q}_{\boldsymbol{\theta\theta}}^{-1/2}\boldsymbol{Q}_{\boldsymbol{\theta x}}\boldsymbol{Q}_{\boldsymbol{xx}}^{T/2} = \boldsymbol{FKG}^T.$$

Consider the SVD $\widetilde{\boldsymbol{W}} = \boldsymbol{U\Sigma V}^T$. Then,

$$\det(\boldsymbol{Q}_{ee}) = \det(\boldsymbol{Q}_{\boldsymbol{\theta\theta}}) \times \det(\boldsymbol{I} - \begin{bmatrix} \boldsymbol{I}_m & 0 \\ 0 & 0 \end{bmatrix} \boldsymbol{V}^T\boldsymbol{GK}^T\boldsymbol{KG}^T\boldsymbol{V})$$

Proceeding as the min-trace problem, one can show that

$$\det(\boldsymbol{Q}_{ee}) \geq \det(\boldsymbol{Q}_{\boldsymbol{\theta\theta}}) \times \prod_{i=1}^{m}(1 - k_i^2).$$

The minimum is attained at $\widetilde{\boldsymbol{W}} = \boldsymbol{G}_m^T$ which yields $\boldsymbol{W}_0^* = \boldsymbol{G}_m^T\boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}$ as the optimal compression matrix minimizing $\det(\boldsymbol{Q}_{ee})$. Note that the minimum can be rewritten as $\det(\boldsymbol{Q}_{\boldsymbol{\theta\theta}}|\boldsymbol{x}) \times \prod_{i=m+1}^{\min\{n,p\}}(1 - k_i^2)^{-1}$, from the fact that $\det(\boldsymbol{Q}_{\boldsymbol{\theta\theta}}|\boldsymbol{x}) = \det(\boldsymbol{Q}_{\boldsymbol{\theta\theta}}) \times \prod_{i=1}^{\min\{n,p\}}(1 - k_i^2)$.

2.7.2   Proof for Lemma 2.1

Let $\boldsymbol{A}$ and $\boldsymbol{B}$ be two $n \times n$ positive semi-definite matrices with eigenvalues $\mu_1 \geq \ldots \mu_n \geq 0$ and $\lambda_1 \geq \ldots \lambda_n \geq 0$. Then,

$$\det(\boldsymbol{I}_n + \boldsymbol{AB}) \geq \sum_{i=1}^{n} \lambda_i\mu_{n-i+1}. \tag{2.7.1}$$

The proof uses the same technique as the proof of Lemma 3 in Witsenhausen [1975] and is omitted here.

By Theorem 2.3, it suffices to consider the compression matrix $\boldsymbol{W}$ with factorization

$$\boldsymbol{W} = \boldsymbol{U}_{\boldsymbol{\omega}}\boldsymbol{\Pi}_m\boldsymbol{\Sigma}\boldsymbol{\Pi}_n^T\boldsymbol{V}^T\boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2} \tag{2.7.2}$$

where $\boldsymbol{V}$ is an orthogonal matrix, and $\boldsymbol{\Pi}_n$ and $\boldsymbol{\Pi}_m$ are permutation matrices. Define $\boldsymbol{\Lambda} = (\boldsymbol{I}_n + \boldsymbol{\Sigma}^T\boldsymbol{\Pi}_m^T\boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1}\boldsymbol{\Pi}_m\boldsymbol{\Sigma})^{-1}$, $\boldsymbol{L} = \boldsymbol{Q}_{\boldsymbol{\theta x}}\boldsymbol{Q}_{\boldsymbol{xx}}^{-T/2}$ and $\boldsymbol{C} = \boldsymbol{Q}_{\boldsymbol{\theta\theta}}^{-1/2}\boldsymbol{L}$. Then,

$$\det(\boldsymbol{Q}_{ee}) = \det\left(\boldsymbol{Q}_{\boldsymbol{\theta\theta}} - \boldsymbol{LL}^T + \boldsymbol{LV}\boldsymbol{\Pi}_n\boldsymbol{\Lambda}\boldsymbol{\Pi}_n^T\boldsymbol{V}^T\boldsymbol{L}^T\right)$$

$$= \det(\boldsymbol{Q}_{\boldsymbol{\theta\theta}})\det\left(\boldsymbol{I}_p - \boldsymbol{CC}^T + \boldsymbol{CV}\boldsymbol{\Pi}_n\boldsymbol{\Lambda}\boldsymbol{\Pi}_n^T\boldsymbol{V}^T\boldsymbol{C}^T\right)$$

$$= \det(\boldsymbol{Q}_{\boldsymbol{\theta\theta}})\det\left(\boldsymbol{I}_p - \boldsymbol{CC}^T\right)\det\left(\boldsymbol{I} + \boldsymbol{\Pi}_n\boldsymbol{\Lambda}\boldsymbol{\Pi}_n^T\boldsymbol{V}^T\boldsymbol{C}^T(\boldsymbol{I}_p - \boldsymbol{CC}^T)^{-1}\boldsymbol{CV}\right)$$

Given the SVD, it can be seen that $\boldsymbol{C} = \boldsymbol{FKG}^T$ in (2.3.3),

$$\boldsymbol{V}^T\boldsymbol{C}^T(\boldsymbol{I}_p - \boldsymbol{CC}^T)^{-1}\boldsymbol{CV} = \boldsymbol{V}^T\boldsymbol{GK}^T(\boldsymbol{I}_p - \boldsymbol{KK}^T)^{-1}\boldsymbol{KG}^T\boldsymbol{V},$$

which has eigenvalues $\gamma_1^2 \geq \ldots \geq \gamma_n^2 \geq 0$ where $\gamma_i^2 = k_i^2/(1 - k_i^2)$ for $i = 1, \ldots, \min\{n, p\}$ and 0 otherwise. Moreover, the matrix $\boldsymbol{\Pi}_n\boldsymbol{\Lambda}\boldsymbol{\Pi}_n^T$ has the same eigenvalues as $\boldsymbol{\Lambda}$. Let $\boldsymbol{W}^0 = \boldsymbol{U_\omega}\boldsymbol{\Pi}_m\boldsymbol{\Sigma}(\boldsymbol{\Pi}_n^0)^T\boldsymbol{G}^T\boldsymbol{Q_{xx}}^{-1/2}$ where $\boldsymbol{\Pi}_n^0$ is an $n \times n$ permutation matrix such that the diagonal elements of the diagonal matrix $\boldsymbol{\Pi}_n^0\boldsymbol{\Lambda}\boldsymbol{\Pi}_n^{0T}$ are increasingly ordered. Then, $\boldsymbol{W}^0 \in \Omega_{\boldsymbol{G}}$ and

$$\det(\boldsymbol{Q}_{ee}(\boldsymbol{W})) \geq \det(\boldsymbol{Q_{\theta\theta}})\det\left(\boldsymbol{I}_p - \boldsymbol{CC}^T\right)\prod_{i=1}^{n}\gamma_i^2\lambda_{(i)} = \det(\boldsymbol{Q}_{ee}(\boldsymbol{W}^0)). \qquad (2.7.3)$$

where the inequality is a direct consequence of (2.7.1). Minimizing both sides of (2.7.3) over $\boldsymbol{W}$ with factorization (2.7.2) and $\boldsymbol{W}^0 \in \Omega_{\boldsymbol{G}}$, respectively, we have

$$\det(\boldsymbol{W}_{\det}^*) \geq \min_{\boldsymbol{W}\in\Omega_{\boldsymbol{G}}}\det(\boldsymbol{Q}_{ee}(\boldsymbol{W})). \qquad (2.7.4)$$

The proof is therefore completed by the fact that $\det(\boldsymbol{W}_{\det}^*) \leq \min_{\boldsymbol{W}\in\Omega_{\boldsymbol{G}}}\det(\boldsymbol{Q}_{ee}(\boldsymbol{W}))$ by definition.

2.7.3 Proof for Theorem 2.2

Given Lemma 2.1, we can restrict $\boldsymbol{W} = \boldsymbol{U_\omega}\boldsymbol{\Pi}_m\boldsymbol{\Sigma}\boldsymbol{\Pi}_n^T\boldsymbol{G}^T\boldsymbol{Q_{xx}}^{-1/2}$ where $\boldsymbol{\Pi}_n$ and $\boldsymbol{\Pi}_m$ are permutation matrices, and $\boldsymbol{\Sigma} \in \mathbb{R}^{m\times n}$ is a diagonal matrix with diagonal elements $\sigma_{11}, \ldots, \sigma_{mm}$. Let $\pi_m(i)$ be the index of the entry equal to unity in the ith column of $\boldsymbol{\Pi}_m$, and $\pi_n(i)$ be the index of the unity entry in the ith column of $\boldsymbol{\Pi}_n$. Then

$$\det(\boldsymbol{Q}_{ee}(\sigma_{11}, \ldots, \sigma_{mm}, \pi_m, \pi_n)) = \prod_{i=1}^{m}\left(1 + \frac{\gamma_{\pi_n(i)}^2}{1 + \lambda_{\pi_m(i)}\sigma_{ii}^2}\right)\prod_{j=m+1}^{n}\log(1 + \gamma_{\pi_n(j)}^2) \quad (2.7.5)$$

where $\lambda_i = \sigma_{\boldsymbol{\omega},i}^{-2}$ for $i = 1, \ldots, m$ with $\lambda_1 \geq \ldots \geq \lambda_m$. We try to minimize $\det(\boldsymbol{Q}_{ee})$ over all possible permutations and $(\sigma_{11}, \ldots, \sigma_{mm})^T$ subject to $\sum_{i=1}^{m}\sigma_{ii}^2 \leq c$.

First, we show that the minimum of $\det(\boldsymbol{Q}_{ee})$ can be achieved when $\pi_m(i) = i$ for $i = 1, \ldots, m$ and $\pi_n(j) = j$ for $j = 1 \ldots, n$, or equivalently, when the optimal permutation matrices $\boldsymbol{\Pi}_m = \boldsymbol{I}_m$ and $\boldsymbol{\Pi}_n = \boldsymbol{I}_n$. For a given permutation, define

$$f(\pi_n, \pi_m) = \min_{\sum_{i=1}^m \sigma_{ii}^2 \leq c} \det(\boldsymbol{Q}_{ee}).$$

The optimal permutation (π_n^*, π_m^*) satisfies

$$f(\pi_n^*, \pi_m^*) \leq f(\pi_n, \pi_m) \tag{2.7.6}$$

for all other permutations (π_n, π_m). It is easy to see that for any $i = 1, \ldots, m$, one must have $\gamma_{\pi_n^*(i)}^2 \geq \max\{\gamma_{\pi_n^*(m+1)}^2, \ldots, \gamma_{\pi_n^*(n)}^2\}$. Moreover, since the orders of $\{\pi_n(j)\}_{j=m+1}^n$ do not affect the value of (2.7.5), we can set WLOG $\pi_n^*(j) = j$ for $j = m + 1, \ldots, n$. For $i = 1, \ldots, m$, it can be seen that $\pi_n(i)$ and $\pi_m(i)$ appear in (2.7.5) pairwise. Therefore, we can set WLOG that $\pi_m^*(i) = i$ for $i = 1, \ldots, \kappa$ and $f(\pi_n) := f(\pi_n, \pi_m^*)$. Then the objective is to search for the optimal permutation $\pi_n^*(i)$ that minimizes

$$\prod_{i=1}^m \left(1 + \frac{\gamma_{\pi_n(i)}^2}{1 + \lambda_i \sigma_{ii}^2} \right). \tag{2.7.7}$$

Let's start from the simple case with $m = 2$. When $\lambda_1 = \lambda_2$ or $\gamma_1^2 = \gamma_2^2$, the function in (2.7.7) is permutation invariant. When $\lambda_1 > \lambda_2$ and $\gamma_1^2 > \gamma_2^2$, only two permutations are possible, $\pi_n^1(i) = i$ or $\pi_n^2(i) = 3 - i$ for $i = 1, 2$. To minimize (2.7.7), consider the following functions

$$h_1(x) = (1 + \frac{\gamma_1^2}{1 + \lambda_1 cx}) \times (1 + \frac{\gamma_2^2}{1 + \lambda_2 c(1-x)});$$
$$h_2(x) = (1 + \frac{\gamma_2^2}{1 + \lambda_1 cx}) \times (1 + \frac{\gamma_1^2}{1 + \lambda_2 c(1-x)}).$$

with $cx = \sigma_{11}^2$ and $c(1-x) = \sigma_{22}^2$, in which case $\sigma_{11}^2 + \sigma_{22}^2 = c$. Then, $f(\pi_n^1) < f(\pi_n^2)$ is equivalent to

$$\min_{x \in [0,1]} h_1(x) < \min_{x \in [0,1]} h_2(x). \tag{2.7.8}$$

31

Straightforward calculation gives that

$$h_2(x) - h_1(x) = \frac{c(\gamma_1^2 - \gamma_2^2)((\lambda_1 + \lambda_2)x - \lambda_2)}{(1 + \lambda_1 cx)(1 + \lambda_2 c(1 - x))}$$

Therefore $h_2(x) > h_1(x)$ for any $x \in [\frac{\lambda_2}{\lambda_1 + \lambda_2}, 1]$. For $x \in [0, \frac{\lambda_2}{\lambda_1 + \lambda_2}]$, one can show that $h_2(x) > h_1(\frac{\lambda_2}{\lambda_1}(1 - x))$ with $\frac{\lambda_2}{\lambda_1}(1 - x) \in [\frac{\lambda_2}{\lambda_1 + \lambda_2}, 1]$. Therefore (2.7.8) holds and $\pi_n^*(i) = i$ for $i = 1, 2$. By checking the first and second derivative of $h_1(x)$, the minimum of h_1 is attained at $\boldsymbol{x} \in [1/2, 1]$. This directly yields that $cx \geq c(1 - x)$ or equivalently, $\sigma_{11}^2 \geq \sigma_{22}^2$, meaning that allocation of power decreases with increasing channel index.

For the general cases where $m \geq 2$, define $(\sigma_{11}^{\pi_n}, \ldots, \sigma_{mm}^{\pi_n})$

$$(\sigma_{11}^{\pi_n}, \ldots, \sigma_{mm}^{\pi_n}) = \underset{\sum_{i=1}^m \sigma_{ii}^2 \leq c}{\arg \min} \; \det(\boldsymbol{Q}_{ee}(\sigma_{11}, \ldots, \sigma_{mm}, \pi_n, \pi_m^*)).$$

Suppose that there exist $1 \leq i < j \leq m$ with $\gamma_{\pi_n(i)}^2 < \gamma_{\pi_n(j)}^2$. Let $\tilde{\pi}_n$ be a new permutation with $\tilde{\pi}_n(i) = \pi_n(j)$, $\tilde{\pi}_n(j) = \pi_n(i)$, and $\tilde{\pi}_n(k) = \pi_n(k)$ for $k \neq i, j$. Define

$$(\tilde{\sigma}_{ii}, \tilde{\sigma}_{jj}) = \underset{\sigma_{ii}^2 + \sigma_{jj}^2 \leq \sigma_{ii}^{\pi_n 2} + \sigma_{jj}^{\pi_n 2}}{\arg \min} \left(1 + \frac{\gamma_{\tilde{\pi}_n(i)}^2}{1 + \lambda_i \sigma_{ii}^2}\right) \left(1 + \frac{\gamma_{\tilde{\pi}_n(j)}^2}{1 + \lambda_j \sigma_{jj}^2}\right). \tag{2.7.9}$$

Given the result for the $m = 2$ case, it is straightforward to see that

$$\det(\boldsymbol{Q}_{ee}(\sigma_{11}^{\pi_n}, \ldots, \tilde{\sigma}_{ii}, \ldots, \tilde{\sigma}_{jj}, \ldots, \sigma_{mm}^{\pi_n}, \tilde{\pi}_n, \pi_m)) < \det(\boldsymbol{Q}_{ee}(\sigma_{11}^{\pi_n}, \ldots, \sigma_{mm}^{\pi_n}, \pi_n, \pi_m))$$

Therefore, the permutation π_n cannot be the optimal permutation. Among all the permutations $\gamma_{\pi_n(1)}^2 \geq \ldots \geq \gamma_{\pi_n(m)}^2$, we can choose WLOG that $\pi_n^*(i) = i$ for $i = 1, \ldots, m$.

Next our problem focuses on finding the sequence

$$\{\sigma_{ii}^*\}_{i=1}^m = \underset{\sum_{i=1}^m \sigma_{ii}^2 \leq c}{\arg \min} \sum_{i=1}^m \log\left(1 + \frac{\gamma_i^2}{1 + \lambda_i \sigma_{ii}^2}\right). \tag{2.7.10}$$

The log operator is implemented to simplify calculation. Note that the objective function in (2.7.10) is a strictly convex function, therefore (2.7.10) is a convex optimization problem with unique minimizer. Moreover, the function is strictly decreasing in σ_{ii}^2. Hence the minimum is attained at $\sum_{i=1}^m \sigma_{ii}^2 = c$. The Lagrangian is

$$L(\sigma_{11}, \ldots, \sigma_{mm}; \mu) = \sum_{i=1}^m \log\left(1 + \frac{\gamma_i^2}{1 + \lambda_i \sigma_{ii}^2}\right) + \mu\left(\sum_{i=1}^m \sigma_{ii}^2 - c\right)$$

By setting the first derivative of the Lagrangian with respect to $\sigma_{11}, \ldots, \sigma_{mm}$, and μ to zero, the necessary conditions for any minimizer of problem (2.7.10) are

$$-\left(1 + \frac{\gamma_i^2}{1 + \lambda_i \sigma_{ii}^2}\right)^{-1} \frac{2\gamma_i^2 \lambda_i \sigma_{ii}}{(1 + \lambda_i \sigma_{ii}^2)^2} + 2\mu\sigma_{ii} = 0 \tag{2.7.11}$$

$$\sum_{i=1}^{m} \sigma_{ii}^2 = c \tag{2.7.12}$$

Equation (2.7.11) yields either $\sigma_{ii} = 0$ or

$$\sigma_{ii} = \sqrt{\frac{1}{2\lambda_i}\left(\sqrt{\gamma_i^4 + 4\lambda_i\gamma_i^2\mu^{-1}} - 2 - \gamma_i^2\right)} \tag{2.7.13}$$

The solution in (2.7.13) provides a feasible solution for σ_{ii} only when $\mu \leq \lambda_i k_i^2 = k_i^2/\sigma_i^2$ or $1/\mu > \sigma_{\boldsymbol{w},i}^2/k_i^2$ where $k_i^2 = \gamma_i^2(1 + \gamma_i^2)^{-1}$ is the squared canonical correlation.

Next we investigate the possible minimizers by checking the second derivative of L, which is

$$\frac{\partial^2 L(\sigma_{11}, \ldots, \sigma_{mm}; \mu)}{\partial \sigma_{ii}^2} = 2\mu - \frac{2\gamma_i^2\lambda_i}{(1 + \gamma_i^2 + \lambda_i\sigma_{ii}^2)(1 + \lambda_i\sigma_{ii}^2)}$$
$$+ \frac{4\gamma_i^2\lambda_i^2\sigma_{ii}^2}{(1 + \gamma_i^2 + \lambda_i\sigma_{ii}^2)^2(1 + \lambda_i\sigma_{ii}^2)} + \frac{4\gamma_i^2\lambda_i^2\sigma_{ii}^2}{(1 + \gamma_i^2 + \lambda_i\sigma_{ii}^2)(1 + \lambda_i\sigma_{ii}^2)^2} \tag{2.7.14}$$

Substituting (2.7.13), the second derivative

$$\frac{\partial^2 L(\sigma_{11}, \ldots, \sigma_{mm}; \mu)}{\partial \sigma_{ii}^2} = 4\mu\lambda_i\sigma_{ii}\left(\frac{1}{1 + \gamma_i^2 + \lambda_i\sigma_{ii}^2} + \frac{1}{1 + \lambda_i\sigma_{ii}^2}\right)$$

is strictly positive when $1/\mu > \sigma_{\boldsymbol{w},i}^2/k_i^2$. When $\sigma_{ii} = 0$,

$$\left.\frac{\partial^2 L(\sigma_{11}, \ldots, \sigma_{mm}; \mu)}{\partial \sigma_{ii}^2}\right|_{\sigma_{ii}=0} = 2\mu - \frac{2\lambda_i\gamma_i^2}{1 + \gamma_i^2}$$

which is negative when $1/\mu \geq \sigma_{\boldsymbol{w},i}^2/k_i^2$ and positive when $1/\mu < \sigma_{\boldsymbol{w},i}^2/k_i^2$.

Let κ be the maximum integer between 1 and m such that $1/\mu > \sigma_{\boldsymbol{w},i}^2/k_i^2$ (or equivalently $\sigma_{ii} > 0$ for $i = 1, \ldots, \kappa$) and $1/\mu \leq \sigma_{\boldsymbol{w},i}^2/k_i^2$ ($\sigma_{ii} = 0$ for $i = \kappa + 1, \ldots, m$), where the value of μ is determined by the power constraint (2.7.12). Then, the Hessian matrix at

$$\sigma_{ii}^* = \begin{cases} \sqrt{\frac{1}{2\lambda_i}\left(\sqrt{\gamma_i^4 + 4\lambda_i\gamma_i^2/\mu} - 2 - \gamma_i^2\right)}, & \text{for } i = 1, \ldots, \kappa \\ 0 & \text{for } i = \kappa + 1, \ldots, m \end{cases} \tag{2.7.15}$$

is strictly positive and (2.7.15) is the minimizer.

As a summary, the optimal compression matrix minimizing $\det(\boldsymbol{Q}_{ee})$ is

$$\boldsymbol{W}_{\det}^* = \boldsymbol{U}_{\boldsymbol{\omega}}\boldsymbol{\Sigma}_{det}^*\boldsymbol{G}^T\boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}.$$

where $\boldsymbol{\Sigma}_{det}^* \in \mathbb{R}^{m \times n}$ is a diagonal scaling matrix with diagonal elements given in (2.7.15).

2.7.4 Proof for Theorem 2.3

Since $\boldsymbol{U}_{\boldsymbol{\omega}}$ and $\boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}$ are invertible, one can uniquely define a matrix $\boldsymbol{\Phi} \in \mathbb{R}^{m \times n}$ such that

$$\boldsymbol{W} = \boldsymbol{U}_{\boldsymbol{\omega}}\boldsymbol{\Phi}\boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}. \tag{2.7.16}$$

The power constraint is equivalent to $\operatorname{tr}(\boldsymbol{\Phi}\boldsymbol{\Phi}^T) \leq c$ since $\operatorname{tr}(\boldsymbol{W}\boldsymbol{Q}_{\boldsymbol{xx}}\boldsymbol{W}^T) = \operatorname{tr}(\boldsymbol{\Phi}\boldsymbol{\Phi}^T)$ for any pair of $(\boldsymbol{W}, \boldsymbol{\Phi})$ satisfying (2.7.16). Moreover, the error covariance $\boldsymbol{Q}_{ee}$ simplifies to

$$\boldsymbol{Q}_{ee} = \boldsymbol{Q}_{\boldsymbol{\theta\theta}} - \boldsymbol{L}\boldsymbol{\Phi}^T(\boldsymbol{\Phi}\boldsymbol{\Phi}^T + \boldsymbol{\Sigma}_{\boldsymbol{\omega}})^{-1}\boldsymbol{\Phi}\boldsymbol{L}^T, \tag{2.7.17}$$

where $\boldsymbol{L} = \boldsymbol{Q}_{\boldsymbol{\theta x}}\boldsymbol{Q}_{\boldsymbol{xx}}^{-T/2}$ is the LMMSE filter for estimating $\boldsymbol{\theta}$ from $\boldsymbol{x}$. Using the matrix inversion lemma $(\boldsymbol{I} + \boldsymbol{\Phi}^T\boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1}\boldsymbol{\Phi})^{-1} = \boldsymbol{I} - \boldsymbol{\Phi}^T(\boldsymbol{\Phi}\boldsymbol{\Phi}^T + \boldsymbol{\Sigma}_{\boldsymbol{\omega}})^{-1}\boldsymbol{\Phi}$, the error covariance can be rewritten as

$$\boldsymbol{Q}_{ee} = \boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}} + \boldsymbol{L}(\boldsymbol{I} + \boldsymbol{\Phi}^T\boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1}\boldsymbol{\Phi})^{-1}\boldsymbol{L}^T, \tag{2.7.18}$$

where $\boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}} = \boldsymbol{Q}_{\boldsymbol{\theta\theta}} - \boldsymbol{L}\boldsymbol{L}^T$ is a constant matrix with respect to $\boldsymbol{\Phi}$.

We define an alternative optimization w.r.t. $\boldsymbol{\Phi}$ as

$$\boldsymbol{\Phi}^* = \arg\min \varphi(\boldsymbol{Q}_{ee}) \quad s.t. \quad \boldsymbol{\Phi} \in \mathbb{R}^{m \times n}, \operatorname{tr}(\boldsymbol{\Phi}\boldsymbol{\Phi}^T) \leq c. \tag{2.7.19}$$

The Lagrangian is $L(\boldsymbol{\Phi}; \mu) = \varphi(\boldsymbol{Q}_{ee}) + \mu(\operatorname{tr}(\boldsymbol{\Phi}\boldsymbol{\Phi}^T) - c)$, and the necessary condition for any optimizer $\boldsymbol{\Phi}$ is

$$\frac{\partial}{\partial\boldsymbol{\Phi}}L(\boldsymbol{\Phi}; \mu) = \frac{\partial\varphi(\boldsymbol{Q}_{ee})}{\partial\boldsymbol{\Phi}} + 2\mu\boldsymbol{\Phi} = \boldsymbol{0}$$

By the matrix derivative chain rule,

$$
\begin{aligned}
\frac{\partial \varphi(\boldsymbol{Q}_{ee})}{\partial \phi_{ij}} &= \operatorname{tr}\left(\left(\frac{\partial \varphi(\boldsymbol{Q}_{ee})}{\partial \boldsymbol{Q}_{ee}}\right)^T \frac{\partial \boldsymbol{Q}_{ee}}{\partial \phi_{ij}}\right) \\
&= \operatorname{tr}(\varphi'(\boldsymbol{Q}_{ee})^T \frac{\partial \boldsymbol{L}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \boldsymbol{L}^T}{\partial \phi_{ij}}) \\
&= -\operatorname{tr}\left((\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \boldsymbol{L}^T \varphi'(\boldsymbol{Q}_{ee})^T \boldsymbol{L}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{J}_{ij}\right) \\
&\quad -\operatorname{tr}\left(\boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \boldsymbol{L}^T \varphi'(\boldsymbol{Q}_{ee})^T \boldsymbol{L}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \boldsymbol{J}_{ij}^T\right)
\end{aligned}
$$

By the fact $\operatorname{tr}(\boldsymbol{A} \boldsymbol{J}_{ij}) = \boldsymbol{a}_{ji}$ and $\operatorname{tr}(\boldsymbol{A} \boldsymbol{J}_{ij}^T) = \boldsymbol{a}_{ij}$, we have

$$
\begin{aligned}
\frac{\partial \varphi(\boldsymbol{Q}_{ee})}{\partial \boldsymbol{\Phi}} &= -\boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \boldsymbol{L}^T \varphi'(\boldsymbol{Q}_{ee})^T \boldsymbol{L}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \\
&\quad -\boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \boldsymbol{L}^T \varphi'(\boldsymbol{Q}_{ee}) \boldsymbol{L}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1}
\end{aligned}
$$

Left multiply $\frac{\partial}{\partial \boldsymbol{\Phi}} L(\boldsymbol{\Phi}; \mu)$ by $\boldsymbol{\Sigma}_{\boldsymbol{\omega}}$ and right multiply by $\boldsymbol{\Phi}$. Then

$$
\begin{aligned}
2\mu \boldsymbol{\Sigma}_{\boldsymbol{\omega}} \boldsymbol{\Phi} \boldsymbol{\Phi}^T &= \boldsymbol{\Phi}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \boldsymbol{L}^T \varphi'(\boldsymbol{Q}_{ee})^T \boldsymbol{L}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \boldsymbol{\Phi}^T \\
&\quad + \boldsymbol{\Phi}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \boldsymbol{L}^T \varphi'(\boldsymbol{Q}_{ee}) \boldsymbol{L}(\boldsymbol{I} + \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1} \boldsymbol{\Phi})^{-1} \boldsymbol{\Phi}^T
\end{aligned}
$$

Since the RHS is a symmetric matrix, $\boldsymbol{\Phi} \boldsymbol{\Phi}^T \boldsymbol{\Sigma}_{\boldsymbol{\omega}} = \boldsymbol{\Sigma}_{\boldsymbol{\omega}} \boldsymbol{\Phi} \boldsymbol{\Phi}^T$. When the diagonal matrix $\boldsymbol{\Sigma}_{\boldsymbol{\omega}}$ has distinct diagonal elements, it can be seen that the symmetric matrix $\boldsymbol{\Phi} \boldsymbol{\Phi}^T$ must be a diagonal matrix.

Given the SVD $\boldsymbol{\Phi} = \boldsymbol{U} \boldsymbol{\Sigma} \boldsymbol{V}^T$ where $\boldsymbol{U}, \boldsymbol{V}$ are orthogonal matrices, and $\boldsymbol{\Sigma} \in \mathbb{R}^{m \times n}$ is diagonal. Then by Lemma 6.8 of Pukelsheim [1993], there exists an $m \times m$ permutation matrix $\boldsymbol{\Pi}_m$ such that $\boldsymbol{\Phi} \boldsymbol{\Phi}^T = \boldsymbol{\Pi}_m \boldsymbol{\Sigma} \boldsymbol{\Sigma}^T \boldsymbol{\Pi}_m^T$. Therefore, $\boldsymbol{\Phi} = \boldsymbol{\Pi}_m \boldsymbol{\Sigma} \boldsymbol{V}^T$. Plugging $\boldsymbol{\Phi}$ in (2.7.16), it can be seen that $\boldsymbol{W}^*$ can be factorized as $\boldsymbol{W}^* = \boldsymbol{U}_{\boldsymbol{\omega}} \boldsymbol{\Pi}_m^* \boldsymbol{\Sigma}^* \boldsymbol{V}^{*T} \boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}$.

2.7.5 Proof for Lemma 2.2

Suppose that $\boldsymbol{W}^* = \boldsymbol{U}_{\boldsymbol{\omega}} \boldsymbol{\Pi}_m^* \boldsymbol{\Sigma}^* \boldsymbol{V}^{*T} \boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}$ is a solution for problem (2.7.19). For an anti-symmetric matrix $\boldsymbol{X} \in \mathbb{R}^{n \times n}$, let $\boldsymbol{W}(t) = \boldsymbol{U}_{\boldsymbol{\omega}} \boldsymbol{\Pi}_m^* \boldsymbol{\Sigma}^* e^{-t\boldsymbol{X}} \boldsymbol{V}^{*T} \boldsymbol{Q}_{\boldsymbol{xx}}^{-1/2}$ and $\boldsymbol{Q}_{ee}(t)$ the

resulting error covariance at compression $\boldsymbol{W}(t)$. Define a function $f(t)$ for $t \in \mathbb{R}^1$ with $f(t) = \varphi(\boldsymbol{Q}_{ee}(t))$. It is easy to show that

$$f(t) = \varphi(\boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}} + \boldsymbol{L}(\boldsymbol{I} + \boldsymbol{V}^* e^{t\boldsymbol{X}} \boldsymbol{\Delta}^* e^{-t\boldsymbol{X}} \boldsymbol{V}^{*T})^{-1} \boldsymbol{L}^T),$$

where $\boldsymbol{\Delta}^* = \boldsymbol{\Sigma}^{*T} \boldsymbol{\Pi}_m^{*T} \boldsymbol{\Sigma}_{\boldsymbol{\omega}} \boldsymbol{\Pi}_m^* \boldsymbol{\Sigma}^*$. Notice that $\boldsymbol{W}(0) = \boldsymbol{W}^*$. Since $\boldsymbol{W}^*$ is a minimizer for problem (2.7.19), we have $f(t) \geq f(0)$ for any $t \in \mathbb{R}$. Therefore, the necessary condition for $\boldsymbol{V}^*$ is $\frac{\partial f}{\partial t}\big|_{t=0} = 0$ for any anti-symmetric matrix $\boldsymbol{X}$. Let $\boldsymbol{X} = [\boldsymbol{V}^{*T} \boldsymbol{L}^T \varphi'(\boldsymbol{Q}_{ee}^*) \boldsymbol{L} \boldsymbol{V}^*, (\boldsymbol{I}_n + \boldsymbol{\Delta}^*)^{-1}]$, one can show that the necessary condition yields $\mathrm{tr}(\boldsymbol{X} \boldsymbol{X}^T) = 0$. Therefore $\boldsymbol{X} = \boldsymbol{0}$, i.e, $\boldsymbol{V}^{*T} \boldsymbol{L}^T \varphi'(\boldsymbol{Q}_{ee}^*) \boldsymbol{L} \boldsymbol{V}^*$ and $(\boldsymbol{I}_n + \boldsymbol{\Delta}^*)^{-1}$ commute.

CHAPTER 3

FUSION INSPIRED CHANNEL DESIGN[2]

3.1 Introduction

Consider the following two-channel system, as illustrated in Fig. 3.1,

$$
\begin{aligned}
\boldsymbol{x} &= \boldsymbol{F\theta} + \boldsymbol{u} \\
\boldsymbol{y} &= \boldsymbol{G\phi} + \boldsymbol{v}.
\end{aligned}
\tag{3.1.1}
$$

The first channel is the primary channel that carries the signal of interest $\boldsymbol{\theta}$. The secondary channel carries a signal $\boldsymbol{\phi}$ that shares a joint distribution with $\boldsymbol{\theta}$. The measurements $\boldsymbol{x}$ and $\boldsymbol{y}$ are linear transformations of the input signals with measurement noises $\boldsymbol{u}$ and $\boldsymbol{v}$, respectively. For example, the elements of the primary signal $\boldsymbol{\theta}$ may be the complex scattering coefficients of several radar-scattering targets and the elements of the secondary signal $\boldsymbol{\phi}$ may be intensities in an optical map of these same optical-scattering targets. The measurement $\boldsymbol{x}$ is then a range-doppler map and the measurement $\boldsymbol{y}$ is an optical image. We assume a known signal model, i.e., the joint distribution of $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. When the signals $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are correlated, the measurements $\boldsymbol{x}$ and $\boldsymbol{y}$ both contain information about $\boldsymbol{\theta}$ and we can combine them to estimate $\boldsymbol{\theta}$. The fused estimate is expected to perform better than the estimate from a single source of measurements. In this chapter, our objective is to design the channel matrix $\boldsymbol{G}$, with the primary channel fixed, such that the fused estimate achieves the best performance.

For a one-channel system $\boldsymbol{x} = \boldsymbol{F\theta} + \boldsymbol{u}$, designing the channel matrix $\boldsymbol{F}$ exhibits parallels to the linear precoding problem for multiple-input-multiple-output (MIMO) communication systems by considering $\boldsymbol{F}$ as the precoder into an identity channel matrix. The linear precoding design for MIMO channels has been studied in the literature,

---

37
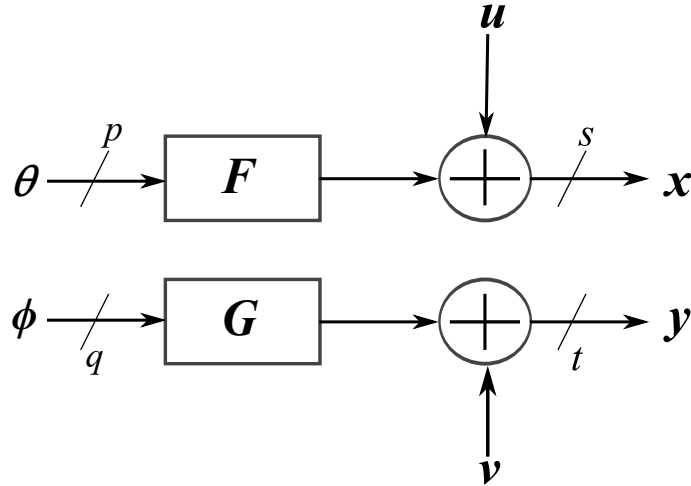
Figure 3.1: A two-channel system with two linear channels.

e.g., Palomar et al. [2003], Scaglione et al. [2002], Pérez-Cruz et al. [2010], Vosoughi and Scaglione [2007], Liu et al. [2012], Lamarca [2009], Sampath et al. [2001], Xiao et al. [2011], and Palomar and Jiang [2007]. The optimal precoding is designed under various criteria, for example, maximizing signal-to-noise ratio (SNR) and signal-to-interference-noise ratio (SINR) (see Palomar et al. [2003] and Scaglione et al. [2002]). Another criterion that has drawn more attention recently is the mutual information between input and output signals (see Pérez-Cruz et al. [2010], Carson et al. [2012], Liu et al. [2012], and Lamarca [2009]). This information-based criterion is connected with estimation theory in a vector Gaussian channel with arbitrary input distribution by linking the mutual information with the minimum mean squared error (MMSE) (see Guo et al. [2005] and Palomar and Verdú [2006]). In Pérez-Cruz et al. [2010], an optimal precoding matrix for the MIMO Gaussian channel with arbitrary input is expressed as the solution of a fixed point equation. When the input signal is Gaussian distributed, the one-channel design problem can be solved as a singular value decomposition (SVD) problem. More specially, the optimal channel matrix has its singular vectors allocated to create non-interfering subchannels and the singular values chosen to solve a generalized waterfilling problem (see Lamarca [2009] and Cover and Thomas [2005]). In Liu et al.

[2012], a greedy adaptive approach is considered to design a channel matrix row by row to maximize information gain.

In Figure 3.1, if both $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are of interest, the two-channel system can be expressed as a one-channel system with a block-diagonal channel matrix. However, due to the "nuisance" signal $\boldsymbol{\phi}$, our two-channel system design problem is fundamentally more difficult than the one-channel system design. We fix the primary channel and design the secondary channel matrix $\boldsymbol{G}$ that maximizes the information gain brought by adding the secondary channel, subject to the total power constraint $\mathrm{tr}(\boldsymbol{G}\boldsymbol{G}^T) \leq c$ with c a pre-determined constant. We call this a *one-channel design problem in a two-channel system.* Analytical solutions are derived for some special cases. In general, this is not a convex problem. Moreover, this problem cannot be formulated as an SVD problem, in contrast to the one-channel system design. Here, we propose two gradient-based algorithms, one extrinsic and the other intrinsic, to approximate the optimal channel matrix. The extrinsic algorithm is a gradient-ascent algorithm with projection onto the constrained space as in Bertsekas [1982]. The intrinsic algorithm, a gradient-ascent algorithm on a manifold, exploits the geometry that codes for the total power constraint by vectorizing the channel matrix. The optimization on manifold has been widely studied in literature, e.g., Absil et al. [2008], Edelman et al. [1998], Smith [1994], and Gabay [1982].

The rest of the chapter is organized as follows. We formulate the channelization problem in Section 3.2 and point out the challenges for design in a two-channel system. In Section 3.3, we give an analytical solution when the conditional covariance of $\boldsymbol{\phi}$ given $\boldsymbol{\theta}$ is the identity matrix. In Section 3.4, we propose two numerical algorithms, extrinsic and intrinsic gradient searches, to approximate the optimal channel matrix for general cases. A simulation study is presented to illustrate the performance of the proposed algorithms in Section 3.4.3. In Section 3.5, we discuss the choice of number of measurements for the secondary channel. Section 3.6 concludes the chapter.

3.2 Overview

3.2.1 Problem Statement

The two channels of the system described in (3.1.1) have input signals $\boldsymbol{\theta} \in \mathbb{R}^p$ and $\boldsymbol{\phi} \in \mathbb{R}^q$, respectively. The signal $\boldsymbol{\theta}$ is of key interest and $\boldsymbol{\phi}$ is a secondary signal that is jointly distributed with $\boldsymbol{\theta}$. The first channel $\boldsymbol{x} \in \mathbb{R}^s$ is a direct measurement of $\boldsymbol{\theta}$, while the secondary channel $\boldsymbol{y} \in \mathbb{R}^t$ is an indirect measurement of $\boldsymbol{\theta}$ through $\boldsymbol{\phi}$. Both $\boldsymbol{x}$ and $\boldsymbol{y}$ contain information about $\boldsymbol{\theta}$, and one can expect that fusing measurements from both channels will provide a better estimate than using a single measurement. The data fusion problem has been widely studied in various areas including sensor networks, image processing, etc. While much of the literature focuses on the methodology of fusion or data integration, we are interested in designing the measurement system. More specifically, our interest is to design the channel matrix $\boldsymbol{G}$ (or the precoding matrix with an identity channel), with the first channel fixed, such that the differential information rate at which $\boldsymbol{x}$ and $\boldsymbol{y}$ bring information about $\boldsymbol{\theta}$ is maximized.

We make the following assumptions:

a1) The signals $\boldsymbol{\theta} \in \mathbb{R}^p$ and $\boldsymbol{\phi} \in \mathbb{R}^q$ are jointly Gaussian distributed as

$$\left(\begin{array}{c} \boldsymbol{\theta} \\ \boldsymbol{\phi} \end{array} \right) \sim N \left(\left(\begin{array}{c} \boldsymbol{\mu_\theta} \\ \boldsymbol{\mu_\phi} \end{array} \right), \left(\begin{array}{cc} \boldsymbol{Q_{\theta\theta}} & \boldsymbol{Q_{\theta\phi}} \\ \boldsymbol{Q_{\phi\theta}} & \boldsymbol{Q_{\phi\phi}} \end{array} \right) \right)$$

with known $\boldsymbol{Q_{\theta\theta}}$, $\boldsymbol{Q_{\theta\phi}}$, $\boldsymbol{Q_{\phi\theta}}$ and $\boldsymbol{Q_{\phi\phi}}$.

a2) The noises $\boldsymbol{u} \in \mathbb{R}^s$ and $\boldsymbol{v} \in \mathbb{R}^t$ are Gaussian distributed with mean zero and known covariance matrices $\boldsymbol{Q_{uu}}$ and $\boldsymbol{Q_{vv}}$, respectively.

a3) The noises $\boldsymbol{u}$ and $\boldsymbol{v}$ are mutually independent, and independent of $(\boldsymbol{\theta}, \boldsymbol{\phi})$.

Based on these assumptions, the mutual information between $\boldsymbol{\theta}$ and $\boldsymbol{x}$ is

$$I(\boldsymbol{\theta}; \boldsymbol{x}) = \frac{1}{2} \log \det(\boldsymbol{Q_{\theta\theta}}) - \frac{1}{2} \log \det(\boldsymbol{Q_{\theta\theta|x}}),$$

where $\boldsymbol{Q_{\theta\theta|x}} = (\boldsymbol{Q_{\theta\theta}^{-1}} + \boldsymbol{F^T Q_{uu}^{-1} F})^{-1}$ is the conditional covariance of $\boldsymbol{\theta}$ given $\boldsymbol{x}$. The mutual information between $\boldsymbol{\theta}$ and $\boldsymbol{x}, \boldsymbol{y}$ is

$$I(\boldsymbol{\theta}; \boldsymbol{x}, \boldsymbol{y}) = \frac{1}{2} \log \det(\boldsymbol{Q_{\theta\theta}}) - \frac{1}{2} \log \det(\boldsymbol{Q_{\theta\theta|x,y}}),$$

where $\boldsymbol{Q_{\theta\theta|x,y}} = [\boldsymbol{Q_{\theta\theta|x}^{-1}} + \boldsymbol{M^T G^T}(\boldsymbol{G Q_{\phi\phi|\theta} G^T} + \boldsymbol{Q_{vv}})^{-1}\boldsymbol{G M}]^{-1}$ with $\boldsymbol{M} = \boldsymbol{Q_{\phi\theta} Q_{\theta\theta}^{-1}}$ and $\boldsymbol{Q_{\phi\phi|\theta}} = \boldsymbol{Q_{\phi\phi}} - \boldsymbol{Q_{\phi\theta} Q_{\theta\theta}^{-1} Q_{\theta\phi}}$. Note that $\boldsymbol{M\theta}$ would be the MMSE estimator of $\boldsymbol{\phi}$ from $\boldsymbol{\theta}$, and $\boldsymbol{Q_{\phi\phi|\theta}}$ would be its error covariance, if $\boldsymbol{\theta}$ could be measured.

The *information gain* is the extra information about $\boldsymbol{\theta}$ brought by $\boldsymbol{y}$, which is defined by

$$D(\boldsymbol{G}) := I(\boldsymbol{\theta}; \boldsymbol{x}, \boldsymbol{y}) - I(\boldsymbol{\theta}; \boldsymbol{x}) = \frac{1}{2} \log \det \boldsymbol{Q_{\theta\theta|x}} - \frac{1}{2} \log \det \boldsymbol{Q_{\theta\theta|x,y}}.$$

By plugging in $\boldsymbol{Q_{\theta\theta|x}}$ and $\boldsymbol{Q_{\theta\theta|x,y}}$, $D(\boldsymbol{G})$ can be written as

$$D(\boldsymbol{G}) = \frac{1}{2} \log \det[\boldsymbol{I_p} + \boldsymbol{M^T G^T}(\boldsymbol{G Q_{\phi\phi|\theta} G^T} + \boldsymbol{Q_{vv}})^{-1}\boldsymbol{G M Q_{\theta\theta|x}}] \qquad (3.2.1)$$

The function $D(\boldsymbol{G})$ is bounded and nonnegative. In fact, one can show that $D(\boldsymbol{G}) \leq I(\boldsymbol{\theta}; \boldsymbol{x}, \boldsymbol{\phi}) - I(\boldsymbol{\theta}; \boldsymbol{x})$, which means the maximum information gain the measurement $\boldsymbol{y}$ can bring is no greater than what could be brought by $\boldsymbol{\phi}$. We further notice that, for any $\boldsymbol{G}$, $D(\lambda\boldsymbol{G})$ is monotone increasing for $\lambda \geq 0$. Therefore, without any constraint, maximization of the information gain in (3.2.1) will lead to a trivial solution by letting the norm of $\boldsymbol{G}$ go to infinity. Here we maximize the information gain subject to the total power constraint $\text{tr}(\boldsymbol{G G^T}) \leq c$. This constraint bounds the total power of $\boldsymbol{G\phi}$ since $\text{tr} \, \mathbb{E}[\boldsymbol{G \phi \phi^T G^T}] \leq c \, \text{tr} \, \mathbb{E}[\boldsymbol{\phi \phi^T}]$. In short, the problem of interest is

$$\boldsymbol{G^*} = \underset{\boldsymbol{G} \in \mathbb{R}^{t \times q}}{\arg \max} \; D(\boldsymbol{G}) \text{ subject to } \text{tr}(\boldsymbol{G G^T}) \leq c. \qquad (3.2.2)$$

Problem (3.2.2) is a one-channel design problem in a two-channel system. In general, the optimization problem cannot be reformulated as an SVD problem in contrast to a one-channel system. The difficulty arises due to the non-degenerate joint distribution of $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. However, when the conditional covariance matrix $\boldsymbol{Q_{\phi\phi|\theta}}$ is zero, i.e., the value of $\boldsymbol{\phi}$ is fixed given $\boldsymbol{\theta}$, the optimal channel matrix $\boldsymbol{G}$ can be solved from an SVD problem, as in a one-channel system.

3.2.2 An Insightful Discussion of the Information Gain

To motivate our discussion, we decompose the secondary channel as follows:

$$\boldsymbol{y} = (\boldsymbol{GM}\mathbb{E}[\boldsymbol{\theta}|\boldsymbol{x}]) + (\boldsymbol{GM}(\boldsymbol{\theta} - \mathbb{E}[\boldsymbol{\theta}|\boldsymbol{x}])) + (\boldsymbol{G}(\boldsymbol{\phi} - \mathbb{E}[\boldsymbol{\phi}|\boldsymbol{\theta}]) + \boldsymbol{v}) , \qquad (3.2.3)$$

where $\boldsymbol{M} = \boldsymbol{Q}_{\phi\theta}\boldsymbol{Q}_{\theta\theta}^{-1}$ and $\boldsymbol{M\theta} = \mathbb{E}[\boldsymbol{\phi}|\boldsymbol{\theta}]$. It can be seen that the secondary channel $\boldsymbol{y}$ is decomposed into three independent components, which are illustrated in Figure 3.2. The first component $\boldsymbol{GM}\mathbb{E}[\boldsymbol{\theta}|\boldsymbol{x}]$ is completely determined by the first channel $\boldsymbol{x}$ and does not contribute to the information gain brought by $\boldsymbol{y}$. The second component $\boldsymbol{GM}(\boldsymbol{\theta} - \mathbb{E}[\boldsymbol{\theta}|\boldsymbol{x}])$, denoted by $\boldsymbol{\omega}$, is (by orthogonality) independent of $\boldsymbol{x}$ and it carries the extra information in channel $\boldsymbol{y}$ about $\boldsymbol{\theta}$. The third component $\boldsymbol{G}(\boldsymbol{\phi} - \mathbb{E}[\boldsymbol{\phi}|\boldsymbol{\theta}]) + \boldsymbol{v}$, denoted by $\boldsymbol{\zeta}$, is independent of both $\boldsymbol{x}$ and $\boldsymbol{\theta}$, and it can be viewed as noise.
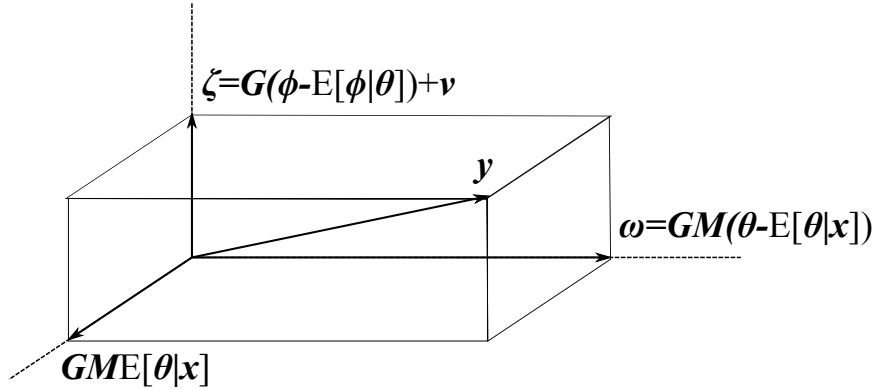
Figure 3.2: Decomposition of the secondary channel.

Notice that the covariance matrices of $\boldsymbol{\omega}$ and $\boldsymbol{\zeta}$ are $\boldsymbol{Q}_{\omega\omega} = \boldsymbol{GMQ}_{\theta\theta|\boldsymbol{x}}\boldsymbol{M}^T\boldsymbol{G}^T$ and $\boldsymbol{Q}_{\zeta\zeta} = \boldsymbol{GQ}_{\phi\phi|\theta}\boldsymbol{G}^T + \boldsymbol{Q}_{vv}$, respectively. By the cyclic property of determinants, i.e., $\det(\boldsymbol{I}_m + \boldsymbol{AB}) = \det(\boldsymbol{I}_n + \boldsymbol{BA})$ for any $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $\boldsymbol{B} \in \mathbb{R}^{n \times m}$, the information gain of (3.2.1) can be re-written as

$$D(\boldsymbol{G}) = \frac{1}{2} \log \det[\boldsymbol{I} + \boldsymbol{Q}_{\zeta\zeta}^{-1/2}\boldsymbol{Q}_{\omega\omega}\boldsymbol{Q}_{\zeta\zeta}^{-1/2}]. \qquad (3.2.4)$$

By viewing $\boldsymbol{\omega}$ as a signal and $\boldsymbol{\zeta}$ as a noise, $\boldsymbol{Q}_{\zeta\zeta}^{-1/2}\boldsymbol{Q}_{\omega\omega}\boldsymbol{Q}_{\zeta\zeta}^{-1/2}$ is a generalized signal-to-noise ratio matrix. Maximizing (3.2.4) essentially balances the tradeoff between the noise

covariance and the signal covariance. As illustrated in Figure 3.2, a good channel design will favor a long parallelepiped with short height. The difficulty of designing the channel matrix $\boldsymbol{G}$ arises because $\boldsymbol{G}$ shapes both $\boldsymbol{Q_{\omega\omega}}$ and $\boldsymbol{Q_{\zeta\zeta}}$. When the secondary channel has a single output, the channel matrix $\boldsymbol{G}$ is a row vector and $\boldsymbol{Q_{\zeta\zeta}}, \boldsymbol{Q_{\omega\omega}}$ are scalars. In this case, the optimal $\boldsymbol{G}$ equivalently maximizes a generalized Rayleigh quotient and the analytical solution can be derived by an eigendecomposition. For a general channel with multiple outputs, designing the matrix $\boldsymbol{G}$ is fundamentally more difficult than the single output case. In Section 3.3, we obtain a closed-form expression for the optimal $\boldsymbol{G}$ in a special case.

3.3 Analytical Solution

In this section, we consider the problem of maximizing the information gain, subject to the total power constraint. The problem can be described as follows:

$$
\begin{aligned}
\underset{\boldsymbol{G}\in\mathbb{R}^{t\times q}}{\text{maximize}} \quad & D(\boldsymbol{G}) = \tfrac{1}{2}\log\det[\boldsymbol{I}_p + \boldsymbol{M}^T\boldsymbol{G}^T(\boldsymbol{G}\boldsymbol{Q}_{\phi\phi|\boldsymbol{\theta}}\boldsymbol{G}^T + \boldsymbol{Q_{vv}})^{-1}\boldsymbol{G}\boldsymbol{M}\boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}}] \\
\text{subject to} \quad & \text{tr}(\boldsymbol{G}\boldsymbol{Q}_{\phi\phi}\boldsymbol{G}^T) \le c.
\end{aligned} \tag{3.3.1}
$$

Note that problem (3.3.1) is not a convex problem since the information gain $D(\boldsymbol{G})$ is not concave in $\boldsymbol{G}$ (see Boyd and Vandenberghe [2004] and Payaró and Palomar [2009]). In Section 3.3.1, we give an analytical solution when the second channel has a single output, i.e., $t = 1$. For the multiple-output channel ($t > 1$), we give an analytical expression for the optimal channel matrix in Section 3.3.2, when the secondary channel has white noise and the conditional covariance matrix $\boldsymbol{Q}_{\phi\phi|\boldsymbol{\theta}}$ is proportional to the identity matrix.

3.3.1 The Case of MISO Channel

Suppose that the second channel has a single output. Then $\boldsymbol{Q_{vv}} = \sigma_v^2 \in \mathbb{R}_+^1$. The channel matrix $\boldsymbol{G}$ is a row vector, and we denote $\boldsymbol{G} = \boldsymbol{g}^T$ for some vector $\boldsymbol{g} \in \mathbb{R}^q$. The

information gain is

$$D(\boldsymbol{g}) = \log\left(1 + \frac{\boldsymbol{g}^T \boldsymbol{M} \boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}} \boldsymbol{M}^T \boldsymbol{g}}{\boldsymbol{g}^T \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}} \boldsymbol{g} + \sigma_v^2}\right). \tag{3.3.2}$$

Notice that $\log(1 + x)$ is strictly increasing for $x \in (0, \infty)$. Therefore, the optimization problem is equivalent to solve the following problem

$$\underset{\boldsymbol{g} \in \mathbb{R}^q}{\text{maximize}} \quad \frac{\boldsymbol{g}^T \boldsymbol{M} \boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}} \boldsymbol{M}^T \boldsymbol{g}}{\boldsymbol{g}^T \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}} \boldsymbol{g} + \sigma_v^2} \tag{3.3.3}$$
$$\text{subject to } \|\boldsymbol{g}\| \le 1.$$

Note that the objective function in (3.3.3) is exactly $\sigma_\omega^2/\sigma_\nu^2$, with ω and ν the signal and noise defined in Section 3.2.2. This means a good channel matrix $\boldsymbol{G} = \boldsymbol{g}^T$ will maximize the ratio of signal power σ_ω^2 to noise power σ_ν^2. It is easy to see that the maximum is attained when $\|\boldsymbol{g}\| = 1$. Therefore we have $\boldsymbol{g}^T \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}} \boldsymbol{g} + \sigma_v^2 = \boldsymbol{g}^T (\boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}} + \sigma_v^2 \boldsymbol{I}_q) \boldsymbol{g}$. Let $\widetilde{\boldsymbol{g}} = (\sigma_v^2 \boldsymbol{I}_q + \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}})^{1/2} \boldsymbol{g}$, where $(\sigma_v^2 \boldsymbol{I}_q + \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}})^{1/2}$ is a matrix square root of $\sigma_v^2 \boldsymbol{I}_q + \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}}$. The objective function in (3.3.3) is reduced to a Rayleigh quotient

$$\frac{\tilde{\boldsymbol{g}}^T \boldsymbol{A} \tilde{\boldsymbol{g}}}{\tilde{\boldsymbol{g}}^T \tilde{\boldsymbol{g}}} \tag{3.3.4}$$

with $\boldsymbol{A} = (\sigma_v^2 \boldsymbol{I}_q + \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}})^{-1/2} \boldsymbol{M} \boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}} \boldsymbol{M}^T (\sigma_v^2 \boldsymbol{I}_q + \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}})^{-1/2}$, which is the signal-to-noise-ratio matrix. The maximum of (3.3.4) is $\lambda_{\max}(\boldsymbol{A})$, the largest eigenvalue of $\boldsymbol{A}$. This maximum is attained when $\widetilde{\boldsymbol{g}}^* = \alpha \boldsymbol{u}_{\max}(\boldsymbol{A})$ where $\boldsymbol{u}_{\max}(\boldsymbol{A})$ is the eigenvector of $\boldsymbol{A}$ corresponding to its largest eigenvalue, and α is a scalar such that $(\sigma_v^2 \boldsymbol{I}_q + \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}})^{-1/2} \widetilde{\boldsymbol{g}}^*$ has norm 1. The optimal channel vector $\boldsymbol{g}^*$ is $\boldsymbol{g}^* = (\sigma_v^2 \boldsymbol{I}_q + \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}})^{-1/2} \widetilde{\boldsymbol{g}}^*$, and the maximum information gain is determined by the maximum eigenvalue of the signal-to-noise ratio matrix:

$$D(\boldsymbol{G}^*) = \log\left(1 + \lambda_{\max}(\boldsymbol{A})\right).$$

3.3.2 An Important Special Case of a MIMO Auxiliary Channel

Suppose that the conditional covariance of $\boldsymbol{\phi}$ given $\boldsymbol{\theta}$ is identity, i.e., $\boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}} = \sigma_{\boldsymbol{\phi}|\boldsymbol{\theta}}^2 \boldsymbol{I}_q$. For example, $\boldsymbol{\phi} = \boldsymbol{M}\boldsymbol{\theta} + \boldsymbol{\tau}$ where $\boldsymbol{M} = \boldsymbol{Q}_{\boldsymbol{\phi\theta}} \boldsymbol{Q}_{\boldsymbol{\theta\theta}}^{-1}$ and $\boldsymbol{\tau} \sim N(\boldsymbol{0}, \sigma_{\boldsymbol{\phi}|\boldsymbol{\theta}}^2 \boldsymbol{I}_q)$. In this case, the

noise $\boldsymbol{\zeta}$ in (3.2.4) has a relatively simple covariance $\boldsymbol{Q}_{\boldsymbol{\zeta\zeta}} = \boldsymbol{G}\boldsymbol{G}^T + \boldsymbol{Q}_{\boldsymbol{vv}}$ and the signal $\boldsymbol{\omega}$ has covariance $\boldsymbol{Q}_{\boldsymbol{\omega\omega}} = \boldsymbol{G}\boldsymbol{M}\boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}}\boldsymbol{M}^T\boldsymbol{G}^T$. While $\boldsymbol{G}$ still affects both covariance matrices, we are able to find the balanced matrix $\boldsymbol{G}$ that maximizes the information gain. Note that we focus on the case $t \leq q$, i.e., the dimension of measurement $\boldsymbol{y}$ is at most the dimension of $\boldsymbol{\phi}$. When $t > q$, the optimization problem can be reformulated and solved as a special case of $t = q$, which will be discussed in Section 3.5.

Begin with the eigendecompositions $\boldsymbol{Q}_{\boldsymbol{vv}} = \boldsymbol{U}_{\boldsymbol{v}}\boldsymbol{\Sigma}_{\boldsymbol{v}}\boldsymbol{U}_{\boldsymbol{v}}$ and $\boldsymbol{M}\boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}}\boldsymbol{M}^T = \boldsymbol{U}_{\boldsymbol{\xi}}\boldsymbol{\Sigma}_{\boldsymbol{\xi}}\boldsymbol{U}_{\boldsymbol{\xi}}^T$ where $\boldsymbol{U}_{\boldsymbol{v}} \in \mathbb{R}^{t \times t}$, $\boldsymbol{U}_{\boldsymbol{\xi}} \in \mathbb{R}^{q \times q}$ are orthogonal matrices, and $\boldsymbol{\Sigma}_{\boldsymbol{v}} \in \mathbb{R}^{t \times t}$ and $\boldsymbol{\Sigma}_{\boldsymbol{\xi}} \in \mathbb{R}^{q \times q}$ are diagonal matrices with diagonal elements $0 < \sigma_{v,1}^2 \leq \ldots \leq \sigma_{v,m}^2$ and $\sigma_{\boldsymbol{\xi},1}^2 \geq \ldots \geq \sigma_{\boldsymbol{\xi},q}^2 \geq 0$, respectively. Because the matrices $\boldsymbol{U}_{\boldsymbol{\xi}}$ and $\boldsymbol{U}_{\boldsymbol{v}}$ are invertible, for each $\boldsymbol{G} \in \mathbb{R}^{t \times q}$, there is a unique matrix $\boldsymbol{\Phi} \in \mathbb{R}^{t \times q}$ such that

$$\boldsymbol{G} = \boldsymbol{U}_{\boldsymbol{v}}\boldsymbol{\Phi}\boldsymbol{U}_{\boldsymbol{\xi}}^T \tag{3.3.5}$$

Then, the information gain $D(\boldsymbol{G})$ in (3.2.1) can be written as

$$D(\boldsymbol{\Phi}) = \frac{1}{2}\log\det[\boldsymbol{I} + \boldsymbol{\Phi}^T(\sigma_{\boldsymbol{\phi}|\boldsymbol{\theta}}^2\boldsymbol{\Phi}\boldsymbol{\Phi}^T + \boldsymbol{\Sigma}_{\boldsymbol{v}})^{-1}\boldsymbol{\Phi}\boldsymbol{\Sigma}_{\boldsymbol{\xi}}]$$

Moreover, the total power constraint is $\mathrm{tr}(\boldsymbol{\Phi}\boldsymbol{\Phi}^T) \leq c$ since $\mathrm{tr}(\boldsymbol{G}\boldsymbol{G}^T) = \mathrm{tr}(\boldsymbol{\Phi}\boldsymbol{\Phi}^T)$. For the given eigendecompositions, the matrices $\boldsymbol{U}_{\boldsymbol{v}}$ and $\boldsymbol{U}_{\boldsymbol{\xi}}$ are fixed. Therefore, the information gain can be maximized with respect to $\boldsymbol{\Phi}$ and the optimal channel matrix $\boldsymbol{G}$ is returned by (3.3.5). WOLG we assume $\sigma_{\boldsymbol{\phi}|\boldsymbol{\theta}}^2 = 1$. The solution for general $\sigma_{\boldsymbol{\phi}|\boldsymbol{\theta}}^2$ is just different by a scaling factor. We give in Lemma 3.1 an important feature of any possible maximizer $\boldsymbol{\Phi}$.

Lemma 3.1. *Suppose that $\boldsymbol{Q}_{\boldsymbol{vv}}$ has distinct eigenvalues, i.e., $0 < \sigma_{v,1}^2 < \ldots < \sigma_{v,m}^2$, and $\boldsymbol{M}\boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}}\boldsymbol{M}^T$ has distinct nonzero eigenvalues, i.e., $\sigma_{\boldsymbol{\xi},1}^2 > \ldots > \sigma_{\boldsymbol{\xi},\rho}^2 > 0$ where $\rho \leq t$ is the rank of $\boldsymbol{\Sigma}_{\boldsymbol{\xi}}$. Then $\boldsymbol{\Phi}$ contains at most one nonzero entry in each row and column and all the nonzero entries are located at the first ρ columns.*

Proof: See Section 3.7.1.

Lemma 3.1 restricts the optimal matrix $\boldsymbol{\Phi}$ within a class of matrices with a special structure. That is, $\boldsymbol{\Phi}$ has at most one nonzero entry in each row and column. Searching within this class, we are able to obtain the closed form expression for the optimal matrix $\boldsymbol{\Phi}$. The corresponding optimal channel matrix $\boldsymbol{G}$ is given in Theorem 3.1.

Theorem 3.1. *Suppose that $\boldsymbol{Q_{vv}}$ and $\boldsymbol{MQ_{\theta\theta|x}M}^T$ have distinct nonzero eigenvalues. Then the optimal secondary channel matrix $\boldsymbol{G}^*$ solving problem (3.2.2) is*

$$\boldsymbol{G}^* = \boldsymbol{U_v}\boldsymbol{\Lambda}^*\boldsymbol{U_\xi}^T. \tag{3.3.6}$$

Here $\boldsymbol{\Lambda}^* \in \mathbb{R}^{t \times q}$ is a diagonal matrix with diagonal elements $\lambda_{11}^*, \ldots, \lambda_{tt}^*$ such that

$$\lambda_{ii}^{*2} = \begin{cases} \frac{\sigma_{v,i}^2}{2(1+\sigma_{\xi,i}^2)}\left(-(2+\sigma_{\xi,i}^2) + \sqrt{\sigma_{\xi,i}^4 + \frac{4(1+\sigma_{\xi,i}^2)\sigma_{\xi,i}^2}{2\mu\sigma_{v,i}^2}}\right) & i = 1, \ldots, \kappa \\ 0 & i = \kappa+1, \ldots, t \end{cases} \tag{3.3.7}$$

where κ is the maximum integer between 1 and $\mathrm{rank}(\boldsymbol{\Sigma_\xi})$ such that $\lambda_{ii}^{*2} > 0$ for $i = 1, \ldots, \kappa$. The value of μ is non-negative and uniquely solves $\sum_{i=1}^{\kappa} \lambda_{ii}^{*2} = c$.

Proof: See Section 3.7.2.

Notice that although Theorem 3.1 requires that $\boldsymbol{Q_{vv}}$ and $\boldsymbol{MQ_{\theta\theta|x}M}^T$ have distinct eigenvalues, the result can be extended to general cases because the solution in (3.3.7) is a continuous function of the eigenvalues of $\boldsymbol{Q_{vv}}$ and $\boldsymbol{MQ_{\theta\theta|x}M}^T$.

Theorem 3.1 factors the optimal channel matrix $\boldsymbol{G}^*$ into the product of three matrices. The first matrix $\boldsymbol{U_\xi}^T$ rotates the signal $\boldsymbol{\phi}$. Given $\boldsymbol{Q_{\phi\phi|\theta}} = \sigma_{\phi|\theta}^2\boldsymbol{I}_q$, the conditional covariance of $\boldsymbol{\phi}$ given $\boldsymbol{x}$ is $\boldsymbol{MQ_{\theta\theta|x}M}^T + \sigma_{\phi|\theta}^2\boldsymbol{I}_q$, which is diagonalized by $\boldsymbol{U_\xi}^T$. Therefore, the components of the rotated signal $\boldsymbol{U_\xi}^T\boldsymbol{\phi}$ are conditionally independent given $\boldsymbol{x}$. The second matrix $\boldsymbol{\Lambda}^* \in \mathbb{R}^{t \times q}$ is a diagonal matrix that extracts the first t components of $\boldsymbol{U_\xi}^T\boldsymbol{\phi}$ and distributes power across the t subchannels optimally. The third matrix $\boldsymbol{U_v}$ rotates the scaled components into the sub-dominant invariant subspace of the noise covariance $\boldsymbol{Q_{vv}}$.

46

The power allocation policy, given by the diagonal elements of $\mathbf{\Lambda}^*$, can be interpreted as a *mercury/waterfilling* algorithm, which is a three-step procedure that has been introduced in Lozano et al. [2006]:

1. For the ith vessel, fill in the solid base with height $2\sigma_{\boldsymbol{v},i}^2/\sigma_{\boldsymbol{\xi},i}^2$, where $\sigma_{\boldsymbol{v},i}^2$ is a noise variance component in the $\boldsymbol{y}$ channel, and $\sigma_{\boldsymbol{\xi},i}^2$ yields a variance component of $\boldsymbol{\phi}$ given $\boldsymbol{x}$.

2. Compute μ. For the vessels with base height less than $1/\mu$, fill in mercury in the vessel until the height reaches

$$
\max\left\{\frac{1}{\mu} - \lambda_{ii}^{*2}, \frac{2\sigma_{\boldsymbol{v},i}^2}{\sigma_{\boldsymbol{\xi},i}^2}\right\}
$$

.

3. Pour water into all the vessels until the height of each vessel reaches $1/\mu$.

The height of the solid base, $2\sigma_{\boldsymbol{v},i}^2/\sigma_{\boldsymbol{\xi},i}^2$, is half of the variance of the ith noise component, weighted by the variance components of $\boldsymbol{M}\boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}}\boldsymbol{M}^T$. A higher solid base means a less informative channel with high channel noise and weak correlation with $\boldsymbol{\theta}$. For any vessel with base height exceeding $1/\mu$, neither mercury nor water will be added, or equivalently, no power will be assigned to the corresponding subchannel. Note that the value of μ is computed by the constraint that the total volume of water equals c. The mercury stage balances the noise contained in $\boldsymbol{\phi}$ and the measurement noise contained in $\boldsymbol{y}$. Without adding mercury, the optimal power allocation would have variable water-plus-solid levels among different vessels. The mercury is added to regulate the water level for each vessel. Given the value of μ, the information gain is maximized when the value of λ_{ii}^{*2} equals the height of water in the corresponding vessel.

From the mercury/waterfilling procedure, it can be seen that the resulting optimal channel matrix $\boldsymbol{G}^*$ may not be full-rank. We will see in Section 3.5 that a rank-reduced
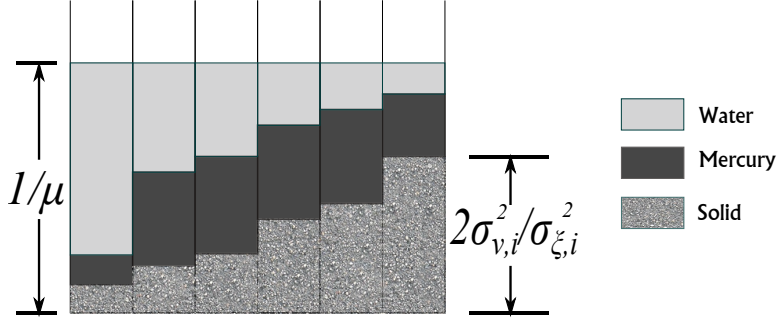
47

Figure 3.3: Mercury/waterfilling. For each vessel, the water height above mercury gives the optimal power allocation for the corresponding subchannel. The total volume of water equals c.

channel matrix can in some cases give a dimension-reduced secondary channel that carries the same information gain as a full-dimensional channel, under the power constraint. To better illustrate the possible rank-reduced optimal channel, we consider the following simple example.

Example 3.1. *Consider a two-channel system in* (3.1.1) *with* $p = q = s = t = 5$. *The primary channel matrix* $\boldsymbol{F} \in \mathbb{R}^{5 \times 5}$ *is set to* $\frac{1}{\sqrt{5}} \boldsymbol{I}_5$. *The covariance matrices* $\boldsymbol{Q_{uu}}$, $\boldsymbol{Q_{vv}}$, *and* $\boldsymbol{Q_{\theta\theta}}$ *are* $\boldsymbol{I}_5$. *We consider three scenarios. In each scenario, we choose* $\boldsymbol{Q_{\phi\phi}}$ *and* $\boldsymbol{Q_{\phi\theta}}$ *such that* $\boldsymbol{Q_{\phi\phi|\theta}} = \boldsymbol{I}_5$ *and the eigenvalues of* $\boldsymbol{M Q_{\theta\theta|x} M}^T$ *have various levels of spread. The corresponding* $\boldsymbol{G}^*$ *is given in Table I.*

In the first scenario, $\boldsymbol{M Q_{\theta\theta|x} M}^T$ *has constant eigenvalues and* $\boldsymbol{G}^*$ *has full-rank and equal singular values. In the second scenario, the eigenvalues of* $\boldsymbol{M Q_{\theta\theta|x} M}^T$ *have moderate spread and the corresponding* $\boldsymbol{G}^*$ *has rank* 4. *In the third scenario, when the spread of eigenvalues of* $\boldsymbol{M Q_{\theta\theta|x} M}^T$ *further increases, the rank of* $\boldsymbol{G}^*$ *is further reduced to* 3.

In the compression design of Chapter 2, we generalized the problem of reduced-rank filtering and precoding/equalizing by designing the matrix $\boldsymbol{G}$ in the bottom channel of Figure 3.4 so that $\boldsymbol{y}$ maximizes the differential rate at which $\boldsymbol{y}$ brings information about $\boldsymbol{\theta}$. The difference between the compression design in Chapter 2 and the fusion design

Table 3.1: The optimal channel matrices $\boldsymbol{G}^*$ for three scenarios.

	(1)	(2)	(3)
$\boldsymbol{Q}_{\phi\phi}$	$2\boldsymbol{I}_5$	$\mathrm{Diag}(26, 17, 10, 5, 2)$	$\mathrm{Diag}(82, 65, 50, 5, 2)$
$\boldsymbol{Q}_{\phi\theta}$	$\boldsymbol{I}_5$	$\mathrm{Diag}(5, 4, 3, 2, 1)$	$\mathrm{Diag}(9, 8, 7, 2, 1)$
$\boldsymbol{M}\boldsymbol{Q}_{\theta\theta|\boldsymbol{x}}\boldsymbol{M}^T$	$\frac{5}{6}\boldsymbol{I}_5$	$\frac{5}{6}\mathrm{Diag}(25, 16, 9, 4, 1)$	$\frac{5}{6}\mathrm{Diag}(81, 64, 49, 4, 1)$
$\boldsymbol{G}^*$	$\frac{1}{\sqrt{5}}\boldsymbol{I}_5$	$\mathrm{Diag}(0.32, 0.30, 0.25, 0.14, 0)$	$\mathrm{Diag}(0.34, 0.33, 0.33, 0, 0)$

here is that there was no existing channel $\boldsymbol{x}$ to be fused with $\boldsymbol{y}$. For compression, the compression $\boldsymbol{G}$ of Figure 3.4 is designed to maximize the differential information rate at which $\boldsymbol{y}$ brings information about $\boldsymbol{\theta}$. For fusion, the compressor $\boldsymbol{G}$ is designed to maximize the differential rate at which $\boldsymbol{y}$ and $\boldsymbol{x}$ bring information about $\boldsymbol{\theta}$.
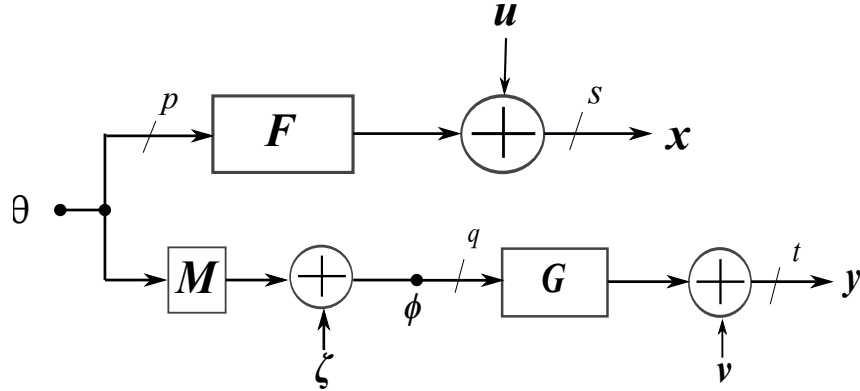
Figure 3.4: An alternative representation for the two-channel system.

3.4   Numerical Algorithms

In general, the constrained optimization problem (3.2.2) is not a convex problem (Boyd and Vandenberghe [2004]) since the information gain $D(\boldsymbol{G})$ is not concave (Payaró and Palomar [2009]). Notice that for any $\boldsymbol{G}$ with $\mathrm{tr}(\boldsymbol{G}\boldsymbol{G}^T) < c$, there exists $\widetilde{\boldsymbol{G}} = \frac{\sqrt{c}}{\|\boldsymbol{G}\|}\boldsymbol{G}$ such that $\mathrm{tr}(\widetilde{\boldsymbol{G}}\widetilde{\boldsymbol{G}}^T) = c$ and $D(\widetilde{\boldsymbol{G}}) \geq D(\boldsymbol{G})$. Therefore, it is sufficient to maximize the information gain on the boundary $\mathrm{tr}(\boldsymbol{G}\boldsymbol{G}^T) = c$. This fact motivates two gradient-based search algorithms, one extrinsic and the other intrinsic, to approximate the optimal

channel matrix. Both algorithms are very general and applicable for arbitrary covariance between $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. In the extrinsic gradient search, the gradient is computed by treating the matrix $\boldsymbol{G}$ as a point in the Euclidean space $\mathbb{R}^{t \times q}$. In the intrinsic gradient search, we consider $\boldsymbol{G}$ to be a point on the unit sphere S^{tq-1}, which is a submanifold of $\mathbb{R}^{tq}$. The intrinsic gradient is computed by taking the geometry of the manifold S^{tq-1} into consideration. WLOG we assume $c = 1$.

3.4.1 Extrinsic Gradient Search Algorithm

Let $\nabla_{\boldsymbol{G}} D$ be the gradient of the information gain w.r.t $\boldsymbol{G}$. We show, in Section 3.7.3, that the gradient can be written as

$$\nabla_{\boldsymbol{G}} D = \boldsymbol{Q}_{vv}^{-1} \boldsymbol{G} [(\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T \boldsymbol{Q}_{vv}^{-1} \boldsymbol{G}) - \boldsymbol{B}]^{-1} \boldsymbol{B} (\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T \boldsymbol{Q}_{vv}^{-1} \boldsymbol{G})^{-1}, \qquad (3.4.1)$$

where $\boldsymbol{B} = \boldsymbol{Q}_{\phi\phi|\theta}^{-1} \boldsymbol{M} \boldsymbol{Q}_{\theta\theta|x} \boldsymbol{M}^T (\boldsymbol{I}_q + \boldsymbol{Q}_{\phi\phi|\theta}^{-1} \boldsymbol{M} \boldsymbol{Q}_{\theta\theta|x} \boldsymbol{M}^T)^{-1} \boldsymbol{Q}_{\phi\phi|\theta}^{-1}$. The gradient $\nabla_{\boldsymbol{G}} D$ points in the direction of greatest increase of the function D in the neighborhood of $\boldsymbol{G}$. However, when moving along this direction, the constraint $\text{tr}(\boldsymbol{G}\boldsymbol{G}^T) = 1$ may be violated. To circumvent this problem, we normalize the updated $\boldsymbol{G}$ at each iteration to meet the unit norm constraint. The table below outlines the proposed extrinsic gradient search algorithm.

Algorithm: Extrinsic Gradient Search

Input: Initial $\boldsymbol{G}_0 \in \mathbb{R}^{t \times q}, \text{tr}(\boldsymbol{G}_0 \boldsymbol{G}_0^T) = 1$.
Output: Sequence of iterates $\{\boldsymbol{G}_k\}$.
for $k = 0, 1, 2, \ldots$ **do**
Select $\boldsymbol{G}_{k+1} = a_k (\boldsymbol{G}_k + \delta_k \nabla_{\boldsymbol{G}_k} D)$ where $a_k = \frac{1}{\|\boldsymbol{G}_k + \delta_k \nabla_{\boldsymbol{G}_k} D\|}$ is a
normalization constant such that $\text{tr}(\boldsymbol{G}_{k+1} \boldsymbol{G}_{k+1}^T) = 1$, δ_k is a small step size.
end for

In this extrinsic algorithm, the gradient of the information gain is computed on the unconstrained Euclidean space $\mathbb{R}^{t \times q}$. Note that $\boldsymbol{G}_k + \delta_k \nabla_{\boldsymbol{G}_k} D$ is the unconstrained update when maximizing D. The normalized update $\boldsymbol{G}_{k+1} = a_k (\boldsymbol{G}_k + \delta_k \nabla_{\boldsymbol{G}_k} D)$ is a projection of $\boldsymbol{G}_k + \delta_k \nabla_{\boldsymbol{G}_k} D$ onto the set of all $\boldsymbol{G} \in \mathbb{R}^{t \times q}$ with unit Frobenius norm. We call it an extrinsic gradient search in contrast to the intrinsic gradient search algorithm, in which the information gain is considered as a function on the manifold S^{tq-1}.

3.4.2   Intrinsic Gradient Search Algorithm

Let $\boldsymbol{g}$ be the vectorization of matrix $\boldsymbol{G}$, denoted $\boldsymbol{g} = \text{vec}(\boldsymbol{G})$. That is,

$$\boldsymbol{g} = [\boldsymbol{G}_{1,1}, \ldots, \boldsymbol{G}_{1,q}, \boldsymbol{G}_{2,1}, \ldots, \boldsymbol{G}_{2,q}, \ldots, \boldsymbol{G}_{t,1}, \ldots, \boldsymbol{G}_{t,q}]^T.$$

This vectorization operation is a one-to-one and onto mapping from $\mathbb{R}^{t \times q}$ to $\mathbb{R}^{tq}$. Thus, for any $\boldsymbol{g} \in \mathbb{R}^{tq}$, there exists a unique matrix $\boldsymbol{G} \in \mathbb{R}^{t \times q}$ such that $\text{vec}(\boldsymbol{G}) = \boldsymbol{g}$. Under the power constraint $\text{tr}(\boldsymbol{G}\boldsymbol{G}^T) = 1$, the corresponding vectorization $\boldsymbol{g}$ lies on the unit sphere $S^{tq-1} = \{\boldsymbol{g} \in \mathbb{R}^{tq} : \sum_{i=1}^{tq} g_i^2 = 1\}$. Therefore, the constrained optimization problem (3.2.2) is an optimization on the manifold S^{tq-1}. Note that S^{tq-1} is an embedded submanifold of $\mathbb{R}^{tq}$, a geometry that has been studied in Absil et al. [2008] and Lee [2000].

Before presenting the algorithm, we briefly introduce the basic terms of a manifold. Readers may refer to Absil et al. [2008] and Lee [2000] for more details.

A set $\mathcal{M}$ is a topological manifold with dimension n if it is a second-countable Hausdorff space and every point of $\mathcal{M}$ has a neighborhood that is homeomorphic to an open subset of $\mathbb{R}^n$. For each point $p \in \mathcal{M}$, there is a open set $U \subset \mathcal{M}$ that contains p and a homeomorphism φ such that $\varphi(U) = \tilde{U}$ for some open subset $\tilde{U} \subset \mathbb{R}^n$. (U, φ) is an n-dimensional *local chart of* $\mathcal{M}$. A smooth *atlas* $\mathcal{A}$ of $\mathcal{M}$ into $\mathbb{R}^n$ is a collection of charts $(U_\alpha, \varphi_\alpha)$ of the set $\mathcal{M}$ such that

1. $\bigcup_\alpha U_\alpha = \mathcal{M}$,

2. for any pair α, β with $U_\alpha \bigcap U_\beta \neq \emptyset$, the sets $\varphi_\alpha(U_\alpha \bigcap U_\beta)$ and $\varphi_\beta(U_\alpha \bigcap U_\beta)$ are open sets in $\mathbb{R}^n$ and the transition $\varphi_\beta \circ \varphi_\alpha^{-1} : \mathbb{R}^n \to \mathbb{R}^n$ is smooth.

An atlas $\mathcal{A}$ is *maximal* if it is not contained in any strictly larger atlas. A smooth maximal atlas of a set $\mathcal{M}$ is called a *smooth manifold structure* on $\mathcal{M}$ and the pair $(\mathcal{M}, A)$ is a smooth manifold with dimension n. Often, we omit mention of the manifold structure $\mathcal{A}$ and simply say "the manifold $\mathcal{M}$". When $\mathcal{M}$ is a smooth manifold, a function

$f : \mathcal{M} \mapsto \mathbb{R}^k$ is *smooth* if the composite function $f \circ \varphi^{-1} : \tilde{U} \mapsto \mathbb{R}^k$ has continuous partial derivatives of all orders. A smooth mapping $\gamma : \mathbb{R} \to \mathcal{M} : t \mapsto \gamma(t)$ is called a *curve* in $\mathcal{M}$. Let $\mathcal{F}(\mathcal{M})$ be the space of smooth scalar functions on $\mathcal{M}$. The tangent vector ξ_x to a manifold $\mathcal{M}$ at a point x then is a mapping from $\mathcal{F}(\mathcal{M})$ to $\mathbb{R}$ such that there exists a curve γ on $\mathcal{M}$ with $\gamma(0) = x$, satisfying

$$\xi_x(f) := \left. \frac{d(f(\gamma(t)))}{dt} \right|_{t=0}, \quad \forall f \in \mathcal{F}(\mathcal{M}).$$

The tangent vector generalizes the notion of a directional derivative on a manifold and it satisfies, for any $a, b \in \mathbb{R}$ and $f, g \in \mathcal{F}(\mathcal{M})$,

1. $\xi_x(af + bg) = a\xi_x(g) + b\xi_x(g)$, and

2. $\xi_x(fg) = \xi_x(g)f + \xi_x(f)g$.

The tangent space to $\mathcal{M}$ at x, denoted $T_x\mathcal{M}$, is the set of all tangent vectors to $\mathcal{M}$ at x. A smooth manifold whose tangent space $T_x\mathcal{M}$ is endowed with an inner product $\langle \cdot, \cdot \rangle_x$ is a Riemannian manifold. Given a smooth scalar function f on the Riemannian manifold $\mathcal{M}$, the *gradient* of f at x, denoted by $\mathrm{grad}(f)_x$, is defined as the unique element of $T_x\mathcal{M}$ that satisfies

$$\langle \mathrm{grad}(f)_x, \xi_x \rangle_x = \xi_x(f) \quad \forall \xi_x \in T_x\mathcal{M}.$$

The direction of $\mathrm{grad}(f)_x$ is the steepest-ascent direction of f at x, which points in the direction of search when maximizing f over $\mathcal{M}$.

Given $\boldsymbol{g} \in S^{tq-1}$, the tangent space to S^{tq-1} at $\boldsymbol{g}$ is

$$T_{\boldsymbol{g}} S^{tq-1} = \{ \boldsymbol{\eta} \in \mathbb{R}^{tq} : \boldsymbol{g}^T \boldsymbol{\eta} = 0 \}.$$

The orthogonal projection of any $\boldsymbol{h} \in \mathbb{R}^{tq}$ onto the tangent space is

$$P_{T_{\boldsymbol{g}} S^{tq-1}} \boldsymbol{h} = (\boldsymbol{I}_{tq} - \boldsymbol{g}\boldsymbol{g}^T)\boldsymbol{h}.$$

Since S^{tq-1} is a submanifold embedded in $\mathbb{R}^{tq}$, the gradient of I_S on S^{tq-1} is

$$\boldsymbol{\eta_g} \triangleq \operatorname{grad} D_{\boldsymbol{g}} = P_{T_{\boldsymbol{g}}S^{tq-1}}(\nabla_{\boldsymbol{g}}D)$$

where $\nabla_{\boldsymbol{g}}D$ is given in (3.4.1). The intrinsic gradient $\boldsymbol{\eta_g}$ points us in the direction of search, and a retraction $R_{\boldsymbol{g}}$ generalizes the notion of moving on a manifold along that direction. The line search on the manifold is then

$$\boldsymbol{g}_{k+1} = R_{\boldsymbol{g}_k}(t\boldsymbol{\eta_{g_k}}).$$

On the manifold S^{tq-1}, one choice of the retraction is

$$R_{\boldsymbol{g}}(\boldsymbol{\eta}) = \boldsymbol{g}\cos(\|\boldsymbol{\eta}\|) + \frac{\boldsymbol{\eta}}{\|\boldsymbol{\eta}\|}\sin(\|\boldsymbol{\eta}\|), \tag{3.4.2}$$

where $\|\boldsymbol{\eta}\|^2 = \sum_{i=1}^{tq} \eta_i^2$. For any tangent vector $\boldsymbol{\eta} \in T_{\boldsymbol{g}}S^{tq-1}$, $R_{\boldsymbol{g}}(\boldsymbol{\eta})$ returns a point on the manifold S^{tq-1}. Notice the normalization in the extrinsic algorithm is also a choice of retraction on S^{tq-1}.

The following algorithm encodes the intrinsic gradient search, which approximates a maximizer of the information gain on the manifold S^{tq-1}. A graphical illustration is depicted in Figure 3.5.

Algorithm: Intrinsic Gradient Search

Input: Initial $\boldsymbol{g}_0 \in S^{tq-1}$
Output: Sequence of iterates $\{\boldsymbol{g}_k\}$.
for $k = 0, 1, 2, \dots$ **do**
Select $\boldsymbol{g}_{k+1} = R_{\boldsymbol{g}_k}(\delta_k\boldsymbol{\eta_{g_k}})$ where $\boldsymbol{\eta_{g_k}} = (\boldsymbol{I}_{tq} - \boldsymbol{g}_k\boldsymbol{g}_k^T)\nabla_{\boldsymbol{g}_k}D$ is the intrinsic gradient and δ_k is the step size.
end for

For any $\boldsymbol{g} \in S^{tq-1}$, the tangent plane $T_{\boldsymbol{g}}S^{tq-1}$ is the subspace orthogonal to $\boldsymbol{g}$. The intrinsic gradient, denoted by $\boldsymbol{\eta_g}$, is the Euclidean gradient $\nabla_{\boldsymbol{g}}D$ projected onto the tangent plane $T_{\boldsymbol{g}}S^{tq-1}$. The function $R_{\boldsymbol{g}}$ is a mapping from the tangent plane $T_{\boldsymbol{g}}S^{tq-1}$ to the manifold S^{tq-1} with

$$R_{\boldsymbol{g}}(\boldsymbol{\eta_g}) = \boldsymbol{g}\cos(\|\boldsymbol{\eta_g}\|) + \frac{\boldsymbol{\eta_g}}{\|\boldsymbol{\eta_g}\|}\sin(\|\boldsymbol{\eta_g}\|) \tag{3.4.3}$$

for any tangent vector $\boldsymbol{\eta_g} \in T_{\boldsymbol{g}} S^{tq-1}$. For $\delta \geq 0$, $R_{\boldsymbol{g}}(\delta \boldsymbol{\eta_g})$ is a curve on the manifold S^{tq-1} starting from $\boldsymbol{g}$. This curve generalizes the idea of straight line in Euclidean space on the manifold S^{tq-1} along the direction $\boldsymbol{\eta_g}$. Given $\boldsymbol{g}_k$, $R_{\boldsymbol{g}}(\delta \boldsymbol{\eta_{g_k}})$ is a periodic function of τ with period $2\pi / \|\boldsymbol{\eta_{g_k}}\|$, thus the step size δ_k can be chosen within the interval $\delta \in [0, 2\pi / \|\boldsymbol{\eta_{g_k}}\|)$ to maximize the information gain $D(R_{\boldsymbol{g}}(\delta \boldsymbol{\eta_{g_k}}))$. By the choice of δ_k, the information gain is non-decreasing, i.e., $D(\boldsymbol{g}_{k+1}) \geq D(\boldsymbol{g}_k)$ for each k.
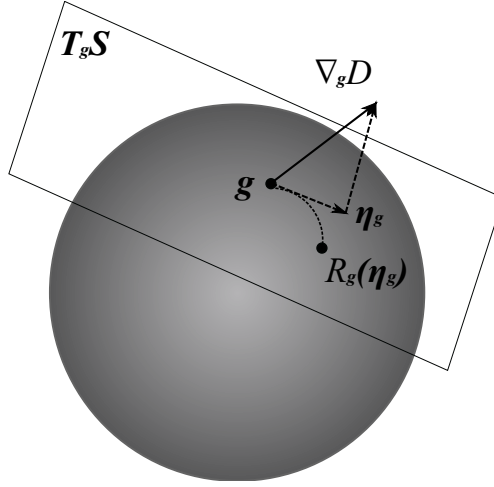
Figure 3.5: Projection of the Euclidean gradient to the tangent plane of the unit sphere.

3.4.3   A Numerical Study

Consider a two-channel system in (3.1.1) with $p = q = 4$ and $s = t = 3$. The input signals $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are characterized recursively as

$$\phi_i = \sum_{j=1}^{i} \rho^{i-j+1} \theta_j + \tau_i,$$

where $\tau_1, \ldots, \tau_4$ are i.i.d. Gaussian random variables with mean 0 and variance 1, and the value of ρ is to be specified. The covariance matrices for the signal $\boldsymbol{\theta}$ and the noises $\boldsymbol{u}, \boldsymbol{v}$ are proportional to the identity matrix with variances 2, 1, 0.1, respectively. The first channel matrix $\boldsymbol{F} \in \mathbb{R}^{3 \times 4}$ is a diagonal matrix with 1 on the diagonal. The initial channel matrix $\boldsymbol{G}_0 \in \mathbb{R}^{3 \times 4}$ is randomly generated with unit norm. For the intrinsic algorithm, the initial value is $\boldsymbol{g}_0 = \text{vec}(\boldsymbol{G}_0)$.

The results are shown in Figure 3.6. Here we set the step size $\delta_k = 0.1$. The x-axis is the index for iterations and the y-axis gives the information gain for the secondary channel returned at step k. First, it can be seen that, as ρ increases, the information gain is increasing as well because the correlation between $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ is increasing. Next, it can also be seen that the performance of the two algorithms is quite comparable and both algorithms converge for each value of ρ. From our empirical evidence, when the step size is constant, both algorithms would perform similarly, and in fact, the extrinsic algorithm converges slightly faster. For more complex problems, we could choose the optimal step size over a finite interval as suggested by the intrinsic algorithm in Section 3.4.2. For extrinsic algorithm, such a strategy for the optimal step size is not available.
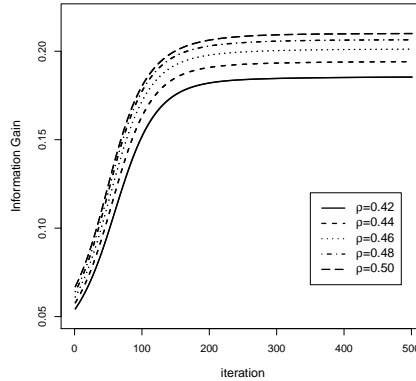
Figure 3.6: A numerical study. The x-axis is the index for iteration and the y-axis is the information gain obtained at each iteration. The solid curve is for the intrinsic algorithm and the dashed curve is for the extrinsic algorithm.

3.5   Discussion on Low-dimensional Channel Design

In the two-channel design problem considered in this chapter, the number of measurements of the secondary channel, i.e., the number of rows of the channel matrix $\boldsymbol{G}$ is an important factor. Ideally we want t to be as small as possible while keeping the information gain as large as possible. More measurements will generally bring more information. However, under the total power constraint, the information a channel carries

is bounded and the upper bound may be attained by a small number of measurements. In fact, for a secondary channel with a q-dimensional input $\boldsymbol{\phi}$, a q-dimensional output $\boldsymbol{y}$ is sufficient to achieve the maximum information gain, which is a consequence of the following lemma. Here we assume that the measurement noise $\boldsymbol{v}$ in the secondary channel is white noise.

Lemma 3.2. *Suppose that the noise covariance $\boldsymbol{Q_{vv}}$ is proportional to the identity matrix. Then, for any channel matrix $\boldsymbol{G} \in \mathbb{R}^{t \times q}$ with rank r, there exists an r-dimensional secondary channel with the same noise variance that achieves the same information gain.*

The proof is given in Section 3.7.4.

Since the maximum rank of $\boldsymbol{G}$ is q, Lemma 3.2 suggests that a q-dimensional $\boldsymbol{y}$ is sufficient to achieve the maximum information gain. Thus, we restrict our attention to the channel matrix with dimension $t \times q$ with $t \leq q$. In some cases, the power constraint will further reduce the dimension of $\boldsymbol{y}$ to $t < q$. For instance, as shown in Example 1, the 5×5 optimal channel matrices can have rank 5, 4, or 3, and the dimension of $\boldsymbol{y}$ may be reduced correspondingly. Denote $\boldsymbol{G}_k^*$ the optimal channel matrix of dimension $k \times q$ for $k = 1, \ldots, q$. The optimal dimension of $\boldsymbol{y}$, denoted by t^*, is defined as the smallest k such that $D(\boldsymbol{G}_q^*) = D(\boldsymbol{G}_k^*)$; that is, $t^* = \min\{k : D(\boldsymbol{G}_q^*) - D(\boldsymbol{G}_k^*) = 0\}$. Note that $D(\boldsymbol{G}_k^*) = D(\boldsymbol{G}_q^*)$ for any $k \geq t^*$, and $D(\boldsymbol{G}_k^*) < D(\boldsymbol{G}_q^*)$ for any $k < t^*$. In general, the values of t^* is unknown since no analytical solution for $\boldsymbol{G}_k^*$ is available. From a practical viewpoint, it is natural to approximate t^* using the approximate optimal channel matrices. Here we consider the following approach to obtain an approximant of t^*.

For $k = 1, \ldots, q$, obtain an approximate optimal channel matrix of dimension $k \times q$, denoted by $\widehat{\boldsymbol{G}}_k^*$, using either the extrinsic or intrinsic algorithm. Denote $\widehat{t}^* = \min\{k : D(\widehat{\boldsymbol{G}}_q^*) - D(\widehat{\boldsymbol{G}}_k^*) \leq c\}$, where c is a predetermined threshold value, and $\widehat{t}^*$ is the proposed dimension of $\boldsymbol{y}$. The following example demonstrates this suggested strategy with more details.

Example 3.2. *Consider a two-channel system* (3.1.1) *with* $p = q = 20$ *and* $s = 10$. *The channel matrix* $\boldsymbol{F} \in \mathbb{R}^{10 \times 20}$ *is randomly generated with Frobenius norm* 1. *The noise covariances* $\boldsymbol{Q_{uu}} = \boldsymbol{Q_{vv}} = \boldsymbol{I}_{20}$. *The covariances* $\boldsymbol{Q_{\theta\theta}}$ *and* $\boldsymbol{Q_{\phi\phi}}$ *are randomly generated positive definite matrices. We consider two different correlation structures between* $\boldsymbol{\theta}$ *and* $\boldsymbol{\phi}$: 1) $\boldsymbol{Q_{\phi\phi|\theta}} = \boldsymbol{I}_{20}$ *(analytical solution available);* 2) $\boldsymbol{Q_{\phi\phi|\theta}}$ *is a banded matrix with* 2 *on the main diagonal line and* 0.2 *on the superdiagonal and subdiagonal lines (analytical solution not available). The results are shown in Figure 3.7 and Figure 3.8, where the x-axis is* k $(k = 1, \ldots, q)$, *the number of rows of the secondary channel matrix* $\boldsymbol{G}$, *the y-axis on the left is the information gain for a k-dimensional secondary channel, and the y-axis on the right is the rank of the channel matrix* $\boldsymbol{G}$ *of dimension* $k \times q$.

In the first scenario (Figure 3.7), we obtain $\boldsymbol{G}_k^*$ *for* $k = 1, \ldots, q$ *analytically (shown in the left panel). It can be seen that the information gain remains constant for all* $k \geq 4$. *Therefore the optimal dimension is* $t^* = 4$. *Moreover, one can see that the rank of all the optimal channel matrices* $\boldsymbol{G}_k^*$ *with* $k \geq 4$ *equal* 4, *which may suggest that the optimal dimension* t^* *equals the maximum rank of the optimal channel matrices. Therefore the curve for the rank of the optimal matrices can be used as an important guidance. The extrinsic (the middle panel) and intrinsic (the right panel) algorithms are implemented, with the initial channel matrices randomly generated. Here we set the constant step size* $\delta_k = 0.1$. *For both algorithms we get* $\widehat{t}^* = t^* = 4$ *for* $c = 10^{-3}$, *and so is the maximum rank.*

In the second scenario (Figure 3.8), we implement the extrinsic and intrinsic algorithms to approximate the optimal channel matrix. Note that the solutions for $k = 4$ *and* $k = 5$ *have similar information gain but different ranks. If the threshold value* $c = 10^{-3}$, *we have* $\widehat{t}^* = 4$ *in both algorithms, while the maximum rank equals* 5. *The difference may be caused by approximation error in the numerical algorithms.*

Figure 3.7: Choice of number of rows of the secondary channel matrix G. The three panels are associated with the channel matrices returned by the analytical solution (top), the extrinsic algorithm (middle) and the intrinsic algorithm (bottom), respectively. In each panel, the x-axis indicates the number of rows of G, the y-axis on the left is the information gain (solid line), and the y-axis on the right is the rank of the channel matrices (star dotted line).
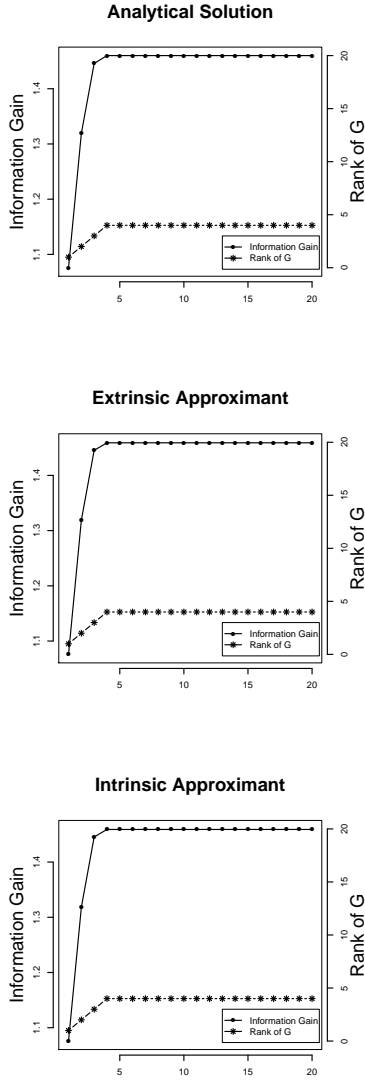
Figure 3.8: Choice of number of rows of the secondary channel matrix $\boldsymbol{G}$. The two panels are associated with the channel matrices returned by the extrinsic algorithm (left) and the intrinsic algorithm (right), respectively. In each panel, the x-axis indicates the number of rows of $\boldsymbol{G}$, the y-axis on the left is the information gain (solid line), and the y-axis on the right is the rank of the channel matrices (star dotted line).
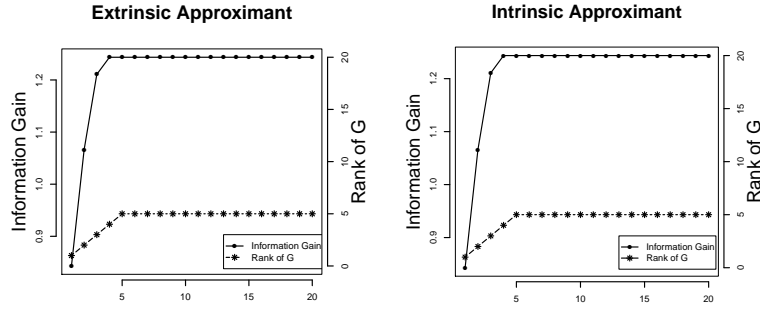
3.6   Summary

In this chapter, we have studied the problem of fusing multiple sources of information. We have modeled the problem as a two-channel system where the signal in the primary channel is of interest, and the signal in the secondary channel is jointly distributed with the signal of interest. The objective is to design the secondary channel to maximize the information gain brought by fusing measurements from the primary and secondary channels. Based on the Gaussian distribution and linear channel assumptions, we obtain a closed-form expression of the information gain. When the input signals have a special covariance structure, we obtain an explicit solution for the optimal channel matrix, where the singular vectors are allocated to create non-interfering subchannels and the singular values solve a generalized water-filling problem. For general cases, we propose two gradient search algorithms, an extrinsic algorithm and an intrinsic algorithm to approximate the optimal channel matrix. Both algorithms can be extended to optimize other design criteria under a power constraint. With the designed secondary channel matrix, combining the measurements of both channels achieves the best information gain. Note that, without the Gaussian assumption, our results maximize the volume of

59

the error concentration ellipsoid of the LMMSE.

3.7 Proofs

3.7.1 Proof of Lemma 3.1

By the matrix inversion lemma $\boldsymbol{\Phi}^T(\boldsymbol{\Phi}\boldsymbol{\Phi}^T + \boldsymbol{\Sigma_v})^{-1}\boldsymbol{\Phi} = \boldsymbol{I} - (\boldsymbol{I} + \boldsymbol{\Phi}\boldsymbol{\Sigma}_v^{-1}\boldsymbol{\Phi}^T)^{-1}$. So $D(\boldsymbol{\Phi})$ may be rewritten as

$$
\begin{aligned}
D(\boldsymbol{\Phi}) &= \frac{1}{2}\log\det[\boldsymbol{I} + \boldsymbol{\Sigma_\xi} - (\boldsymbol{I} + \boldsymbol{\Phi}\boldsymbol{\Sigma}_v^{-1}\boldsymbol{\Phi}^T)^{-1}\boldsymbol{\Sigma_\xi}] \\
&= \frac{1}{2}\log\det[\boldsymbol{I} + \boldsymbol{\Sigma_\xi}] + \frac{1}{2}\log\det[\boldsymbol{I} - (\boldsymbol{I} + \boldsymbol{\Phi}\boldsymbol{\Sigma}_v^{-1}\boldsymbol{\Phi}^T)^{-1}\boldsymbol{\Sigma_\xi}(\boldsymbol{I} + \boldsymbol{\Sigma_\xi})^{-1}]
\end{aligned}
$$

Define $\boldsymbol{\Lambda} := \boldsymbol{\Sigma}_v^{-1}$, a $t \times t$ diagonal matrix with diagonal elements $\lambda_i = \sigma_{v,i}^{-2}$, and $\boldsymbol{\Gamma} := \boldsymbol{\Sigma_\xi}(\boldsymbol{I} + \boldsymbol{\Sigma_\xi})^{-1}$, a $q \times q$ diagonal matrix with diagonal elements $\gamma_i = \sigma_{\xi,i}^2/(1 + \sigma_{\xi,i}^2)$. Let ρ be the rank of $\boldsymbol{\Sigma_\xi}$. Then $1 > \gamma_1 \geq \ldots \geq \gamma_\rho > 0$ and $\gamma_{\rho+1} = \ldots = \gamma_q = 0$. The Lagrangian is

$$
L(\boldsymbol{\Phi}; \mu) = \frac{1}{2}\log\det(\boldsymbol{I}_n - (\boldsymbol{I}_n + \boldsymbol{\Phi}^T\boldsymbol{\Lambda}\boldsymbol{\Phi})^{-1}\boldsymbol{\Gamma}) + \mu(\text{tr}(\boldsymbol{\Phi}\boldsymbol{\Phi}^T) - c) + \frac{1}{2}\log\det[\boldsymbol{I} + \boldsymbol{\Sigma_\xi}]
$$

(3.7.1)

where μ is the Lagrangian multiplier. The partial derivative of $L(\boldsymbol{\Phi}; \mu)$ with respect to the elements of $\boldsymbol{\Phi}$ is

$$
\nabla_{\boldsymbol{\Phi}} L(\boldsymbol{\Phi}; \mu) = \boldsymbol{\Lambda}\boldsymbol{\Phi}(\boldsymbol{I}_n - \boldsymbol{\Gamma} + \boldsymbol{\Phi}^T\boldsymbol{\Lambda}\boldsymbol{\Phi})^{-1} - \boldsymbol{\Lambda}\boldsymbol{\Phi}(\boldsymbol{I}_n + \boldsymbol{\Phi}^T\boldsymbol{\Lambda}\boldsymbol{\Phi})^{-1} + 2\mu\boldsymbol{\Phi}.
$$

Left multiply the gradient by $\boldsymbol{\Lambda}^{-1}$ and right multiply by $\boldsymbol{\Phi}^T$:

$$
-\boldsymbol{\Phi}(\boldsymbol{I}_n - \boldsymbol{\Gamma} + \boldsymbol{\Phi}^T\boldsymbol{\Lambda}\boldsymbol{\Phi})^{-1}\boldsymbol{\Phi}^T + \boldsymbol{\Phi}(\boldsymbol{I}_n + \boldsymbol{\Phi}^T\boldsymbol{\Lambda}\boldsymbol{\Phi})^{-1}\boldsymbol{\Phi}^T = 2\mu\boldsymbol{\Lambda}^{-1}\boldsymbol{\Phi}\boldsymbol{\Phi}^T
$$

Since the LHS is symmetric, $\boldsymbol{\Phi}\boldsymbol{\Phi}^T\boldsymbol{\Lambda}^{-1} = \boldsymbol{\Lambda}^{-1}\boldsymbol{\Phi}\boldsymbol{\Phi}^T$. Therefore, when $\boldsymbol{\Lambda}$ has distinct diagonal elements, the symmetric matrix $\boldsymbol{\Phi}\boldsymbol{\Phi}^T$ must be diagonal. Next we show that $\boldsymbol{\Phi}(\boldsymbol{I}_n - \boldsymbol{\Gamma})^{-1}\boldsymbol{\Phi}^T$ is diagonal. Notice that

$$
\boldsymbol{\Phi}(\boldsymbol{I}_n - \boldsymbol{\Gamma} + \boldsymbol{\Phi}^T\boldsymbol{\Lambda}\boldsymbol{\Phi})^{-1}\boldsymbol{\Phi}^T = \boldsymbol{\Lambda}^{-1} - \boldsymbol{\Lambda}^{-1}(\boldsymbol{\Lambda}^{-1} + \boldsymbol{\Phi}(\boldsymbol{I}_n - \boldsymbol{\Gamma})^{-1}\boldsymbol{\Phi}^T)\boldsymbol{\Lambda}^{-1}
$$

$$
\boldsymbol{\Phi}(\boldsymbol{I}_n + \boldsymbol{\Phi}^T\boldsymbol{\Lambda}\boldsymbol{\Phi})^{-1}\boldsymbol{\Phi}^T = \boldsymbol{\Lambda}^{-1} - \boldsymbol{\Lambda}^{-1}(\boldsymbol{\Lambda}^{-1} + \boldsymbol{\Phi}\boldsymbol{\Phi}^T)\boldsymbol{\Lambda}^{-1}
$$

Then, right multiply the gradient by $\boldsymbol{\Phi}^T$:

$$\boldsymbol{\Lambda}^{-1}\boldsymbol{\Phi}(\boldsymbol{I}_n - \boldsymbol{\Gamma})^{-1}\boldsymbol{\Phi}^T = \boldsymbol{\Lambda}^{-1}\boldsymbol{\Phi}\boldsymbol{\Phi}^T - \mu\boldsymbol{\Phi}\boldsymbol{\Phi}^T$$

The RHS is symmetric since $\boldsymbol{\Phi}\boldsymbol{\Phi}^T$ is diagonal. Therefore we have $\boldsymbol{\Lambda}^{-1}\boldsymbol{\Phi}(\boldsymbol{I}_n - \boldsymbol{\Gamma})^{-1}\boldsymbol{\Phi}^T = \boldsymbol{\Phi}(\boldsymbol{I}_n - \boldsymbol{\Gamma})^{-1}\boldsymbol{\Phi}^T\boldsymbol{\Lambda}^{-1}$, which implies that $\boldsymbol{\Phi}(\boldsymbol{I}_n - \boldsymbol{\Gamma})^{-1}\boldsymbol{\Phi}^T$ is diagonal.

Denote $\boldsymbol{\Phi} := [\phi_{ij}]$. Given the fact that $\boldsymbol{\Phi}(\boldsymbol{I}_n - \boldsymbol{\Gamma})^{-1}\boldsymbol{\Phi}^T$ and $\boldsymbol{\Phi}\boldsymbol{\Phi}^T$ are diagonal, (3.7.1) can be rewritten as

$$\begin{aligned}
L(\boldsymbol{\Phi}; \mu) &= \frac{1}{2}\log\det(\boldsymbol{I}_m + \boldsymbol{\Lambda}\boldsymbol{\Phi}(\boldsymbol{I}_n - \boldsymbol{\Gamma})^{-1}\boldsymbol{\Phi}^T) - \frac{1}{2}\log\det(\boldsymbol{I}_m + \boldsymbol{\Lambda}\boldsymbol{\Phi}\boldsymbol{\Phi}^T) \\
&\quad + \frac{1}{2}\log\det(\boldsymbol{I}_n - \boldsymbol{\Gamma}) + \frac{1}{2}\log\det[\boldsymbol{I} + \boldsymbol{\Sigma}_{\boldsymbol{\xi}}] + \mu(\mathrm{tr}(\boldsymbol{\Phi}\boldsymbol{\Phi}^T) - c) \\
&= \sum_{i=1}^{m}\frac{1}{2}\log(1 + \lambda_i\sum_{j=1}^{n}\frac{\phi_{ij}^2}{1 - \gamma_j}) - \sum_{i=1}^{m}\frac{1}{2}\log(1 + \lambda_i\sum_{j=1}^{n}\phi_{ij}^2) \\
&\quad + \frac{1}{2}\log\det(\boldsymbol{I}_n - \boldsymbol{\Gamma}) + \frac{1}{2}\log\det[\boldsymbol{I} + \boldsymbol{\Sigma}_{\boldsymbol{\xi}}] + \mu(\sum_{i=1}^{m}\sum_{j=1}^{n}\phi_{ij}^2 - c)
\end{aligned}$$

Notice that $L(\boldsymbol{\Phi}; \mu)$ is quadratic in each ϕ_{ij}. Therefore, we can assume WLOG $\phi_{ij} \geq 0$. The partial derivative of $L(\boldsymbol{\Phi}; \mu)$ w.r.t ϕ_{ij} is

$$\frac{\partial L(\boldsymbol{\Phi}; \mu)}{\partial \phi_{ij}} = \phi_{ij}\left[\frac{\lambda_i(1 - \gamma_j)^{-1}}{1 + \sum_{j=1}^{n}\phi_{ij}^2(1 - \gamma_j)^{-1}} - \frac{\lambda_i}{1 + \sum_{j=1}^{n}\phi_{ij}^2} + 2\mu\right]$$

For $j > \rho$, we have $\gamma_j = 0$, and $L(\boldsymbol{\Phi}; \mu)$ is monotone decreasing in ϕ_{ij} since $\mu \leq 0$. Hence for any minimizer $\boldsymbol{\Phi}$, $\phi_{ij} = 0$ for any $j > \rho$.

For the ith row, suppose that there exist two non-zero elements ϕ_{ij_1} and ϕ_{ij_2}. Then the partial derivative $\frac{\partial L(\boldsymbol{\Phi}; \mu)}{\partial \phi_{ij_1}} = \frac{\partial L(\boldsymbol{\Phi}; \mu)}{\partial \phi_{ij_2}} = 0$ yields

$$\frac{\lambda_i(1 - \gamma_{j_1})^{-1}}{1 + \sum_{j=1}^{n}\phi_{ij}^2(1 - \gamma_j)^{-1}} = \frac{\lambda_i(1 - \gamma_{j_2})^{-1}}{1 + \sum_{j=1}^{n}\phi_{ij}^2(1 - \gamma_j)^{-1}}$$

which contradicts the assumption $\gamma_{j_1} \neq \gamma_{j_2}$. For the jth column, if there are two non-zero elements ϕ_{i_1j} and ϕ_{i_2j}, then $\phi_{i_1k} = \phi_{i_2k} = 0$ for any $k \neq j$ since each row of $\boldsymbol{\Phi}$ has at most one non-zero entry. Hence, $[\boldsymbol{\Phi}\boldsymbol{\Phi}^T]_{i_1i_2} = \sum_{k=1}^{n}\phi_{i_1k}\phi_{i_2k} = \phi_{i_1j}\phi_{i_2j} \neq 0$, which contradicts diagonal $\boldsymbol{\Phi}\boldsymbol{\Phi}^T$.

3.7.2 Proof of Theorem 3.1

Restricting the matrix $\mathbf{\Phi}$ within the class of matrices satisfying Lemma 3.1, $\mathbf{\Phi}$ can be written as $\mathbf{\Phi} = \mathbf{\Pi}_2 \mathbf{\Lambda} \mathbf{\Pi}_1^T$ where $\mathbf{\Pi}_1 \in \mathbb{R}^{q \times q}$ and $\mathbf{\Pi}_2 \in \mathbb{R}^{t \times t}$ are permutation matrices and $\mathbf{\Lambda}$ is a $t \times q$ diagonal matrix with diagonal elements $\lambda_{11}, \dots, \lambda_{tt}$. The maximum information gain is taken over the permutations $\mathbf{\Pi}_1$ $\mathbf{\Pi}_2$ and $\lambda_{11}, \dots, \lambda_{tt}$ subject to $\sum_{i=1}^{t} \lambda_{ii}^2 \leq c$.

First of all, we show the optimal permutation matrices are $\mathbf{\Pi}_1 = \mathbf{I}_q$ and $\mathbf{\Pi}_2 = \mathbf{I}_t$. Denote

$$f(\mathbf{\Pi}_1, \mathbf{\Pi}_2) = \max D(\mathbf{\Phi}) \text{ subject to } \mathbf{\Phi} = \mathbf{\Pi}_2 \mathbf{\Lambda} \mathbf{\Pi}_1^T \text{ and } \sum_{i=1}^{t} \lambda_{ii}^2 \leq c.$$

The objective is to show

$$f(\mathbf{I}_q, \mathbf{I}_t) \geq f(\mathbf{\Pi}_1, \mathbf{\Pi}_2)$$

for all the possible permutations $\mathbf{\Pi}_1$ and $\mathbf{\Pi}_2$.

Let $\pi_1(i)$ be the index of the entry equal to unity in the ith column of $\mathbf{\Pi}_1$, and $\pi_2(i)$ the index of the unity entry in the ith column of $\mathbf{\Pi}_2$. Then the information gain $D(\mathbf{\Phi})$ can be written as

$$D(\mathbf{\Phi}|\mathbf{\Pi}_1, \mathbf{\Pi}_2) = \frac{1}{2} \sum_{i=1}^{t} \log \left(1 + \frac{\sigma_{\boldsymbol{\xi}, \pi_1(i)}^2 \lambda_{ii}^2}{\lambda_{ii}^2 + \sigma_{\boldsymbol{v}, \pi_2(i)}^2} \right)$$

It is easy to see that for any $i = 1, \dots, t$, one must have

$$\sigma_{\boldsymbol{\xi}, \pi_1(i)}^2 \geq \max\{\sigma_{\boldsymbol{\xi}, \pi_1(t+1)}^2, \dots, \sigma_{\boldsymbol{\xi}, \pi_1(q)}^2\}.$$

Moreover, since the orders of $\{\pi_1(j)\}_{j=t+1}^{q}$ do not affect the value of $D(\mathbf{\Phi})$, we can set WLOG $\pi_1(j) = j$ for $j = t+1, \dots, q$. For $i = 1, \dots, t$, it can be seen that $\pi_1(i)$ and $\pi_2(i)$ appear pairwise in $D(\mathbf{\Phi}|\mathbf{\Pi}_1, \mathbf{\Pi}_2)$. Therefore, we can set WLOG that $\pi_2(i) = i$ for $i = 1, \dots, t$ and then search for the optimal permutation $\pi_1(i)$ to maximize

$$D(\mathbf{\Phi}|\mathbf{\Pi}_1, \mathbf{I}_t) = \frac{1}{2} \sum_{i=1}^{t} \log \left(1 + \frac{\sigma_{\boldsymbol{\xi}, \pi_1(i)}^2 \lambda_{ii}^2}{\lambda_{ii}^2 + \sigma_{\boldsymbol{v}, i}^2} \right) \tag{3.7.2}$$

The proof that $\mathbf{\Pi}_1 = \boldsymbol{I}_q$ is the optimal permutation matrix is similar to the proof for Theorem 2 in Wang et al. [2013]. First prove the case $t = 2$ and generalize the results to $t \geq 2$. The details are omitted.

Next, the objective is to solve a simpler optimization problem:

$$\{\lambda_{ii}^*\}_{i=1}^t = \arg\max \frac{1}{2} \sum_{i=1}^t \log\left(1 + \frac{\sigma_{\boldsymbol{\xi},i}^2 \lambda_{ii}^2}{\lambda_{ii}^2 + \sigma_{\boldsymbol{v},i}^2}\right) \text{ subject to } \sum_{i=1}^t \lambda_{ii}^2 \leq c. \qquad (3.7.3)$$

The Lagrangian is

$$L(\lambda_{11}, \ldots, \lambda_{tt}, ; \mu) = \frac{1}{2} \sum_{i=1}^t \log\left(1 + \frac{\sigma_{\boldsymbol{\xi},\pi_1(i)}^2 \lambda_{ii}^2}{\lambda_{ii}^2 + \sigma_{\boldsymbol{v},\pi_2(i)}^2}\right) - \mu\left(\sum_{i=1}^t \lambda_{ii}^2 - c\right) \qquad (3.7.4)$$

where $\mu \geq 0$ is the Lagrange multiplier. Setting the first derivative of L w.r.t. λ_{ii} equal to zero, we have either $\lambda_{ii} = 0$ or

$$\lambda_{ii} = \sqrt{\frac{b_i\left(-(2 + a_i) + \sqrt{(2 + a_i)^2 - 4(1 + a_i)(1 - a_i/(2\mu b_i))}\right)}{2(1 + a_i)}} \qquad (3.7.5)$$

where $a_i = \sigma_{\boldsymbol{\xi},i}^2$ and $b_i = \sigma_{\boldsymbol{v},i}^2$. Equation (3.7.5) provides a feasible solution for λ_{ii} when $\mu \leq a_i/(2b_i)$. To see whether the solution is the maximizer for (3.7.3) , we check the Hessian matrix. The second derivative of L w.r.t. λ_{ii} is

$$\frac{\partial^2 L(\boldsymbol{\Phi}; \mu)}{\partial \lambda_{ii}^2} = -2\mu + \frac{a_i b_i}{(\lambda_{ii}^2 + b_i)(\lambda_{ii}^2 + b_i + a_i \lambda_{ii}^2)} - \frac{2a_i b_i \lambda_{ii}^2((\lambda_{ii}^2 + b_i)(2 + a_i) + a_i \lambda_{ii}^2)}{(\lambda_{ii}^2 + b_i)^2(\lambda_{ii}^2 + b_i + a_i \lambda_{ii}^2)^2}$$
$$(3.7.6)$$

For $i = 1 \ldots, \kappa$, upon substituting (3.7.5),

$$\frac{\partial^2 L(\boldsymbol{\Phi}; \mu)}{\partial \lambda_{ii}^2} = -\frac{8\mu^2 \lambda_{ii}^2}{a_i} \sqrt{(2 + a_i)^2 - 4(1 + a_i)(1 - \frac{a_i}{2\mu b_i})}.$$

which is negative when $\mu < a_i/(2b_i)$. For $\lambda_{ii} = 0$,

$$\left.\frac{\partial^2 L(\boldsymbol{\Phi}; \mu)}{\partial \lambda_{ii}^2}\right|_{\lambda_{ii}=0} = -2\mu + \frac{a_i}{b_i},$$

is negative when $\mu > a_i/(2b_i)$. Let κ be the maximum integer such that $\mu < a_i/(2b_i)$ for $i = 1, \ldots, \kappa$ with μ uniquely solves that $\sum_{i=1}^\kappa \lambda_{ii}^2 = c$. Then, the maximizer of (3.7.3) is $\lambda_{11}^*, \ldots, \lambda_{tt}^*$ where

$$\lambda_{ii}^{*2} = \begin{cases} \frac{b_i\left(\sqrt{(2+a_i)^2 - 4(1+a_i)(1-a_i/(2\mu b_i))} - (2+a_i)\right)}{2(1+a_i)}, & \text{for } i = 1, \ldots, \kappa \\ 0 & \text{for } i = \kappa + 1, \ldots, t. \end{cases} \qquad (3.7.7)$$

3.7.3   Proof of Equation 3.4.1

Applying the matrix inversion lemma yields

$$\boldsymbol{G}^T(\boldsymbol{G}\boldsymbol{Q}_{\phi\phi|\theta}\boldsymbol{G}^T + \boldsymbol{Q}_{vv})^{-1}\boldsymbol{G} = \boldsymbol{Q}_{\phi\phi|\theta}^{-1} - \boldsymbol{Q}_{\phi\phi|\theta}^{-1}(\boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G} + \boldsymbol{Q}_{\phi\phi|\theta}^{-1})^{-1}\boldsymbol{Q}_{\phi\phi|\theta}^{-1}.$$

Therefore, the information gain is

$$
\begin{aligned}
D(\boldsymbol{G}) &= \frac{1}{2}\log\det[\boldsymbol{I}_q + (\boldsymbol{Q}_{\phi\phi|\theta}^{-1} - \boldsymbol{Q}_{\phi\phi|\theta}^{-1}(\boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G} + \boldsymbol{Q}_{\phi\phi|\theta}^{-1})^{-1} \\
&\quad \times \boldsymbol{Q}_{\phi\phi|\theta}^{-1})\boldsymbol{M}\boldsymbol{Q}_{\theta\theta|x}\boldsymbol{M}^T] \\
&= \frac{1}{2}\log\det[\boldsymbol{I}_q + \boldsymbol{Q}_{\phi\phi|\theta}^{-1}\boldsymbol{M}\boldsymbol{Q}_{\theta\theta|x}\boldsymbol{M}^T] + \frac{1}{2}\log\det[\boldsymbol{I} - \boldsymbol{B}(\boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1} \\
&\quad \times \boldsymbol{G} + \boldsymbol{Q}_{\phi\phi|\theta}^{-1})^{-1}]
\end{aligned}
$$

where $\boldsymbol{B} = \boldsymbol{Q}_{\phi\phi|\theta}^{-1}\boldsymbol{M}\boldsymbol{Q}_{\theta\theta|x}\boldsymbol{M}^T(\boldsymbol{I}_q + \boldsymbol{Q}_{\phi\phi|\theta}^{-1}\boldsymbol{M}\boldsymbol{Q}_{\theta\theta|x}\boldsymbol{M}^T)^{-1}\boldsymbol{Q}_{\phi\phi|\theta}^{-1}$.

Let $\boldsymbol{J}_{i,j}$ be a $t \times q$ matrix with value 1 at element (i,j) and 0 elsewhere. From Petersen and Pedersen [2008], for a matrix $\boldsymbol{X}$, we have the partial derivatives

$$\partial\boldsymbol{X}^{-1} = -\boldsymbol{X}^{-1}(\partial\boldsymbol{X})\boldsymbol{X}^{-1}, \quad \partial\log\det\boldsymbol{X} = \mathrm{tr}(\boldsymbol{X}^{-1}\partial\boldsymbol{X}).$$

Let $\boldsymbol{C} = (\boldsymbol{I}_q - (\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G})^{-1}\boldsymbol{B})$. Then $D(\boldsymbol{G}) = \frac{1}{2}\log\det\boldsymbol{C}$ and we have

$$
\begin{aligned}
\frac{\partial D}{\partial \boldsymbol{G}_{i,j}} &= \frac{1}{2}\mathrm{tr}\{\boldsymbol{C}^{-1}\frac{\partial\boldsymbol{C}}{\partial\boldsymbol{G}_{ij}}\} \\
&= -\frac{1}{2}\mathrm{tr}\{\boldsymbol{C}^{-1}\frac{\partial(\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G})^{-1}}{\partial\boldsymbol{G}_{ij}}\boldsymbol{B}\} \\
&= \frac{1}{2}\mathrm{tr}\{\boldsymbol{C}^{-1}(\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G})^{-1}\frac{\partial(\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G})}{\partial\boldsymbol{G}_{i,j}} \\
&\quad \times (\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G})^{-1}\boldsymbol{B}\} \\
&= \frac{1}{2}\mathrm{tr}\{\boldsymbol{C}^{-1}(\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G})^{-1}(\boldsymbol{J}_{i,j}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{J}_{i,j}) \\
&\quad \times (\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G})^{-1}\boldsymbol{B}\} \\
&= \left\{(\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G})^{-1}\boldsymbol{B}\boldsymbol{C}^{-1}(\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G})^{-1}\boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\right\}_{j,i}
\end{aligned}
$$

where the last equality follows from $\mathrm{tr}(\boldsymbol{A}\boldsymbol{J}_{ij}) = \boldsymbol{A}_{j,i} = \mathrm{tr}(\boldsymbol{J}_{ij}^T\boldsymbol{A}^T)$. Hence, the gradient of function D with respect to $\boldsymbol{G}$ is

$$\nabla_{\boldsymbol{G}}D = \boldsymbol{Q}_{vv}^{-1}\boldsymbol{G}[(\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G}) - \boldsymbol{B}]^{-1}\boldsymbol{B}(\boldsymbol{Q}_{\phi\phi|\theta}^{-1} + \boldsymbol{G}^T\boldsymbol{Q}_{vv}^{-1}\boldsymbol{G})^{-1}.$$

3.7.4 Proof of Lemma 3.2

Suppose that the noise $\boldsymbol{v}$ has covariance $\boldsymbol{Q_{vv}} = \sigma_{\boldsymbol{v}}^2 \boldsymbol{I}_t$. Then for any $t \times t$ orthogonal matrix $\boldsymbol{U}$,

$$
\begin{aligned}
D(\boldsymbol{UG}) &= \frac{1}{2} \log \det[\boldsymbol{I}_p + \boldsymbol{M}^T \boldsymbol{G}^T \boldsymbol{U}^T (\boldsymbol{UG} \boldsymbol{Q_{\phi\phi|\theta}} \boldsymbol{G}^T \boldsymbol{U}^T + \sigma_{\boldsymbol{v}}^2 \boldsymbol{I}_t)^{-1} \boldsymbol{UG} \boldsymbol{M} \boldsymbol{Q_{\theta\theta|x}}] \\
&= \frac{1}{2} \log \det[\boldsymbol{I}_p + \boldsymbol{M}^T \boldsymbol{G}^T (\boldsymbol{G} \boldsymbol{Q_{\phi\phi|\theta}} \boldsymbol{G}^T + \sigma_{\boldsymbol{v}}^2 \boldsymbol{I}_t)^{-1} \boldsymbol{G} \boldsymbol{M} \boldsymbol{Q_{\theta\theta|x}}] \\
&= D(\boldsymbol{G})
\end{aligned}
$$

Therefore, the information gain $D(\boldsymbol{G})$ is invariant to left unitary multiplication of $\boldsymbol{G}$. For any $\boldsymbol{G} \in \mathbb{R}^{t \times q}$ with rank r, $\boldsymbol{G}$ has the singular value decomposition $\boldsymbol{G} = \boldsymbol{U \Delta V}^T$ where $\boldsymbol{U}$ and $\boldsymbol{V}$ are orthogonal matrices, and $\boldsymbol{\Delta} \in \mathbb{R}^{t \times q}$ is a diagonal matrix with diagonal elements $\boldsymbol{\Delta}_{1,1} \geq \ldots \geq \boldsymbol{\Delta}_{r,r} > 0$ and $\boldsymbol{\Delta}_{i,i} = 0$ for any $i \geq r$. By the invariance property, we can assume WLOG that $\boldsymbol{U} = \boldsymbol{I}_t$. Let $\widetilde{\boldsymbol{G}} = \text{Diag}(\boldsymbol{\Delta}_{1,1}, \ldots, \boldsymbol{\Delta}_{r,r}) \boldsymbol{V}_r^T \in \mathbb{R}^{r \times q}$ where $\boldsymbol{V}_r \in \mathbb{R}^{q \times r}$ contains the first r columns of $\boldsymbol{V}$. It can be seen that $\boldsymbol{G} = [\widetilde{\boldsymbol{G}}^T, \boldsymbol{0}_{q \times (t-r)}]^T$ and $\text{tr}(\boldsymbol{GG}^T) = \text{tr}(\widetilde{\boldsymbol{G}} \widetilde{\boldsymbol{G}}^T)$. Moreover, one can easily check that

$$
D(\boldsymbol{G}) = \frac{1}{2} \log \det[\boldsymbol{I}_p + \boldsymbol{M}^T \widetilde{\boldsymbol{G}}^T (\widetilde{\boldsymbol{G}} \boldsymbol{Q_{\phi\phi|\theta}} \widetilde{\boldsymbol{G}}^T + \sigma_{\boldsymbol{v}}^2 \boldsymbol{I}_r)^{-1} \widetilde{\boldsymbol{G}} \boldsymbol{M} \boldsymbol{Q_{\theta\theta|x}}]. \tag{3.7.8}
$$

The RHS of (3.7.8) is the information gain brought by an r-dimensional channel $\widetilde{\boldsymbol{y}}$,

$$
\widetilde{\boldsymbol{y}} = \widetilde{\boldsymbol{G}} \boldsymbol{\phi} + \widetilde{\boldsymbol{v}}
$$

where $\widetilde{\boldsymbol{v}}$ is r-dimensional white noise with variance $\sigma_{\boldsymbol{v}}^2$. The new channel $\widetilde{\boldsymbol{y}}$ brings the same information gain, that is $I(\boldsymbol{\theta}; \boldsymbol{x}, \boldsymbol{y}) - I(\boldsymbol{\theta}; \boldsymbol{x}) = I(\boldsymbol{\theta}; \boldsymbol{x}, \widetilde{\boldsymbol{y}}) - I(\boldsymbol{\theta}; \boldsymbol{x})$.

CHAPTER 4

Conclusion and Future Work

4.1 Conclusion

In this dissertation, we have considered two design problems arising in MIMO channel signal processing. In Chapter 2, we have considered the linear compression or dimension reduction of a noisy measurement $\boldsymbol{x}$, which is then transmitted over a noisy channel. The final dimension-reduced measurement $\boldsymbol{z}$ is used to recover the signal of interest $\boldsymbol{\theta}$. Over the processes of compression and noisy transmission, we lose some information that $\boldsymbol{x}$ contains about $\boldsymbol{\theta}$. The optimal compression matrices that minimize the trace or determinant of the error covariance matrix, subject to the power constraint, were derived. Both analytical solutions return the scaled and rotated canonical coordinates of $\boldsymbol{x}$, while the choice of the canonical coordinate system depends on the criterion implemented. The scaling coefficients are determined by the canonical correlations between $\boldsymbol{x}$ and $\boldsymbol{\theta}$, and the eigenvalues of the noise covariance. The rotation sends the scaled canonical coordinates into the subdominant invariant space of the noise covariance matrix. We further extend the discussion to a general criterion and show that the general solution also returns the scaled and rotated canonical coordinates, with a particular choice of the coordinate system and scaling coefficients.

In Chapter 3, we have considered a system with multiple sources of measurements and correlated input signals. More specifically, a secondary channel is added to an existing primary channel. The objective is to design the optimal secondary channel to maximize the mutual information between the signal of interest and the measurements from both channels, subject to a total power constraint. In this problem, the input signals and the channel noises are Gaussian distributed, which allows an explicit expression for the mutual information. When the conditional covariance of the secondary input signal

66

given the primary input signal is proportional to the identity matrix, we have obtained an analytical solution for the optimal channel. For general cases, we have proposed two numerical algorithms to approximate the optimal secondary channel. The first extrinsic algorithm implements a projection onto the constraint space given by the power constraint. The second intrinsic algorithm exploits the geometry of the power constraint and restricts the search to a manifold. Both algorithms converge to a local optimal channel. We also show that the optimal secondary matrix can in some cases carry a compression of the secondary input signal.

4.2   Future Work

The work so far focuses on linear systems. Moreover, for compression and for fusion, the optimal solutions only exploit the second moments of the signals. In Chapter 2, the optimal compression matrix is derived in a system of canonical coordinates. These coordinates are given by the SVD of the coherence matrix, which is fully determined by the second moments of the signals. In Chapter 3, the Gaussian distribution is fully characterized by the first and second moments as well. Future work would seek design for compressing and fusing nonlinear features of the measurements, based on higher-order correlations between signals.

The idea of using nonlinear maps prior to linear processing has been exploited in the theory of support vector machine, where the data are mapped by a nonlinear mapping into a high-dimensional feature space, in which the features are linearly separable (see Vapnik [1995] and Vapnik [1998]). The idea of kernel methods avoids the high-dimensional nonlinear mapping and allows all computations to be carried out in the original low-dimensional space. In short, the kernel function is defined on the input space and returns the inner products in the feature space. Since the development of the support vector machine, numerous results have been reported on kernel nonlinear counterparts of standard information processing techniques including principal component

analysis, Fisher discriminant analysis, linear least squares estimation, etc. See Scholkopf et al. [1998], Ruiz and de Teruel [2001], and Scholkopf and Smola [2002] for more details. In a coming study, we will explore compression and fusion in system of featured, extracted from kernel pre-processing.

CHAPTER 5

BACKGROUND MATERIALS

5.1 Lagrange Duality Theory and KKT Optimality Conditions

Consider the optimization problem

$$
\begin{array}{ll}
\text{minimize} & f_0(\boldsymbol{x}) \\
\text{subject to} & f_i(\boldsymbol{x}) \leq u_i, \quad i = 1, \dots, m \\
& h_i(\boldsymbol{x}) = v_i, \quad i = 1, \dots, n
\end{array}
\tag{5.1.1}
$$

with variable $\boldsymbol{x} \in \mathbb{R}^p$. This is a constrained optimization problem with inequality constraints and equality constraints. We assume the domain of the functions $f_i (i = 0, 1, \dots, m)$ and $h_i (i = 1, \dots, n)$ have nonempty intersection.

The Lagrange duality theory is based on a dual optimization problem associated with problem (5.1.1). To initiate the discussion, we first introduce the Lagrangian and the Lagrange dual function. The Lagrangian L is a mapping from $\mathbb{R}^p \times \mathbb{R}^m \times \mathbb{R}^n$ to $\mathbb{R}^1$ as

$$
L(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = f_0(\boldsymbol{x}) + \sum_{i=1}^{m} \lambda_i f_i(\boldsymbol{x}) + \sum_{i=1}^{n} \nu_i h_i(\boldsymbol{x}).
\tag{5.1.2}
$$

The vectors $\boldsymbol{\lambda}$ and $\boldsymbol{\nu}$ are the Lagrange multipliers associated with the problem (5.1.1). The Lagrange dual function, denotes g, is a mapping from $\mathbb{R}^m \times \mathbb{R}^n$ to $\mathbb{R}^1$ that returns the minimum of the Lagrangian over $\boldsymbol{x}$:

$$
g(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \inf_{\boldsymbol{x}} \left(f_0(\boldsymbol{x}) + \sum_{i=1}^{m} \lambda_i f_i(\boldsymbol{x}) + \sum_{i=1}^{n} \nu_i h_i(\boldsymbol{x}) \right).
\tag{5.1.3}
$$

For any $\boldsymbol{\lambda} \geq 0$ and $\boldsymbol{\nu}$, the Lagrange dual function $g(\boldsymbol{\lambda}, \boldsymbol{\nu})$ gives a lower bound on the minimum of problem (5.1.1). This fact is easy to verify and omitted here. Readers may refer to Boyd and Vandenberghe [2004] for more details. The Lagrange dual problem is then the maximization of the Lagrange dual function, which is to

$$
\begin{array}{ll}
\text{maximize} & g(\boldsymbol{\lambda}, \boldsymbol{\nu}) \\
\text{subject to} & \boldsymbol{\lambda} \geq 0
\end{array}
\tag{5.1.4}
$$

The problem (5.1.1) is sometimes referred to the primary problem. The Lagrange dual problem (5.1.4) is a convex optimization problem, which is true no matter whether the primal problem (5.1.1) is convex or not.

Let $\boldsymbol{x}^*$ and $(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$ be the optimal solutions of the primary problem and the dual problem, respectively. Moreover, define $p^* = f_0(\boldsymbol{x}^*)$ and $d^* = g(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$. It can be seen that

$$d^* \leq p^*$$

The difference between d^* and p^* is called the duality gap. When the duality gap is zero, we say that strong duality holds. In this case, any primal optimal point is also a minimizer of $L(\boldsymbol{x}, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$, where $(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$ is a dual optimal solution. In other words, a primal optimal solution is also a solution of the following unconstrained problem:

$$\text{minimize } L(\boldsymbol{x}, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) = f_0(\boldsymbol{x}) + \sum_{i=1}^{m} \lambda_i^* f_i(\boldsymbol{x}) + \sum_{i=1}^{n} \nu_i^* h_i(\boldsymbol{x}) \qquad (5.1.5)$$

This is the motivation for Karush-Kuhn-Tucker (KKT) conditions:

$$
\begin{aligned}
f_i(\boldsymbol{x}^*) &\leq 0, \quad i = 1, \ldots, m \\
h_i(\boldsymbol{x}^*) &= 0, \quad i = 1, \ldots, n \\
\lambda_i^* &\geq 0, \quad i = 1, \ldots, m \\
\lambda_i^* f_i(\boldsymbol{x}^*) &= 0, \quad i = 1, \ldots, m \\
\nabla f_0(\boldsymbol{x}^*) + \sum_{i=1}^{m} \lambda_i \nabla f_i(\boldsymbol{x}) + \sum_{i=1}^{n} \nu_i \nabla h_i(\boldsymbol{x}) &= 0
\end{aligned}
\qquad (5.1.6)
$$

The first three conditions are necessary for any feasible solutions. The last condition is a direct consequence of problem (5.1.5). The fourth condition is called complementary slackness. Note that

$$
\begin{aligned}
f_0(\boldsymbol{x}^*) &= g(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) \\
&= \inf_{\boldsymbol{x}} \left(f_0(\boldsymbol{x}) + \sum_{i=1}^{m} \lambda_i^* f_i(\boldsymbol{x}) + \sum_{i=1}^{n} \nu_i^* h_i(\boldsymbol{x}) \right) \\
&\leq f_0(\boldsymbol{x}^*) + \sum_{i=1}^{m} \lambda_i^* f_i(\boldsymbol{x}^*) + \sum_{i=1}^{n} \nu_i^* h_i(\boldsymbol{x}^*) \\
&\leq f_0(\boldsymbol{x}^*)
\end{aligned}
$$

where the first equation follows the strong duality and the last inequality follows from $\boldsymbol{\lambda}^* \geq 0, f_i(\boldsymbol{x}^*) \leq 0$ and $h_i(\boldsymbol{x}^*) = 0$. This directly yields that $\sum_{i=1}^{m} \lambda_i^* f_i(\boldsymbol{x}^*) = 0$ and $\lambda_i^* f_i(\boldsymbol{x}^*) = 0$ since each term is nonnegative.

In summary, for problem (5.1.1) with differentiable objective and constraint functions, any optimal solution must satisfy the KKT conditions (5.1.6).

5.2 Miscellaneous Matrix Results

When solving the matrix variate optimization problems in this dissertation, simplifying the objective function reduces the complexity of the problem. In this section, we introduce several useful results in matrix analysis.

First, we show a fundamental equation for the determinant function, which is known as Sylvester's Determinant Theorem: For any $m \times n$ matrix $\boldsymbol{A}$ and $n \times m$ matrix $\boldsymbol{B}$,

$$\det(\boldsymbol{I}_m + \boldsymbol{AB}) = \det(\boldsymbol{I}_n + \boldsymbol{BA})$$

More generally, for any invertible $m \times m$ matrix $\boldsymbol{X}$,

$$\det(\boldsymbol{X} + \boldsymbol{AB}) = \det(\boldsymbol{I} + \boldsymbol{AB}\boldsymbol{X}^{-1}) \det(\boldsymbol{X})$$

The matrix inversion formula, or Woodbury matrix identity, gives a useful transformation of matrix inverse: Suppose that the matrices $\boldsymbol{A}$ and $\boldsymbol{C}$ are invertible. Then,

$$(\boldsymbol{A} + \boldsymbol{UCV})^{-1} = \boldsymbol{A}^{-1} - \boldsymbol{A}^{-1}\boldsymbol{U} \left(\boldsymbol{C}^{-1} + \boldsymbol{V}\boldsymbol{A}^{-1}\boldsymbol{U}\right)^{-1} \boldsymbol{V}\boldsymbol{A}^{-1}.$$

When the matrix $\boldsymbol{A}$ is the identity matrix and $\boldsymbol{V} = \boldsymbol{U}^T$, we have

$$(\boldsymbol{I} + \boldsymbol{UCU}^T)^{-1} = \boldsymbol{I} - \boldsymbol{U} \left(\boldsymbol{C}^{-1} + \boldsymbol{U}^T\boldsymbol{U}\right)^{-1} \boldsymbol{U}^T.$$

Next, we introduce differentiation of a scalar function with respect to a matrix $\boldsymbol{X}$. Given a scalar function $y = f(\boldsymbol{X})$, the matrix derivative of f to $\boldsymbol{X}$ is defined as

$$\frac{\partial y}{\partial \mathbf{X}} = \begin{bmatrix} \frac{\partial y}{\partial x_{11}} & \frac{\partial y}{\partial x_{21}} & \cdots & \frac{\partial y}{\partial x_{p1}} \\ \frac{\partial y}{\partial x_{12}} & \frac{\partial y}{\partial x_{22}} & \cdots & \frac{\partial y}{\partial x_{p2}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial y}{\partial x_{1q}} & \frac{\partial y}{\partial x_{2q}} & \cdots & \frac{\partial y}{\partial x_{pq}} \end{bmatrix}.$$

A list of the derivative of the trace and determinant function is

$$\frac{\partial \operatorname{tr}(\boldsymbol{X})}{\partial \boldsymbol{X}} = \boldsymbol{I}$$
$$\frac{\partial \det \boldsymbol{X}}{\partial \boldsymbol{X}} = \det \boldsymbol{X}(\boldsymbol{X}^{-1})$$
$$\frac{\partial \log \det \boldsymbol{X}}{\partial \boldsymbol{X}} = \boldsymbol{X}^{-1}$$
$$\frac{\partial \operatorname{tr}(\boldsymbol{X}^{-1})}{\partial \boldsymbol{X}} = -\boldsymbol{X}^{-2}$$
$$\frac{\partial \det(\boldsymbol{X}^{-1})}{\partial \boldsymbol{X}} = -\det \boldsymbol{X}^{-1}(\boldsymbol{X}^{-1})$$

In our optimization problem, computing the first-order derivative provides a necessary condition for the optimal solution. Solving the first-order derivative often gives a reduced candidate set of the solutions.

5.3 Differential Entropy and Information

In Chapter 3, the criterion we used is based on information theory. In this section, we will briefly introduce the concept of entropy and mutual information. For a discrete random variable Z with probability mass function $p(z) = \operatorname{Prob}(Z = z)$, the entropy for Z is

$$H(Z) = E[-\log_b(p(Z))] = -\sum_i p(z_i) \log_b p(z_i)$$

Here b is the base of the logarithm used. When $b = 2$, then entropy is measured in bits. The entropy $H(\boldsymbol{z})$ is a measure of uncertainty of $\boldsymbol{z}$. For example, the random experiment of flipping a fair coin has entropy 1 bits. Through this work, we will use Euler's base where $\log_b$ is the natural logarithm. The continuous version of the entropy is called the differential entropy. The differential entropy for a random variable Z with density function f is

$$H(Z) = \mathbb{E}\left[\log \frac{1}{f(Z)}\right].$$

A Gaussian distributed random vector $\boldsymbol{z} \in \mathbb{R}^k$ with mean $\boldsymbol{\mu_z}$ and covariance $\boldsymbol{Q_{zz}}$ has differential entropy

$$H(\boldsymbol{z}) = \frac{1}{2} \log \det \boldsymbol{Q_{zz}} + \frac{k}{2} \log(2\pi).$$

For two random variables Z_1 and Z_2, the equivocation or conditional entropy for Z_1 given Z_2 is

$$H(Z_1|Z_2) = \mathbb{E}\left[\log \frac{1}{f_{Z_1|Z_2}(Z_1|Z_2)}\right],$$

where $f_{Z_1|Z_2}(\cdot)$ is the conditional density of Z_1 given Z_2. The expectation is taken over the joint distribution of (Z_1, Z_2). The equivocation $H(Z_1|Z_2)$ measures the average remaining information in Z_1 after revealing the value of Z_2. $H(Z_1|Z_2) = 0$ if and only if the value of Z_1 is determined by the value of Z_2. For $\boldsymbol{z}_1 \in \mathbb{R}^{k_1}$ and $\boldsymbol{z}_2 \in \mathbb{R}^{k_2}$ with a joint Gaussian distribution,

$$H(\boldsymbol{z}_1|\boldsymbol{z}_2) = \frac{1}{2} \log \det \boldsymbol{Q}_{z_1 z_1 | z_2} + \frac{k_1}{2} \log(2\pi),$$

where $\boldsymbol{Q}_{z_1 z_1 | z_2}$ is the conditional covariance matrix of $\boldsymbol{z}_1$ given $\boldsymbol{z}_2$. The mutual information between two random variables Z_1 and Z_2 is

$$I(Z_1; Z_2) = H(Z_1) - H(Z_1|Z_2).$$

The mutual information $I(Z_1; Z_2)$ measures the amount of information Z_2 contains about Z_1, or equivalently, the amount of information Z_1 contains about Z_2. $I(Z_1; Z_2) = 0$ if and only if Z_1 and Z_2 are independent random variables. For jointly Gaussian distributed random vectors $\boldsymbol{z}_1$ and $\boldsymbol{z}_2$, the mutual information is

$$I(\boldsymbol{z}_1; \boldsymbol{z}_2) = \frac{1}{2} \log \det \boldsymbol{Q}_{z_1 z_1} - \frac{1}{2} \log \det \boldsymbol{Q}_{z_1 z_1 | z_2} \tag{5.3.1}$$

with $\boldsymbol{Q}_{z_1 z_1}$ the covariance matrix of $\boldsymbol{z}_1$. In Chapter 3, the explicit expression (5.3.1) is used to derive the information gain criterion, when the input signals are Gaussian distributed.

Bibliography

Absil, P. A., Mahony, R., and Sepulchre, R. (2008). *Optimization algorithms on matrix manifolds*. Princeton University Press.

Bertsekas, D. P. (1982). Projected Newton methods for optimization problems with simple constraints. *SIAM Journal on Control and Optimization*, 20:221–246.

Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.

Candès, E. J., Romberg, J. K., and Tao, T. (2006). Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59:1207–1223.

Carson, W. R., Rodrigues, M. R. D., Chen, M., Carin, L., and Calderbank, R. (2012). Communications-inspired projection design with application to compressive sensing. *SIAM Journal on Imaging Sciences*, 5:1185–1212.

Cover, T. M. and Thomas, J. A. (2005). *Elements of information theory, 2nd ed.* Wiley-Interscience.

Ding, Y., Davidson, T., Luo, Z., and Wong, K. (2003). Minimum BER block precoders for zero-forcing equalization. *IEEE Transactions on Signal Processing*, 51:2410–2423.

Donoho, D. L. (2006). Compressed sensing. *IEEE Transactions on Information Theory*, 52:1289–1306.

Edelman, A., Arias, T. A., and Smith, S. T. (1998). The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, 20:303–353.

Foschini, G. and Gans, M. (1998). On limits of wireless communications in a fading environment when using multiple antennas. *Wireless Personal Communications*, 6:311–335.

Gabay, D. (1982). *Minimizing a differentiable function over a differential manifold.* INRIA Rocquencourt.

Ginis, G. and Cioffi, J. (2002). Vectored transmission for digital subscriber line systems. *IEEE Journal on Selected Areas in Communications*, 20:1085–1104.

Guo, D., Shamai, S., and Verdù, S. (2005). Mutual information and minimum mean-square error in gaussian channels. *IEEE Transactions on Information Theory*, 51:1261–1282.

Honig, M. L., Steiglitz, K., and Gopinath, B. (1990). Multichannel signal processing for data communications in the presence of crosstalk. *IEEE Transactions on Communications*, 38:551–558.

Hotelling, H. (1936). Relations between two sets of variates. *Biometrika*, 28:321–377.

Hua, Y., Nikpour, M., and Stocia, P. (2001). Optimal reduced-rank estimation and filtering. *IEEE Transactions on Signal Processing*, 49:457–469.

Johnson, D. H. and Dudgeon, D. E. (1993). *Array signal processing: concepts and techniques.* Prentice Hall.

Komaroff, N. (1990). Bounds on eigenvalues of matrix products with an application to the algebraic riccati equation. *IEEE Transactions on Automatic Control*, 35:348–350.

Krim, H. and Viberg, M. (1996). Two decades of array signal processing research. *IEEE Signal Processing Magazine*, 13:67–94.

Lamarca, M. (2009). Linear precoding for mutual information maximization in mimo systems. In *Intl. Symp. on Wireless Communication Systems (ISWCS '09)*.

Lee, J. M. (2000). *Introduction to smooth manifolds*. Springer.

Lee, K. H. and Petersen, D. P. (1976). Optimal linear coding for vector channels. *IEEE Transactions on Communications*, COM-24:1283–1290.

Liu, E., Chong, E. K. P., and Scharf, L. L. (2012). Greedy adaptive compression in signal-plus-noise models. available at arXiv:1202.3913.

Lozano, A., Tulino, A., and Verdú, S. (2006). Optimum power allocation for parallel Gaussian channels with arbitrary input distributions. *IEEE Transactions on Information Theory*, 52:3033–3051.

Monzingo, R. A. and Miller, T. W. (1980). *Introduction to adaptive arrays*. New York: Wiley.

Palomar, D. and Jiang, Y. (2007). *MIMO transceiver design via majorization theory*. Now Publishers Inc.

Palomar, D. and Verdú, S. (2006). Gradient of mutual information in linear vector Gaussian channels. *IEEE Transactions on Information Theory*, 52:141–154.

Palomar, D. P., Cioffi, J., and Lagunas, M. A. (2003). Joint TxRx beamforming design for multcarrier MIMO channels: A unifed framework for convex optimization. *IEEE Transactions on Signal Processing*, 51:2381–2401.

Payaró, M. and Palomar, D. P. (2009). Hessian and concavity of mutual information, entropy, and entropy power in linear vector Gaussian channels. *IEEE Transactions on Information Theory*, 55:3613–3628.

Pérez-Cruz, F., Rodrigues, M. R. D., and Verdú, S. (2010). MIMO Gaussian channels with arbitrary inputs: Optimal precoding and power allocation. *IEEE Transactions on Information Theory*, 56:1070–1084.

Petersen, K. B. and Pedersen, M. S. (2008). *The matrix cookbook.* available online.

Pukelsheim, F. (1993). *Optimal design of experiments.* Wiley-Interscience.

Ruiz, A. and de Teruel, P. E. L. (2001). Nonlinear kernel-based statistical pattern analysis. *IEEE Trans. Neural Networks*, 12:16–32.

Salz, J. (1985). Digital transmission over cross-coupled linear channels. *AT&T Technical Journal*, 64:1147–1159.

Sampath, H., Stoica, P., and Paulraj, A. (2001). Generalized linear precoder and decoder design for MIMO channels using the weighted MMSE criterion. *IEEE Transactions on Communications*, 49:2198–2206.

Scaglione, A., Stoica, P., Barbarossa, S., Giannakis, G. B., and Sampath, H. (2002). Optimal designs for space-time linear precoders and decoders. *IEEE Transactions on Signal Processing*, 50:1051–1064.

Scharf, L. L. (1990). *Statistical signal processing: Detection, estimation, and time series analysis.* Prentice Hall.

Scharf, L. L. and Mullis, C. T. (1998). Canonical coordinates and the geometry of inference, rate, and capacity. *IEEE Transactions on Signal Processing*, 48:824–831.

Scharf, L. L. and Thomas, J. K. (1998). Wiener filters in canonical coordinates for transform coding, filtering, and quantizing. *IEEE Transactions on Signal Processing*, 46:647–654.

Schizas, I. D., Giannakis, G. B., and Luo, Z. (2007). Distributed estimation using reduced-dimensionality sensor observations. *IEEE Transactions on Signal Processing*, 55:4284–4299.

Scholkopf, B., Smola, A., and Müller, K. R. (1998). Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10:1299–1319.

Scholkopf, B. and Smola, A. J. (2002). *Learning with kernels, support vector machines, regularization, optimization, and beyond.* Cambridge, MA: MIT Press.

Schreier, P. J. and Scharf, L. L. (2006). Canonical coordinates for transform coding of noisy sources. *IEEE Transactions on Signal Processing*, 54:235–243.

Skoglund, M. and Jöngren, G. (2003). On the capacity of a multiple-antenna communication link with channel side information. *IEEE Journal on Selected Areas in Communications*, 21:395–405.

Smith, S. T. (1994). Optimization techniques on Riemannian manifolds. *Hamiltonian and Gradient Flows, Algorithms, and Control American Mathematical Soc.*, 1:113–136.

Tarokh, V., Seshadri, N., and Calderbank, A. R. (1998). Space-time codes for high data rate wireless communication: Performance criterion and code construction. *IEEE Transactions on Information Theory*, 44:744–765.

Telatar, I. E. (1999). Capacity of multi-antenna Gaussian channels. *European Transactions on Telecommunications*, 10:585–595.

Van Trees, H. (2002). *Detection, estimation, and modulation theory. Part IV: Optimum array processing.* Wiley.

Van Veen, B. D. and Buckley, K. M. (1988). Beamforming: a versatile approach to spatial filtering. *IEEE ASSP Magazine*, 5:4–24.

Vapnik, V. N. (1995). *The nature of statistical learning theory.* Springer.

Vapnik, V. N. (1998). *Statistical learning theory.* Wiley.

Visotsky, E. and Madhow, U. (2001). Space-time transmit precoding with imperfect feedback. *IEEE Transactions on Information Theory*, 47:2632–2639.

Vojcic, B. R. and Jang, W. M. (1998). Transmitter precoding in synchronous multiuser communication. *IEEE Transactions on Communications*, 46:1346–1355.

Vosoughi, A. and Scaglione, A. (2007). Precoding and decoding paradigms for distributed vector data compression. *IEEE Transactions on Signal Processing*, 55:1445–1460.

Vu, M. and Paulraj, A. (2007). MIMO wireless linear precoding. *IEEE Signal Processing Magazine*, 24:86–105.

Wang, Y., Wang, H., and Scharf, L. L. (2013). Optimum compression of a noisy measurement for transmission over a noisy channel. *submitted.*

Wiesel, A., Eldar, Y. C., and Shamai, S. (2006). Linear precoding via conic optimization for fixed MIMO receivers. *IEEE Transactions on Signal Processing*, 54:161–176.

Witsenhausen, H. (1975). A determinant maximization problem occuring in the theory of data communication. *Jounal of Applied Mathematics*, 29:515–522.

Xiao, C., Zheng, Y. R., and Ding, Z. (2011). Globally optimal linear precoders for finite alphabet signals over complex vector gaussian channels. *IEEE Transactions on Signal Processing*, 59:3301–3314.