The Geometry of Fusion Inspired Channel Design

Yuan Wang[a,∗], Haonan Wang[a], Louis L. Scharf[a,b]

[a]*Department of Statistics, Colorado State University, Fort Collins, CO 80523, USA.*
[b]*Department of Mathematics, Colorado State University, Fort Collins, CO 80523, USA.*

Abstract

This paper is motivated by the 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 matrix, when the primary channel matrix is fixed. We formulate the problem as an optimization problem, in which the optimal secondary channel matrix 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.

Keywords: Embedded submanifold, information fusion, MIMO channel design,

---

∗Corresponding author. Tel.: +1 970-491-3778.
Email addresses: `wangy@stat.colostate.edu` (Yuan Wang),
`wanghn@stat.colostate.edu` (Haonan Wang), `scharf@engr.colostate.edu`
(Louis L. Scharf)

1. Introduction

Consider the following two-channel system, as illustrated in Fig. 1,

$$\boldsymbol{x} = \boldsymbol{F}\boldsymbol{\theta} + \boldsymbol{u}$$
$$\boldsymbol{y} = \boldsymbol{G}\boldsymbol{\phi} + \boldsymbol{v}. \tag{1}$$

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 paper, 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 $\boldsymbol{x} = \boldsymbol{F}\boldsymbol{\theta} + \boldsymbol{u}$, designing the channel matrix $\boldsymbol{F}$ exhibits parallels to the linear precoding design 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 [1]-[10]. The optimal precoding is designed under various
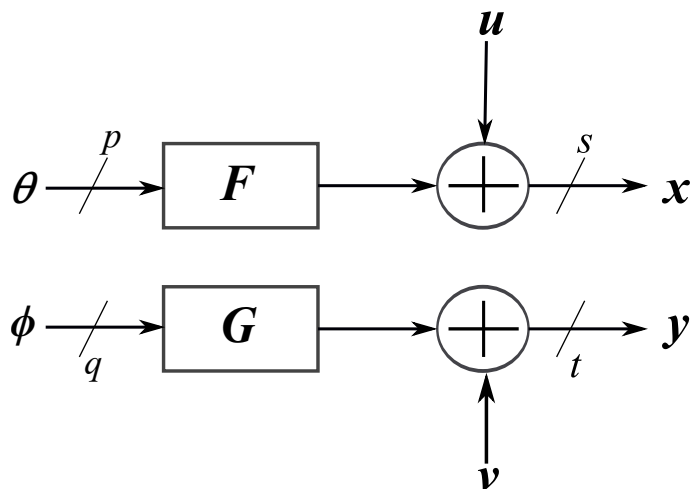
Figure 1: A two-channel system with two linear channels.

criteria, for example, signal-to-noise ratio (SNR) and signal-to-interference-noise ratio (SINR), [1]-[2]. Another criterion that has drawn more attention recently is the mutual information between input and output signals, [3]-[4], [6]-[7]. 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) [11]-[12]. In [3], 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 [7], [23]. In [6], a greedy adaptive approach is considered to design a channel matrix row by row to maximize information gain.

In Fig. 1, if both θ and ϕ are of interest, the two-channel system can be ex-

pressed as a one-channel system with a block-diagonal channel matrix. However, due to the "nuisance" signal ϕ, our two-channel system design problem is fundamentally more difficult than the one-channel system design. In this paper, 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 $\mathrm{tr}(\boldsymbol{G}\boldsymbol{G}^T) \leq P$ with P 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 to the constrained space [14]. The intrinsic algorithm, a gradient-ascent algorithm of manifold, exploits the geometry that codes for the total power constraint by vectorizing the channel matrix, [15], [17]-[19].

The rest of the paper is organized as follows. We formulate the channelization problem in Section 2 and point out the challenges for design in a two-channel system. In Section 3, we give an analytical solution when the conditional covariance of ϕ given $\boldsymbol{\theta}$ is the identity matrix. In Section 4, we propose two numerical algorithms, the 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 4.3. In Section 5, we discuss the choice of number of measurements for the secondary channel. Section 6 concludes the paper.

Notation: The set of length m real vectors is denoted by $\mathbb{R}^m$ and the set of

4

$m \times n$ real matrices is denoted $\mathbb{R}^{m \times n}$. Bold upper case letters denote matrices, bold 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, trace and determinant of a matrix are denoted by $(\cdot)^T$, $(\cdot)^{-1}$, $\mathrm{tr}(\cdot)$ and $\det(\cdot)$, respectively.

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

2. Overview

2.1. Problem Statement

The two channels of the system described in (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 would 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 $\boldsymbol{G}$, with the first channel fixed, such that the 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

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

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 all the 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 $\boldsymbol{Q_{\theta\theta|x}} = (\boldsymbol{Q_{\theta\theta}^{-1}} + \boldsymbol{F}^T \boldsymbol{Q_{uu}^{-1}} \boldsymbol{F})^{-1}$ 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 $\boldsymbol{Q_{\theta\theta|x,y}} = [\boldsymbol{Q_{\theta\theta|x}^{-1}} + \boldsymbol{M}^T \boldsymbol{G}^T (\boldsymbol{G} \boldsymbol{Q_{\phi\phi|\theta}} \boldsymbol{G}^T + \boldsymbol{Q_{vv}})^{-1} \boldsymbol{G} \boldsymbol{M}]^{-1}$ with $\boldsymbol{M} = \boldsymbol{Q_{\phi\theta}} \boldsymbol{Q_{\theta\theta}^{-1}}$ and $\boldsymbol{Q_{\phi\phi|\theta}} = \boldsymbol{Q_{\phi\phi}} - \boldsymbol{Q_{\phi\theta}} \boldsymbol{Q_{\theta\theta}^{-1}} \boldsymbol{Q_{\theta\phi}}$. Note that $\boldsymbol{M}\boldsymbol{\theta}$ 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 $\boldsymbol{y}$, defined as

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

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

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

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 that 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 (2) 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 $\text{tr}(\boldsymbol{G}\boldsymbol{G}^T) \leq P$. This constraint bounds the total power of $\boldsymbol{G}\boldsymbol{\phi}$ since $\text{tr}\,\mathbb{E}[\boldsymbol{G}\boldsymbol{\phi}\boldsymbol{\phi}^T\boldsymbol{G}^T] \leq P\,\text{tr}\,\mathbb{E}[\boldsymbol{\phi}\boldsymbol{\phi}^T]$. In short, the problem of interest is

$$\boldsymbol{G}^* = \arg\max_{\boldsymbol{G}\in\mathbb{R}^{t\times q}} D(\boldsymbol{G}) \text{ subject to } \text{tr}(\boldsymbol{G}\boldsymbol{G}^T) \leq P. \qquad (3)$$

Problem (3) 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 $\boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}}$ 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.

2.2. An Insightful Discussion of the Information Gain

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

$$\boldsymbol{y} = (\boldsymbol{G}\boldsymbol{M}\mathbb{E}[\boldsymbol{\theta}|\boldsymbol{x}]) + (\boldsymbol{G}\boldsymbol{M}(\boldsymbol{\theta} - \mathbb{E}[\boldsymbol{\theta}|\boldsymbol{x}])) + (\boldsymbol{G}(\boldsymbol{\phi} - \mathbb{E}[\boldsymbol{\phi}|\boldsymbol{\theta}]) + \boldsymbol{v}), \qquad (4)$$

7

where $M = Q_{\phi\theta}Q_{\theta\theta}^{-1}$ and $M\theta = \mathbb{E}[\phi|\theta]$. It can be seen that the secondary channel y is decomposed into three independent components, which are illustrated in Fig. 2. The first component $GM\mathbb{E}[\theta|x]$ is completely determined by the first channel x and does not contribute to the information gain brought by y. The second component $GM(\theta - \mathbb{E}[\theta|x])$, denoted by ω, is (by orthogonality) independent of x and it carries the extra information in channel y about θ. The third component $G(\phi - \mathbb{E}[\phi|\theta]) + v$, denoted by ζ, is independent of both x and θ, and it can be viewed as noise.
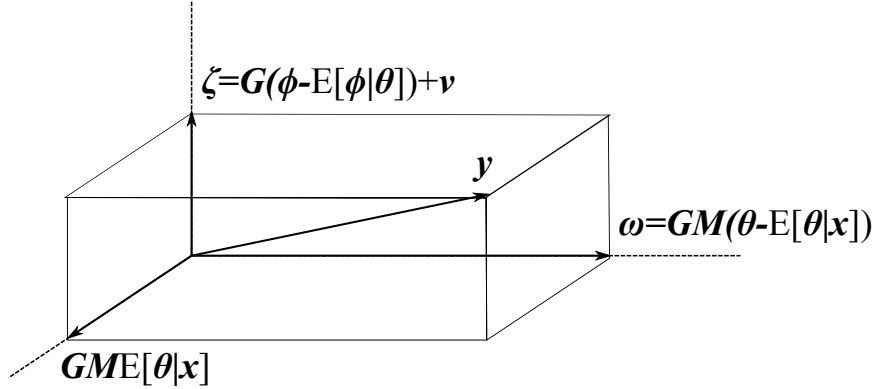
Figure 2: Decomposition of the secondary channel.

Notice that the covariance matrices of ω and ζ are $Q_{\omega\omega} = GMQ_{\theta\theta|x}M^TG^T$ and $Q_{\zeta\zeta} = GQ_{\phi\phi|\theta}G^T + Q_{vv}$, respectively. By the cyclic property of determinants, $\det(I_m + AB) = \det(I_n + BA)$ for any $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{n \times m}$, the information gain of (2) can be re-written as

$$D(G) = \frac{1}{2}\log\det[I + Q_{\zeta\zeta}^{-1/2}Q_{\omega\omega}Q_{\zeta\zeta}^{-1/2}]. \tag{5}$$

By viewing ω as a signal and ζ as a noise, $Q_{\zeta\zeta}^{-1/2}Q_{\omega\omega}Q_{\zeta\zeta}^{-1/2}$ is a generalized signal-to-noise ratio matrix. Maximizing (5) essentially balances the tradeoff between the noise covariance and the signal covariance. As illustrated in Fig. 2, a

8

good channel design will favor a long parallelepiped with short height. The difficulty of designing the channel matrix G arises because G shapes both $Q_{\omega\omega}$ and $Q_{\zeta\zeta}$. When the secondary channel has a single output, the channel matrix G is a row vector and $Q_{\zeta\zeta}, Q_{\omega\omega}$ are scalars. In this case, the optimal 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 G is fundamentally more difficult than the single output case. In Section 3, we obtain a closed-form expression for the optimal G in a special case.

3. Analytical Solution

Suppose that the conditional covariance of ϕ given θ is identity, i.e., $Q_{\phi\phi|\theta} = \sigma^2_{\phi|\theta}I_q$. For example, $\phi = M\theta + \tau$ where $\tau \sim N(0, \sigma^2_{\phi|\theta}I_q)$. In this case, the noise ζ in (5) has a relatively simple covariance $Q_{\zeta\zeta} = GG^T + Q_{vv}$ and the signal ω has covariance $Q_{\omega\omega} = GMQ_{\theta\theta|x}M^TG^T$. While G still affects both covariance matrices, we are able to find the balanced matrix G that maximizes the information gain. Note that we focus on the case $t \leq q$, i.e., the dimension of measurement y is at most the dimension of ϕ. When $t > q$, the optimization problem can be reformulated and solved as a special case of $t = q$, which will be discussed in Section 5.

Given the eigendecompositions $Q_{vv} = U_v\Sigma_vU_v$ and $MQ_{\theta\theta|x}M^T = U_\xi\Sigma_\xi U_\xi^T$ where $U_v \in \mathbb{R}^{t\times t}$, $U_\xi \in \mathbb{R}^{q\times q}$ are orthogonal matrices, and $\Sigma_v \in \mathbb{R}^{t\times t}$ and $\Sigma_\xi \in \mathbb{R}^{q\times q}$ are diagonal matrices with diagonal elements $0 < \sigma^2_{v,1} \leq \ldots \leq \sigma^2_{v,m}$ and $\sigma^2_{\xi,1} \geq \ldots \geq \sigma^2_{\xi,q} \geq 0$, respectively. Because the matrices U_ξ and U_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

$$\boldsymbol{G} = \boldsymbol{U_v} \boldsymbol{\Phi} \boldsymbol{U_\xi}^T \tag{6}$$

Then, the information gain $D(\boldsymbol{G})$ in (2) can be written as

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

Moreover, the total power constraint is $\text{tr}(\boldsymbol{\Phi} \boldsymbol{\Phi}^T) \leq P$ since $\text{tr}(\boldsymbol{G} \boldsymbol{G}^T) = \text{tr}(\boldsymbol{\Phi} \boldsymbol{\Phi}^T)$. For the given eigendecompositions, the matrices $\boldsymbol{U_v}$ and $\boldsymbol{U_\xi}$ are fixed. Therefore, the information gain can be maximized with respect to $\boldsymbol{\Phi}$ and the optimal channel matrix $\boldsymbol{G}$ is returned by (6). WOLG we assume $\sigma_{\phi|\boldsymbol{\theta}}^2 = 1$. The solution for general $\sigma_{\phi|\boldsymbol{\theta}}^2$ is just different by a scaling factor. We give in Lemma 1 an important feature of any possible maximizer $\boldsymbol{\Phi}$.

Lemma 1. *Suppose that $\boldsymbol{Q_{vv}}$ has distinct eigenvalues, i.e., $0 < \sigma_{v,1}^2 < \ldots < \sigma_{v,m}^2$, and $\boldsymbol{M} \boldsymbol{Q_{\theta\theta|x}} \boldsymbol{M}^T$ has distinct nonzero eigenvalues, i.e., $\sigma_{\xi,1}^2 > \ldots > \sigma_{\xi,\rho}^2 > 0$ where $\rho \leq t$ is the rank of $\boldsymbol{\Sigma_\xi}$. 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 Appendix Appendix A.

Lemma 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 1.

Theorem 1. *Suppose that $\boldsymbol{Q_{vv}}$ and $\boldsymbol{M} \boldsymbol{Q_{\theta\theta|x}} \boldsymbol{M}^T$ have distinct nonzero eigenvalues. Then the optimal secondary channel matrix $\boldsymbol{G}^*$ solving problem (3) is*

$$\boldsymbol{G}^* = \boldsymbol{U_v} \boldsymbol{\Lambda}^* \boldsymbol{U_\xi}^T. \tag{7}$$

10

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

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

$$(8)$$

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

Proof: See Appendix Appendix B.

Notice that although Theorem 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 (8) is a continuous function of the eigenvalues of $\boldsymbol{Q_{vv}}$ and $\boldsymbol{MQ_{\theta\theta|x}M}^T$.

Theorem 1 factors the optimal channel matrix $\boldsymbol{G}^*$ into the product of three matrices. The first matrix $\boldsymbol{U}^T_{\boldsymbol{\xi}}$ rotates the signal $\boldsymbol{\phi}$. Given $\boldsymbol{Q_{\phi\phi|\theta}} = \sigma^2_{\phi|\boldsymbol{\theta}}\boldsymbol{I}_q$, the conditional covariance of $\boldsymbol{\phi}$ given $\boldsymbol{x}$ is $\boldsymbol{MQ_{\theta\theta|x}M}^T + \sigma^2_{\phi|\boldsymbol{\theta}}\boldsymbol{I}_q$, which is diagonalized by $\boldsymbol{U}^T_{\boldsymbol{\xi}}$. Therefore, the components of the rotated signal $\boldsymbol{U}^T_{\boldsymbol{\xi}}\boldsymbol{\phi}$ 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 $\boldsymbol{U}^T_{\boldsymbol{\xi}}\boldsymbol{\phi}$ and distributes power across the t sub-channels optimally. The third matrix $\boldsymbol{U_v}$ rotates the scaled components into the sub-dominant invariant subspace of the noise covariance $\boldsymbol{Q_{vv}}$.

The power allocation policy, given by the diagonal elements of $\boldsymbol{\Lambda}^*$, can be interpreted as a *mercury/waterfilling* algorithm, which is a three-step procedure that has been introduced in [25]:

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

11

component of ϕ given x.

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

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

.

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

The height of the solid base, $2\sigma_{v,i}^2/\sigma_{\xi,i}^2$, is half of the variance of the ith noise weighted by the variance components of $MQ_{\theta\theta|x}M^T$. A higher solid base means a less informative channel with high channel noise and weak correlation with θ. 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 P. The mercury stage balances the noise contained in ϕ and the measurement noise contained in y. Without adding mercury, the optimal power allocation will 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 G^* may not be full-rank. We will see in Section 5 that a rank-reduced 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
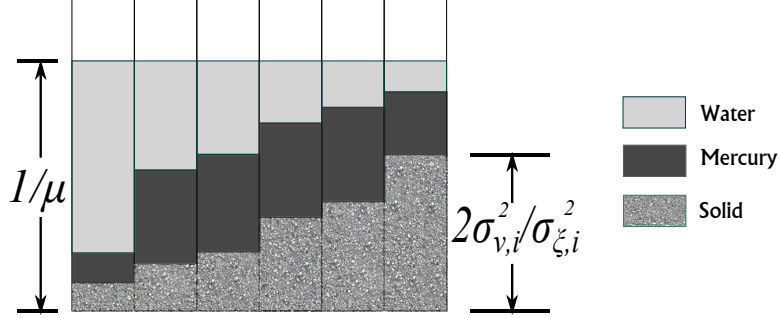
12

Figure 3: Mercury/waterfilling. For each vessel, the water height above mercury gives the optimal power allocation to the corresponding subchannel. The total volume of water equals P.

channel, we consider the following simple example.

Example 1. *Consider a two-channel system in* (1) *with $p = q = s = t = 5$. The primary channel matrix $\boldsymbol{F} \in \mathbb{R}^{5 \times 5}$ is set as $\frac{1}{\sqrt{5}} \boldsymbol{I}_5$. 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 [26], Wang et al. generalized the problem of reduced-rank filtering and precoding/equalizing be designing the matrix $\boldsymbol{G}$ in the bottom channel of Fig. 4 so that $\boldsymbol{y}$ maximizes the differential rate at which $\boldsymbol{y}$ brings information about $\boldsymbol{\theta}$. The difference between that design and the fusion design of this paper is that there was no existing channel $\boldsymbol{x}$ to be fused with $\boldsymbol{y}$. Thus, the optimal channel matrix

	$MQ_{\theta\theta\|x}M^T$	G^*
(1)	$\frac{5}{6}I_5$	$\frac{1}{\sqrt{5}}I_5$
(2)	$\frac{5}{6}\mathrm{Diag}(25,16,9,4,1)$	$\mathrm{Diag}(0.316,0.294,0.249,0.141,0)$
(3)	$\frac{5}{6}\mathrm{Diag}(81,64,49,4,1)$	$\mathrm{Diag}(0.338,0.334,0.328,0,0)$

Table 1: The optimal channel matrices G^* for three scenarios.

G needs to maximize the extra information y contains about θ, taking account of the primary channel x.
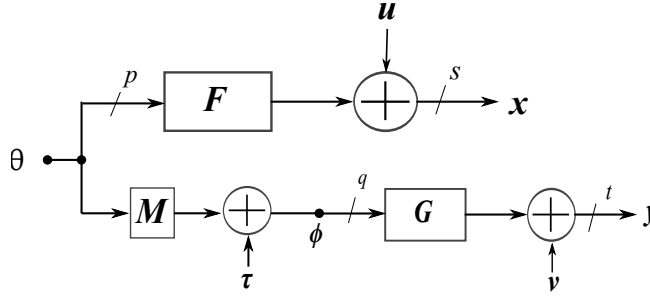

Figure 4: An alternative representation for the two-channel system.

4. Numerical Algorithms

In general, the constrained optimization problem (3) is not a convex problem [24] since the information gain $D(G)$ is not concave [13]. Notice that for any G with $\mathrm{tr}(GG^T) < P$, there exists $\widetilde{G} = \frac{\sqrt{P}}{\|G\|}G$ such that $\mathrm{tr}(\widetilde{G}\widetilde{G}^T) = P$ and $D(\widetilde{G}) \geq D(G)$. Therefore, it is sufficient to maximize the information gain on the boundary $\mathrm{tr}(GG^T) = P$. 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

14

between $\boldsymbol{\theta}$ and ϕ. 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 $\boldsymbol{G}$ as 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 $P = 1$.

4.1. Extrinsic Gradient Search Algorithm

Let $\nabla_{\boldsymbol{G}} D$ be the gradient of the information gain w.r.t $\boldsymbol{G}$, and it can be written as

$$\nabla_{\boldsymbol{G}} D = \boldsymbol{Q}_{\boldsymbol{vv}}^{-1} \boldsymbol{G}[(\boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}}^{-1} + \boldsymbol{G}^T \boldsymbol{Q}_{\boldsymbol{vv}}^{-1} \boldsymbol{G}) - \boldsymbol{B}]^{-1} \boldsymbol{B}(\boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}}^{-1} + \boldsymbol{G}^T \boldsymbol{Q}_{\boldsymbol{vv}}^{-1} \boldsymbol{G})^{-1}, \quad (9)$$

where $\boldsymbol{B} = \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}}^{-1} \boldsymbol{M} \boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}} \boldsymbol{M}^T (\boldsymbol{I}_q + \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}}^{-1} \boldsymbol{M} \boldsymbol{Q}_{\boldsymbol{\theta\theta}|\boldsymbol{x}} \boldsymbol{M}^T)^{-1} \boldsymbol{Q}_{\boldsymbol{\phi\phi}|\boldsymbol{\theta}}^{-1}$. See Appendix Appendix C for details. 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 $\mathrm{tr}(\boldsymbol{G}\boldsymbol{G}^T) = 1$ 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 $\boldsymbol{G}_0 \in \mathbb{R}^{t \times q}, \mathrm{tr}(\boldsymbol{G}_0 \boldsymbol{G}_0^T) = 1$.

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 $\mathrm{tr}(\boldsymbol{G}_{k+1} \boldsymbol{G}_{k+1}^T) = 1$, 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}.

4.2. Intrinsic Gradient Search Algorithm

Let $\boldsymbol{g}$ be the vectorization of matrix $\boldsymbol{G}$, denoted $\boldsymbol{g} = \mathrm{vec}(\boldsymbol{G})$. 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 $\mathrm{vec}(\boldsymbol{G}) = \boldsymbol{g}$. Under the power constraint $\mathrm{tr}(\boldsymbol{G}\boldsymbol{G}^T) = 1$, 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) is an optimization on the manifold S^{tq-1}. Note that S^{tq-1} is an embedded submanifold of $\mathbb{R}^{tq}$, and its geometry has been studied in [15]-[16].

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 Fig. 5.

16

Algorithm: Intrinsic Gradient Search

Input: Initial $\boldsymbol{g}_0 \in S^{tq-1}$

Output: Sequence of iterates $\{\boldsymbol{g}_k\}$.

for $k = 0, 1, 2, \ldots$ **do**

Select $\boldsymbol{g}_{k+1} = R_{\boldsymbol{g}_k}(\delta_k \boldsymbol{\eta}_{\boldsymbol{g}_k})$ where $\boldsymbol{\eta}_{\boldsymbol{g}_k} = (\boldsymbol{I}_{tq} - \boldsymbol{g}_k \boldsymbol{g}_k^T) \nabla_{\boldsymbol{g}_k} D$ 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 $\boldsymbol{\eta}_{\boldsymbol{g}}$, 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

$$R_{\boldsymbol{g}}(\boldsymbol{\eta}_{\boldsymbol{g}}) = \boldsymbol{g} \cos(\|\boldsymbol{\eta}_{\boldsymbol{g}}\|) + \frac{\boldsymbol{\eta}_{\boldsymbol{g}}}{\|\boldsymbol{\eta}_{\boldsymbol{g}}\|} \sin(\|\boldsymbol{\eta}_{\boldsymbol{g}}\|) \tag{10}$$

for any tangent vector $\boldsymbol{\eta}_{\boldsymbol{g}} \in T_{\boldsymbol{g}} S^{tq-1}$. For $\delta \geq 0$, $R_{\boldsymbol{g}}(\delta \boldsymbol{\eta}_{\boldsymbol{g}})$ 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 $\boldsymbol{\eta}_{\boldsymbol{g}}$. Given $\boldsymbol{g}_k$, $R_{\boldsymbol{g}}(\delta \boldsymbol{\eta}_{\boldsymbol{g}_k})$ is a periodic function of τ with period $2\pi/\|\boldsymbol{\eta}_{\boldsymbol{g}_k}\|$, thus the step size δ_k can be chosen within the interval $\delta \in [0, 2\pi/\|\boldsymbol{\eta}_{\boldsymbol{g}_k}\|)$ to maximize the information gain $D(R_{\boldsymbol{g}}(\delta \boldsymbol{\eta}_{\boldsymbol{g}_k}))$. 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.

4.3. A Numerical Study

Consider a two-channel system in (1) with $p = q = 4$ and $s = t = 3$. The input signals $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are characterized recursively as

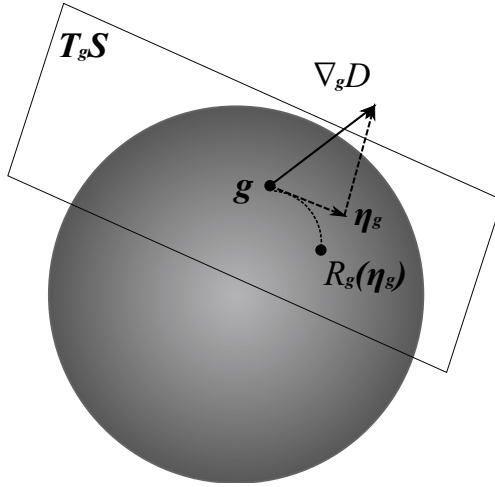$$\phi_i = \sum_{j=1}^{i} \rho^{i-j+1} \theta_j + \tau_i,$$

17

Figure 5: Projection of the Euclidean gradient to the tangent plane of unit sphere.

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 variance 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 Fig. 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 ϕ is increasing. Next, it can also be seen that the performance of the two algorithms are 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

18

suggested by the intrinsic algorithm in Section 4.2. While, for extrinsic algorithm, such strategy for the optimal step size is not available.
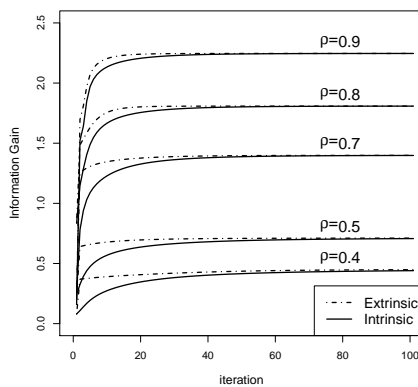

Figure 6: A numerical study. The x-axis is the index for iteration and the y-axis is the information gain obtained at each iteration. The solid curve is for the intrinsic algorithm and the dashed curve is for the extrinsic algorithm.

5. Discussion on Low-dimensional Channel Design

In the two-channel design problem considered in this paper, the number of measurements of the secondary channel, i.e., the number of rows of the channel matrix 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 ϕ, a q-dimensional output y is sufficient to achieve the maximum information gain, which is a consequence of the following lemma. Here we assume that the measurement noise v in the secondary channel is white noise.

19

Lemma 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 Appendix Appendix D.

Since the maximum rank of $\boldsymbol{G}$ is q, Lemma 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 algorithms. 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 2. *Consider a two-channel system* (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 Fig. 7 and Fig. 8, where the x-axis is k ($k = 1, \ldots, q$), number of rows of the secondary channel matrix $\boldsymbol{G}$, the y-axis on the left is the information gain for an k-dimensional secondary channel, and the y-axis on the right is the rank of the channel matrix with dimension $k \times q$.

In the first scenario (Fig. 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 (Fig. 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. Such difference may be caused by approximation error of the numerical algorithms.
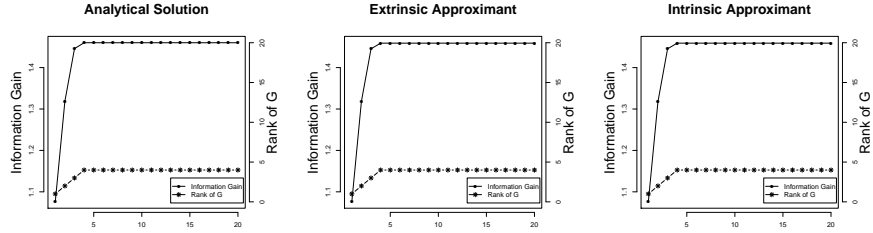
Figure 7: Choice of number of rows of the secondary channel matrix G. The three panels are associated with the channel matrices returned by the analytical solution (left), the extrinsic algorithm (middle) and the intrinsic algorithm (right), respectively. In each panel, the x-axis indicates the number 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 8: Choice of number of rows of the secondary channel matrix 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 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).
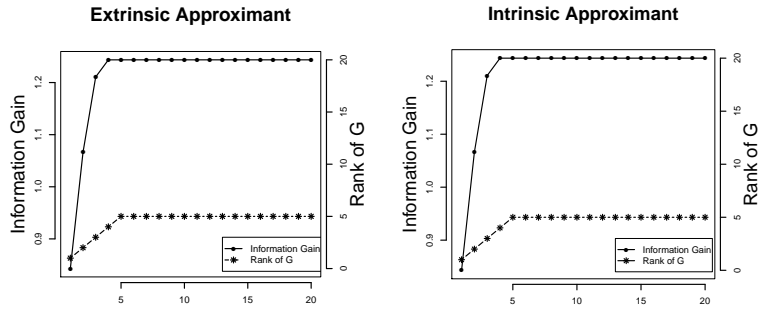
6. Conclusions

In this paper, 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.

Acknowledgements

Appendix A. Proof of Lemma 1

By the matrix inversion lemma $\boldsymbol{\Phi}^T(\boldsymbol{\Phi}\boldsymbol{\Phi}^T + \boldsymbol{\Sigma}_{\boldsymbol{v}})^{-1}\boldsymbol{\Phi} = \boldsymbol{I} - (\boldsymbol{I} + \boldsymbol{\Phi}\boldsymbol{\Sigma}_{\boldsymbol{v}}^{-1}\boldsymbol{\Phi}^T)^{-1}$, $D(\boldsymbol{\Phi})$ may be rewritten as

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

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

$$
\begin{aligned}
L(\boldsymbol{\Phi}; \mu) &= \frac{1}{2}\log\det(\boldsymbol{I}_n - (\boldsymbol{I}_n + \boldsymbol{\Phi}^T\boldsymbol{\Lambda}\boldsymbol{\Phi})^{-1}\boldsymbol{\Gamma}) + \mu(\mathrm{tr}(\boldsymbol{\Phi}\boldsymbol{\Phi}^T) - P) \\
&\quad + \frac{1}{2}\log\det[\boldsymbol{I} + \boldsymbol{\Sigma}_{\boldsymbol{\xi}}] \qquad\qquad\qquad\qquad\qquad\qquad\text{(A.1)}
\end{aligned}
$$

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

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

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

24

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, (A.1) can be rewritten as

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

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 $\phi_{i_1 j}$ and $\phi_{i_2 j}$, then $\phi_{i_1 k} = \phi_{i_2 k} = 0$ for any $k \neq j$ since each

25

row of $\mathbf{\Phi}$ has at most one non-zero entry. Hence, $[\mathbf{\Phi}\mathbf{\Phi}^T]_{i_1 i_2} = \sum_{k=1}^{n} \phi_{i_1 k}\phi_{i_2 k} = \phi_{i_1 j}\phi_{i_2 j} \neq 0$, which contradicts diagonal $\mathbf{\Phi}\mathbf{\Phi}^T$.

Appendix B. Proof of Theorem 1

Restricting the matrix $\mathbf{\Phi}$ within the class of matrices satisfying Lemma 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 $\lambda_{11}, \ldots, \lambda_{tt}$. The maximum information gain is taken over the permutations $\mathbf{\Pi}_1 \, \mathbf{\Pi}_2$ and $\lambda_{11}, \ldots, \lambda_{tt}$ subject to $\sum_{i=1}^{t} \lambda_{ii}^2 \leq P$.

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

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

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 $i = 1, \ldots, t$, one must have

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

26

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 $j = t+1, \ldots, q$. For $i = 1, \ldots, t$, it can be seen that $\pi_1(i)$ and $\pi_2(i)$ are pairwise. Therefore, we can set WLOG that $\pi_2(i) = i$ for $i = 1, \ldots, t$ and then search for the optimal permutation $\pi_1(i)$ to maximize

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

The proof that $\mathbf{\Pi}_1 = \mathbf{I}_q$ is the optimal permutation matrix is similar to the proof for Theorem 2 in [26]. 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:

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

The Lagrangian is

$$L(\lambda_{11}, \ldots, \lambda_{tt}, ; \mu) = \frac{1}{2}\sum_{i=1}^t \log\left(1 + \frac{\sigma_{\boldsymbol{\xi},\pi_1(i)}^2 \lambda_{ii}^2}{\lambda_{ii}^2 + \sigma_{\boldsymbol{v},\pi_2(i)}^2}\right) - \mu\left(\sum_{i=1}^t \lambda_{ii}^2 - P\right) \tag{B.3}$$

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

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

where $a_i = \sigma_{\boldsymbol{\xi},i}^2$ and $b_i = \sigma_{\boldsymbol{v},i}^2$. Equation (B.4) provides a feasible solution for λ_{ii} when $\mu \leq a_i/(2b_i)$. To see whether the solution is the maximizer for (B.2), we check the Hessian matrix. The second derivative of $\boldsymbol{L}$ w.r.t. λ_{ii} is

$$\frac{\partial^2 L(\mathbf{\Phi}; \mu)}{\partial \lambda_{ii}^2} = -2\mu + \frac{a_i b_i}{(\lambda_{ii}^2 + b_i)(\lambda_{ii}^2 + b_i + a_i\lambda_{ii}^2)}$$
$$- \frac{2a_i b_i \lambda_{ii}^2((\lambda_{ii}^2 + b_i)(2 + a_i) + a_i\lambda_{ii}^2)}{(\lambda_{ii}^2 + b_i)^2(\lambda_{ii}^2 + b_i + a_i\lambda_{ii}^2)^2} \tag{B.5}$$

For $i = 1 \ldots, \kappa$, upon substituting (B.4),

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

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

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

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 $\sum_{i=1}^{\kappa} \lambda_{ii}^2 = P$. Then, the maximizer of (B.2) is $\lambda_{11}^*, \ldots, \lambda_{tt}^*$ where

$$\lambda_{ii}^{*2} = \begin{cases} \frac{b_i \left(\sqrt{(2+a_i)^2 - 4(1+a_i)(1 - a_i/(2\mu b_i))} - (2+a_i) \right)}{2(1+a_i)}, & \text{for } i = 1, \ldots, \kappa \\ 0 & \text{for } i = \kappa + 1, \ldots, t. \end{cases} \quad \text{(B.6)}$$

Appendix C. Gradient of the Information Gain

Applying the matrix inversion lemma yields

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

Therefore, the information gain $D(\mathbf{G})$ satisfies

$$2D(\mathbf{G})$$
$$= \log \det[\mathbf{I}_q + (\mathbf{Q}_{\phi\phi|\theta}^{-1} - \mathbf{Q}_{\phi\phi|\theta}^{-1} (\mathbf{G}^T \mathbf{Q}_{vv}^{-1} \mathbf{G} + \mathbf{Q}_{\phi\phi|\theta}^{-1})^{-1} \mathbf{Q}_{\phi\phi|\theta}^{-1}) \mathbf{M} \mathbf{Q}_{\theta\theta|x} \mathbf{M}^T]$$
$$= \log \det[\mathbf{I}_q + \mathbf{Q}_{\phi\phi|\theta}^{-1} \mathbf{M} \mathbf{Q}_{\theta\theta|x} \mathbf{M}^T] + \log \det[\mathbf{I} - \mathbf{B}(\mathbf{G}^T \mathbf{Q}_{vv}^{-1} \mathbf{G} + \mathbf{Q}_{\phi\phi|\theta}^{-1})^{-1}]$$

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

Let $\mathbf{J}_{i,j}$ be a $t \times q$ matrix with value 1 at element (i, j) and 0 elsewhere. From [22], for a matrix $\mathbf{X}$, we have the partial derivatives

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

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

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

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

$$
\nabla_G D = Q_{vv}^{-1} G[(Q_{\phi\phi|\theta}^{-1} + G^T Q_{vv}^{-1} G) - B]^{-1} B (Q_{\phi\phi|\theta}^{-1} + G^T Q_{vv}^{-1} G)^{-1}.
$$

Appendix D. Proof of Lemma 2

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

$$
\begin{aligned}
&D(UG) \\
&= \frac{1}{2}\log\det[I_p + M^T G^T U^T (UGQ_{\phi\phi|\theta}G^T U^T + \sigma_v^2 I_t)^{-1} UGMQ_{\theta\theta|x}] \\
&= \frac{1}{2}\log\det[I_p + M^T G^T (GQ_{\phi\phi|\theta}G^T + \sigma_v^2 I_t)^{-1} GMQ_{\theta\theta|x}] \\
&= D(G)
\end{aligned}
$$

Therefore, the information gain $D(G)$ is invariant to left unitary multiplication of G. For any $G \in \mathbb{R}^{t\times q}$ with rank r, G has the singular value decomposition $G = U\Delta V^T$ where U and V are orthogonal matrices, and $\Delta \in \mathbb{R}^{t\times q}$ 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 WOLG that $\boldsymbol{U} = \boldsymbol{I}_t$. Let $\widetilde{\boldsymbol{G}} = \mathrm{Diag}(\boldsymbol{\Delta}_{1,1}, \ldots, \boldsymbol{\Delta}_{r,r})\boldsymbol{V}_r^T \in \mathbb{R}^{r \times q}$ 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 $\mathrm{tr}(\boldsymbol{G}\boldsymbol{G}^T) = \mathrm{tr}(\widetilde{\boldsymbol{G}}\widetilde{\boldsymbol{G}}^T)$. Moreover, one can easily check that

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

The RHS of (D.1) is the information gain brought by a r-dimensional channel $\widetilde{\boldsymbol{y}}$ as

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

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 $I(\boldsymbol{\theta}; \boldsymbol{x}, \boldsymbol{y}) - I(\boldsymbol{\theta}; \boldsymbol{x}) = I(\boldsymbol{\theta}; \boldsymbol{x}, \widetilde{\boldsymbol{y}}) - I(\boldsymbol{\theta}; \boldsymbol{x})$.

References

[1] D. P. Palomar, J. Cioffi, and M. A. Lagunas, "Joint TxRx beamforming design for multcarrier MIMO channels: A unifed framework for convex optimization," *IEEE Trans. Signal Process.*, vol. 51, no. 9, pp. 2381–2401, Sept. 2003.

[2] A. Scaglione, P. Stoica, S. Barbarossa, G. B. Giannakis, and H. Sampath, "Optimal designs for space-time linear precoders and decoders," *IEEE Trans. Signal Process.*, vol. 50, no. 5, pp. 1051–1064, May 2002.

[3] F. Pèrez-Cruz, M. R. D. Rodrigues, and S. Verdù, "MIMO Gaussian channels with arbitrary inputs: Optimal precoding and power allocation," in *IEEE Trans. Inf. Theory*, vol. 56, no. 3, pp. 1070–1084, Mar. 2010.

[4] W. R. Carson, M. R. D. Rodrigues, M. Chen, L. Carin, and R. Calder-bank, "Communications inspired compressive sensing," availabel at arXiv:1206.1973.

[5] A. Vosoughi and A. Scaglione, "Precoding and decoding paradigms for distributed vector data compression," *IEEE Trans. on Signal Process.*, vol. 55, no. 4, pp. 1445–1460, Apr. 2007.

[6] E. Liu, E. K. P. Chong, and L. L. Scharf, "Greedy adaptive compression in signal-plus-noise models," availabel at arXiv:1202.3913.

[7] M. Lamarca, "Linear precoding for mutual information maximization in MIMO systems," in *Intl. Symp. on Wireless Communication Systems (ISWCS '09),* pp. 26–30, Sept. 2009.

[8] H. Sampath, P. Stoica, and A. Paulraj, "Generalized linear precoder and decoder design for MIMO channels using the weighted MMSE criterion," *IEEE Trans. Communication*, vol. 49, no.12, pp. 2198–2206, Dec. 2001.

[9] C. Xiao, Y. R. Zheng, and Z. Ding, "Globally optimal linear precoders for finite alphabet signals over complex vector Gaussian channels," *IEEE Trans. Signal Process.,* vol. 59, no. 7, pp. 3301–3314, July 2011.

[10] D. Palomar and Y. Jiang, "MIMO transceiver design via majorization the-ory," *Now Publishers Inc*, 2007.

[11] D. Guo, S. Shamai, and S. Verdù, "Mutual information and minimum mean-square error in Gaussian channels," *IEEE Trans. Inf. Theory*, vol. 51, no. 4, pp. 1261–1282, Apr. 2005.

[12] D. Palomar and S. Verdù, "Gradient of mutual information in linear vector Gaussian channels," *IEEE Trans. Inf. Theory*, vol. 52, no. 1, pp. 141–154, Jan. 2006.

[13] M. Payarò and D. P. Palomar, "Hessian and concavity of mutual information, entropy, and entropy power in linear vector Gaussian channels," *IEEE Trans. Inf. Theory*, vol. 55, no. 8, pp. 3613–3628, Aug. 2009.

[14] D. P. Bertsekas "Projected Newton methods for optimization problems with simple constraints," *SIAM J. Control and Optimization*, vol. 20, no. 2, pp. 221–246, Mar. 1982.

[15] P. A. Absil, R. Mahony and R. Sepulchre, "Optimization algorithms on matrix manifolds," *Princeton University Press*, 2008.

[16] J. M. Lee, "Introduction to Smooth Manifolds," *Springer*, 2000.

[17] A. Edelman, T. A. Arias, and S. T. Smith, "The geometry of algorithms with orthogonality constraints," *SIAM J. Matrix Anal. Appl.*, vol. 20, no. 2, pp. 303–353, Oct. 1998.

[18] S. T. Smith, "Optimization techniques on Riemannian manifolds," *Hamiltonian and Gradient Flows, Algorithms, and Control*, American Mathematical Soc., pp. 113–136, 1994.

[19] D. Gabay, "Minimizing a differentiable function over a differential manifold," *INRIA Rocquencourt*, 1982.

[20] H. Witsenhausen, "A determinant maximization problem occuring in the

theory of data communication," *Jounal of Applied Mathematics*, vol. 29, no. 3, pp. 515–522, Nov. 1975.

[21] R. A. Horn and C. R. Johnson, "Matrix analysis," *Cambridge University Press*, 1985.

[22] K. B. Petersen and M. S. Pedersen, "The matrix cookbook," Nov., 2008.

[23] T. M. Cover and J. A. Thomas, "Elements of information theory," 2nd ed. *Wiley-Interscience*, 2005.

[24] S. Boyd and L. Vandenberghe, "Convex optimization", *Cambridge University Press*, 2004.

[25] A. Lozano, A.M. Tulino, and S. Verdú, "Optimum Power Allocation for Parallel Gaussian Channels with Arbitrary Input Distributions", *IEEE Trans. Inf. Theory*, vol. 52, no. 7, pp. 3033–3051, Jul. 2006.

[26] Y. Wang, H. Wang and L. L. Scharf, "Optimum Compression of a Noisy Measurement for Transmission over a Noisy Channel", *submitted*.