A Framework for Joint Design of Pilot Sequence and Linear Precoder

Adriano Pastore, *Member, IEEE*, Michael Joham, *Member, IEEE*,
and Javier Rodríguez Fonollosa, *Senior Member, IEEE*

***Abstract*—Most performance measures of pilot-assisted multiple-input multiple-output (MIMO) systems are functions of the linear precoder and the pilot sequence. A framework for the optimization of these two parameters is proposed, based on a matrix-valued generalization of the concept of effective signal-to-noise ratio (SNR) introduced in the famous work by Hassibi and Hochwald [1]. Our framework aims to extend [1] by allowing for transmit-side fading correlations, and by considering a class of utility functions of said effective SNR matrix, most notably including the well-known capacity lower bound used in [1]. We tackle the joint optimization problem by recasting the optimization of the precoder (resp. pilot sequence) subject to a fixed pilot sequence (resp. precoder) into a convex problem. Furthermore, we prove that joint optimality requires that the eigenbases of the precoder and pilot sequence be both aligned along the eigenbasis of the channel correlation matrix. We finally describe how to wrap all studied subproblems into an iteration that converges to a local optimum of the joint optimization.**

***Index Terms*—Channel estimation, mutual information, Rayleigh fading, block fading, wireless communications, MIMO precoding, pilot/training sequence optimization, power boost, joint pilot/precoder design.**

I. Introduction

WHEN the receiver has no *a priori* knowledge of the fading gains, a common approach is to incorporate a pre-agreed pattern of *training* or *pilot* symbols into the transmitted signal. The receiver first exploits the training observation to generate an estimate of the fading gains, and then uses this channel estimate to decode the transmitted message coherently (i.e., using the channel estimate as if it were the true channel). In terms of achievable rates, this two-stage approach is suboptimal compared to maximum-likelihood decoding, but it drastically reduces decoding complexity while maintaining near-optimal performance results. We shall consider a narrowband block-fading MIMO channel with time-duplexed training, that is, certain time slots are reserved exclusively for transmitting pilot symbols, while other time slots are reserved for data symbols. However, much of what is described in this article could be applied as well to wideband channels with frequency-duplexed training (pilot tones).

In our discrete-time Rayleigh block-fading model, the statistics of the channel gains are fully described by their second-order moments (the covariance of the fading coefficients) and by the fading-block length, also called the *coherence time*. Therefore, when the feedback is limited to being statistical (as in the scenario we shall consider), the pilot sequence and the precoder can only be designed based on these two channel parameters.

A frequent yet suboptimal choice for the pilot sequence is that of orthonormal pilot symbols. Furthermore, some publications such as [3]–[5] focus on minimizing distortion measures like the mean-square error when designing the pilot sequence (as in [6], [7]), while focusing on other measures such as bit-error rate or mutual information when designing the precoder. For instance in [4], [5], in an admittedly more complex system model, the authors propose two alternative figures of merit for optimizing the training sequence: the variance of the channel estimation error norm, and the average volume in which said error vector is concentrated. Though insightful, these are heuristic approaches. In light of this situation, it is of both practical and theoretical interest to examine how the pilot sequence and the precoder can be *jointly* designed with respect to the *same* performance metric.

In the present article, the metric of choice will be a well-known lower bound on the noncoherent capacity, based on a training routine combined with a lower bound on the Gaussian-input mutual information conditioned on the training observation. Said mutual information bound was originally proposed by Médard [8] for single-antenna channels and later generalized to MIMO systems by Hassibi and Hochwald [1]. In numerous variations and different settings, the bounding technique proposed in [1], [8] has been extensively used in subsequent work on transmission with imperfect channel-state information (e.g. in [3]–[5], [9]–[13], to cite only a few) as a performance metric for system design.

In [1], the problem was considered of finding the optimal time share between training and transmission, as well as optimally balancing the power levels of pilot and data symbol. Assuming independent and identically distributed (i.i.d.) coefficients for the channel estimate and the channel estimation error, [10] proves that the optimal transmit covariance is diagonal, and its eigenvalues are solutions to a convex problem. However, in both [1] and [10], all results were derived exclusively for i.i.d. fading. When facing the more difficult—

This work is supported by the "Ministerio de Economía y Competitividad" of the Spanish Government and ERDF funds (TEC2013-41315-R DISNET and TEC2015-69648-REDC Red COMONSENS) and the Catalan Government (2014 SGR 60 AGAUR). Parts of this work were published in [2].

A. Pastore was with the Department of Signal Processing and Communications, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain. He is now with the Laboratory of Information in Networked Systems, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland. M. Joham is with the Associate Institute for Signal Processing, Technical University of Munich, 80290 Munich, Germany. J.R. Fonollosa is with the Department of Signal Processing and Communications, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain.

yet more realistic—situation of *correlated* fading, the question of joint optimality of pilot sequence and precoder is much more involved. For example, one can intuit that the optimal number of pilot symbols and data streams will depend, among other things, on the conditioning of the channel's correlation structure. One central result of the present article is the proof that the respective left singular bases of the precoder and of the pilot sequence must be equal (up to column permutations) to the eigenvectors of the channel correlation matrix. While there is no reason to believe that this alignment property will still hold for arbitrary fading distributions (cf. Myth 1 in [14]), it is nonetheless expected in the case of correlated Rayleigh fading channels. In fact, our result is in line with similar results for Rayleigh fading found in [15]–[17], in which the optimal transmit directions were shown to be aligned with the channel eigenmodes.

We consider a single-user multiple-input multiple-output (MIMO) link and assume a highly scattering environment at the receiver—as is the case in many downlink scenarios—so that the fading is correlated only at the transmitter side. This encompasses the important special case of an arbitrarily correlated multiple-input single-output (MISO) link. In this setting with antenna correlations, we propose to revisit the problem treated in [1] for uncorrelated fading. In [1], the concept of *effective signal-to-noise ratio* (SNR) was introduced to designate an SNR that accounts for the imperfection of channel state information (CSI) at the receiver and serves as the utility to be maximized. In the present work, due to the presence of antenna correlation, the effective SNR needs to be extended from a scalar to a matrix-valued quantity, and the entire spatial structures of the pilot sequence and the linear precoder need to be jointly optimized. Thus, we are facing a multidimensional optimization problem which entails a non-trivial extension of the approach from [1].

A. Main Contributions

1) Regarding the two *separate* problems that consist in optimizing the precoder for a prescribed pilot sequence, and the pilot sequence for a prescribed precoder, it is shown that for the former problem, the precoded streams should not outnumber the pilot symbols [Theorem V.1,(21)], while for the latter problem, the pilot symbols should not outnumber the precoded streams [Theorem VI.1]. Furthermore, it is shown that both problems can be cast into (quasi-)convex problems, provided that the utility is (quasi-)concave [Theorem V.2 and Section VI]. The main utilities of interest exhibit this concavity [Table I in Appendix A].
2) Regarding the *joint* optimization problem, we prove that the number of streams should equal the number of pilot symbols [Section VII-B] and that the eigenbases of the transmit covariance and Gram form[1] of the pilot sequence should be equal to the eigenbasis of the channel's transmit correlation matrix [Theorem VII.1]. Loosely speaking, the pilot symbols and the transmit beamforming vectors should all be aligned in direction of the channel eigenmodes. This alignment significantly reduces the dimensionality of the problem.
3) To further reduce the remaining optimization of the eigenvalues of the transmit covariance and pilot Gram, a method is proposed in Section VII for computing eigenvalue vectors that are *jointly Pareto optimal* [in a sense explained in detail in Subsection VII-A].
4) All the above findings are combined constructively to yield an iterative procedure which converges towards a local optimum of the joint optimization problem.

[1]The Gram form of a matrix $\boldsymbol{A}$ shall be defined as $\boldsymbol{A}\boldsymbol{A}^\dagger$.

B. Outline

The article is structured as follows: Section II defines notation; Section III describes the system model; Section IV defines and motivates the class of utility functions considered in the optimization framework and states the optimization problem to consider; Sections V and VI describe how the optimization of the precoder (resp. pilot sequence) subject to a fixed pilot sequence (resp. precoder) is cast into a convex problem; Section VII specifies the jointly optimal training and transmit directions and shows how the residual problem of computing pilot-precoder power loading vectors that are Pareto optimal in terms of the effective SNR, can be formulated as a quasi-convex problem; Section VIII compiles prior results from Sections V, VI and VII into an iterative design algorithm that converges to a local joint optimum; Section IX presents simulations; the concluding Section X closes with a brief discussion on the value of our findings.

II. Notation

The operators $(\bullet)^{\mathrm{T}}$, $(\bullet)^*$ and $(\bullet)^\dagger$ denote the transpose, the complex conjugate, and the conjugate transpose (Hermitian adjoint) of a matrix, respectively. Matrix square roots are denoted as $(\bullet)^{\frac{1}{2}}$ and are assumed to be Hermitian. The Moore-Penrose pseudoinverse $\boldsymbol{A}^+$ of a matrix $\boldsymbol{A}$ is uniquely defined by the four identities

$$\boldsymbol{A}\boldsymbol{A}^+\boldsymbol{A} = \boldsymbol{A} \qquad \boldsymbol{A}^+\boldsymbol{A}\boldsymbol{A}^+ = \boldsymbol{A}^+$$
$$(\boldsymbol{A}\boldsymbol{A}^+)^\dagger = \boldsymbol{A}\boldsymbol{A}^+ \qquad (\boldsymbol{A}^+\boldsymbol{A})^\dagger = \boldsymbol{A}^+\boldsymbol{A}.$$

The range of a matrix $\boldsymbol{A}$, denoted as $\mathrm{range}(\boldsymbol{A})$, is the linear space spanned by its columns.

We will occasionally use the entrywise comparison $\boldsymbol{a} \leq \boldsymbol{b}$ between two real-valued vectors $\boldsymbol{a}$ and $\boldsymbol{b}$, of i-th entries a_i and b_i, defined as $\boldsymbol{a} \leq \boldsymbol{b} \Leftrightarrow \forall i\colon a_i \leq b_i$. If $\boldsymbol{A}$ and $\boldsymbol{B}$ denote Hermitian matrices, $\boldsymbol{A} \preceq \boldsymbol{B}$ (resp. $\boldsymbol{A} \prec \boldsymbol{B}$) means that $\boldsymbol{A} - \boldsymbol{B}$ is positive semidefinite (resp. positive definite). The cone of positive definite (resp. positive semidefinite) matrices from $\mathbb{C}^{n\times n}$ is denoted $\mathbb{C}^{n\times n}_{++}$ (resp. $\mathbb{C}^{n\times n}_{+}$).

The vectorization operator $\mathrm{vec}(\boldsymbol{A})$ takes a matrix $\boldsymbol{A} = [\boldsymbol{a}_1, \boldsymbol{a}_2, \dots]$ as argument, and returns a vector $\mathrm{vec}(\boldsymbol{A}) = [\boldsymbol{a}_1^{\mathrm{T}}, \boldsymbol{a}_2^{\mathrm{T}}, \dots]^{\mathrm{T}}$ containing the columns of $\boldsymbol{A}$ stacked on top of each other.

III. System model

Our discrete-time system consists of a standard single-user MIMO link with an $N_{\mathsf{R}} \times N_{\mathsf{T}}$ random channel matrix $\boldsymbol{H}$ expressible as

$$\boldsymbol{H} = \boldsymbol{W}\boldsymbol{R}^{\frac{1}{2}}, \tag{1}$$

where the entries of $\boldsymbol{W} \in \mathbb{C}^{N_\mathsf{R} \times N_\mathsf{T}}$ are independent and identically distributed (i.i.d.) zero-mean circularly-symmetric unit-variance complex Gaussian, i.e., $\mathrm{vec}(\boldsymbol{W}) \sim \mathcal{N}_\mathbb{C}(\mathbf{0}, \mathbf{I})$. The deterministic matrix $\boldsymbol{R} = \frac{1}{N_\mathsf{R}} \mathsf{E}[\boldsymbol{H}^\dagger \boldsymbol{H}]$ characterizes the transmit-side correlation, and is assumed to be full-rank, since we shall ignore keyhole effects. Equivalently, we can write the distribution of $\boldsymbol{H}$ as

$$\boldsymbol{h} = \mathrm{vec}(\boldsymbol{H}) \sim \mathcal{N}_\mathbb{C}(\mathbf{0}, \boldsymbol{R}^\mathrm{T} \otimes \mathbf{I}). \tag{2}$$

This correlation model is valid in setups where numerous scatterers are located in the vicinity of the receiver, and notably subsumes the case of arbitrarily correlated multiple-input single-output (MISO) channels, which are especially relevant in wireless downlinks.

The channel $\boldsymbol{H}$ remains constant for a duration T called the channel coherence time, after which it changes to a new realization that is independent of all previous ones (block-fading). Within every such fading block, we reserve T_p time slots to transmit a sequence of pilot symbols known at the receiver, while the data is transmitted during the remaining $T_\mathsf{d} = T - T_\mathsf{p}$ time slots. For example (and without loss of generality), we can accommodate the pilot symbols into the first T_p time slots of each fading block. The resulting training observation is

$$\boldsymbol{Y}_\mathsf{p} = \boldsymbol{H}\boldsymbol{X}_\mathsf{p} + \boldsymbol{Z}_\mathsf{p}, \tag{3}$$

where $\boldsymbol{X}_\mathsf{p} \in \mathbb{C}^{N_\mathsf{T} \times T_\mathsf{p}}$ is a matrix whose columns are the pilot symbols, and the noise matrix $\boldsymbol{Z}_\mathsf{p} \in \mathbb{C}^{N_\mathsf{R} \times T_\mathsf{p}}$ is distributed as $\mathrm{vec}(\boldsymbol{Z}_\mathsf{p}) \sim \mathcal{N}_\mathbb{C}(\mathbf{0}, \mathbf{I})$ and independent of $\boldsymbol{H}$. During the data transmission phase of one fading block, the received signal is

$$\boldsymbol{Y}_\mathsf{d} = \boldsymbol{H}\boldsymbol{F}\boldsymbol{X}_\mathsf{d} + \boldsymbol{Z}_\mathsf{d} \tag{4}$$

where $\boldsymbol{X}_\mathsf{d} \in \mathbb{C}^{r \times T_\mathsf{d}}$ (with $r \leq N_\mathsf{T}$) and $\boldsymbol{Z}_\mathsf{d} \in \mathbb{C}^{N_\mathsf{R} \times T_\mathsf{p}}$ are respectively the data symbol matrix and the additive noise matrix, with respective distributions $\mathrm{vec}(\boldsymbol{X}_\mathsf{d}) \sim \mathcal{N}_\mathbb{C}(\mathbf{0}, \mathbf{I})$ and $\mathrm{vec}(\boldsymbol{Z}_\mathsf{d}) \sim \mathcal{N}_\mathbb{C}(\mathbf{0}, \mathbf{I})$, whereas $\boldsymbol{F} \in \mathbb{C}^{N_\mathsf{T} \times r}$ is the linear precoder, and is assumed to have full column rank. The number r of rows of $\boldsymbol{X}_\mathsf{d}$ (and of columns of $\boldsymbol{F}$) represents the number of streams into which the channel input is multiplexed.

It can be shown that the minimum mean-square error (MMSE) channel estimate $\hat{\boldsymbol{H}} = \mathsf{E}[\boldsymbol{H}|\boldsymbol{Y}_\mathsf{p}]$ is obtained from the training observation $\boldsymbol{Y}_\mathsf{p}$ by right-multiplying it with the estimator matrix $\boldsymbol{G} = (\boldsymbol{X}_\mathsf{p}^\dagger \boldsymbol{R} \boldsymbol{X}_\mathsf{p} + \mathbf{I})^{-1} \boldsymbol{X}_\mathsf{p}^\dagger \boldsymbol{R}$ as

$$\hat{\boldsymbol{H}} = \boldsymbol{Y}_\mathsf{p}\boldsymbol{G}. \tag{5}$$

As a consequence of the distributions and correlation models for $\boldsymbol{H}$ and $\boldsymbol{Z}_\mathsf{p}$, the respective marginal distributions of the estimate $\hat{\boldsymbol{H}}$ and of the estimation error $\tilde{\boldsymbol{H}} = \boldsymbol{H} - \hat{\boldsymbol{H}}$ turn out to be

$$\hat{\boldsymbol{h}} = \mathrm{vec}(\hat{\boldsymbol{H}}) \sim \mathcal{N}_\mathbb{C}(\mathbf{0}, \hat{\boldsymbol{R}}^\mathrm{T} \otimes \mathbf{I}) \tag{6a}$$

$$\tilde{\boldsymbol{h}} = \mathrm{vec}(\tilde{\boldsymbol{H}}) \sim \mathcal{N}_\mathbb{C}(\mathbf{0}, \tilde{\boldsymbol{R}}^\mathrm{T} \otimes \mathbf{I}) \tag{6b}$$

with respective transmit-side covariances

$$\hat{\boldsymbol{R}} = \tfrac{1}{N_\mathsf{R}} \mathsf{E}[\hat{\boldsymbol{H}}^\dagger \hat{\boldsymbol{H}}] = \boldsymbol{R} - \tilde{\boldsymbol{R}} \tag{7a}$$

$$\tilde{\boldsymbol{R}} = \tfrac{1}{N_\mathsf{R}} \mathsf{E}[\tilde{\boldsymbol{H}}^\dagger \tilde{\boldsymbol{H}}] = (\boldsymbol{R}^{-1} + \boldsymbol{X}_\mathsf{d}\boldsymbol{X}_\mathsf{d}^\dagger)^{-1}. \tag{7b}$$

We further define the two $N_\mathsf{T} \times N_\mathsf{T}$ Gram matrices

$$\boldsymbol{Q} = \boldsymbol{F}\boldsymbol{F}^\dagger \qquad\qquad \boldsymbol{P} = \boldsymbol{X}_\mathsf{p}\boldsymbol{X}_\mathsf{p}^\dagger \tag{8}$$

which are called the *transmit covariance*[2] and the *pilot Gram*, respectively. These two matrices will be subject to optimization.

IV. Problem statement

A capacity lower bound (in bits per channel use) of the communication system is given by $\frac{1}{T} I(\boldsymbol{X}_\mathsf{d}; \boldsymbol{Y}_\mathsf{d} | \boldsymbol{Y}_\mathsf{p})$, where $I(\boldsymbol{X}_\mathsf{d}; \boldsymbol{Y}_\mathsf{d} | \boldsymbol{Y}_\mathsf{p})$ stands for the mutual information between the block of input symbols $\boldsymbol{X}_\mathsf{d}$ and their corresponding outputs $\boldsymbol{Y}_\mathsf{d}$, conditioned on the side information $\boldsymbol{Y}_\mathsf{p}$ (training observation). Using a well-known lower bound on this mutual information, we get the achievable rate expression

$$\frac{T_\mathsf{d}}{T} \mathsf{E}\left[\log\det\left(\mathbf{I} + \frac{\hat{\boldsymbol{H}}\boldsymbol{Q}\hat{\boldsymbol{H}}^\dagger}{1 + \mathrm{tr}(\tilde{\boldsymbol{R}}\boldsymbol{Q})}\right)\right] \leq \frac{1}{T} I(\boldsymbol{X}_\mathsf{d}; \boldsymbol{Y}_\mathsf{d} | \boldsymbol{Y}_\mathsf{p}), \tag{9}$$

where $\hat{\boldsymbol{H}}$ is given in (5). This bound is based on a worst-case noise approach [1] and has been widely used and studied in the literature, e.g., [8], [1], [18], [5] and many more. In [3], the bound (9) is used in the exact same system setup as here. Using (6a), we can write $\hat{\boldsymbol{H}} = \hat{\boldsymbol{W}}\hat{\boldsymbol{R}}^{\frac{1}{2}}$ with $\mathrm{vec}(\hat{\boldsymbol{W}}) \sim \mathcal{N}_\mathbb{C}(\mathbf{0}, \mathbf{I})$ so that the expectation in (9) reads simply

$$I(\boldsymbol{S}) \triangleq \mathsf{E}\left[\log\det\left(\mathbf{I} + \hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger\right)\right]. \tag{10}$$

The latter is a function of the matrix argument

$$\boldsymbol{S} = \boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q}) = \frac{\hat{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{Q}\hat{\boldsymbol{R}}^{\frac{1}{2}}}{1 + \mathrm{tr}(\tilde{\boldsymbol{R}}\boldsymbol{Q})}, \tag{11}$$

which constitutes a matrix-valued *effective SNR* generalization of the scalar effective SNR introduced by Hassibi and Hochwald in [1]. We write $\boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q})$ instead of $\boldsymbol{S}$ whenever we wish to emphasize the dependency on $(\boldsymbol{P}, \boldsymbol{Q})$. Let $\boldsymbol{s}$ denote the *(eigenvalue) profile* of $\boldsymbol{S}$, i.e., the vector of non-increasingly ordered eigenvalues of $\boldsymbol{S}$. Since $\hat{\boldsymbol{W}}$ and $\hat{\boldsymbol{W}}\boldsymbol{U}$ have the same marginal distribution for any unitary $\boldsymbol{U}$, the function $I(\boldsymbol{S})$ is in fact a function of the unordered eigenvalues of $\boldsymbol{S}$, i.e., a symmetric function of $\boldsymbol{s}$. We denote the latter as $\breve{I}(\boldsymbol{s})$. We thus have

$$\breve{I}(\boldsymbol{s}) \triangleq \mathsf{E} \log\det\big(\mathbf{I} + \hat{\boldsymbol{W}}\,\mathrm{diag}(\boldsymbol{s})\hat{\boldsymbol{W}}^\dagger\big) \tag{12}$$

with $I(\boldsymbol{S}) = \breve{I}(\boldsymbol{s})$. We will prefer one notation over the other depending on the situation.

To keep derivations as general as possible, we consider a general class $\mathcal{F}$ of utility functions sharing similar properties with I (resp. $\breve{I}$). The class $\mathcal{F}$ (resp. $\bar{\mathcal{F}}$) shall be the set of functions which, like I (resp. $\breve{I}$), are unitarily invariant and matrix-monotonic (resp. symmetric and vector-monotonic). That is, for $f \in \mathcal{F}$ (resp. $\bar{f} \in \bar{\mathcal{F}}$) we have

- $f(\boldsymbol{S}) = f(\boldsymbol{U}\boldsymbol{S}\boldsymbol{U}^\dagger)$ for any unitary matrix $\boldsymbol{U}$
- $\mathbf{0} \preceq \boldsymbol{S} \preceq \breve{\boldsymbol{S}} \Rightarrow f(\boldsymbol{S}) \leq f(\breve{\boldsymbol{S}})$

and

[2] in reference to the fact that $\boldsymbol{F}\boldsymbol{F}^\dagger$ is the covariance of the transmitted data signal

- $\bar{f}(\boldsymbol{\Pi s}) = \bar{f}(\boldsymbol{s})$ for any permutation $\boldsymbol{\Pi}$
- $\mathbf{0} \leq \boldsymbol{s} \leq \boldsymbol{s}' \Rightarrow \bar{f}(\boldsymbol{s}) \leq \bar{f}(\boldsymbol{s}')$.

We define the trace-constrained sets

$$\mathcal{P}(\mu_{\mathcal{P}}) = \left\{ \boldsymbol{P} \in \mathbb{C}_{+}^{N_{\mathsf{T}} \times N_{\mathsf{T}}} \colon \operatorname{tr}(\boldsymbol{P}) \leq \mu_{\mathcal{P}} \right\} \tag{13a}$$

$$\mathcal{Q}(\mu_{\mathcal{Q}}) = \left\{ \boldsymbol{Q} \in \mathbb{C}_{+}^{N_{\mathsf{T}} \times N_{\mathsf{T}}} \colon \operatorname{tr}(\boldsymbol{Q}) \leq \mu_{\mathcal{Q}} \right\} \tag{13b}$$

where $\mu_{\mathcal{P}}$ and $\mu_{\mathcal{Q}}$ stand for the maximum overall pilot symbol energy, and the average data symbol power (energy per time unit), respectively.

The pilot-assisted system shall be constrained by a maximum average energy consumption per time unit, denoted as μ. Consequently, the set of admissible values of the pilot-precoder pair $(\boldsymbol{P}, \boldsymbol{Q})$ is

$$\mathcal{PQ}(T_{\mathsf{p}}) \triangleq \bigcup_{\substack{\mu_{\mathcal{P}}, \mu_{\mathcal{Q}} \geq 0 \\ \mu_{\mathcal{P}} + (T - T_{\mathsf{p}})\mu_{\mathcal{Q}} = T\mu}} \mathcal{P}(\mu_{\mathcal{P}}) \times \mathcal{Q}(\mu_{\mathcal{Q}}). \tag{14}$$

For future reference, we define five different pilot/precoder optimization problems (of increasing number of variables to be optimized) for some given utility $F \in \mathcal{F}$: [3]

$$\text{(P.1.a)} \quad \max_{\boldsymbol{Q} \in \mathcal{Q}(\mu_{\mathcal{Q}})} F(\boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q})) \tag{15a}$$

$$\text{(P.1.b)} \quad \max_{\boldsymbol{P} \in \mathcal{P}(\mu_{\mathcal{P}})} F(\boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q})) \tag{15b}$$

$$\text{(P.2)} \quad \max_{(\boldsymbol{P}, \boldsymbol{Q}) \in \mathcal{P}(\mu_{\mathcal{P}}) \times \mathcal{Q}(\mu_{\mathcal{Q}})} F(\boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q})) \tag{15c}$$

$$\text{(P.3)} \quad \max_{(\boldsymbol{P}, \boldsymbol{Q}) \in \mathcal{PQ}(T_{\mathsf{p}})} F(\boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q})) \tag{15d}$$

$$\text{(P.4)} \quad \max_{T_{\mathsf{p}} \in \{1, \ldots, T-1\}} \max_{(\boldsymbol{P}, \boldsymbol{Q}) \in \mathcal{PQ}(T_{\mathsf{p}})} F(\boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q})) \tag{15e}$$

Problems (P.1.a) and (P.1.b) are the partial problems that consist in optimizing one among the two variables $\boldsymbol{P} \in \mathcal{P}(\mu_{\mathcal{P}})$ and $\boldsymbol{Q} \in \mathcal{Q}(\mu_{\mathcal{Q}})$, while the other variable has a fixed value. The parameters $\mu_{\mathcal{P}}$ and $\mu_{\mathcal{Q}}$ can be considered as arbitrary constants. Problem (P.2) is the simplest joint optimization problem, having two independent trace constraints on $\boldsymbol{P}$ and $\boldsymbol{Q}$. Adding on (P.2) an outer optimization over pairs $(\mu_{\mathcal{P}}, \mu_{\mathcal{Q}})$ fulfilling the weighted-sum constraint $\mu_{\mathcal{P}} + (T - T_{\mathsf{p}})\mu_{\mathcal{Q}} = T\mu$, we obtain Problem (P.3). This balancing of pilot symbol and data symbol energies is what we call *energy boost*[4]. Adding yet another optimization over the training duration T_{p}, and incorporating an overhead factor $(T - T_{\mathsf{p}})/T$ which accounts for the loss of spectral efficiency due to the pilot symbols, we obtain the full-fledged Problem (P.4). Step by step, the present article builds up a procedure for tackling (P.4).

The optimization of T_{p} in (P.4) is over a finite set and is solved by an exhaustive search, thus we will leave it aside until Section VIII. Notice that in [1] the authors posit for the same (though uncorrelated) channel model (and utility $f = I$) that the receiver should have a representative estimate of the *complete* channel state, described by $N_{\mathsf{T}} N_{\mathsf{R}}$ fading coefficients. Therefore, they assume that the training duration T_{p} should be at least the number of transmit antennas N_{T}, so as to generate at least as many observables as there are coefficients to estimate. However, in the case where only a limited number of data streams are to be precoded, it might be more economic to only estimate a properly chosen subspace of the channel covariance spanned by the stronger eigenmodes. In fact, since T_{p} is defined as the number of columns of the pilot matrix $\boldsymbol{X}_{\mathsf{p}}$, and given that the utility function and power constraint [cf. (14)–(15e)] depend on $\boldsymbol{X}_{\mathsf{p}}$ only via its Gram matrix $\boldsymbol{P} = \boldsymbol{X}_{\mathsf{p}} \boldsymbol{X}_{\mathsf{p}}^{\dagger}$, we can assume that $\boldsymbol{X}_{\mathsf{p}}$ has full column rank and set the training duration equal to the rank of $\boldsymbol{P}$, i.e., $T_{\mathsf{p}} = \operatorname{rank}(\boldsymbol{P}) \leq N_{\mathsf{T}}$, and accordingly reduce the search interval in (15e) from $\{1, \ldots, T-1\}$ down to $\{1, \ldots, \min(T-1, N_{\mathsf{T}})\}$.

We will start by studying Problems (P.1.a) and (P.1.b) in the next two sections. These individual optimizations will form two building blocks of an algorithmic approach that aims to solve (P.3), and eventually (P.4). However, they may also be considered as two stand-alone problems in their own right.

V. Precoder design for prescribed pilots

In this section, we consider the optimization of the transmit covariance $\boldsymbol{Q}$ alone, while the pilot Gram $\boldsymbol{P}$ has a fixed value [Problem (P.1.a)]. The optimal transmit covariance is

$$\boldsymbol{Q}^{\star}(\boldsymbol{P}) = \operatorname*{argmax}_{\boldsymbol{Q} \in \mathcal{Q}(\mu_{\mathcal{Q}})} F(\boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q})). \tag{16}$$

A. Number of streams and pilot symbols

Recall that the ranks of $\boldsymbol{P}$ and $\boldsymbol{Q}$ represent the number of pilot symbols and precoded streams, respectively. To establish a relation between them, we first need to uncover an important property of the range space of $\boldsymbol{Q}^{\star}(\boldsymbol{P})$.

Theorem V.1. *For any utility $F \in \mathcal{F}$ and a prescribed pilot Gram $\boldsymbol{P}$, the range space of the optimal transmit covariance $\boldsymbol{Q}^{\star}(\boldsymbol{P})$ must be contained in the range space of the channel estimate covariance $\hat{\boldsymbol{R}}$:*

$$\operatorname{range}(\boldsymbol{Q}^{\star}(\boldsymbol{P})) \subseteq \operatorname{range}(\hat{\boldsymbol{R}}). \tag{17}$$

Proof: See Appendix B. ∎

Complementing Theorem V.1, notice that the rank equality

$$\operatorname{rank}(\hat{\boldsymbol{R}}) = \operatorname{rank}(\boldsymbol{P}) \tag{18}$$

always holds. This is easily seen by application of the matrix inversion lemma:

$$\begin{aligned} \hat{\boldsymbol{R}} &= \boldsymbol{R} - (\boldsymbol{R}^{-1} + \boldsymbol{X}_{\mathsf{p}} \boldsymbol{X}_{\mathsf{p}}^{\dagger})^{-1} \\ &= \boldsymbol{R} \boldsymbol{X}_{\mathsf{p}} (\mathbf{I} + \boldsymbol{X}_{\mathsf{p}}^{\dagger} \boldsymbol{R} \boldsymbol{X}_{\mathsf{p}})^{-1} \boldsymbol{X}_{\mathsf{p}}^{\dagger} \boldsymbol{R}. \end{aligned} \tag{19}$$

Since $\boldsymbol{X}_{\mathsf{p}}$ has full column rank by assumption [cf. Section IV], it becomes manifest that the rank of $\hat{\boldsymbol{R}}$ equals the number of columns of $\boldsymbol{X}_{\mathsf{p}}$, which is equal to $T_{\mathsf{p}} = \operatorname{rank}(\boldsymbol{P})$, hence (18).

Combining (18) with Theorem V.1 directly implies the rank inequality

$$\operatorname{rank}(\boldsymbol{Q}^{\star}(\boldsymbol{P})) \leq \operatorname{rank}(\hat{\boldsymbol{R}}) = \operatorname{rank}(\boldsymbol{P}), \tag{20}$$

[3] Obviously, these problems could be equivalently stated in terms of $\bar{f} \in \bar{\mathcal{F}}$.

[4] In the literature, the balancing between pilot/data symbol *powers* under an overall average power constraint and for fixed time fractions assigned to training and data transmission, is sometimes referred to as *power boost* (e.g., [12]). Since in our setup, $\mu_{\mathcal{P}}$ represents a pilot *energy* budget, and given that the training duration T_{p} is not fixed (but subject to an outer optimization), we prefer the term *energy boost*.

or in words,

$$\textbf{number of streams} \leq \textbf{number of pilot symbols} \tag{21}$$

The idea behind the proof of Theorem V.1 is that, if $\boldsymbol{Q}^\star(\boldsymbol{P})$ had eigenvectors (transmit directions) lying outside the range space of the estimate covariance $\hat{\boldsymbol{R}}$, then the transmitter would be radiating some of its transmit power into channel directions of which the receiver has no estimate (and thus cannot detect coherently), thus incurring a waste of power. As a particular consequence, (20) tells us that the number of precoded streams should never exceed the number of training symbols.

B. Feasible effective SNR matrices

It is convenient to reformulate Problem (P.1.a) as

$$\boldsymbol{S}^\star(\boldsymbol{P}) = \underset{\boldsymbol{S}\in\boldsymbol{S}(\boldsymbol{P},\mathcal{Q}(\mu_\mathcal{Q}))}{\operatorname{argmax}} F(\boldsymbol{S}) \tag{22}$$

in terms of effective SNR matrices, which belong to a feasible set

$$\boldsymbol{S}(\boldsymbol{P},\mathcal{Q}(\mu_\mathcal{Q})) = \big\{\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q}) \;\big|\; \boldsymbol{Q}\in\mathcal{Q}(\mu_\mathcal{Q})\big\}. \tag{23}$$

We will show that $\boldsymbol{S}(\boldsymbol{P},\mathcal{Q}(\mu_\mathcal{Q}))$ is a convex set. For this purpose, observe that in the expression of the function

$$\boldsymbol{Q}\mapsto\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q}) = \frac{\hat{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{Q}\hat{\boldsymbol{R}}^{\frac{1}{2}}}{1+\operatorname{tr}(\boldsymbol{Q}\tilde{\boldsymbol{R}})}, \tag{24}$$

the argument $\boldsymbol{Q}$ appears in the matrix-valued numerator, and inside a trace operator in the denominator. This function $\boldsymbol{Q}\mapsto\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q})$ is thus reminiscent of fractions of monomials such as $q\mapsto\frac{aq}{1+bq}$, except that it is defined for matrices. In fact, the function $\boldsymbol{Q}\mapsto\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q})$ pertains to what can be defined in the following Definition V.1 as a generalization of linear fractional functions. The latter are commonly defined for the scalar case (e.g., [19, Sec. 2.3.3]).

Definition V.1. *Let $\mathcal{X}\subset\mathbb{C}^{n\times n}$ denote a set of Hermitian matrices of size $n\times n$ whose elements $\boldsymbol{X}\in\mathcal{X}$ satisfy $\operatorname{tr}(\boldsymbol{B}\boldsymbol{X})\neq-1$ with some given Hermitian matrix $\boldsymbol{B}\in\mathbb{C}^{n\times n}$. A function $\boldsymbol{X}\mapsto\phi(\boldsymbol{X};\boldsymbol{A},\boldsymbol{B})$ that is defined as*

$$\mathcal{X}\;\to\;\mathbb{C}^{m\times m},\;\boldsymbol{X}\;\mapsto\;\phi(\boldsymbol{X};\boldsymbol{A},\boldsymbol{B}) = \frac{\boldsymbol{A}\boldsymbol{X}\boldsymbol{A}^\dagger}{1+\operatorname{tr}(\boldsymbol{B}\boldsymbol{X})} \tag{25}$$

shall be called a linear fractional *function with parameters $\boldsymbol{A}\in\mathbb{C}^{m\times n}$ and $\boldsymbol{B}\in\mathbb{C}^{n\times n}$.*

Note that the Hermitianity of $\boldsymbol{B}$ and of the argument $\boldsymbol{X}$ ensures the Hermitianity of the image $\phi(\boldsymbol{X};\boldsymbol{A},\boldsymbol{B})$. Linear fractional functions may or may not be injective functions, depending on the properties of the parameter $\boldsymbol{A}$. Let $\boldsymbol{A}^\sharp = (\boldsymbol{A}^\dagger\boldsymbol{A})^{-1}\boldsymbol{A}^\dagger$ denote the left pseudoinverse of $\boldsymbol{A}$, and $\boldsymbol{A}^\flat = \boldsymbol{A}^\dagger(\boldsymbol{A}\boldsymbol{A}^\dagger)^{-1}$ denote the right pseudoinverse of $\boldsymbol{A}$.

Lemma V.1. *The linear fractional function $\boldsymbol{X}\mapsto\phi(\boldsymbol{X};\boldsymbol{A},\boldsymbol{B})$ from Definition V.1 is injective (one-to-one) if one at least of the following two conditions apply:*

1) *The parameter $\boldsymbol{A}$ has full column rank*
2) *The parameter $\boldsymbol{A}$ has full row rank and the domain $\mathcal{X}$ is such that $\forall\boldsymbol{X}\in\mathcal{X}\colon\ \operatorname{range}(\boldsymbol{X}) = \operatorname{range}(\boldsymbol{A}^\dagger)$.*[5]

In these two respective cases, its inverse function $\phi^{-1}\colon\phi(\mathcal{X};\boldsymbol{A},\boldsymbol{B})\to\mathcal{X},\boldsymbol{Y}\mapsto\phi^{-1}(\boldsymbol{Y};\boldsymbol{A},\boldsymbol{B})$ is

1) *linear fractional with parameters $\boldsymbol{A}^\sharp$ and $-\boldsymbol{A}^{\sharp\dagger}\boldsymbol{B}\boldsymbol{A}^\sharp$, i.e., $\phi^{-1}(\bullet;\boldsymbol{A},\boldsymbol{B}) = \phi(\bullet;\boldsymbol{A}^\sharp,-\boldsymbol{A}^{\sharp\dagger}\boldsymbol{B}\boldsymbol{A}^\sharp)$.*
2) *linear fractional with parameters $\boldsymbol{A}^\flat$ and $-\boldsymbol{A}^{\flat\dagger}\boldsymbol{B}\boldsymbol{A}^\flat$, i.e., $\phi^{-1}(\bullet;\boldsymbol{A},\boldsymbol{B}) = \phi(\bullet;\boldsymbol{A}^\flat,-\boldsymbol{A}^{\flat\dagger}\boldsymbol{B}\boldsymbol{A}^\flat)$.*

Proof: See Appendix C. ∎

In the following we will optimize $\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q})$ rather than $\boldsymbol{Q}$, and consider that Lemma V.1 can be used to recover the optimal transmit covariance $\boldsymbol{Q}^\star(\boldsymbol{P})$ from the optimal $\boldsymbol{S}$, by means of the appropriate inverse linear fractional function.

Prescribing the pilot Gram $\boldsymbol{P}$ means that the matrices $\tilde{\boldsymbol{R}} = (\boldsymbol{R}^{-1}+\boldsymbol{P})^{-1}$ and $\hat{\boldsymbol{R}} = \boldsymbol{R}-\tilde{\boldsymbol{R}}$ are prescribed. Therefore, the function $\boldsymbol{Q}\mapsto\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q})$ as given in (24) is linear fractional with parameters $\boldsymbol{A} = \hat{\boldsymbol{R}}^{\frac{1}{2}}$ and $\boldsymbol{B} = \tilde{\boldsymbol{R}}$, i.e.,

$$\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q}) = \phi(\boldsymbol{Q};\hat{\boldsymbol{R}}^{\frac{1}{2}},\tilde{\boldsymbol{R}}). \tag{26}$$

The key property of linear fractional functions that we need for understanding Problem (P.1.a) is that they preserve the linearity of segments.

Lemma V.2. *An injective linear fractional function $\varphi(\bullet) = \phi(\bullet;\boldsymbol{A},\boldsymbol{B})$ with some given parameters $\boldsymbol{A}$ and $\boldsymbol{B}$ uniquely maps linear segments onto linear segments in a one-to-one manner, i.e.,*

$$\begin{aligned}&\forall(\boldsymbol{X}_1,\boldsymbol{X}_2,\alpha)\in\mathcal{X}^2\times[0;1],\exists\beta\in[0;1]\colon\\ &\varphi(\alpha\boldsymbol{X}_1+(1-\alpha)\boldsymbol{X}_2) = \beta\varphi(\boldsymbol{X}_1)+(1-\beta)\varphi(\boldsymbol{X}_2).\end{aligned} \tag{27}$$

Proof: This is readily verified by inserting

$$\beta = \frac{\alpha(1+\operatorname{tr}(\boldsymbol{B}\boldsymbol{X}_1))}{1+\alpha\operatorname{tr}(\boldsymbol{B}\boldsymbol{X}_1)+(1-\alpha)\operatorname{tr}(\boldsymbol{B}\boldsymbol{X}_2)} \tag{28}$$

into the equality (27). ∎

Figure 1 symbolically depicts the behavior of linear fractional functions: a convex combination of two points is mapped onto a convex combination of the respective images of said points, thus preserving segments. They are not linear functions though, because α and β can be different.

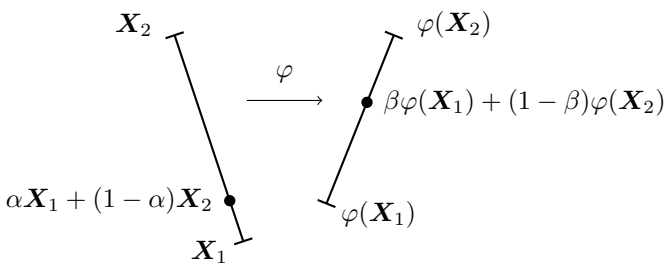

Fig. 1. Linear fractional functions preserve segments

[5] In case $\boldsymbol{A}$ has neither full column nor full row rank, one can bring the problem back to one of the two considered cases by an appropriate rank reduction.

Corollary V.1. *Linear fractional mappings preserve set convexity.*

Proof: Take a pair $(\boldsymbol{X}_1, \boldsymbol{X}_2) \in \mathcal{X}^2$ with a convex $\mathcal{X}$. According to Lemma V.2, any convex combination of $\boldsymbol{X}_1$ and $\boldsymbol{X}_2$ is mapped onto a convex combination of $\varphi(\boldsymbol{X}_1)$ and $\varphi(\boldsymbol{X}_2)$. Therefore, the codomain $\varphi(\mathcal{X})$ is convex. ∎

As a consequence, $\boldsymbol{S}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))$ is a convex set because $\mathcal{Q}(\mu_{\mathcal{Q}})$ is convex [cf. (13b)]. So if a utility F is concave in $\boldsymbol{S}$, then Problem (P.1.a) in its formulation (22) is convex. The optimal transmit covariance $\boldsymbol{Q}^\star(\boldsymbol{P})$ is then computed from $\boldsymbol{S}^\star(\boldsymbol{P})$ by means of the appropriate inverse linear fractional function (cf. Lemma V.1). More generally speaking, if F is quasi-concave in $\boldsymbol{S}$, then the problem (16) can be recast into a convex problem by an appropriate transformation. The mutual information I is one example of a concave utility. Other examples of concave or log-concave (quasi-concave) utilities are given in Table I in Appendix A.

C. Feasible effective SNR eigenvalues

It is convenient to rewrite Problem (P.1.a) as

$$\boldsymbol{s}^\star(\boldsymbol{P}) = \operatorname*{argmax}_{\boldsymbol{s} \in \boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))} \bar{f}(\boldsymbol{s}), \tag{29}$$

in terms of effective SNR eigenvalue profiles, which are to be searched in a feasible set

$$\boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}})) = \big\{ \boldsymbol{s}(\boldsymbol{P}, \boldsymbol{Q}) \;\big|\; \boldsymbol{Q} \in \mathcal{Q}(\mu_{\mathcal{Q}}) \big\}. \tag{30}$$

In the previous subsection, we have shown that the set $\boldsymbol{S}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))$ is convex. Note that this convexity, however, does not generally imply (nor is implied by) the convexity of the set $\boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))$, hence establishing the convexity of $\boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))$ requires a separate proof. Indeed, we show in the following that $\boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))$ is also convex and has a simplex shape, whose vertices are characterized by Theorem V.2 below.

Let ω_i denote the non-increasingly ordered eigenvalues of the generalized eigenvalue problem

$$\hat{\boldsymbol{R}} \boldsymbol{v}_i = \omega_i \big(\mu_{\mathcal{Q}}^{-1} \mathbf{I} + \tilde{\boldsymbol{R}} \big) \boldsymbol{v}_i. \tag{31}$$

Due to $\operatorname{rank}(\hat{\boldsymbol{R}}) = \operatorname{rank}(\boldsymbol{P})$ [cf. (18)], only the first $r_{\boldsymbol{P}} = \operatorname{rank}(\boldsymbol{P})$ eigenvalues ω_i are different from zero. Let $\boldsymbol{\lambda}(\boldsymbol{A})$ denote the vector of non-increasingly ordered eigenvalues of a Hermitian matrix $\boldsymbol{A}$.

Theorem V.2. *The set [cf.* (13b)*,* (24)*]*

$$\boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}})) = \left\{ \boldsymbol{\lambda} \left(\frac{\hat{\boldsymbol{R}}^{\frac{1}{2}} \boldsymbol{Q} \hat{\boldsymbol{R}}^{\frac{1}{2}}}{1 + \operatorname{tr}(\boldsymbol{Q} \tilde{\boldsymbol{R}})} \right) \;\middle|\; \boldsymbol{Q} \in \mathcal{Q}(\mu_{\mathcal{Q}}) \right\} \tag{32}$$

is a simplex given by the convex hull of the origin $\boldsymbol{\sigma}^{(0)} \triangleq \mathbf{0}$ *and of the* $r_{\boldsymbol{P}}$ *linearly independent points*

$$\boldsymbol{\sigma}^{(n)} = \mathcal{H}(\omega_1, \ldots, \omega_n) \sum_{j=1}^{n} \boldsymbol{e}_j, \quad n \in \{1, \ldots, r_{\boldsymbol{P}}\} \tag{33}$$

where $[\boldsymbol{e}_1, \ldots, \boldsymbol{e}_{N_{\mathsf{T}}}] = \mathbf{I}$ *is the canonical basis, and* $\mathcal{H}(x_1, \ldots, x_n) = (\sum_{i=1}^{n} x_i^{-1})^{-1}$ *with* n *arguments* $x_1, \ldots, x_n$ *denotes the harmonic mean thereof, divided by* n*.*

Proof: See Appendix D. ∎

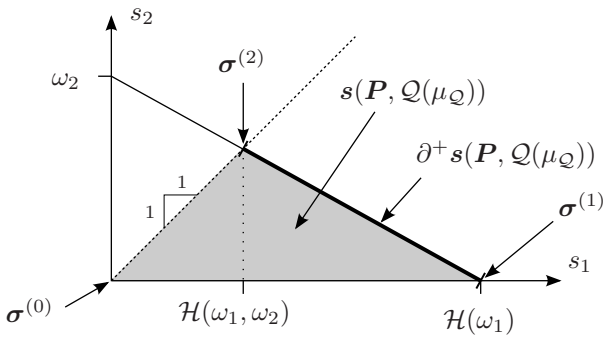

Fig. 2. Sketch of a simplex set $\boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))$ (shaded region). The so-called Pareto border $\partial^+ \boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))$ contains those points from $\boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))$ that are not dominated by any other point from $\boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))$, and is the convex hull of $\boldsymbol{\sigma}^{(n)}$ for $n \in \{1, \ldots, r_{\boldsymbol{P}}\}$ (excluding the origin).

As a byproduct, the proof of Theorem V.2 reveals that if the set of eigenvectors of $\hat{\boldsymbol{R}}$ is contained in the set of eigenvectors of $\tilde{\boldsymbol{R}}$, i.e., $\mathcal{C}(\boldsymbol{U}_{\hat{\boldsymbol{R}}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\tilde{\boldsymbol{R}}})$, then it is optimal with respect to any utility $F \in \mathcal{F}$ that the eigenbasis $\boldsymbol{U}_{\boldsymbol{Q}^\star(\boldsymbol{P})}$ of the optimal matrix $\boldsymbol{Q}^\star(\boldsymbol{P})$ be chosen such that as

$$\mathcal{C}(\boldsymbol{U}_{\boldsymbol{Q}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\hat{\boldsymbol{R}}}). \tag{34}$$

Note that this requirement is stronger than the range space inclusion property of Theorem V.1 [cf. (17)]. This particular situation of eigenbasis alignment $\mathcal{C}(\boldsymbol{U}_{\hat{\boldsymbol{R}}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\tilde{\boldsymbol{R}}})$ occurs, for example, when

- using N_{T} unitary pilots (i.e., $\boldsymbol{P} = \frac{\operatorname{tr}(\boldsymbol{P})}{N_{\mathsf{T}}} \mathbf{I}_{N_{\mathsf{T}}}$ is a scaled identity matrix)
- the channel gains are independently and identically distributed ($\boldsymbol{R} = \mathbf{I}$)
- the channel estimation error vanishes ($\tilde{\boldsymbol{R}} = \mathbf{0}$, $\hat{\boldsymbol{R}} = \boldsymbol{R}$)
- the pilots are aligned with the channel covariance, i.e., $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}})$.

As we shall see later in Section VII, the latter condition $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}})$ is in fact necessary for joint optimality of $\boldsymbol{P}$ and $\boldsymbol{Q}$.

VI. Pilot design for a prescribed precoder

To complement the previous Section V, we will now swap the roles of $\boldsymbol{P}$ and $\boldsymbol{Q}$ so as to consider the optimization of the pilot Gram $\boldsymbol{P}$ under a trace constraint, while the transmit covariance $\boldsymbol{Q}$ has a fixed value [Problem (P.1.b)]. The optimal pilot sequence reads as

$$\boldsymbol{P}^\star(\boldsymbol{Q}) = \operatorname*{argmax}_{\boldsymbol{P} \in \mathcal{P}(\mu_{\mathcal{P}})} F(\boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q})). \tag{35}$$

A. Number of streams and pilot symbols

In analogy to the inequality (20) relating the ranks of $\boldsymbol{P}$ and $\boldsymbol{Q}^\star(\boldsymbol{P})$, we have a similar rank inequality for Problem (P.1.b) too.

Theorem VI.1. *For any utility* $f \in \mathcal{F}$ *and a prescribed transmit covariance* $\boldsymbol{Q}$*, the rank of the optimal pilot Gram*

$\boldsymbol{P}^\star(\boldsymbol{Q})$ under a trace constraint is not larger than the rank of $\boldsymbol{Q}$:

$$\operatorname{rank}(\boldsymbol{P}^\star(\boldsymbol{Q})) \leq \operatorname{rank}(\boldsymbol{Q}). \tag{36}$$

Proof: See Appendix E. ∎

In words, we can state this as [compare with (21)]

$$\textbf{number of pilot symbols} \leq \textbf{number of streams} \tag{37}$$

The interpretation behind this rank inequality is that, if there were more orthogonal training directions than there are data streams precoded, we would necessarily be wasting some pilot energy into directions that are not used for transmission anyway.

B. Feasible set of effective SNR matrices

Let us rewrite Problem (P.1.b) as

$$\boldsymbol{S}^\star(\boldsymbol{Q}) = \underset{\boldsymbol{S}\in\boldsymbol{S}(\mathcal{P}(\mu_\mathcal{P}),\boldsymbol{Q})}{\operatorname{argmax}} F(\boldsymbol{S}) \tag{38}$$

in terms of effective SNR matrices, which belong to a feasible set

$$\boldsymbol{S}(\mathcal{P}(\mu_\mathcal{P}),\boldsymbol{Q}) = \big\{\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q}) \;\big|\; \boldsymbol{P}\in\mathcal{P}(\mu_\mathcal{P})\big\}. \tag{39}$$

We will show that this set is convex. To this end, we write out $\tilde{\boldsymbol{R}}$ as $\boldsymbol{R}-\hat{\boldsymbol{R}}$, then $\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q})$ reads as [cf. (11)]

$$\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q}) = \frac{\hat{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{Q}\hat{\boldsymbol{R}}^{\frac{1}{2}}}{1+\operatorname{tr}(\boldsymbol{Q}\boldsymbol{R})-\operatorname{tr}(\boldsymbol{Q}\hat{\boldsymbol{R}})}, \tag{40}$$

which is unitarily equivalent to

$$\begin{aligned}\breve{\boldsymbol{S}}(\boldsymbol{P},\boldsymbol{Q}) &= \frac{\boldsymbol{Q}^{\frac{1}{2}}\hat{\boldsymbol{R}}\boldsymbol{Q}^{\frac{1}{2}}}{1+\operatorname{tr}(\boldsymbol{Q}\boldsymbol{R})-\operatorname{tr}(\boldsymbol{Q}\hat{\boldsymbol{R}})}\\ &= \frac{\boldsymbol{Q}^{\frac{1}{2}}\hat{\boldsymbol{R}}\boldsymbol{Q}^{\frac{1}{2}}}{\tau-\operatorname{tr}(\boldsymbol{Q}\hat{\boldsymbol{R}})},\end{aligned} \tag{41}$$

where $\tau = 1+\operatorname{tr}(\boldsymbol{Q}\boldsymbol{R})$. The unitary equivalence is due the Hermitian matrices $\hat{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{Q}\hat{\boldsymbol{R}}^{\frac{1}{2}}$ and $\boldsymbol{Q}^{\frac{1}{2}}\hat{\boldsymbol{R}}\boldsymbol{Q}^{\frac{1}{2}}$ having the same eigenvalues because of the identity $\boldsymbol{\lambda}(\boldsymbol{A}\boldsymbol{B}) = \boldsymbol{\lambda}(\boldsymbol{B}\boldsymbol{A})$. As a consequence, $F(\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q})) = F(\breve{\boldsymbol{S}}(\boldsymbol{P},\boldsymbol{Q}))$ for any $F\in\mathcal{F}$, so $\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q})$ and $\breve{\boldsymbol{S}}(\boldsymbol{P},\boldsymbol{Q})$ can be used interchangeably. By comparing Expression (41) with the definition of linear fractional functions (cf. Definition V.1), we identify $\hat{\boldsymbol{R}}\mapsto\breve{\boldsymbol{S}}(\boldsymbol{P},\boldsymbol{Q})$ as a linear fractional function with parameters $\boldsymbol{A} = \frac{1}{\sqrt{\tau}}\boldsymbol{Q}^{\frac{1}{2}}$ and $\boldsymbol{B} = -\frac{1}{\tau}\boldsymbol{Q}$, i.e.,

$$\breve{\boldsymbol{S}}(\boldsymbol{P},\boldsymbol{Q}) = \phi\Big(\hat{\boldsymbol{R}};\tfrac{1}{\sqrt{\tau}}\boldsymbol{Q}^{\frac{1}{2}},-\tfrac{1}{\tau}\boldsymbol{Q}\Big). \tag{42}$$

In Appendix F, we show that the set of feasible $\hat{\boldsymbol{R}}$, namely

$$\big\{\boldsymbol{R}-(\boldsymbol{R}^{-1}+\boldsymbol{P})^{-1} \;\big|\; \boldsymbol{P}\in\mathcal{P}(\mu_\mathcal{P})\big\}, \tag{43}$$

is convex, from which follows immediately with Corollary V.1 that $\breve{\boldsymbol{S}}(\mathcal{P}(\mu_\mathcal{P}),\boldsymbol{Q})$ is a convex set. For solving Problem (P.1.b), it now suffices to replace $\boldsymbol{S}(\mathcal{P}(\mu_\mathcal{P}),\boldsymbol{Q})$ with $\breve{\boldsymbol{S}}(\mathcal{P}(\mu_\mathcal{P}),\boldsymbol{Q})$ in formulation (38) of Problem (P.1.b). Granted that the utility F is concave or quasi-concave, Problem (P.1.b) can be solved efficiently.

VII. Jointly Pareto optimal pilot-precoder pairs

A. Problem statement

We move on to study a subproblem of (P.3). For this purpose, we restate Problem (P.3) in the vector domain of feasible profiles $\boldsymbol{s}$ so that it reads

$$\max_{\boldsymbol{s}\in\boldsymbol{s}(\mathcal{PQ}(T_\mathsf{p}))} \bar{f}(\boldsymbol{s}), \tag{44}$$

where the search set is

$$\boldsymbol{s}(\mathcal{PQ}(T_\mathsf{p})) = \Big\{\boldsymbol{s}(\boldsymbol{P},\boldsymbol{Q}) \;\Big|\; (\boldsymbol{P},\boldsymbol{Q})\in\mathcal{PQ}(T_\mathsf{p})\Big\} \tag{45}$$

and $\mathcal{PQ}(T_\mathsf{p})$ was defined in (14). Exploiting the monotonicity of utilities $\bar{f}\in\bar{\mathcal{F}}$, we can restrict the search set $\boldsymbol{s}(\mathcal{PQ}(T_\mathsf{p}))$ to its Pareto border $\partial^+\boldsymbol{s}(\mathcal{PQ}(T_\mathsf{p}))$. To be precise, the Pareto border of a set $\mathcal{A}\subset\mathbb{R}^N$ is the subset

$$\partial^+\mathcal{A} = \Big\{\boldsymbol{a}\in\mathcal{A} \;\Big|\; \nexists\boldsymbol{a}'\in\mathcal{A}\colon \boldsymbol{a}'\geq\boldsymbol{a} \text{ with } \boldsymbol{a}'\neq\boldsymbol{a}\Big\}. \tag{46}$$

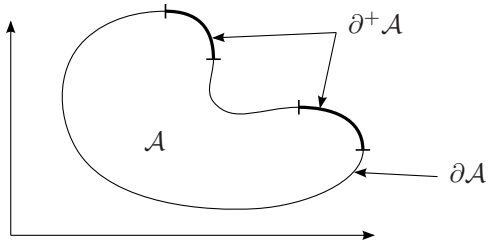

Fig. 3. Pareto border $\partial^+\mathcal{A}$ of a compact set $\mathcal{A}\subset\mathbb{R}^2$

Regardless of which utility function $\bar{f}\in\bar{\mathcal{F}}$ we are considering, there exists an important subproblem of (P.3) that is common to all utility functions: the characterization of the search set $\partial^+\boldsymbol{s}(\mathcal{PQ}(T_\mathsf{p}))$. The following subsections will build up this characterization in several steps.

B. Number of streams and pilot symbols

Recall that (P.3) can be expressed in terms of (P.2) by adding an outer optimization of the energy boost [cf. (15)]:

$$\max_{\substack{\mu_\mathcal{P},\mu_\mathcal{Q}\geq 0\\ \mu_\mathcal{P}+(T-T_\mathsf{p})\mu_\mathcal{Q}=T\mu}} \left\{\max_{(\boldsymbol{P},\boldsymbol{Q})\in\mathcal{P}(\mu_\mathcal{P})\times\mathcal{Q}(\mu_\mathcal{Q})} F(\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q}))\right\}. \tag{47}$$

Inside the curly braces is Problem (P.2), which can be equivalently written as

$$\max_{\boldsymbol{P}\in\mathcal{P}(\mu_\mathcal{P})} F\big(\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q}^\star(\boldsymbol{P}))\big) = \max_{\boldsymbol{Q}\in\mathcal{Q}(\mu_\mathcal{Q})} F\big(\boldsymbol{S}(\boldsymbol{P}^\star(\boldsymbol{Q}),\boldsymbol{Q})\big) \tag{48}$$

with $\boldsymbol{Q}^\star(\boldsymbol{P})$ and $\boldsymbol{P}^\star(\boldsymbol{Q})$ defined in (16) and (35), respectively. We infer that the jointly optimal pilot-precoder pair for Problem (P.2) must simultaneously fulfill the rank inequalities (20) and (36). This means that the pilot Gram and transmit covariance have equal ranks. Since this rank equality holds regardless of the value of $(\mu_\mathcal{P},\mu_\mathcal{Q})$, it also holds for Problem (P.3). Since it also holds regardless of the value of T_p, it also holds for Problem (P.4). Hence, we can state generally that, for Problems (P.2), (P.3), (P.4),

$$r^\star \triangleq \textbf{number of streams} = \textbf{number of pilot symbols} \tag{49}$$

is a necessary condition for a pilot-precoder pair $(\boldsymbol{P}, \boldsymbol{Q})$ to be jointly optimal. Note that, in addition, Theorem V.1 requires the optimal transmit covariance to lie in the range space of $\hat{\boldsymbol{R}}$ and as a consequence, the rank of $\boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q})$, which is the number of non-zero entries in $\boldsymbol{s}(\boldsymbol{P}, \boldsymbol{Q})$, is equal to $r^\star$ as well.

C. Jointly optimal transmit and training directions

A fortunate circumstance when treating the joint problems (P.2), (P.3), (P.4), is that the jointly optimal transmit and training directions have a very simple and intuitive characterization, enunciated in Theorem VII.1 below. As in the previous Section, we set our focus on Problem (P.2), since the property will extend immediately to Problems (P.3) and (P.4).

Let the channel covariance $\boldsymbol{R}$, the pilot Gram $\boldsymbol{P}$ and the transmit covariance $\boldsymbol{Q}$ have the following (reduced) eigendecompositions:

$$\boldsymbol{R} = \boldsymbol{U}_{\boldsymbol{R}} \boldsymbol{\Lambda}_{\boldsymbol{R}} \boldsymbol{U}_{\boldsymbol{R}}^\dagger, \quad \boldsymbol{P} = \boldsymbol{U}_{\boldsymbol{P}} \boldsymbol{\Lambda}_{\boldsymbol{P}} \boldsymbol{U}_{\boldsymbol{P}}^\dagger, \quad \boldsymbol{Q} = \boldsymbol{U}_{\boldsymbol{Q}} \boldsymbol{\Lambda}_{\boldsymbol{Q}} \boldsymbol{U}_{\boldsymbol{Q}}^\dagger.$$

Without loss of generality, we assume that the eigenvalues of $\boldsymbol{R}$ are arranged in non-increasing order on the diagonal positions of $\boldsymbol{\Lambda}_{\boldsymbol{R}}$, whereas the eigenvalues of $\boldsymbol{\Lambda}_{\boldsymbol{P}}$ and $\boldsymbol{\Lambda}_{\boldsymbol{Q}}$ are not sorted in any specific order. Let the set of columns of a matrix $\boldsymbol{A}$ be denoted as $\mathcal{C}(\boldsymbol{A})$.

Theorem VII.1. *For any utility $\bar{f} \in \bar{\mathcal{F}}$, in the joint optimization problem (P.2), there is no loss of optimality in restricting the eigenvectors of the pilot Gram $\boldsymbol{P}$ (i.e., the left singular vectors of the pilot sequence $\boldsymbol{X}_{\mathsf{p}}$) and the eigenvectors of the transmit covariance $\boldsymbol{Q}$ (i.e., the left singular vectors of the precoder $\boldsymbol{F}$) to be a common subset of the eigenvectors of the channel covariance $\boldsymbol{R}$ corresponding to the largest eigenvalues of $\boldsymbol{R}$. Formally, this is to say that the (reduced) eigenbases $\boldsymbol{U}_{\boldsymbol{P}}$ and $\boldsymbol{U}_{\boldsymbol{Q}}$ should satisfy*

$$\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) = \mathcal{C}(\boldsymbol{U}_{\boldsymbol{Q}}) = \{\boldsymbol{u}_{\boldsymbol{R},1}, \ldots, \boldsymbol{u}_{\boldsymbol{R},r^\star}\} \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}}), \quad (50)$$

where $\boldsymbol{U}_{\boldsymbol{R}} \triangleq [\boldsymbol{u}_{\boldsymbol{R},1}, \ldots, \boldsymbol{u}_{\boldsymbol{R},N_{\mathsf{T}}}]$, and $r^\star = \mathrm{rank}(\boldsymbol{P}^\star) = \mathrm{rank}(\boldsymbol{Q}^\star)$ denotes the pilot/precoder rank at the joint optimum $(\boldsymbol{P}^\star, \boldsymbol{Q}^\star)$ of Problem (P.2). [6]

Proof: See Appendix G. ∎

Since Theorem VII.1 holds irrespective of the values of $(\mu_{\mathcal{P}}, \mu_{\mathcal{Q}})$ and of T_{p}, we infer that it also holds for Problems (P.3) and (P.4).

This Theorem echoes similar results from previous publications. For example, in [3], the authors find that the optimal pilot symbols are, as in the present case, scaled eigenvectors of the channel's transmit correlation matrix. However, they optimize the pilot sequence with respect to the Frobenius norm of the channel estimation's mean-square error matrix instead of the achievable rate. Similarly, in [5] it is proven for a multiple-access setup that, under the assumption that the channel estimate follows an UIU model (in the terminology of [20]) slightly more general than ours, the transmit covariances are aligned with the channel correlation as well. As to the pilot sequence in [5], it is optimized with respect to different objectives: the trace and the determinant of the mean-square error matrix of the channel estimation. The corresponding optimal eigenbases for the pilot sequences are similarly aligned. The main contribution of Theorem VII.1 is that of establishing the jointly optimal eigenbases of pilot Gram and transmit covariance with respect to a single utility function.

[6] Obviously, the rank $r^\star$ is not known *a priori* before solving the problem. The notation in (50) is merely to indicate that $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}})$ and $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{Q}})$ should contain eigenvectors of $\boldsymbol{R}$ corresponding to the *largest* eigenvalues of $\boldsymbol{R}$.

D. Pareto optimal allocation

Consequently, and without loss of optimality, we will align the eigenbases of $\boldsymbol{P}$ and $\boldsymbol{Q}$ in conformity with (50). The scalars $[\boldsymbol{r}]_i = r_i$, $[\boldsymbol{p}]_i = p_i$, and $[\boldsymbol{q}]_i = q_i$ shall denote the eigenvalues of $\boldsymbol{R}$, $\boldsymbol{P}$, and $\boldsymbol{Q}$, respectively. Under such assumptions, all matrices involved in the expression of the effective SNR (11), namely $\hat{\boldsymbol{R}}$ and $\tilde{\boldsymbol{R}}$ [cf. (7)], as well as $\boldsymbol{Q}$, acquire the same eigenbasis $\boldsymbol{U}_{\boldsymbol{R}}$. We can readily see from Expression (11) that $\boldsymbol{S}$ then inherits the (common) eigenvectors of $\boldsymbol{P}$ and $\boldsymbol{Q}$, i.e., $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{S}}) = \mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) = \mathcal{C}(\boldsymbol{U}_{\boldsymbol{Q}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}})$, so that the profile $\boldsymbol{s}$ is given by [cf. (11)]

$$\boldsymbol{s} = \frac{\hat{\boldsymbol{r}} \odot \boldsymbol{q}}{1 + \boldsymbol{q}^{\mathrm{T}} \tilde{\boldsymbol{r}}} \quad (51)$$

and '$\odot$' denotes the componentwise product. Here, the eigenvalue vectors $\tilde{\boldsymbol{r}} = \tilde{\boldsymbol{r}}(\boldsymbol{p})$ and $\hat{\boldsymbol{r}} = \hat{\boldsymbol{r}}(\boldsymbol{p}) = \boldsymbol{r} - \tilde{\boldsymbol{r}}(\boldsymbol{p})$ are functions of $\boldsymbol{p}$ and respectively have entries

$$\tilde{r}_i(p_i) = \frac{r_i}{1 + r_i p_i}, \qquad \hat{r}_i(p_i) = \frac{r_i^2 p_i}{1 + r_i p_i}. \quad (52)$$

Hereinforth, we will write $\boldsymbol{s}(\boldsymbol{p}, \boldsymbol{q})$ instead of $\boldsymbol{s}(\boldsymbol{P}, \boldsymbol{Q})$ whenever we implicitly assume that the eigenbases are optimally aligned according to (50). We do not impose any ordering of the eigenvalues p_i, q_i, and r_i. Instead we assume, without loss of generality, that they are arranged in such way that the s_i are non-increasingly ordered.

Upon optimally aligning the eigenbases as according to Theorem VII.1, we now consider the remaining problem that consists in jointly optimizing the allocation vector pair $(\boldsymbol{p}, \boldsymbol{q})$, which belongs to a set that constrains the average power radiated by the transmitter array:

$$\mathit{\Gamma} = \left\{ (\boldsymbol{p}, \boldsymbol{q}) \in \mathbb{R}_+^{2N_{\mathsf{T}}} \middle| \mathbf{1}^{\mathrm{T}} \boldsymbol{p} + (T - T_{\mathsf{p}}) \mathbf{1}^{\mathrm{T}} \boldsymbol{q} \leq T\mu \right\}. \quad (53)$$

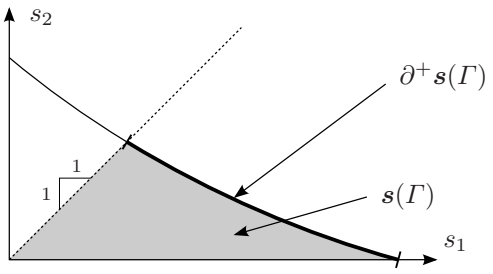

Fig. 4. Sketch of the typical shape of a set $\boldsymbol{s}(\mathit{\Gamma})$ and its Pareto border $\partial^+ \boldsymbol{s}(\mathit{\Gamma})$ for $N_{\mathsf{T}} = 2$.

By virtue of Theorem VII.1, we have $\boldsymbol{s}(\mathcal{PQ}(T_{\mathsf{p}})) = \boldsymbol{s}(\mathit{\Gamma})$. In the following, we will devise a procedure for computing the set of all allocations $(\boldsymbol{p}, \boldsymbol{q})$ that yield points located

on the Pareto border $\partial^+\boldsymbol{s}(\mathcal{PQ}(T_\mathsf{p})) = \partial^+\boldsymbol{s}(\mathit{\Gamma})$. Given the monotonicity of the entries of $\boldsymbol{s}(\boldsymbol{p},\boldsymbol{q})$ in p_i and q_i, we are certain that any Pareto optimal allocation $(\boldsymbol{p},\boldsymbol{q})$ will expend the full power budget, and thus belong to

$$\partial^+\mathit{\Gamma} = \left\{(\boldsymbol{p},\boldsymbol{q}) \in \mathit{\Gamma} \middle| \mathbf{1}^\mathrm{T}\boldsymbol{p} + (T - T_\mathsf{p})\mathbf{1}^\mathrm{T}\boldsymbol{q} = T\mu\right\}. \quad (54)$$

Hence, $\partial^+\boldsymbol{s}(\mathit{\Gamma}) = \partial^+\boldsymbol{s}(\partial^+\mathit{\Gamma})$. Now note that the search set $\partial^+\boldsymbol{s}(\partial^+\mathit{\Gamma})$ is *not* equal to the set $\boldsymbol{s}(\partial^+\mathit{\Gamma})$, meaning that it is not sufficient to simply choose some full-power allocation $(\boldsymbol{p},\boldsymbol{q}) \in \partial^+\mathit{\Gamma}$ in order to obtain a *Pareto optimal* allocation. Instead, we have the proper inclusion

$$\partial^+\boldsymbol{s}(\mathit{\Gamma}) = \partial^+\boldsymbol{s}(\partial^+\mathit{\Gamma}) \subsetneq \boldsymbol{s}(\partial^+\mathit{\Gamma}). \quad (55)$$

In fact, any Pareto optimal allocation is a full-power allocation, but the converse is not true. This becomes clear when counting dimensions: the vector $\boldsymbol{s}$ has N_T real entries, so any parametrization of the feasible set $\boldsymbol{s}(\mathit{\Gamma})$ with minimal number of parameters will require at most N_T real parameters. A parametrization of the Pareto border $\partial^+\boldsymbol{s}(\mathit{\Gamma})$ will require $N_\mathsf{T} - 1$ parameters. However, the entries of the vector pair $(\boldsymbol{p},\boldsymbol{q})$ represent $2N_\mathsf{T}$ parameters. Thus, to obtain a minimal parametrization of $\partial^+\boldsymbol{s}(\mathit{\Gamma})$ there are at least $N_\mathsf{T} - 1$ excess parameters to be eliminated. A direct elimination by working off the explicit expression of $\boldsymbol{s}(\boldsymbol{p},\boldsymbol{q})$ in (51) does not seem possible. Even replacing $\mathit{\Gamma}$ with $\partial^+\mathit{\Gamma}$ only saves one parameter.

The idea for reducing the parameter set so as to efficiently compute Pareto optimal allocations $(\boldsymbol{p},\boldsymbol{q})$ will be as follows: we choose some vector norm $\|\cdot\|$, then fix a non-negative direction vector $\boldsymbol{e} \geq \mathbf{0}$ that is normalized as $\|\boldsymbol{e}\| = 1$. This normalized vector points into the positive orthant of the $\boldsymbol{s}$ domain and defines a half-line departing from the origin. We then maximize the norm $\|\boldsymbol{s}(\boldsymbol{p},\boldsymbol{q})\|$ with respect to the allocation $(\boldsymbol{p},\boldsymbol{q})$ under the constraint that $\boldsymbol{s}(\boldsymbol{p},\boldsymbol{q})$ points into the direction of $\boldsymbol{e}$. In other terms, we determine the point from the set $\boldsymbol{s}(\mathit{\Gamma})$ which lies farthest away from the origin, and is located on the line running along $\boldsymbol{e}$.

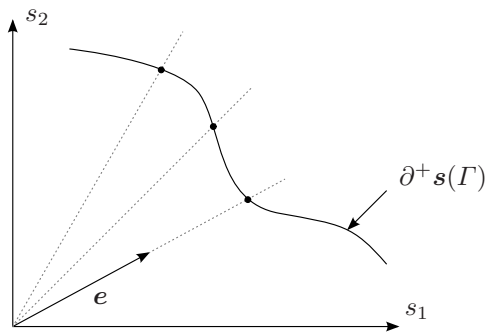

Fig. 5. Symbolic sketch of the procedure for computing Pareto border points from $\partial^+\boldsymbol{s}(\mathit{\Gamma})$. Said points are parametrized by a unit-norm direction vector $\boldsymbol{e}$

Formally, the problem at hand can be stated as:

$$\max_{(\boldsymbol{p},\boldsymbol{q})\in\mathit{\Gamma}} \nu \qquad \text{s.t.} \quad \boldsymbol{s}(\boldsymbol{p},\boldsymbol{q}) = \nu\boldsymbol{e} \quad (56)$$

where $\nu = \|\boldsymbol{s}(\boldsymbol{p},\boldsymbol{q})\|$ stands for the norm of $\boldsymbol{s}$, while the function $\boldsymbol{s}(\boldsymbol{p},\boldsymbol{q})$ is given by (51) as

$$\boldsymbol{s}(\boldsymbol{p},\boldsymbol{q}) = \frac{\hat{\boldsymbol{r}}\odot\boldsymbol{q}}{1+\tilde{\boldsymbol{r}}^\mathrm{T}\boldsymbol{q}} = \frac{\hat{\boldsymbol{r}}\odot\boldsymbol{q}}{1+\boldsymbol{r}^\mathrm{T}\boldsymbol{q}-\hat{\boldsymbol{r}}^\mathrm{T}\boldsymbol{q}} \quad (57)$$

and $\boldsymbol{e}$ is some normalized direction vector pointing into the positive orthant, i.e., $\boldsymbol{e} \geq \mathbf{0}$ and $\|\boldsymbol{e}\| = 1$. As usual, the search set $\mathit{\Gamma}$ can be reduced to $\partial^+\mathit{\Gamma}$.

When we vary $\boldsymbol{e}$, the set of all points $\nu_{\max}\boldsymbol{e}$ that are determined by this maximization procedure constitute what we shall call a *front border*.

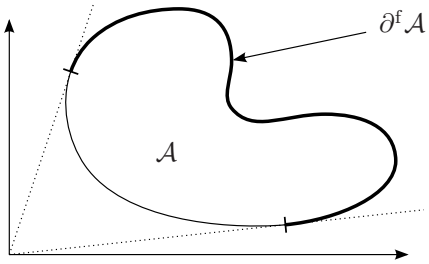

Fig. 6. Front border $\partial^\mathrm{f}\mathcal{A}$ of a closed set $\mathcal{A} \subset \mathbb{R}^2$

The front border of a compact set $\mathcal{A} \subseteq \mathbb{R}^n_+$ shall be denoted as $\partial^\mathrm{f}\mathcal{A}$ and be formally defined as

$$\partial^\mathrm{f}\mathcal{A} = \bigcup_{\substack{\boldsymbol{e}\geq\mathbf{0}\\ \|\boldsymbol{e}\|=1}} \operatorname*{argmax}_{\substack{\boldsymbol{a}\in\mathcal{A}\\ \boldsymbol{a}=\nu\boldsymbol{e}}} \nu. \quad (58)$$

Note that certain directions $\boldsymbol{e}$ may yield empty sets $\{\boldsymbol{a} \in \mathcal{A} | \boldsymbol{a} = \nu\boldsymbol{e}\} = \emptyset$, so only non-trivial contributions (non-empty sets) should be retained when taking the union (58). As we easily intuit from comparing Figures 3 and 6, the *Pareto border* and *front border* of a compact set are not generally identical. However, according to the next Lemma, identity holds for the set $\boldsymbol{s}(\mathit{\Gamma})$.

Lemma VII.1. *The Pareto border and the front border of the set $\boldsymbol{s}(\mathit{\Gamma})$ coincide.*

Proof: See Appendix H. ∎

As a consequence, we can compute the Pareto border by the above-mentioned technique. Let us choose the norm $\|\cdot\|$ to be the 1-norm $\|\boldsymbol{s}\|_1 = \sum_i s_i$, as this will turn out to be a convenient choice. The quantity ν that is maximized in (56) is the 1-norm of the vector $\boldsymbol{s}(\boldsymbol{p},\boldsymbol{q})$, constrained to being colinear with $\boldsymbol{e}$, i.e.,

$$\boldsymbol{s}(\boldsymbol{p},\boldsymbol{q}) = \nu\boldsymbol{e} \qquad \|\boldsymbol{s}(\boldsymbol{p},\boldsymbol{q})\|_1 = \frac{\eta}{1+\boldsymbol{r}^\mathrm{T}\boldsymbol{q}-\eta} = \nu \quad (59)$$

where η stands for [cf. (57)]

$$\eta = \|\hat{\boldsymbol{r}}\odot\boldsymbol{q}\|_1 = \hat{\boldsymbol{r}}^\mathrm{T}\boldsymbol{q}. \quad (60)$$

Note that the colinearity constraint $\boldsymbol{s}(\boldsymbol{p},\boldsymbol{q}) = \nu\boldsymbol{e}$ implies the colinearity $\hat{\boldsymbol{r}}\odot\boldsymbol{q} = \eta\boldsymbol{e}$. Componentwise, the latter reads as [cf. (52)]

$$\frac{r_i^2 p_i q_i}{1+r_i p_i} = \eta e_i \quad (61)$$

Consider $\boldsymbol{e}$ to be fixed. Then we see from (61) that, once η is given, p_i and q_i are entirely determined from one another: given any value of $q_i \geq 0$, the corresponding value of $p_i \geq 0$ is uniquely determined (as long as $\frac{\eta e_i}{q_i} < r_i$), and conversely, given any value of $p_i \geq 0$, the value of $q_i \geq 0$ is uniquely determined. This allows us to effectuate a (one-to-one) change of parameters: we drop the q_i and replace them by e_i, thus effectively replacing the parameter pair $(\boldsymbol{p}, \boldsymbol{q}) \in \mathit{\Gamma}$ by the new pair $(\boldsymbol{p}, \boldsymbol{e}) \in \mathcal{D}(T\mu) \times \mathcal{D}(1)$, where $\mathcal{D}(\cdot)$ is defined as

$$\mathcal{D}(\alpha) = \big\{ \boldsymbol{d} \in \mathbb{R}_+^{N_\mathsf{T}} \;\big|\; \mathbf{1}^\mathrm{T} \boldsymbol{d} \leq \alpha \big\}. \tag{62}$$

From (61), the q_i can now be expressed in terms of p_i and e_i as

$$q_i(p_i, e_i) = \eta e_i \frac{1 + r_i p_i}{r_i^2 p_i}. \tag{63}$$

By summing (63) up over i, and taking into account the energy conservation $\sum_i p_i + (T - T_\mathsf{p}) \sum_i q_i = T\mu$, we obtain expressions of η and of q_i which are functions of $(\boldsymbol{p}, \boldsymbol{e})$:

$$\eta(\boldsymbol{p}, \boldsymbol{e}) = \frac{T\mu - \mathbf{1}^\mathrm{T} \boldsymbol{p}}{T - T_\mathsf{p}} \left(\sum_{i=1}^{N_\mathsf{T}} e_i \frac{1 + r_i p_i}{r_i^2 p_i} \right)^{-1} \tag{64}$$

$$q_i(\boldsymbol{p}, \boldsymbol{e}) = \frac{T\mu - \mathbf{1}^\mathrm{T} \boldsymbol{p}}{T - T_\mathsf{p}} \frac{e_i \frac{1 + r_i p_i}{r_i^2 p_i}}{\sum_j e_j \frac{1 + r_j p_j}{r_j^2 p_j}}. \tag{65}$$

Consequently, ν can itself be expressed as a function of $(\boldsymbol{p}, \boldsymbol{e})$ too [cf. (56)]:

$$\nu(\boldsymbol{p}, \boldsymbol{e}) = \frac{\eta(\boldsymbol{p}, \boldsymbol{e})}{1 + \boldsymbol{r}^\mathrm{T} \boldsymbol{q}(\boldsymbol{p}, \boldsymbol{e}) - \eta(\boldsymbol{p}, \boldsymbol{e})}. \tag{66}$$

We can now dismiss the initial problem formulation (56) in favor of the equivalent formulation

$$\boldsymbol{p}^\star(\boldsymbol{e}) = \operatorname*{argmax}_{\boldsymbol{p} \in \mathcal{D}(T\mu)} \nu(\boldsymbol{p}, \boldsymbol{e}) \tag{67}$$

with $\nu(\boldsymbol{p}, \boldsymbol{e})$ as given in (66). Once the maximizer $\boldsymbol{p}^\star(\boldsymbol{e})$ is determined, we compute the corresponding $\boldsymbol{q}^\star(\boldsymbol{e})$ via (65) as

$$\boldsymbol{q}^\star(\boldsymbol{e}) = \begin{bmatrix} q_1(\boldsymbol{p}^\star(\boldsymbol{e}), \boldsymbol{e}) \\ \vdots \\ q_{N_\mathsf{T}}(\boldsymbol{p}^\star(\boldsymbol{e}), \boldsymbol{e}) \end{bmatrix}. \tag{68}$$

The Pareto border $\partial^+ \boldsymbol{s}(\mathcal{PQ}(T_\mathsf{p})) = \partial^+ \boldsymbol{s}(\mathit{\Gamma})$ is described in its entirety by the union (see Figure 5)

$$\partial^+ \boldsymbol{s}(\mathit{\Gamma}) = \bigcup_{\substack{\boldsymbol{e} \geq \mathbf{0} \\ \|\boldsymbol{e}\|_1 = 1}} \boldsymbol{s}(\boldsymbol{p}^\star(\boldsymbol{e}), \boldsymbol{q}^\star(\boldsymbol{e})). \tag{69}$$

Definition VII.1. *A function $f \colon \mathcal{X} \mapsto \mathbb{R}$ is quasi-concave (resp. quasi-convex) on a convex and compact set $\mathcal{X} \subset \mathbb{R}^n$ if it can be represented as a concatenation*

$$f(\boldsymbol{x}) = (g \circ h)(\boldsymbol{x}) \tag{70}$$

of a concave (resp. convex) function $h \colon \mathcal{X} \to \mathbb{R}$ and a non-decreasing function $g \colon \mathbb{R} \to \mathbb{R}$.

Lemma VII.2. *The function $\nu(\boldsymbol{p}, \boldsymbol{e})$ is quasi-concave in $\boldsymbol{p}$.*

Proof: See Appendix I. ∎

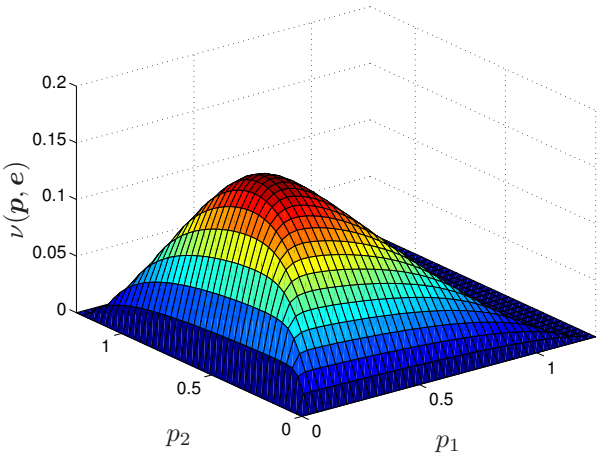

(a) $\nu(\boldsymbol{p}, \boldsymbol{e})$ as a function of p_1 and p_2

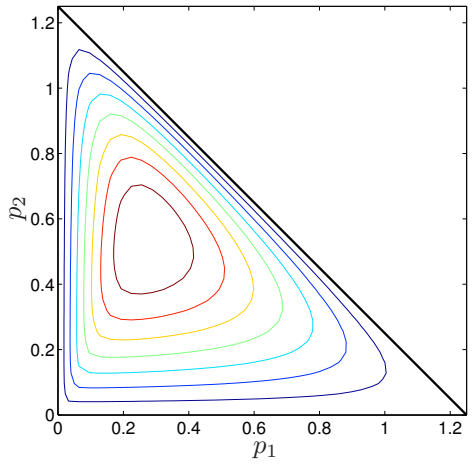

(b) Contour plot of the function $\nu(\boldsymbol{p}, \boldsymbol{e})$ from Figure 7(a)

Fig. 7. Three-dimensional representation and corresponding contour plot of a function $\nu(\boldsymbol{p}, \boldsymbol{e})$

This lemma renders (67) a quasi-convex problem, which can be solved efficiently.

Figures 7(a) and 7(b) illustrate an example of a function $\nu(\boldsymbol{p}, \boldsymbol{e})$ for $N_\mathsf{T} = 2$ transmit antennas, channel coherence $T = 10$ and SNR $\mu = 1$, $\boldsymbol{r} = \begin{bmatrix} 2/3 & 1/3 \end{bmatrix}^\mathrm{T}$ and $\boldsymbol{e} = \begin{bmatrix} 1/2 & 1/2 \end{bmatrix}^\mathrm{T}$. The quasi-concavity (but non-concavity) can be well appreciated in said plot, since $\nu(\boldsymbol{p}, \boldsymbol{e})$ appears to be convex in $\boldsymbol{p}$ near the borders of its triangular domain $\mathcal{D}(T\mu)$, while it is concave in an inner region. Notwithstanding this change of curvature, the function is globally quasi-concave in $\boldsymbol{p}$, since all upper contour sets, as illustrated in Figure 7(b), are convex.

VIII. Iterative joint design

Having studied Problems (P.1.a) and (P.1.b) as well as a subproblem (Pareto optimal allocations) of Problem (P.3), we now propose an iterative approach to solving Problems (P.3) and (P.4).

Using Theorem VII.1, we align the eigenbases as $\mathcal{C}(\boldsymbol{U_P}) = \mathcal{C}(\boldsymbol{U_Q}) \subseteq \mathcal{C}(\boldsymbol{U_R})$. Since we know from Sections V and VI how to efficiently solve Problems (P.1.a) and (P.1.b) for

(quasi-)concave utilities F, a natural way of tackling the joint problem (P.2) is by alternating between the two problems in the fashion of a block gradient ascent:

$$\left\{ \begin{array}{l} \boldsymbol{p}_{n+1} = \boldsymbol{p}^\star(\boldsymbol{q}_n) \\ \boldsymbol{q}_{n+1} = \boldsymbol{q}^\star(\boldsymbol{p}_{n+1}) \end{array} \right. \quad \text{or} \quad \left\{ \begin{array}{l} \boldsymbol{q}_{n+1} = \boldsymbol{q}^\star(\boldsymbol{p}_n) \\ \boldsymbol{p}_{n+1} = \boldsymbol{p}^\star(\boldsymbol{q}_{n+1}) \end{array} \right. \tag{71}$$

This procedure converges monotonically towards a fixed-point of the iteration $(\boldsymbol{p}_{n+1}, \boldsymbol{q}_{n+1}) = (\boldsymbol{p}^\star(\boldsymbol{q}_n), \boldsymbol{q}^\star(\boldsymbol{p}_n))$, which is a local optimum for Problem (P.2). However, it yields no local optimum of Problem (P.3), the reason being that no step in the above iteration ever changes the balance between pilot energy $\mathbf{1}^{\mathrm{T}}\boldsymbol{p}$ and transmit power $\mathbf{1}^{\mathrm{T}}\boldsymbol{q}$ (energy boost). To accomodate the energy boost, an additional step needs to be inserted in the above iteration in order to readjust the allocation $(\boldsymbol{p}, \boldsymbol{q})$ so as to remain Pareto optimal. This step can be performed with the methods for computing the Pareto border, developed in Subsection VII-D. Therefore, the proposed algorithm for solving (P.3) should cycle through the following three steps:

1) Optimize $\boldsymbol{p}$ for a prescribed $\boldsymbol{q}$
2) Optimize $\boldsymbol{q}$ for a prescribed $\boldsymbol{p}$
3) Adjust $(\boldsymbol{p}, \boldsymbol{q})$ to be Pareto optimal

Algorithm 1 Iteration for solving (P.3)

1: $\boldsymbol{p}_0 \leftarrow \frac{T_{\mathsf{p}}\mu}{N_{\mathsf{T}}}\mathbf{1}_{N_{\mathsf{T}}}$
2: $\boldsymbol{q}_0 \leftarrow \frac{(T-T_{\mathsf{p}})\mu}{N_{\mathsf{T}}}\mathbf{1}_{N_{\mathsf{T}}}$
3: $n \leftarrow 0$
4: **repeat**
5: $\quad \boldsymbol{p}' \leftarrow \operatorname{argmax}_{\boldsymbol{p}\in\mathcal{D}(\mathbf{1}^{\mathrm{T}}\boldsymbol{p}_n)} f(\boldsymbol{s}(\boldsymbol{p}, \boldsymbol{q}))$
6: $\quad \boldsymbol{q}' \leftarrow \operatorname{argmax}_{\boldsymbol{q}\in\mathcal{D}(\mathbf{1}^{\mathrm{T}}\boldsymbol{q}_n)} f(\boldsymbol{s}(\boldsymbol{p}, \boldsymbol{q}))$
7: $\quad \boldsymbol{e}_{n+1} \leftarrow \frac{\boldsymbol{s}(\boldsymbol{p}', \boldsymbol{q}')}{\|\boldsymbol{s}(\boldsymbol{p}', \boldsymbol{q}')\|_1}$
8: $\quad \boldsymbol{p}_{n+1} \leftarrow \operatorname{argmax}_{\boldsymbol{p}\in\mathcal{D}(T\mu)} \nu(\boldsymbol{p}, \boldsymbol{e}_{n+1})$
9: $\quad \boldsymbol{q}_{n+1} \leftarrow \boldsymbol{q}(\boldsymbol{p}_{n+1}, \boldsymbol{e}_{n+1})$
10: $\quad \boldsymbol{s}_{n+1} \leftarrow \boldsymbol{e}_{n+1}\nu(\boldsymbol{p}_{n+1}, \boldsymbol{e}_{n+1})$
11: $\quad n \leftarrow n+1$
12: **until** $f(\boldsymbol{s}_n) - f(\boldsymbol{s}_{n-1}) \leq \epsilon$

In pseudocode, the algorithm is written out in Algorithm 1. For (quasi-)concave utilities $\bar{f} \in \bar{\mathcal{F}}$, Steps 5 and 6 were shown to be (quasi-)convex optimizations in Sections V and VI, respectively. The computation of Steps 7 through 10 has been explained in detail in Subsection VII-D, wherein Step 8 was shown to be a quasi-convex optimization.

We already mentioned in Section IV that the problem of optimally tuning the training duration length T_{p} could be tackled by an exhaustive search over the set $\{1, \ldots, \min(T-1, N_{\mathsf{T}})\}$. We thus simply need to wrap Algorithm 1 into an outer loop. Note that the function $\nu(\cdot, \cdot)$ is dependent on the parameter T_{p} (though this is not reflected in notation), and should be updated accordingly with the loop count. The algorithm for (P.4) in pseudocode is written out in Algorithm 2 below:

The proposed algorithm may still suffer from efficiency issues due, for instance, to the precision parameter ϵ [Step 12 in Algorithm 1] being left as an arbitrary choice. A fast single-loop iteration in the spirit of Algorithm 1 in [20] would be preferable, but such a reformulation of our algorithm following the ideas of [20] does not seem straightforward. For one thing,

Algorithm 2 Iteration for solving (P.4)

1: $T_{\mathsf{p}} \leftarrow 1$
2: **repeat**
3: $\quad \ldots$ Algorithm 1 $\ldots$
4: $\quad (\boldsymbol{p}^\star(T_{\mathsf{p}}), \boldsymbol{q}^\star(T_{\mathsf{p}})) \leftarrow (\boldsymbol{p}_n, \boldsymbol{q}_n)$
5: $\quad T_{\mathsf{p}} \leftarrow T_{\mathsf{p}} + 1$
6: **until** $T_{\mathsf{p}} = \min(T-1, N_{\mathsf{T}})$
7: $(\boldsymbol{p}^\star, \boldsymbol{q}^\star) = \max_i f(\boldsymbol{s}(\boldsymbol{p}^\star(i), \boldsymbol{q}^\star(i)))$

because the constraints on the feasible values of the effective SNR eigenvalues $\boldsymbol{s}$ are far more intricate as is the simple trace constraint on the transmit covariance in [20].

IX. Simulations

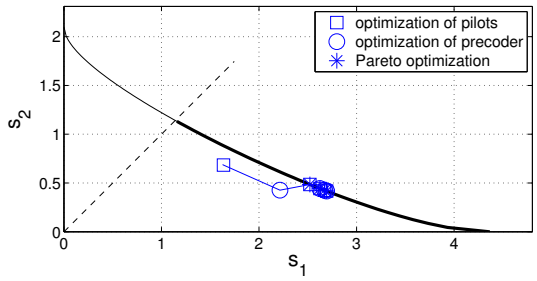

(a) convergence of the profile $\boldsymbol{s}$ towards the optimizer, located on the Pareto border (thick black line) of the feasible set $\boldsymbol{s}(\Gamma)$

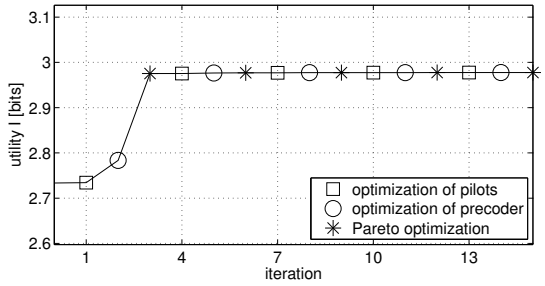

(b) convergence of the utility function $I(\boldsymbol{s})$ towards the optimum

Fig. 8. Convergence of Algorithm 1 at an SNR of 10dB ($\mu = 10$) for an exemplary 2×2 MIMO channel, both in the $\boldsymbol{s}$-domain [Fig. 8(a)] and in terms of the utility value $I(\boldsymbol{s})$ [Fig. 8(b)]

Figure 8 shows how Algorithm 1 (for fixed $T_{\mathsf{p}} = 2$) converges to the jointly optimal solution for the utility function $f(\boldsymbol{s}) = I(\boldsymbol{s})$. The parameters chosen in this simulation are $T = 10$, $\mu = 10$ (i.e., 10dB), $(N_{\mathsf{T}}, N_{\mathsf{R}}) = (2, 2)$, and $(r_1, r_2) = (\frac{2}{3}, \frac{1}{3})$.

Figures 9(a) and 9(b) respectively show the quantities

$$\bar{f}(\boldsymbol{s}) = \frac{T - T_{\mathsf{p}}}{T} I(\boldsymbol{s}) \quad -\bar{f}(\boldsymbol{s}) = \operatorname{tr} \mathsf{E}\big[(\mathbf{I} + \hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger)^{-1}\big] \tag{72}$$

optimized using Algorithm 2 and plotted against the SNR μ (in decibels) for the same 2×2 system as for Figure 8, i.e., $(r_1, r_2) = (\frac{2}{3}, \frac{1}{3})$ and $T = 10$. The former utility represents an achievable rate [including an overhead factor as in (9)]. The latter utility is the negative of the per-symbol mean-square error achieved by a linear minimum mean-square error

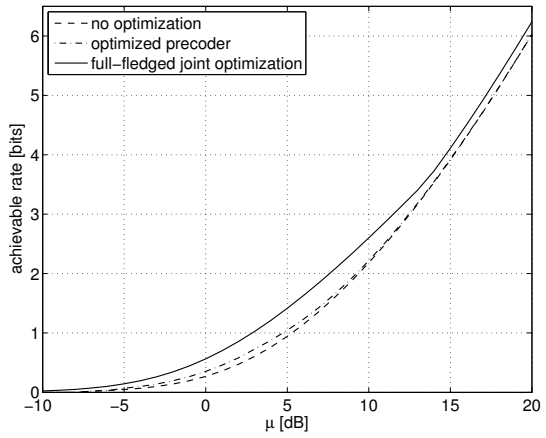

(a) achievable rate as a function of the SNR

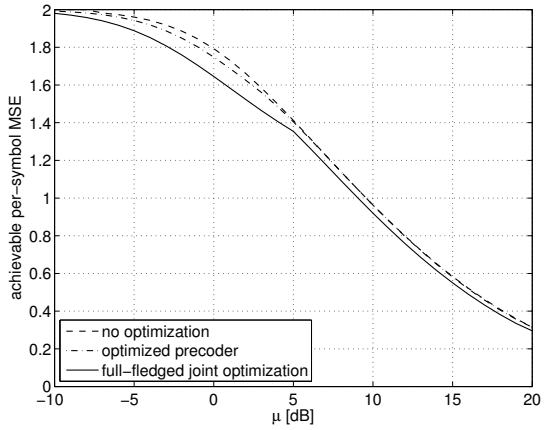

(b) achievable mean-square error as a function of the SNR

Fig. 9. Two different utility functions against the SNR parameter μ for an examplary 2×2 MIMO system. For both utilities, the performance of full-fledged optimization is compared to a partial optimization (precoder only) and no optimization at all.

symbol estimator [see also Utility 9 in Table I]. In each of the two figures, three curves are plotted for comparison: the utility achieved with full-fledged joint optimization (P.4) using Algorithm 2; the utility achieved for precoder optimization (P.1.a) alone; the utility achieved when no optimization is performed, i.e., $\boldsymbol{P}_0 = \frac{T_\mathsf{p}\mu}{N_\mathsf{T}}\mathbf{I}_{N_\mathsf{T}}$ and $\boldsymbol{Q}_0 = \frac{(T-T_\mathsf{p})\mu}{N_\mathsf{T}}\mathbf{I}_{N_\mathsf{T}}$. The relative gains in mutual information are well noticeable especially for low and moderate SNR values. For higher SNR instead, these gains are far less significant.

X. Conclusion

We have presented an in-depth study of a joint pilot-precoder optimization for a general class of utility functions of the effective SNR matrix.

Upon analyzing the two separate problems of pilot and precoder optimization, both of which can be cast into (quasi-)convex problems in the effective SNR domain provided that the utility is itself (quasi-)concave, we have shown that joint optimization requires the eigenbases of both the pilot Gram and the transmit covariance to be aligned with the channel's transmit correlation matrix. This allows to simplify the joint problem significantly. Furthermore, by deriving a method for computing energy allocations which are Pareto optimal in terms of the effective SNR eigenvalues, we managed to further reduce the dimensionality of the problem. By combining these approaches into a single iteration, we have proposed an algorithm for computing local optima of the joint optimization problem.

It should be noted that the main contributions of this work consist in offering insights and theorems that allow us to reduce a multidimensional non-convex optimization problem so as to render it tractable using common optimization tools. As pointed out in Section VIII, the focus has not been on an efficient and practical implementation, though there may be room for further improvement. Future work may thus look at possible refinements of the iterative algorithm, or extensions of the system model, for instance to include more general fading models.

Appendix

A. Examples of utilities

TABLE I
Examples of utilities from the class $\mathcal{F}$

	Utility	Curvature in $\boldsymbol{S}$
1	$I(\boldsymbol{S})$	concave
2	$\mathrm{tr}(\boldsymbol{S})$	linear
3	$\det(\boldsymbol{S})$	log-concave
4	$\mathrm{tr}(\boldsymbol{S}^{-1})^{-1}$	log-concave
5	$\det(\mathbf{I}+\nu\boldsymbol{S})$ with $\nu \geq 0$	log-concave
6	$\mathsf{E}\det(\mathbf{I}+\hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger)$	log-concave
7	$\mathsf{E}\log\det(\hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger)$ for $N_\mathsf{T} \geq N_\mathsf{R}$	concave
8	$\mathsf{E}\det(\hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger)$ for $N_\mathsf{T} \geq N_\mathsf{R}$	log-concave
9	$-\mathrm{tr}\,\mathsf{E}\left\{(\mathbf{I}+\hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger)^{-1}\right\}$	concave
10	$\mathrm{tr}\,\mathsf{E}\left\{(\boldsymbol{S}^{-1}+\hat{\boldsymbol{W}}^\dagger\hat{\boldsymbol{W}})^{-1}\right\}$ for $\det(\boldsymbol{S}) \neq 0$	concave
11	$\mathsf{Pr}\big(\det(\mathbf{I}+\hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger) \geq \eta\big)$	–/–
12	$\mathsf{Pr}\big(\log\det(\hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger) \geq \eta\big)$ for $N_\mathsf{T} \geq N_\mathsf{R}$	–/–
13	$\mathsf{Pr}\big(\det(\hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger) \geq \eta\big)$ for $N_\mathsf{T} \geq N_\mathsf{R}$	–/–
14	$\mathsf{Pr}\big(-\mathrm{tr}\left\{(\mathbf{I}+\hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger)^{-1}\right\} \geq \eta\big)$	–/–
15	$\mathsf{Pr}\big(\mathrm{tr}\left\{(\boldsymbol{S}^{-1}+\hat{\boldsymbol{W}}^\dagger\hat{\boldsymbol{W}})^{-1}\right\} \geq \eta\big)$	–/–
16	$\|\boldsymbol{S}\|_\mathsf{F}^2$	convex
17	$\lambda_{\max}(\boldsymbol{S})$	convex

In the following, we provide a few examples illustrating from which bounds or approximations of the mutual information $I(\boldsymbol{S})$ the above utilities may arise.

a) Utility 2: A simple upper bound on $I(\boldsymbol{S})$ is obtained using the fact that $\mathsf{E}[\hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger] = \mathrm{tr}(\boldsymbol{S})N_\mathsf{T}\mathbf{I}$ and by applying Jensen's inequality to the concave log-determinant:

$$I(\boldsymbol{S}) \leq \sum_{i=1}^{N_\mathsf{R}} \log(1+\mathrm{tr}(\boldsymbol{S})N_\mathsf{T}). \tag{73}$$

b) Utility 5 with $\nu = N_T N_R$: Using the determinant identity $\det(\mathbf{I}+\boldsymbol{A}\boldsymbol{B}) = \det(\mathbf{I}+\boldsymbol{B}\boldsymbol{A})$ to write $I(\boldsymbol{S}) = \mathsf{E}\log\det(\mathbf{I}+\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger\hat{\boldsymbol{W}})$, and applying Jensen's inequality, we get another upper bound:

$$I(\boldsymbol{S}) \leq \log\det\left(\mathbf{I}+N_\mathsf{T}N_\mathsf{R}\boldsymbol{S}\right). \tag{74}$$

Here, we have used $\mathsf{E}[\hat{\boldsymbol{W}}^\dagger\hat{\boldsymbol{W}}] = N_\mathsf{T}N_\mathsf{R}$.

c) Utilities 6 and 11: By applying Jensen's inequality to the concave $\log$ function, we get the upper bound

$$I(\boldsymbol{S}) \leq \log \mathsf{E}[\det(\mathbf{I} + \hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger)]. \tag{75}$$

d) Utilities 3, 7 and 12: We can lower bound $I(\boldsymbol{S})$ by removing the identity matrix inside the $\log$-determinant. Depending on the sizes of antenna arrays, this gives us a bound $I(\boldsymbol{S}) \geq \underline{I}(\boldsymbol{S})$ with

$$\underline{I}(\boldsymbol{S}) = \begin{cases} \mathsf{E}\log\det\big(\hat{\boldsymbol{W}}\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger\big) & \text{for } N_{\mathsf{T}} \geq N_{\mathsf{R}} \\ \mathsf{E}\log\det\big(\boldsymbol{S}\hat{\boldsymbol{W}}^\dagger\hat{\boldsymbol{W}}\big) & \text{for } N_{\mathsf{T}} \leq N_{\mathsf{R}}, \end{cases} \tag{76}$$

The former case (i.e., $N_{\mathsf{T}} \geq N_{\mathsf{R}}$) justifies utilities 7 and 12. In the latter case (i.e., $N_{\mathsf{T}} \leq N_{\mathsf{R}}$), note that

$$\underline{I}(\boldsymbol{S}) = \log\det(\boldsymbol{S}) + \mathsf{E}\log\det\big(\hat{\boldsymbol{W}}^\dagger\hat{\boldsymbol{W}}\big) \tag{77}$$

leads to utility 3. Clearly, $\underline{I}(\boldsymbol{S})$ is good as an approximation of $I(\boldsymbol{S})$ at high SNR, and was used as such in [21], [22]. Let us also mention the tighter lower bound [23]

$$I(\boldsymbol{S}) \geq N_{\mathsf{R}} \log\left(1 + \exp\left(\frac{\log(\mathrm{e})}{N_{\mathsf{R}}}\underline{I}(\boldsymbol{S})\right)\right), \tag{78}$$

the derivation of which makes use of the Minkowski inequality for determinants.

B. Proof of Theorem V.1

Let $\boldsymbol{P}$ and $\boldsymbol{Q}$ have $\mathrm{rank}(\boldsymbol{P}) = r_{\boldsymbol{P}}$ and $\mathrm{rank}(\boldsymbol{Q}) = r_{\boldsymbol{Q}}$ respectively, with $r_{\boldsymbol{P}}, r_{\boldsymbol{Q}} \in \{1, \ldots, N_{\mathsf{T}}\}$. The absolute difference of ranks be $d = |r_{\boldsymbol{P}} - r_{\boldsymbol{Q}}|$. The pilot matrix and precoder have reduced spectral decompositions $\boldsymbol{P} = \boldsymbol{U}_{\boldsymbol{P}}\boldsymbol{\Lambda}_{\boldsymbol{P}}\boldsymbol{U}_{\boldsymbol{P}}^\dagger$ and $\boldsymbol{Q} = \boldsymbol{U}_{\boldsymbol{Q}}\boldsymbol{\Lambda}_{\boldsymbol{Q}}\boldsymbol{U}_{\boldsymbol{Q}}^\dagger$, respectively, where the eigenbases $\boldsymbol{U}_{\boldsymbol{P}} \in \mathbb{U}^{N_{\mathsf{T}} \times r_{\boldsymbol{P}}}$ and $\boldsymbol{U}_{\boldsymbol{Q}} \in \mathbb{U}^{N_{\mathsf{T}} \times r_{\boldsymbol{Q}}}$ are tall or square, whereas $\boldsymbol{\Lambda}_{\boldsymbol{P}}$ and $\boldsymbol{\Lambda}_{\boldsymbol{Q}}$ are diagonal and positive definite. Let $\boldsymbol{U}_{\boldsymbol{P}^\perp} \in \mathbb{U}^{N_{\mathsf{T}} \times (N_{\mathsf{T}} - r_{\boldsymbol{P}})}$ and $\boldsymbol{U}_{\boldsymbol{Q}^\perp} \in \mathbb{U}^{N_{\mathsf{T}} \times (N_{\mathsf{T}} - r_{\boldsymbol{Q}})}$ denote orthonormal bases of the nullspaces of $\boldsymbol{P}$ and $\boldsymbol{Q}$, respectively, so that $\boldsymbol{U}_{\boldsymbol{P}}^\dagger\boldsymbol{U}_{\boldsymbol{P}^\perp} = \mathbf{0}$ and $\boldsymbol{U}_{\boldsymbol{Q}}^\dagger\boldsymbol{U}_{\boldsymbol{Q}^\perp} = \mathbf{0}$.

The reduced eigendecomposition of $\hat{\boldsymbol{R}}$ is consistently denoted as $\boldsymbol{U}_{\hat{\boldsymbol{R}}}\boldsymbol{\Lambda}_{\hat{\boldsymbol{R}}}\boldsymbol{U}_{\hat{\boldsymbol{R}}}^\dagger$, where $\boldsymbol{\Lambda}_{\hat{\boldsymbol{R}}} \in \mathbb{R}_+^{r_{\boldsymbol{P}} \times r_{\boldsymbol{P}}}$ is diagonal positive definite and of size $r_{\boldsymbol{P}} \times r_{\boldsymbol{P}}$, due to the rank equality (18) which states that $\mathrm{rank}(\hat{\boldsymbol{R}}) = \mathrm{rank}(\boldsymbol{P})$. The orthonormal nullspace of $\hat{\boldsymbol{R}}$ is denoted as $\boldsymbol{U}_{\hat{\boldsymbol{R}}^\perp} \in \mathbb{U}^{N_{\mathsf{T}} \times (N_{\mathsf{T}} - r_{\boldsymbol{P}})}$. We introduce the notation $\boldsymbol{L}_{\boldsymbol{A} \cap \boldsymbol{B}}$ to designate an orthonormal basis of the intersection of range spaces $\mathrm{range}(\boldsymbol{A})$ and $\mathrm{range}(\boldsymbol{B})$. If it exists, $\boldsymbol{L}_{\boldsymbol{A} \cap \boldsymbol{B}}$ is a matrix with the maximum number of columns defined as

$$\boldsymbol{L}_{\boldsymbol{A} \cap \boldsymbol{B}} = \left\{ \boldsymbol{L} \;\middle|\; \begin{array}{l} \boldsymbol{L}^\dagger\boldsymbol{L} = \mathbf{I}, \\ \forall \boldsymbol{x} \neq \mathbf{0} \colon \boldsymbol{A}\boldsymbol{L}\boldsymbol{x} \neq \mathbf{0} \text{ and } \boldsymbol{B}\boldsymbol{L}\boldsymbol{x} \neq \mathbf{0} \end{array} \right\}. \tag{79}$$

Assume that $\mathrm{range}(\boldsymbol{Q}) \nsubseteq \mathrm{range}(\hat{\boldsymbol{R}})$, so the matrix $\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}$ is defined and has at least one column. We define a new precoder $\boldsymbol{Q}' \in \mathcal{Q}(\mu_{\mathcal{Q}})$ as

$$\boldsymbol{Q}' = \boldsymbol{Q} - \lambda_{r_{\boldsymbol{Q}}}(\boldsymbol{Q})\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}^\dagger, \tag{80}$$

where $\lambda_{r_{\boldsymbol{Q}}}(\boldsymbol{Q})$ is the smallest non-zero eigenvalue of $\boldsymbol{Q}$. First, we verify that $\boldsymbol{Q}' \in \mathcal{Q}(\mu_{\mathcal{Q}})$. Clearly, since $\lambda_{r_{\boldsymbol{Q}}}(\boldsymbol{Q}) \geq 0$, we have $\boldsymbol{Q}' \preceq \boldsymbol{Q}$, and therefore $\mathrm{tr}(\boldsymbol{Q}') \leq \mathrm{tr}(\boldsymbol{Q})$. What remains to prove is that $\boldsymbol{Q}' \succeq \mathbf{0}$. The smallest eigenvalue of $\boldsymbol{Q}'$ is

$$\lambda_{\min}(\boldsymbol{Q}') = \min_{\|\boldsymbol{w}\|=1} \boldsymbol{w}^\dagger\boldsymbol{Q}'\boldsymbol{w}.$$

But since by definition of $\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}}$, the range space of $\boldsymbol{Q}$ contains the range space of $\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}}\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}}^\dagger$, we have that $\boldsymbol{w}^\dagger\boldsymbol{Q}'\boldsymbol{w}$ is equal to $\boldsymbol{w}^\dagger\boldsymbol{\Pi}_{\boldsymbol{Q}}^\dagger\boldsymbol{Q}'\boldsymbol{\Pi}_{\boldsymbol{Q}}\boldsymbol{w}$, where $\boldsymbol{\Pi}_{\boldsymbol{Q}} = \boldsymbol{U}_{\boldsymbol{Q}}(\boldsymbol{U}_{\boldsymbol{Q}}^\dagger\boldsymbol{U}_{\boldsymbol{Q}})^{-1}\boldsymbol{U}_{\boldsymbol{Q}}^\dagger$ is the projector from $\mathbb{C}^{N_{\mathsf{T}} \times N_{\mathsf{T}}}$ onto the basis $\boldsymbol{U}_{\boldsymbol{Q}}$. We thus have

$$\begin{aligned} &\lambda_{\min}(\boldsymbol{Q}') \\ &\quad = \lambda_{\min}\left(\boldsymbol{\Pi}_{\boldsymbol{Q}}^\dagger\Big(\boldsymbol{Q} - \lambda_{r_{\boldsymbol{Q}}}(\boldsymbol{Q})\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}^\dagger\Big)\boldsymbol{\Pi}_{\boldsymbol{Q}}\right) \\ &\quad \geq \lambda_{\min}\Big(\boldsymbol{\Pi}_{\boldsymbol{Q}}^\dagger\boldsymbol{Q}\boldsymbol{\Pi}_{\boldsymbol{Q}}\Big) \\ &\qquad\qquad - \lambda_{r_{\boldsymbol{Q}}}(\boldsymbol{Q})\lambda_{\max}\Big(\boldsymbol{\Pi}_{\boldsymbol{Q}}^\dagger\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}^\dagger\boldsymbol{\Pi}_{\boldsymbol{Q}}\Big) \\ &\quad \geq \lambda_{r_{\boldsymbol{Q}}}(\boldsymbol{Q})\Big(1 - \lambda_{\max}(\boldsymbol{\Pi}_{\boldsymbol{Q}}^\dagger\boldsymbol{\Pi}_{\boldsymbol{Q}})\lambda_{\max}\big(\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}^\dagger\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}\big)\Big) \\ &\quad = 0. \end{aligned} \tag{81}$$

The second inequality holds because the spectral radius norm $\lambda_{\max}(\cdot)$ is sub-multiplicative, while the last equality holds because the projector $\boldsymbol{\Pi}_{\boldsymbol{Q}}$ and the (sub)unitary $\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}$ have a largest singular value of at most 1. We infer that $\boldsymbol{Q}' \in \mathcal{Q}(\mu_{\mathcal{Q}})$.

Notice that $\boldsymbol{Q}'$ is purposely constructed so that $\boldsymbol{Q}'\hat{\boldsymbol{R}} = \boldsymbol{Q}\hat{\boldsymbol{R}}$. As compared to the matrix $\boldsymbol{S} = \boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q})$ obtained with the precoder $\boldsymbol{Q}$, the new matrix $\boldsymbol{S}' = \boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q}')$ thus reads as

$$\begin{aligned} \boldsymbol{S}' &= \frac{\hat{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{Q}'\hat{\boldsymbol{R}}^{\frac{1}{2}}}{1 + \mathrm{tr}(\boldsymbol{Q}'\tilde{\boldsymbol{R}})} \\ &= \frac{\hat{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{Q}\hat{\boldsymbol{R}}^{\frac{1}{2}}}{1 + \mathrm{tr}(\boldsymbol{Q}\tilde{\boldsymbol{R}}) - \lambda_{r_{\boldsymbol{Q}}}(\boldsymbol{Q})\,\mathrm{tr}\big(\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}^\dagger\tilde{\boldsymbol{R}}\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}\big)} \\ &= k\boldsymbol{S} \end{aligned} \tag{82}$$

and thus turns out to be a scaled version of $\boldsymbol{S}$, where the positive scalar k is

$$k = \frac{1 + \mathrm{tr}(\boldsymbol{Q}\tilde{\boldsymbol{R}})}{1 + \mathrm{tr}(\boldsymbol{Q}\tilde{\boldsymbol{R}}) - \lambda_{r_{\boldsymbol{Q}}}(\boldsymbol{Q})\,\mathrm{tr}\big(\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}^\dagger\tilde{\boldsymbol{R}}\boldsymbol{L}_{\boldsymbol{Q} \cap \hat{\boldsymbol{R}}^\perp}\big)} > 1.$$

Therefore, we have $\boldsymbol{S}' \succ \boldsymbol{S}$, so the precoder $\boldsymbol{Q}$ is necessarily suboptimal, which means that $\mathrm{range}(\boldsymbol{Q}) \nsubseteq \mathrm{range}(\hat{\boldsymbol{R}})$ cannot hold for optimal $\boldsymbol{Q}$. Instead, we must have $\mathrm{range}(\boldsymbol{Q}) \subseteq \mathrm{range}(\hat{\boldsymbol{R}})$ for optimality, which concludes the proof.

C. Proof of Lemma V.1

Supposing we are in the first situation, i.e., $\boldsymbol{A}$ has full column rank, then $\boldsymbol{A}$ has a left pseudoinverse $\boldsymbol{A}^\sharp = (\boldsymbol{A}^\dagger\boldsymbol{A})^{-1}\boldsymbol{A}^\dagger$ which can be used to define the inverse function ϱ^{-1}. Let $\boldsymbol{Z} = \varrho(\boldsymbol{X})$ be the image of $\boldsymbol{X}$. Given $\boldsymbol{Z}$, one obtains $\boldsymbol{X}$ by insulating it via left-multiplication with $\boldsymbol{A}^\sharp$ and right-multiplication with $\boldsymbol{A}^{\sharp\dagger}$, and appropriate scaling:

$$\boldsymbol{A}^\sharp\boldsymbol{Z}\boldsymbol{A}^{\sharp\dagger}(1 + \mathrm{tr}(\boldsymbol{B}\boldsymbol{X})) = \boldsymbol{X} \tag{83}$$

Left-multiplying (83) with $\boldsymbol{B}$ and taking the trace yields

$$\mathrm{tr}(\boldsymbol{B}\boldsymbol{A}^\sharp\boldsymbol{Z}\boldsymbol{A}^{\sharp\dagger}) = \frac{\mathrm{tr}(\boldsymbol{B}\boldsymbol{X})}{1 + \mathrm{tr}(\boldsymbol{B}\boldsymbol{X})}, \tag{84}$$

or equivalently,

$$1+\operatorname{tr}(\boldsymbol{B}\boldsymbol{X})=\frac{1}{1-\operatorname{tr}(\boldsymbol{B}\boldsymbol{A}^{\sharp}\boldsymbol{Z}\boldsymbol{A}^{\sharp\dagger})} \tag{85}$$

By combining (83) with (85), we recover the pre-image $\boldsymbol{X}=\varrho^{-1}(\boldsymbol{Z})$, and see that the inverse function ϱ^{-1} is linear fractional with parameters $\boldsymbol{A}^{\sharp}$ and $-\boldsymbol{A}^{\sharp\dagger}\boldsymbol{B}\boldsymbol{A}^{\sharp}$.

We now suppose that we are in the second situation, i.e., $\boldsymbol{A}$ has full row rank and $\mathcal{X}$ is such that $\operatorname{range}(\boldsymbol{X})=\operatorname{range}(\boldsymbol{A}^{\dagger})$ for every element of $\mathcal{X}$. Due to the latter constraint on the span of $\boldsymbol{X}$, we can write any $\boldsymbol{X}\in\mathcal{X}$ as $\boldsymbol{X}=\boldsymbol{A}^{\dagger}\hat{\boldsymbol{X}}\boldsymbol{A}$, with $\hat{\boldsymbol{X}}$ given by the inverse relation $\hat{\boldsymbol{X}}=\boldsymbol{A}^{\flat\dagger}\boldsymbol{X}\boldsymbol{A}^{\flat}$, where $\boldsymbol{A}^{\flat}=\boldsymbol{A}^{\dagger}(\boldsymbol{A}\boldsymbol{A}^{\dagger})^{-1}$ denotes the right pseudoinverse of $\boldsymbol{A}$. The function ϱ can be represented as

$$\varrho\colon\ \boldsymbol{X}\ \mapsto\ \frac{\boldsymbol{A}\boldsymbol{X}\boldsymbol{A}^{\dagger}}{1+\operatorname{tr}(\boldsymbol{B}\boldsymbol{X})}=\frac{\hat{\boldsymbol{A}}\hat{\boldsymbol{X}}\hat{\boldsymbol{A}}^{\dagger}}{1+\operatorname{tr}(\hat{\boldsymbol{B}}\hat{\boldsymbol{X}})} \tag{86}$$

with abbreviations $\boldsymbol{A}\boldsymbol{A}^{\dagger}\triangleq\hat{\boldsymbol{A}}$ and $\hat{\boldsymbol{B}}\triangleq\boldsymbol{A}\boldsymbol{B}\boldsymbol{A}^{\dagger}$. Since $\hat{\boldsymbol{A}}=\boldsymbol{A}\boldsymbol{A}^{\dagger}$ has full rank (because $\boldsymbol{A}$ has full row rank), the function ϱ appears as an injective linear fractional function of $\hat{\boldsymbol{X}}$ with parameters $\hat{\boldsymbol{A}}$ and $\hat{\boldsymbol{B}}$, whose inverse, according to findings above, is linear fractional with parameters $\hat{\boldsymbol{A}}^{\sharp}$ and $-\hat{\boldsymbol{A}}^{\sharp\dagger}\hat{\boldsymbol{B}}\hat{\boldsymbol{A}}^{\sharp}$. Denoting as $\boldsymbol{Z}=\varrho(\boldsymbol{X})$ the image of $\boldsymbol{X}$ under the function ϱ, we can thus recover the pre-image $\boldsymbol{X}$ from $\boldsymbol{Z}$ as

$$\begin{aligned}\boldsymbol{X}=\boldsymbol{A}^{\dagger}\hat{\boldsymbol{X}}\boldsymbol{A}&=\boldsymbol{A}^{\dagger}\frac{\hat{\boldsymbol{A}}^{\sharp}\boldsymbol{Z}\hat{\boldsymbol{A}}^{\sharp\dagger}}{1-\operatorname{tr}(\hat{\boldsymbol{A}}^{\sharp\dagger}\boldsymbol{B}\hat{\boldsymbol{A}}^{\sharp}\boldsymbol{Z})}\boldsymbol{A}\\&=\frac{\boldsymbol{A}^{\flat}\boldsymbol{Z}\boldsymbol{A}^{\flat\dagger}}{1-\operatorname{tr}(\boldsymbol{A}^{\flat\dagger}\boldsymbol{B}\boldsymbol{A}^{\flat}\boldsymbol{Z})}.\end{aligned} \tag{87}$$

Consequently, the inverse ϱ^{-1} is linear fractional with parameters $\boldsymbol{A}^{\flat}$ and $-\boldsymbol{A}^{\flat\dagger}\boldsymbol{B}\boldsymbol{A}^{\flat}$.

D. Proof of Lemma V.2

We first establish that $\boldsymbol{S}(\boldsymbol{P},\boldsymbol{Q})$ is monotonic in the transmit power, in the sense that

$$0\leq k<k'\Rightarrow\boldsymbol{S}(\boldsymbol{P},k\boldsymbol{Q})\prec\boldsymbol{S}(\boldsymbol{P},k'\boldsymbol{Q}). \tag{88}$$

This monotonicity holds because

$$\boldsymbol{S}(\boldsymbol{P},k\boldsymbol{Q})=\frac{k}{1+k\operatorname{tr}(\boldsymbol{Q}\tilde{\boldsymbol{R}})}\hat{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{Q}\hat{\boldsymbol{R}}^{\frac{1}{2}} \tag{89}$$

$$<\frac{k'}{1+k'\operatorname{tr}(\boldsymbol{Q}\tilde{\boldsymbol{R}})}\hat{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{Q}\hat{\boldsymbol{R}}^{\frac{1}{2}}=\boldsymbol{S}(\boldsymbol{P},k'\boldsymbol{Q}) \tag{90}$$

owing to the fact that $k\mapsto\frac{k}{1+k\operatorname{tr}(\boldsymbol{Q}\tilde{\boldsymbol{R}})}$ is monotonically increasing in k.

Now, as a consequence of Theorem V.1 and of the power monotonicity (88), optimal precoders $\boldsymbol{Q}$ will be elements of $\partial^{+}\mathcal{Q}(\mu_{\mathcal{Q}})\cap\operatorname{range}(\hat{\boldsymbol{R}})$. By definition, this set can be parametrized by $r_{\boldsymbol{P}}$ non-negative coefficients $\boldsymbol{\psi}=[\psi_1,\ldots,\psi_{r_{\boldsymbol{P}}}]^{\mathrm{T}}\in\mathbb{R}_{+}^{r_{\boldsymbol{P}}}$ stored in a diagonal matrix $\boldsymbol{\Psi}=\operatorname{diag}(\boldsymbol{\psi})$, and a tall or square (sub)unitary basis $\boldsymbol{\Upsilon}\in\mathbb{U}^{N_{\mathrm{T}}\times r_{\boldsymbol{P}}}$ as follows:

$$\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}}=\big(\hat{\boldsymbol{R}}^{\frac{1}{2}}\big)^{+}\boldsymbol{\Upsilon}\boldsymbol{\Psi}\boldsymbol{\Upsilon}^{\dagger}\big(\hat{\boldsymbol{R}}^{\frac{1}{2}}\big)^{+}, \tag{91}$$

where $(\bullet)^{+}$ denotes the Moore-Penrose pseudoinverse (cf. Section II), and where the parameter pair $(\boldsymbol{\Psi},\boldsymbol{\Upsilon})$ shall be subject to the four constraints

$$\boldsymbol{\psi}\geq\mathbf{0}, \tag{92a}$$

$$\operatorname{tr}(\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}})=\mu_{\mathcal{Q}}, \tag{92b}$$

$$\boldsymbol{\Upsilon}^{\dagger}\boldsymbol{\Upsilon}=\mathbf{I}, \tag{92c}$$

$$\operatorname{range}(\boldsymbol{\Upsilon})=\operatorname{range}(\hat{\boldsymbol{R}}). \tag{92d}$$

The first two constraints ensure that $\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}}$ belongs to $\partial^{+}\mathcal{Q}(\mu_{\mathcal{Q}})$, while the structure of Expression (91) ensures that $\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}}$ belongs to $\operatorname{range}(\hat{\boldsymbol{R}})$. The third and fourth constraints (92c)–(92d) are clearly not necessary to fulfill $\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}}\in\operatorname{range}(\hat{\boldsymbol{R}})\cap\partial^{+}\mathcal{Q}(\mu_{\mathcal{Q}})$, but they induce no loss of generality either and will turn out helpful later. The set $\partial^{+}\mathcal{Q}(\mu_{\mathcal{Q}})$ is thus entirely parametrized by the parameter pair $(\boldsymbol{\Psi},\boldsymbol{\Upsilon})$ subject to the constraints (92). Consider now the feasible vectors

$$\begin{aligned}\boldsymbol{s}(\boldsymbol{P},\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}})&=\boldsymbol{\lambda}\left(\frac{\hat{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}}\hat{\boldsymbol{R}}^{\frac{1}{2}}}{1+\operatorname{tr}(\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}}\tilde{\boldsymbol{R}})}\right)\\&=\boldsymbol{\lambda}\left(\frac{\boldsymbol{U}_{\hat{\boldsymbol{R}}}\boldsymbol{U}_{\hat{\boldsymbol{R}}}^{\dagger}\boldsymbol{\Upsilon}\boldsymbol{\Psi}\boldsymbol{\Upsilon}^{\dagger}\boldsymbol{U}_{\hat{\boldsymbol{R}}}\boldsymbol{U}_{\hat{\boldsymbol{R}}}^{\dagger}}{1+\operatorname{tr}(\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}}\tilde{\boldsymbol{R}})}\right)\in\mathbb{R}_{+}^{N_{\mathrm{T}}}.\end{aligned} \tag{93}$$

This vector has at most $r_{\boldsymbol{P}}$ non-zero entries because $\boldsymbol{\Psi}$ is $r_{\boldsymbol{P}}\times r_{\boldsymbol{P}}$. Therefore, we define a vector $\bar{\boldsymbol{s}}\in\mathbb{R}_{+}^{r_{\boldsymbol{P}}}$ of reduced dimension, which contains the $r_{\boldsymbol{P}}$ topmost (i.e., largest) entries of $\boldsymbol{s}$. Since $\operatorname{range}(\boldsymbol{\Upsilon})=\operatorname{range}(\hat{\boldsymbol{R}})$ [cf. (92d)], the matrix $\boldsymbol{U}_{\hat{\boldsymbol{R}}}^{\dagger}\boldsymbol{\Upsilon}$ is unitary, so we have

$$\begin{aligned}\bar{\boldsymbol{s}}(\boldsymbol{P},\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}})&=\frac{\boldsymbol{\psi}}{1+\operatorname{tr}(\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}}\tilde{\boldsymbol{R}})}\\&=\frac{\boldsymbol{\psi}}{1+\operatorname{tr}\big(\boldsymbol{\Upsilon}^{\dagger}\big(\hat{\boldsymbol{R}}^{\frac{1}{2}}\big)^{+}\tilde{\boldsymbol{R}}\big(\hat{\boldsymbol{R}}^{\frac{1}{2}}\big)^{+}\boldsymbol{\Upsilon}\boldsymbol{\Psi}\big)}.\end{aligned} \tag{94}$$

Note that we have not assumed so far that the entries of $\boldsymbol{\psi}$ or $\bar{\boldsymbol{s}}$ are sorted in any specific order. For notational brevity, call $\boldsymbol{\alpha}$ the vector of entries $\alpha_i=\big[\boldsymbol{\Upsilon}^{\dagger}\big(\hat{\boldsymbol{R}}^{\frac{1}{2}}\big)^{+}\tilde{\boldsymbol{R}}\big(\hat{\boldsymbol{R}}^{\frac{1}{2}}\big)^{+}\boldsymbol{\Upsilon}\big]_{i,i}$, then

$$\bar{\boldsymbol{s}}(\boldsymbol{P},\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}})=\frac{\boldsymbol{\psi}}{1+\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{\psi}}. \tag{95}$$

On the other hand, the second constraint (92b) translates to $\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{\psi}=\mu_{\mathcal{Q}}$, where $\boldsymbol{\beta}$ denotes the vector of diagonal entries of $\boldsymbol{\Upsilon}^{\dagger}\hat{\boldsymbol{R}}^{+}\boldsymbol{\Upsilon}$, i.e., $\beta_i=[\boldsymbol{\Upsilon}^{\dagger}\hat{\boldsymbol{R}}^{+}\boldsymbol{\Upsilon}]_{i,i}$. Together, this constraint and equation (95) describe an affine plane of dimension $r_{\boldsymbol{P}}-1$, because left-multiplying (95) with $\frac{1}{\mu_{\mathcal{Q}}}\boldsymbol{\beta}^{\mathrm{T}}+\boldsymbol{\alpha}^{\mathrm{T}}$ leads to the affine equation

$$\left(\frac{1}{\mu_{\mathcal{Q}}}\boldsymbol{\beta}+\boldsymbol{\alpha}\right)^{\mathrm{T}}\bar{\boldsymbol{s}}(\boldsymbol{P},\boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}})=1. \tag{96}$$

This affine equation, together with the non-negativity constraint $\boldsymbol{\psi}\geq\mathbf{0}$ [cf. (92a)], thus delimit a $(r_{\boldsymbol{P}}-1)$-dimensional simplex, whose elements fulfill

$$\begin{cases}\displaystyle\sum_{i}\frac{\bar{s}_i}{\omega_i}=1\\ \bar{\boldsymbol{s}}\geq\mathbf{0}\end{cases} \tag{97}$$

where

$$\omega_i = \frac{1}{\frac{1}{\mu_{\mathcal{Q}}}\beta_i + \alpha_i}. \tag{98}$$

The $r_{\boldsymbol{P}}$ vertices of the simplex described by (97) are the axis points $\omega_i \boldsymbol{e}_i$.

Due to the symmetry property of utilities from the class $\mathcal{F}$, the ordering of the $\bar{s}_i$ does not influence the utility value. Assume that, for a given $\bar{\boldsymbol{s}} \geq \boldsymbol{0}$ fulfilling (97), there exists an index permutation π such that $\sum_i \frac{\bar{s}_{\pi(i)}}{\omega_i} < 1$, then $\bar{\boldsymbol{s}}$ is suboptimal, since there exists an $\bar{\boldsymbol{s}}' = k\boldsymbol{\Pi}\bar{\boldsymbol{s}}$ with $k > 1$ which also fulfills (97) and yields a larger utility value $f(\bar{\boldsymbol{s}}') = f(k\boldsymbol{\Pi}\bar{\boldsymbol{s}}) = f(k\bar{\boldsymbol{s}}) > f(\bar{\boldsymbol{s}})$. Therefore, we can discard all $\bar{\boldsymbol{s}}$ for which some permutation π yields $\sum_i \frac{\bar{s}_{\pi(i)}}{\omega_i} < 1$. This is equivalent to the requirement that the $\bar{s}_i$ and ω_i be ordered in the same way, i.e.,

$$\omega_i \leq \omega_j \Rightarrow \bar{s}_i \leq \bar{s}_j. \tag{99}$$

Hence, without loss of optimality, we will restrain the set of admissible $\bar{\boldsymbol{s}}$ to the following convex set, called $\mathcal{S}$:

$$\mathcal{S} = \left\{ \bar{\boldsymbol{s}} \in \mathbb{R}_+^{r_{\boldsymbol{P}}} \;\middle|\; \sum_{i=1}^{r_{\boldsymbol{P}}} \frac{\bar{s}_i}{\bar{\omega}_i} = 1, \forall j \colon \bar{s}_j \geq \bar{s}_{j+1} \right\}, \tag{100}$$

where $\bar{\boldsymbol{\omega}} = [\bar{\omega}_1, \ldots, \bar{\omega}_{r_{\boldsymbol{P}}}]^{\mathrm{T}}$ contains the entries of $\boldsymbol{\omega}$ arranged in non-increasing order, i.e., $\bar{\omega}_1 \geq \ldots \geq \bar{\omega}_{r_{\boldsymbol{P}}}$. Let us define $r_{\boldsymbol{P}}$ special points pertaining to $\mathcal{S}$, which we shall denote as $\boldsymbol{\sigma}^{(n)}$, and define as

$$\forall n = 1, \ldots, r_{\boldsymbol{P}} \colon \quad \boldsymbol{\sigma}^{(n)} = \mathcal{H}(\bar{\omega}_1, \ldots, \bar{\omega}_n) \sum_{j=1}^{n} \boldsymbol{e}_j, \tag{101}$$

where $\mathcal{H}(\cdot, \ldots, \cdot)$ and $\boldsymbol{e}_j$ are defined in the statement of Lemma V.2. In fact, it is easy to see that the $\boldsymbol{\sigma}^{(n)}$ have non-increasing entries and fulfill $\sum_{i=1}^{r_{\boldsymbol{P}}} \frac{\sigma_i^{(n)}}{\bar{\omega}_i} = 1$, and thus belong to $\mathcal{S}$. Now, we will show that the set $\mathcal{S}'$ of all convex combinations of the $\boldsymbol{\sigma}^{(n)}$, i.e,

$$\mathcal{S}' = \left\{ \sum_{n=1}^{r_{\boldsymbol{P}}} \nu_n \boldsymbol{\sigma}^{(n)} \middle| \sum_n \nu_n = 1, \forall n \colon \nu_n \geq 0 \right\}, \tag{102}$$

is the same as the set $\mathcal{S}$. We know that $\mathcal{S}'$ is a subset of the convex set $\mathcal{S}$, for being a convex combination of a collection of points from $\mathcal{S}$, hence $\mathcal{S}' \subseteq \mathcal{S}$. Next, we argue that, if we assume that some particular point $\tilde{\boldsymbol{\sigma}}$ belongs to $\mathcal{S} \setminus \mathcal{S}'$, this implies that $\tilde{\boldsymbol{\sigma}}$ does not lie in $\mathcal{S}$ because it would fail to comply with some constraint from the definition (100) of $\mathcal{S}$. Therefrom, it will follow that $\mathcal{S} = \mathcal{S}'$.

Since the $\boldsymbol{\sigma}^{(n)}$ pertain to $\mathcal{S}$ and are $r_{\boldsymbol{P}}$ linearly independent vectors, they define the $(r_{\boldsymbol{P}} - 1)$-dimensional affine plane described by $\sum_{i=1}^{r_{\boldsymbol{P}}} \frac{\bar{s}_i}{\bar{\omega}_i} = 1$. Therefore, to prove the equality $\mathcal{S} = \mathcal{S}'$, it will be sufficient to take some point $\tilde{\boldsymbol{\sigma}} = [\tilde{\sigma}_1, \ldots, \tilde{\sigma}_{r_{\boldsymbol{P}}}]^{\mathrm{T}}$ to lie on said plane, and show that an infringement of an inequality $\tilde{\sigma}_i \geq \tilde{\sigma}_{i+1}$ implies that $\tilde{\boldsymbol{\sigma}} = \sum_{n=1}^{r_{\boldsymbol{P}}} \tilde{\nu}_n \boldsymbol{\sigma}^{(n)}$ with coefficients $\tilde{\nu}_n$ such that either $\sum_n \tilde{\nu}_n \neq 1$ or $\tilde{\nu}_n < 0$ for some index n. So, assume that $\tilde{\sigma}_i < \tilde{\sigma}_{i+1}$ for a given i. There exist unique coefficients $\tilde{\nu}_n$ such that $\tilde{\boldsymbol{\sigma}} = \sum_{n=1}^{r_{\boldsymbol{P}}} \tilde{\nu}_n \boldsymbol{\sigma}^{(n)}$. The inequality $\tilde{\sigma}_i < \tilde{\sigma}_{i+1}$ can thus be written as

$$\sum_{n=1}^{r_{\boldsymbol{P}}} \tilde{\nu}_n \boldsymbol{e}_i^{\mathrm{T}} \boldsymbol{\sigma}^{(n)} < \sum_{n=1}^{r_{\boldsymbol{P}}} \tilde{\nu}_n \boldsymbol{e}_{i+1}^{\mathrm{T}} \boldsymbol{\sigma}^{(n)}. \tag{103}$$

By inserting (101) into the latter inequality, we get

$$\sum_{n=i}^{r_{\boldsymbol{P}}} \tilde{\nu}_n \mathcal{H}(\bar{\omega}_1, \ldots, \bar{\omega}_n) < \sum_{n=i+1}^{r_{\boldsymbol{P}}} \tilde{\nu}_n \mathcal{H}(\bar{\omega}_1, \ldots, \bar{\omega}_n), \tag{104}$$

which boils down to $\tilde{\nu}_i < 0$. This concludes the proof that $\mathcal{S} = \mathcal{S}'$. Also, this simplex $\mathcal{S}$ contains only Pareto border points, in the sense that $\mathcal{S} = \partial^+ \mathcal{S}$. In fact, any point $\bar{\boldsymbol{s}}''$ dominating some point $\bar{\boldsymbol{s}}' \in \mathcal{S}$ would fulfill $\sum_{i=1}^{r_{\boldsymbol{P}}} \bar{s}_i''/\bar{\omega}_i > 1$ and thus lie outside $\mathcal{S}$.

Now that we have fully characterized the set of Pareto border points $\bar{\boldsymbol{s}} = \bar{\boldsymbol{s}}(\boldsymbol{P}, \boldsymbol{Q}_{\boldsymbol{\Psi},\boldsymbol{\Upsilon}})$ for a fixed $\boldsymbol{\Upsilon}$ as a simplex set $\mathcal{S}$, we ask what the best choice for $\boldsymbol{\Upsilon}$ is under the constraints (92c)–(92d). Clearly, if there exists one single $\boldsymbol{\Upsilon}^\star$ that simultaneously maximizes all vertices $\boldsymbol{\sigma}^{(n)}$ in the sense that for any $\boldsymbol{\Upsilon}$, we have

$$\boldsymbol{\sigma}^{(n)}(\boldsymbol{\Upsilon}^\star) \geq \boldsymbol{\sigma}^{(n)}(\boldsymbol{\Upsilon}), \qquad n = 1, \ldots, r_{\boldsymbol{P}} \tag{105}$$

then this $\boldsymbol{\Upsilon}^\star$ is optimal. Here, $\boldsymbol{\sigma}^{(n)}(\boldsymbol{\Upsilon})$ denotes the value of $\boldsymbol{\sigma}^{(n)}$, as defined in (101), with the $\bar{\omega}_i$ interpreted as functions of $\boldsymbol{\Upsilon}$. Next, we show that such $\boldsymbol{\Upsilon}^\star$ is well-defined and characterize it.

We state the multiobjective optimization problem

$$\forall n = 1, \ldots, r_{\boldsymbol{P}} \colon \max_{\substack{\boldsymbol{\Upsilon} \in \mathbb{U}^{N_{\mathrm{T}} \times r_{\boldsymbol{P}}} \\ \mathrm{range}(\boldsymbol{\Upsilon}) = \mathrm{range}(\hat{\boldsymbol{R}})}} \mathcal{H}(\bar{\omega}_1, \ldots, \bar{\omega}_n). \tag{106}$$

Omitting the range space constraint on $\boldsymbol{\Upsilon}$, we have that, with the definition (98) of the coefficients ω_i together with the definitions of α_i and β_i, this multiobjective problem reads as

$$\forall n \colon \min_{\boldsymbol{\Upsilon} \in \mathbb{U}^{N_{\mathrm{T}} \times r_{\boldsymbol{P}}}} \sum_{i=1}^{n} \Big[\boldsymbol{\Upsilon}^\dagger \big(\hat{\boldsymbol{R}}^{\frac{1}{2}}\big)^+ \big(\tfrac{\mathbf{I}}{\mu_{\mathcal{Q}}} + \tilde{\boldsymbol{R}}\big) \big(\hat{\boldsymbol{R}}^{\frac{1}{2}}\big)^+ \boldsymbol{\Upsilon} \Big]_{\pi(i),\pi(i)}, \tag{107}$$

where π denotes the permutation which orders the diagonal entries of the matrix between square brackets so as to be non-decreasingly ordered. If $\boldsymbol{W}\boldsymbol{D}\boldsymbol{W}^\dagger$ denotes the spectral decomposition of $\big(\hat{\boldsymbol{R}}^{\frac{1}{2}}\big)^+ \big(\frac{\mathbf{I}}{\mu_{\mathcal{Q}}} + \tilde{\boldsymbol{R}}\big) \big(\hat{\boldsymbol{R}}^{\frac{1}{2}}\big)^+$ where $\boldsymbol{W} \in \mathbb{U}^{N_{\mathrm{T}} \times r_{\boldsymbol{P}}}$, and $\boldsymbol{D}$ has non-increasingly ordered, positive diagonal entries, then it is well known from majorization theory that the solution of (107) is $\boldsymbol{\Upsilon}^\star = \boldsymbol{W}$, up to a column permutation (e.g., [24, Theorem 4.3.26]). It turns out as well that $\mathrm{range}(\boldsymbol{\Upsilon}^\star) = \mathrm{range}(\boldsymbol{W}) = \mathrm{range}(\hat{\boldsymbol{R}})$, so the range space constraint is systematically fulfilled. The columns of $\boldsymbol{\Upsilon}^\star$ contain the eigenvectors $\boldsymbol{v}_i$ of the generalized eigenvalue problem

$$\hat{\boldsymbol{R}}\boldsymbol{v}_i = \omega_i \big(\tfrac{1}{\mu_{\mathcal{Q}}}\mathbf{I} + \tilde{\boldsymbol{R}}\big)\boldsymbol{v}_i, \tag{108}$$

corresponding to the $r_{\boldsymbol{P}}$ largest generalized eigenvalues ω_i.

E. Proof of Theorem VI.1

Assume that $r_{\boldsymbol{P}} > r_{\boldsymbol{Q}}$. Similarly as for the proof of Theorem V.1 in Appendix B, we will proceed by constructing another pilot matrix in $\mathcal{P}(\mu_{\mathcal{P}})$ which strictly outperforms $\boldsymbol{P}$. Recall that the covariance of the channel estimate is $\hat{\boldsymbol{R}} = \boldsymbol{R} - \tilde{\boldsymbol{R}}$, as usual, and $\tilde{\boldsymbol{R}} = (\boldsymbol{R}^{-1} + \boldsymbol{P})^{-1}$ is the estimation error covariance. We construct $\boldsymbol{P}'$ as

$$\boldsymbol{P}' = \left[\tilde{\boldsymbol{R}} + \lambda_{r_{\boldsymbol{P}}}(\hat{\boldsymbol{R}}) \boldsymbol{L}_{\boldsymbol{Q}^\perp \cap \hat{\boldsymbol{R}}} \boldsymbol{L}_{\boldsymbol{Q}^\perp \cap \hat{\boldsymbol{R}}}^\dagger\right]^{-1} - \boldsymbol{R}^{-1}. \quad (109)$$

The subunitary matrix $\boldsymbol{L}_{\boldsymbol{Q}^\perp \cap \hat{\boldsymbol{R}}}$ is defined as in the proof of Theorem V.1. It exists and has at least $d = r_{\boldsymbol{P}} - r_{\boldsymbol{Q}}$ columns. First, we verify that $\boldsymbol{P}' \in \mathcal{P}(\mu_{\mathcal{P}})$. In fact, $\boldsymbol{P}'$ can be written out as

$$\boldsymbol{P}' = \left[\boldsymbol{R} - \hat{\boldsymbol{R}} + \lambda_{r_{\boldsymbol{P}}}(\hat{\boldsymbol{R}}) \boldsymbol{L}_{\boldsymbol{Q}^\perp \cap \hat{\boldsymbol{R}}} \boldsymbol{L}_{\boldsymbol{Q}^\perp \cap \hat{\boldsymbol{R}}}^\dagger\right]^{-1} - \boldsymbol{R}^{-1}, \quad (110)$$

where it becomes clear that $\boldsymbol{P}' \succeq \boldsymbol{0}$, because $\hat{\boldsymbol{R}} - \lambda_{r_{\boldsymbol{P}}}(\hat{\boldsymbol{R}}) \boldsymbol{L}_{\boldsymbol{Q}^\perp \cap \hat{\boldsymbol{R}}} \boldsymbol{L}_{\boldsymbol{Q}^\perp \cap \hat{\boldsymbol{R}}}^\dagger \succeq \boldsymbol{0}$. On the other hand, the trace of $\boldsymbol{P}'$ is upper-bounded as

$$\begin{aligned} \operatorname{tr}(\boldsymbol{P}') &= \operatorname{tr}\left(\left[\tilde{\boldsymbol{R}} + \lambda_{r_{\boldsymbol{P}}}(\hat{\boldsymbol{R}}) \boldsymbol{L}_{\boldsymbol{Q}^\perp \cap \hat{\boldsymbol{R}}} \boldsymbol{L}_{\boldsymbol{Q}^\perp \cap \hat{\boldsymbol{R}}}^\dagger\right]^{-1}\right) - \operatorname{tr}(\boldsymbol{R}^{-1}) \\ &< \operatorname{tr}(\tilde{\boldsymbol{R}}^{-1}) - \operatorname{tr}(\boldsymbol{R}^{-1}) \\ &= \operatorname{tr}(\boldsymbol{P}). \end{aligned} \quad (111)$$

If we write $\tilde{\boldsymbol{R}} = \tilde{\boldsymbol{R}}(\boldsymbol{P})$ to stress that it is essentially a function of $\boldsymbol{P}$, then we notice that $\boldsymbol{P}'$ is designed so as to leave the product

$$\begin{aligned} \boldsymbol{Q}\tilde{\boldsymbol{R}}(\boldsymbol{P}') &= \boldsymbol{Q}\big(\tilde{\boldsymbol{R}} + \lambda_{r_{\boldsymbol{P}}}(\hat{\boldsymbol{R}}) \boldsymbol{L}_{\boldsymbol{Q}^\perp \cap \hat{\boldsymbol{R}}} \boldsymbol{L}_{\boldsymbol{Q}^\perp \cap \hat{\boldsymbol{R}}}^\dagger\big) \\ &= \boldsymbol{Q}\tilde{\boldsymbol{R}}(\boldsymbol{P}) \end{aligned} \quad (112)$$

unchanged, irrespective of whether the pilots are $\boldsymbol{P}$ or $\boldsymbol{P}'$. The same is true for $\boldsymbol{s} = \boldsymbol{\lambda}(\boldsymbol{S})$, which is left unchanged when replacing $\boldsymbol{P}$ by $\boldsymbol{P}'$, because $\boldsymbol{s}$ depends on $\boldsymbol{P}$ only via the product $\boldsymbol{Q}\tilde{\boldsymbol{R}}(\boldsymbol{P})$, as seen from the relationship

$$\boldsymbol{s} = \frac{\boldsymbol{\lambda}\big(\hat{\boldsymbol{R}}^{\frac{1}{2}} \boldsymbol{Q} \hat{\boldsymbol{R}}^{\frac{1}{2}}\big)}{1 + \operatorname{tr}(\boldsymbol{Q}\tilde{\boldsymbol{R}})} = \frac{\boldsymbol{\lambda}\big(\boldsymbol{Q}\boldsymbol{R} - \boldsymbol{Q}\tilde{\boldsymbol{R}}\big)}{1 + \operatorname{tr}(\boldsymbol{Q}\tilde{\boldsymbol{R}})}. \quad (113)$$

We have thus constructed alternative pilots $\boldsymbol{P}'$ which yield the same utility value $f(\boldsymbol{s})$, yet saving on the training energy, since $\operatorname{tr}(\boldsymbol{P}') < \operatorname{tr}(\boldsymbol{P})$. We generate another pilot matrix $\boldsymbol{P}'' = \kappa \boldsymbol{P}'$ with $\kappa = \operatorname{tr}(\boldsymbol{P})/\operatorname{tr}(\boldsymbol{P}')$. The new pilots $\boldsymbol{P}''$ spend the same amount of training energy as $\boldsymbol{P}$, but yield a strictly larger $\boldsymbol{S}'' = \boldsymbol{S}(\boldsymbol{P}'', \boldsymbol{Q}) \succ \boldsymbol{S}(\boldsymbol{P}, \boldsymbol{Q})$. Hence, $\boldsymbol{P}$ is suboptimal.

F. Convexity of the set of feasible $\hat{\boldsymbol{R}}$

Showing the convexity of the set of feasible $\hat{\boldsymbol{R}}$ is equivalent to showing the convexity of the set of feasible $\tilde{\boldsymbol{R}}$, because $\hat{\boldsymbol{R}} = \boldsymbol{R} - \tilde{\boldsymbol{R}}$ is merely $\tilde{\boldsymbol{R}}$ scaled with -1 and summed with a constant matrix $\boldsymbol{R}$. Therefore, we show that the set

$$\left\{\tilde{\boldsymbol{R}} = (\boldsymbol{R}^{-1} + \boldsymbol{P})^{-1} \middle| \boldsymbol{P} \in \mathcal{P}(\mu_{\mathcal{P}})\right\} \quad (114)$$

is convex. For any pair $(\boldsymbol{P}_1, \boldsymbol{P}_2) \in \mathcal{P}(\mu_{\mathcal{P}})^2$, there exists a $\boldsymbol{P}_3 \in \mathcal{P}(\mu_{\mathcal{P}})$ and a $\alpha \in [0; 1]$ such that

$$\alpha \tilde{\boldsymbol{R}}_1 + (1 - \alpha)\tilde{\boldsymbol{R}}_2 = \tilde{\boldsymbol{R}}_3, \quad (115)$$

where $\tilde{\boldsymbol{R}}_i = (\boldsymbol{R}^{-1} + \boldsymbol{P}_i)^{-1}$ for $i = 1, 2, 3$. By isolating $\boldsymbol{P}_3$ in (115), the pilot Gram $\boldsymbol{P}_3$ is given by

$$\boldsymbol{P}_3 = \left[\alpha \tilde{\boldsymbol{R}}_1 + (1 - \alpha)\tilde{\boldsymbol{R}}_2\right]^{-1} - \boldsymbol{R}^{-1}. \quad (116)$$

Obviously, since $\tilde{\boldsymbol{R}}_i \preceq \boldsymbol{R}$ for $i = 1, 2$, we have $\boldsymbol{P}_3 \succeq \boldsymbol{0}$. What remains to prove is that $\operatorname{tr}(\boldsymbol{P}_3) \leq \mu_{\mathcal{P}}$. Knowing that the function $\boldsymbol{X} \mapsto \operatorname{tr}(\boldsymbol{X}^{-1})$ is convex on the positive cone $\boldsymbol{X} \succ \boldsymbol{0}$, we have

$$\begin{aligned} \operatorname{tr}(\boldsymbol{P}_3) &\leq \alpha \operatorname{tr}(\tilde{\boldsymbol{R}}_1^{-1}) + (1 - \alpha) \operatorname{tr}(\tilde{\boldsymbol{R}}_2^{-1}) - \operatorname{tr}(\boldsymbol{R}^{-1}) \\ &= \alpha \operatorname{tr}(\boldsymbol{P}_1) + (1 - \alpha) \operatorname{tr}(\boldsymbol{P}_2) \\ &\leq \mu_{\mathcal{P}}. \end{aligned} \quad (117)$$

Hence, the set of feasible $\hat{\boldsymbol{R}}$ is convex, and so is Problem (P.1.b).

G. Proof of Theorem VII.1

We will proceed by showing that, in Problem (P.2), for any given value of the pair $(\mu_{\mathcal{P}}, \mu_{\mathcal{Q}})$, the search set $\boldsymbol{s}(\mathcal{P}(\mu_{\mathcal{P}}), \mathcal{Q}(\mu_{\mathcal{Q}}))$—and thus its Pareto border $\partial^+ \boldsymbol{s}(\mathcal{P}(\mu_{\mathcal{P}}), \mathcal{Q}(\mu_{\mathcal{Q}}))$—is left unchanged whether we allow $(\boldsymbol{P}, \boldsymbol{Q})$ to take *any* value within $\mathcal{P}(\mu_{\mathcal{P}}) \times \mathcal{Q}(\mu_{\mathcal{Q}})$, or whether we restrict the choice of the basis $\boldsymbol{U}_{\boldsymbol{P}}$ such that $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) = \{\boldsymbol{u}_{\boldsymbol{R},1}, \ldots, \boldsymbol{u}_{\boldsymbol{R},r^\star}\}$, where $r^\star$ denotes the number of non-zero entries of the $\boldsymbol{s}^\star$. With a consequence of Theorem V.2, we will eventually conclude on the desired result $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) = \mathcal{C}(\boldsymbol{U}_{\boldsymbol{Q}}) = \{\boldsymbol{u}_{\boldsymbol{R},1}, \ldots, \boldsymbol{u}_{\boldsymbol{R},r^\star}\}$.

To begin with, note that the set $\boldsymbol{s}(\mathcal{P}(\mu_{\mathcal{P}}), \mathcal{Q}(\mu_{\mathcal{Q}}))$ can be represented as the union

$$\boldsymbol{s}(\mathcal{P}(\mu_{\mathcal{P}}), \mathcal{Q}(\mu_{\mathcal{Q}})) = \bigcup_{\boldsymbol{P} \in \mathcal{P}(\mu_{\mathcal{P}})} \boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}})). \quad (118)$$

As a consequence of the rank equality (18), the elements of $\boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))$ have at most $r_{\boldsymbol{P}}$ non-zero entries, because $\operatorname{rank}(\boldsymbol{S}) = \operatorname{rank}(\hat{\boldsymbol{R}}^{\frac{1}{2}} \boldsymbol{Q} \hat{\boldsymbol{R}}^{\frac{1}{2}}) \leq \operatorname{rank}(\hat{\boldsymbol{R}}) = \operatorname{rank}(\boldsymbol{P}) = r_{\boldsymbol{P}}$. They can thus be written as $\boldsymbol{s}(\boldsymbol{P}, \boldsymbol{Q}) = \left[\bar{\boldsymbol{s}}(\boldsymbol{P}, \boldsymbol{Q})^{\mathrm{T}} \ \boldsymbol{0}^{\mathrm{T}}\right]^{\mathrm{T}}$ with $\bar{\boldsymbol{s}}(\boldsymbol{P}, \boldsymbol{Q}) \in \mathbb{R}_+^{r_{\boldsymbol{P}}}$ of reduced size. According to Theorem V.2, this set of reduced-size vectors $\bar{\boldsymbol{s}}(\boldsymbol{P}, \mathcal{Q}(\mu_{\mathcal{Q}}))$ is the simplex given by the convex hull of the points

$$\boldsymbol{\sigma}^{(0)} = \boldsymbol{0} \quad \boldsymbol{\sigma}^{(n)} = \mathcal{H}(\omega_1, \ldots, \omega_n) \sum_{j=1}^{n} \boldsymbol{e}_j, \quad n \in \{1, \ldots, r_{\boldsymbol{P}}\}. \quad (119)$$

Every such simplex is entirely described by $\boldsymbol{\omega} = [\omega_1, \ldots, \omega_{r_{\boldsymbol{P}}}]^{\mathrm{T}}$, the vector non-increasingly ordered eigenvalues of the matrix $\hat{\boldsymbol{R}}\big(\mu_{\mathcal{Q}}^{-1}\mathbf{I} + \tilde{\boldsymbol{R}}\big)^{-1}$, which is a function of $\boldsymbol{P}$ alone (not of $\boldsymbol{Q}$). Consistently with the notation used so far, $\boldsymbol{\omega}(\mathcal{P}(\mu_{\mathcal{P}}))$ shall denote the set of feasible $\boldsymbol{\omega}$ given that $\boldsymbol{P}$ belongs to $\mathcal{P}(\mu_{\mathcal{P}})$. To prove Theorem VII.1, we will first show that the set of Pareto border points $\partial^+ \boldsymbol{\omega}(\mathcal{P}(\mu_{\mathcal{P}}))$ is still achievable under the restriction $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) \subseteq \{\boldsymbol{u}_{\boldsymbol{R},1}, \ldots, \boldsymbol{u}_{\boldsymbol{R},r^\star}\}$. Recalling that $\boldsymbol{R} = \hat{\boldsymbol{R}} + \tilde{\boldsymbol{R}}$ and $\tilde{\boldsymbol{R}} = (\boldsymbol{R}^{-1} + \boldsymbol{P})^{-1}$, we write out $\boldsymbol{\omega}$ as

$$\begin{aligned} \boldsymbol{\omega} &= \boldsymbol{\lambda}\Big(\hat{\boldsymbol{R}}\big(\tfrac{1}{\mu_{\mathcal{Q}}}\mathbf{I} + \tilde{\boldsymbol{R}}\big)^{-1}\Big) \\ &= \boldsymbol{\lambda}\Big(\big(\boldsymbol{R} - (\boldsymbol{R}^{-1} + \boldsymbol{P})^{-1}\big)\big(\tfrac{1}{\mu_{\mathcal{Q}}}\mathbf{I} + (\boldsymbol{R}^{-1} + \boldsymbol{P})^{-1}\big)^{-1}\Big). \end{aligned}$$

Let us denote $\boldsymbol{P}' = \boldsymbol{R}^{\frac{1}{2}}\boldsymbol{P}\boldsymbol{R}^{\frac{1}{2}}$, then using the property $\boldsymbol{\lambda}(\boldsymbol{A}\boldsymbol{B}) = \boldsymbol{\lambda}(\boldsymbol{B}\boldsymbol{A})$, the last expression can be rewritten as

$$\boldsymbol{\omega} = \boldsymbol{\lambda}\left(\left(\mathbf{I} - (\mathbf{I}+\boldsymbol{P}')^{-1}\right)\left((\mu_{\mathcal{Q}}\boldsymbol{R})^{-1} + (\mathbf{I}+\boldsymbol{P}')^{-1}\right)^{-1}\right).$$

Let us now denote $\boldsymbol{P}'' = \boldsymbol{U}_{\boldsymbol{R}}^{\dagger}\boldsymbol{P}'\boldsymbol{U}_{\boldsymbol{R}}$, so that the last expression becomes

$$\boldsymbol{\omega} = \boldsymbol{\lambda}\Big(\left(\mathbf{I} - (\mathbf{I}+\boldsymbol{P}'')^{-1}\right)\left((\mu_{\mathcal{Q}}\boldsymbol{\Lambda}_{\boldsymbol{R}})^{-1} + (\mathbf{I}+\boldsymbol{P}'')^{-1}\right)^{-1}\Big). \tag{120}$$

Let us write out the mutual relations linking $\boldsymbol{P}$ and $\boldsymbol{P}''$ in full:

$$\boldsymbol{P}'' = \boldsymbol{\Lambda}_{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{U}_{\boldsymbol{R}}^{\dagger}\boldsymbol{U}_{\boldsymbol{P}}\operatorname{diag}(\boldsymbol{p})\boldsymbol{U}_{\boldsymbol{P}}^{\dagger}\boldsymbol{U}_{\boldsymbol{R}}\boldsymbol{\Lambda}_{\boldsymbol{R}}^{\frac{1}{2}} \tag{121a}$$

$$\boldsymbol{P} = \boldsymbol{U}_{\boldsymbol{R}}\boldsymbol{\Lambda}_{\boldsymbol{R}}^{-\frac{1}{2}}\boldsymbol{U}_{\boldsymbol{P}''}\operatorname{diag}(\boldsymbol{p}'')\boldsymbol{U}_{\boldsymbol{P}''}^{\dagger}\boldsymbol{\Lambda}_{\boldsymbol{R}}^{-\frac{1}{2}}\boldsymbol{U}_{\boldsymbol{R}}^{\dagger}. \tag{121b}$$

Regarding the (non-reduced) eigendecomposition $\boldsymbol{P}'' = \boldsymbol{U}_{\boldsymbol{P}''}\boldsymbol{\Lambda}_{\boldsymbol{P}''}\boldsymbol{U}_{\boldsymbol{P}''}^{\dagger}$ with $\boldsymbol{U}_{\boldsymbol{P}''} \in \mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}$ and $\boldsymbol{\Lambda}_{\boldsymbol{P}''} = \operatorname{diag}(\boldsymbol{p}'') = \operatorname{diag}(p_1'', \ldots, p_{N_{\mathsf{T}}}'')$, one can say that, if $\boldsymbol{P}$ is drawn from $\mathcal{P}(\mu_{\mathcal{P}})$, then the corresponding eigenvalue profile $\boldsymbol{p}'' = \boldsymbol{\lambda}(\boldsymbol{P}'') = \boldsymbol{\lambda}\big(\boldsymbol{\Lambda}_{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{U}_{\boldsymbol{R}}^{\dagger}\boldsymbol{P}\boldsymbol{U}_{\boldsymbol{R}}\boldsymbol{\Lambda}_{\boldsymbol{R}}^{\frac{1}{2}}\big) = \boldsymbol{\lambda}(\boldsymbol{P}\boldsymbol{R})$ [cf. (121a)] is drawn from a feasible set which we shall call $\boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}}))$, a notation which emphasizes its direct dependence on the domain $\mathcal{P}(\mu_{\mathcal{P}})$. As to the eigenbasis $\boldsymbol{U}_{\boldsymbol{P}''}$, it obviously belongs to $\mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}$ by definition, yet in general, we must presume that not all pairs $(\boldsymbol{p}'', \boldsymbol{U}_{\boldsymbol{P}''}) \in \boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}})) \times \mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}$ are jointly feasible, since the eigenbasis $\boldsymbol{U}_{\boldsymbol{P}''}$ and the eigenvalues $\boldsymbol{p}''$ cannot be chosen independently of each other, due to the special structure of Expression (121a). Instead, $\boldsymbol{U}_{\boldsymbol{P}}''$ belongs to a feasible set $\boldsymbol{U}_{\boldsymbol{P}''}(\boldsymbol{p}'') \subseteq \mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}$ (which depends on $\boldsymbol{p}''$), so the overall set of feasible pairs $(\boldsymbol{p}'', \boldsymbol{U}_{\boldsymbol{P}''})$ forms a *subset* of the Cartesian product $\boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}})) \times \mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}$.

However, suppose for a while that $\boldsymbol{p}''$ and $\boldsymbol{U}_{\boldsymbol{P}''}$ can be drawn *independently* of each other from their respective domains $\boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}}))$ and $\mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}$. This assumption then corresponds to a *relaxation* of the original problem, as it possibly extends the overall set of feasible $\boldsymbol{P}''$, and consequently, of feasible $\boldsymbol{\omega}$. The resulting set of achievable vectors $\boldsymbol{\omega}$ under this relaxation shall be denoted $\bar{\boldsymbol{\omega}}(\mathcal{P}(\mu_{\mathcal{P}})) \supseteq \boldsymbol{\omega}(\mathcal{P}(\mu_{\mathcal{P}}))$ and is formally defined as

$$\bar{\boldsymbol{\omega}}(\mathcal{P}(\mu_{\mathcal{P}})) = \Big\{\boldsymbol{\omega}(\boldsymbol{p}'', \boldsymbol{U}_{\boldsymbol{P}''}) \;\Big|\; (\boldsymbol{p}'', \boldsymbol{U}_{\boldsymbol{P}''}) \in \boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}})) \times \mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}\Big\}, \tag{122}$$

wherein the two-argument notation $\boldsymbol{\omega}(\bullet, \bullet)$ is defined as [cf. (121b)]

$$\boldsymbol{\omega}(\boldsymbol{p}'', \boldsymbol{U}_{\boldsymbol{P}''}) \triangleq \boldsymbol{\omega}\Big(\boldsymbol{U}_{\boldsymbol{R}}\boldsymbol{\Lambda}_{\boldsymbol{R}}^{-\frac{1}{2}}\boldsymbol{U}_{\boldsymbol{P}''}\operatorname{diag}(\boldsymbol{p}'')\boldsymbol{U}_{\boldsymbol{P}''}^{\dagger}\boldsymbol{\Lambda}_{\boldsymbol{R}}^{-\frac{1}{2}}\boldsymbol{U}_{\boldsymbol{R}}^{\dagger}\Big). \tag{123}$$

The set $\bar{\boldsymbol{\omega}}(\mathcal{P}(\mu_{\mathcal{P}}))$ can be represented as a double union

$$\begin{aligned}\bar{\boldsymbol{\omega}}(\mathcal{P}(\mu_{\mathcal{P}})) &= \bigcup_{\boldsymbol{p}''\in\boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}}))} \bigcup_{\boldsymbol{U}_{\boldsymbol{P}''}\in\mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}} \boldsymbol{\omega}(\boldsymbol{p}'', \boldsymbol{U}_{\boldsymbol{P}''}) \\ &= \bigcup_{\boldsymbol{p}''\in\boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}}))} \boldsymbol{\omega}\big(\boldsymbol{p}'', \mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}\big).\end{aligned} \tag{124}$$

As seen from expression (120), $\boldsymbol{\omega}(\boldsymbol{p}'', \boldsymbol{U}_{\boldsymbol{P}''})$ is monotonic in the eigenvalues p_i'', meaning that

$$\forall \boldsymbol{d} \geq \mathbf{0}\colon \boldsymbol{\omega}(\boldsymbol{p}'' + \boldsymbol{d}, \boldsymbol{U}_{\boldsymbol{P}''}) \geq \boldsymbol{\omega}(\boldsymbol{p}'', \boldsymbol{U}_{\boldsymbol{P}''}).$$

Hence, since we are essentially interested in the Pareto border $\partial^+\bar{\boldsymbol{\omega}}(\mathcal{P}(\mu_{\mathcal{P}}))$ of the set $\bar{\boldsymbol{\omega}}(\mathcal{P}(\mu_{\mathcal{P}}))$, we can restrict our further analysis to the set[7]

$$\bar{\boldsymbol{\omega}}^{+}(\mathcal{P}(\mu_{\mathcal{P}})) = \bigcup_{\boldsymbol{p}''\in\partial^+\boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}}))} \boldsymbol{\omega}\big(\boldsymbol{p}'', \mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}\big). \tag{125}$$

The remainder of the proof of Theorem VII.1 is completed in four successive steps, each of which is detailed in a separate paragraph, for the sake of a clearer structure: first, we specify a method for constructing a particular Pareto border point of the set $\boldsymbol{\omega}(\boldsymbol{p}'', \mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}})$ given a particular value of the vector $\boldsymbol{p}''$, where we show that this construction requires the alignment $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}})$; second, we show that the point constructed this way, besides yielding a Pareto border point of the relaxed set $\bar{\boldsymbol{\omega}}(\mathcal{P}(\mu_{\mathcal{P}}))$, is also contained in the smaller (non-relaxed) set $\boldsymbol{\omega}(\mathcal{P}(\mu_{\mathcal{P}}))$, so it must be a Pareto border point of $\boldsymbol{\omega}(\mathcal{P}(\mu_{\mathcal{P}}))$ as well; third, we show that, by varying the eigenvalues $\boldsymbol{p}''$ over the feasible set $\boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}}))$, with the aforementioned method of constructing particular Pareto border points, we reach the whole Pareto border of $\boldsymbol{\omega}(\mathcal{P}(\mu_{\mathcal{P}}))$; fourth, we show that the alignment $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}})$ implies that $\boldsymbol{U}_{\boldsymbol{Q}}$ must as well be aligned such that $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{Q}}) = \mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}})$ to reach the whole feasible set $\boldsymbol{s}(\mathcal{P}(\mu_{\mathcal{P}}), \mathcal{Q}(\mu_{\mathcal{Q}}))$, and conclude.

1) Let the orthonormal eigenbasis $\boldsymbol{U}_{\boldsymbol{P}''} = [\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{N_{\mathsf{T}}}]$ be spanned by unit vectors $\boldsymbol{u}_i$, where the i-th vector $\boldsymbol{u}_i$ is associated to the i-th largest eigenvalue p_i''. Given a fixed value of $\boldsymbol{p}'' \in \boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}}))$, we construct a particular point of the Pareto border $\partial^+\boldsymbol{\omega}(\boldsymbol{p}'', \mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}})$ by solving the sequence of optimization problems:

$$\begin{aligned}&\forall i \in \{1, \ldots, N_{\mathsf{T}}\}\colon \quad \boldsymbol{U}^{(i)} = \operatorname*{argmax}_{\boldsymbol{U}\in\mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}} \omega_i\big(\boldsymbol{p}'', \boldsymbol{U}\big) \\ &\text{s.t.}\;\; \forall \ell \in \{1, \ldots, i-1\}\colon \quad \omega_\ell = \omega_\ell\big(\boldsymbol{p}'', \boldsymbol{U}^{(i-1)}\big).\end{aligned} \tag{126}$$

Clearly, $\boldsymbol{U}^{(N_{\mathsf{T}})}$ will yield a Pareto optimal point, that is,

$$\boldsymbol{\omega}(\boldsymbol{p}'', \boldsymbol{U}^{(N_{\mathsf{T}})}) \in \partial^+\boldsymbol{\omega}(\boldsymbol{p}'', \mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}). \tag{127}$$

Next, we will show by induction that $\boldsymbol{U}^{(N_{\mathsf{T}})} = \mathbf{I}$. For this purpose, let us explicitly solve the first problem ($i = 1$) of (126), i.e.,

$$\boldsymbol{U}^{(1)} = \operatorname*{argmax}_{\boldsymbol{U}\in\mathbb{U}^{N_{\mathsf{T}}\times N_{\mathsf{T}}}} \omega_1\big(\boldsymbol{p}'', \boldsymbol{U}\big) \tag{128}$$

[7]Note that $\bar{\boldsymbol{\omega}}^{+}(\mathcal{P}(\mu_{\mathcal{P}}))$ is generally not the Pareto border of $\bar{\boldsymbol{\omega}}(\mathcal{P}(\mu_{\mathcal{P}}))$, but rather a superset thereof.

With Expression (120), this reads as

$$\begin{aligned}
&\max_{\boldsymbol{U}\in\mathbb{U}^{N_\mathsf{T}\times N_\mathsf{T}}}\max_{\|\boldsymbol{v}_1\|=1}\left[\frac{\boldsymbol{v}_1^\dagger\big(\mathbf{I}-(\mathbf{I}+\boldsymbol{U}\boldsymbol{\Lambda}_{\boldsymbol{P}''}\boldsymbol{U}^\dagger)^{-1}\big)\boldsymbol{v}_1}{\boldsymbol{v}_1^\dagger\big((\mu_{\mathcal{Q}}\boldsymbol{\Lambda}_{\boldsymbol{R}})^{-1}+(\mathbf{I}+\boldsymbol{U}\boldsymbol{\Lambda}_{\boldsymbol{P}''}\boldsymbol{U}^\dagger)^{-1}\big)\boldsymbol{v}_1}\right]\\
&\leq \max_{\boldsymbol{U}\in\mathbb{U}^{N_\mathsf{T}\times N_\mathsf{T}}}\left[\frac{1-\lambda_{\min}\big((\mathbf{I}+\boldsymbol{U}\boldsymbol{\Lambda}_{\boldsymbol{P}''}\boldsymbol{U}^\dagger)^{-1}\big)}{\lambda_{\min}\big(\mu_{\mathcal{Q}}^{-1}\boldsymbol{\Lambda}_{\boldsymbol{R}}^{-1}\big)+\lambda_{\min}\big((\mathbf{I}+\boldsymbol{U}\boldsymbol{\Lambda}_{\boldsymbol{P}''}\boldsymbol{U}^\dagger)^{-1}\big)}\right]\\
&=\frac{1-\frac{1}{1+\lambda_{\max}(\boldsymbol{\Lambda}_{\boldsymbol{P}''})}}{(\mu_{\mathcal{Q}}\lambda_{\max}(\boldsymbol{\Lambda}_{\boldsymbol{R}}))^{-1}+\frac{1}{1+\lambda_{\max}(\boldsymbol{\Lambda}_{\boldsymbol{P}''})}}\\
&=\frac{1-\frac{1}{1+p_1''}}{(\mu_{\mathcal{Q}}r_1)^{-1}+\frac{1}{1+p_1''}}.
\end{aligned}\tag{129}$$

This upper bound is tight and achieved if and only if $\boldsymbol{v}_1=\boldsymbol{e}_1$, and when $\boldsymbol{U}$ is of the form

$$\boldsymbol{U}^{(1)}=\begin{bmatrix}1 & \mathbf{0}\\ \mathbf{0} & \boldsymbol{W}^{(1)}\end{bmatrix}\tag{130}$$

with some arbitrary $\boldsymbol{W}^{(1)}\in\mathbb{U}^{(N_\mathsf{T}-1)\times(N_\mathsf{T}-1)}$. To prove the induction step, we will show that if for a certain $i\geq 1$, all maximizers $\boldsymbol{U}^{(i)}$ are of the form

$$\boldsymbol{U}^{(i)}=\begin{bmatrix}\mathbf{I}_i & \mathbf{0}\\ \mathbf{0} & \boldsymbol{W}^{(i)}\end{bmatrix}\tag{131}$$

with some arbitrary $\boldsymbol{W}^{(i)}\in\mathbb{U}^{(N_\mathsf{T}-i)\times(N_\mathsf{T}-i)}$, and $\forall\ell=1,\dots,i\colon \boldsymbol{v}_\ell=\boldsymbol{e}_\ell$, then $\boldsymbol{U}^{(i+1)}$ also has the above block structure (131), with an identity matrix $\mathbf{I}_{i+1}$ top left and an arbitrary rotation matrix $\boldsymbol{W}^{(i+1)}$ bottom right. After solving the i-th problem, we know that all solutions thereof are of the form (131), which implies that the equality constraints for the $(i+1)$-th problem [cf. (126)] can only be fulfilled if $\boldsymbol{U}^{(i+1)}$ has the same structure as $\boldsymbol{U}^{(i)}$, i.e.,

$$\boldsymbol{U}^{(i+1)}=\begin{bmatrix}\mathbf{I}_i & \mathbf{0}\\ \mathbf{0} & \tilde{\boldsymbol{W}}^{(i)}\end{bmatrix}\tag{132}$$

with some unitary matrix $\tilde{\boldsymbol{W}}^{(i)}\in\mathbb{U}^{(N_\mathsf{T}-i)\times(N_\mathsf{T}-i)}$ to be determined. According to a straightforward adaptation of the Courant-Fisher Theorem [24, Theorem 4.2.11], the non-increasingly ordered eigenvalues $\lambda_i(\boldsymbol{A}\boldsymbol{B}^{-1})$ with corresponding eigenvectors $\boldsymbol{v}_i$ of a product of two Hermitian matrices $\boldsymbol{A}$ and $\boldsymbol{B}^{-1}$ can be expressed as

$$\lambda_i(\boldsymbol{A}\boldsymbol{B}^{-1})=\max_{\boldsymbol{v}\perp\boldsymbol{B}\boldsymbol{v}_{i-1},\dots,\boldsymbol{B}\boldsymbol{v}_1}\frac{\boldsymbol{v}^\dagger\boldsymbol{A}\boldsymbol{v}}{\boldsymbol{v}^\dagger\boldsymbol{B}\boldsymbol{v}}.\tag{133}$$

The $(i+1)$-th optimization problem reads as

$$\begin{aligned}
\boldsymbol{U}^{(i+1)}=\operatorname*{argmax}_{\boldsymbol{U}\in\mathbb{U}^{N_\mathsf{T}\times N_\mathsf{T}}}&\left\{\max_{\boldsymbol{v}_{i+1}\perp\boldsymbol{B}\boldsymbol{v}_i,\dots,\boldsymbol{B}\boldsymbol{v}_1}\frac{\boldsymbol{v}_{i+1}^\dagger\boldsymbol{A}\boldsymbol{v}_{i+1}}{\boldsymbol{v}_{i+1}^\dagger\boldsymbol{B}\boldsymbol{v}_{i+1}}\right\}\\
&\text{s.t.}\quad \boldsymbol{U}=\begin{bmatrix}\mathbf{I}_i & \mathbf{0}\\ \mathbf{0} & \tilde{\boldsymbol{W}}^{(i)}\end{bmatrix}
\end{aligned}\tag{134}$$

with $\boldsymbol{A}=\mathbf{I}-(\mathbf{I}+\boldsymbol{U}\boldsymbol{\Lambda}_{\boldsymbol{P}''}\boldsymbol{U}^\dagger)^{-1}$ and $\boldsymbol{B}=(\mu_{\mathcal{Q}}\boldsymbol{\Lambda}_{\boldsymbol{R}})^{-1}+(\mathbf{I}+\boldsymbol{U}\boldsymbol{\Lambda}_{\boldsymbol{P}''}\boldsymbol{U}^\dagger)^{-1}$ [cf. (120)], and $\forall\ell=1,\dots,i\colon \boldsymbol{v}_\ell=\boldsymbol{e}_\ell$. When writing out $\boldsymbol{B}$, the vectors involved in the orthogonality constraints $\boldsymbol{v}_i\perp\boldsymbol{B}\boldsymbol{e}_{i-1},\dots,\boldsymbol{B}\boldsymbol{e}_1$ read as

$$\begin{aligned}
\boldsymbol{B}\boldsymbol{e}_\ell&=\Big[(\mu_{\mathcal{Q}}\boldsymbol{\Lambda}_{\boldsymbol{R}})^{-1}+(\mathbf{I}+\boldsymbol{U}\boldsymbol{\Lambda}_{\boldsymbol{P}''}\boldsymbol{U}^\dagger)^{-1}\Big]\boldsymbol{e}_\ell\\
&=\frac{\boldsymbol{e}_\ell}{\mu_{\mathcal{Q}}r_\ell}+\begin{bmatrix}\big(\mathbf{I}+\boldsymbol{\Lambda}_{\boldsymbol{P}''}^{[i]}\big)^{-1} & \mathbf{0}\\ \mathbf{0} & \tilde{\boldsymbol{W}}^{(i)}\big(\mathbf{I}+\bar{\boldsymbol{\Lambda}}_{\boldsymbol{P}''}^{[i]}\big)^{-1}\big(\tilde{\boldsymbol{W}}^{(i)}\big)^\dagger\end{bmatrix}\boldsymbol{e}_\ell\\
&=\Big[\frac{1}{\mu_{\mathcal{Q}}r_\ell}+\frac{1}{1+p_\ell''}\Big]\boldsymbol{e}_\ell\qquad\forall\ell=1,\dots,i
\end{aligned}\tag{135}$$

where $\boldsymbol{\Lambda}_{\boldsymbol{P}''}^{[i]}=\operatorname{diag}(p_1'',\dots,p_i'')$ and $\bar{\boldsymbol{\Lambda}}_{\boldsymbol{P}''}^{[i]}=\operatorname{diag}(p_{i+1}'',\dots,p_{N_\mathsf{T}}'')$. Thus, the orthogonality constraints simply translate into $\boldsymbol{v}_{i+1}\perp\boldsymbol{e}_i,\dots,\boldsymbol{e}_1$. In other terms, the first i entries of $\boldsymbol{v}_{i+1}$ must be zero. Thus, we can define matrices $(N_\mathsf{T}-i)\times(N_\mathsf{T}-i)$ matrices $\breve{\boldsymbol{A}}$ and $\breve{\boldsymbol{B}}$ as

$$\begin{aligned}
\breve{\boldsymbol{A}}&=\mathbf{I}-\big(\mathbf{I}+\tilde{\boldsymbol{W}}^{(i)}\bar{\boldsymbol{\Lambda}}_{\boldsymbol{P}''}^{[i]}\big(\tilde{\boldsymbol{W}}^{(i)}\big)^\dagger\big)^{-1}\\
\breve{\boldsymbol{B}}&=(\mu_{\mathcal{Q}}\bar{\boldsymbol{\Lambda}}_{\boldsymbol{R}}^{(i)})^{-1}+\big(\mathbf{I}+\tilde{\boldsymbol{W}}^{(i)}\bar{\boldsymbol{\Lambda}}_{\boldsymbol{P}''}^{[i]}\big(\tilde{\boldsymbol{W}}^{(i)}\big)^\dagger\big)^{-1},
\end{aligned}\tag{136}$$

so that the optimization problem (134) boils down to solving

$$\tilde{\boldsymbol{W}}^{(i)}=\operatorname*{argmax}_{\boldsymbol{W}\in\mathbb{U}^{(N_\mathsf{T}-i)\times(N_\mathsf{T}-i)}}\left\{\max_{\breve{\boldsymbol{v}}_{i+1}}\frac{\breve{\boldsymbol{v}}_{i+1}^\dagger\breve{\boldsymbol{A}}\breve{\boldsymbol{v}}_{i+1}}{\breve{\boldsymbol{v}}_{i+1}^\dagger\breve{\boldsymbol{B}}\breve{\boldsymbol{v}}_{i+1}}\right\}.\tag{137}$$

This problem is fully equivalent in structure to the first optimization problem ($i=1$) as written out in Equation (129) and has the same solution, i.e., [cf. (130)]

$$\tilde{\boldsymbol{W}}^{(i)}=\begin{bmatrix}1 & \mathbf{0}\\ \mathbf{0} & \boldsymbol{W}^{(i+1)}\end{bmatrix}.\tag{138}$$

Consequently, $\boldsymbol{U}^{(i+1)}$ has indeed the structure (131), which concludes the induction proof. We infer that $\boldsymbol{U}^{(N_\mathsf{T})}=\mathbf{I}$, and thus

$$\boldsymbol{\omega}(\boldsymbol{p}'',\mathbf{I})\in\partial^+\boldsymbol{\omega}(\boldsymbol{p}'',\mathbb{U}^{N_\mathsf{T}\times N_\mathsf{T}})\subset\bar{\boldsymbol{\omega}}^+(\mathcal{P}(\mu_{\mathcal{P}})).\tag{139}$$

Thus, we have specified a method to construct specific Pareto optimal points of the inner union in (124).

2) Recalling how $\boldsymbol{P}''$ is obtained from $\boldsymbol{P}\in\mathcal{P}(\mu_{\mathcal{P}})$, namely [cf. (121a)]

$$\boldsymbol{P}''=\boldsymbol{\Lambda}_{\boldsymbol{R}}^{\frac{1}{2}}\boldsymbol{U}_{\boldsymbol{R}}^\dagger\boldsymbol{U}_{\boldsymbol{P}}\operatorname{diag}(\boldsymbol{p})\boldsymbol{U}_{\boldsymbol{P}}^\dagger\boldsymbol{U}_{\boldsymbol{R}}\boldsymbol{\Lambda}_{\boldsymbol{R}}^{\frac{1}{2}},$$

we can leverage Theorem V.2 (although with other variables) to characterize the set $\boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}}))$ of vectors of feasible, non-increasingly sorted eigenvalues of the above matrix. First note that $\boldsymbol{P}''$ has the same eigenvalues as $\boldsymbol{U}_{\boldsymbol{R}}\boldsymbol{P}''\boldsymbol{U}_{\boldsymbol{R}}^\dagger$, so that the set $\boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}}))$ may be defined as [compare with (32)]

$$\boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}}))=\Big\{\boldsymbol{\lambda}\big(\boldsymbol{R}^{\frac{1}{2}}\boldsymbol{P}\boldsymbol{R}^{\frac{1}{2}}\big)\;\Big|\;\boldsymbol{P}\in\mathbb{C}_+^{N_\mathsf{T}\times N_\mathsf{T}},\operatorname{tr}(\boldsymbol{P})\leq\mu_{\mathcal{P}}\Big\}\tag{140}$$

Now, Theorem V.2 can be applied upon replacing $\hat{\boldsymbol{R}}$, $\tilde{\boldsymbol{R}}$, $\boldsymbol{Q}$, $\mathcal{Q}(\mu_{\mathcal{Q}})$ and $\mu_{\mathcal{Q}}$ (as they appear in the formulation of said theorem) with $\boldsymbol{R}$, $\mathbf{0}$, $\boldsymbol{P}$, $\mathcal{P}(\mu_{\mathcal{P}})$ and $\mu_{\mathcal{P}}$ respectively. This leads to $\boldsymbol{p}''(\mathcal{P}(\mu_{\mathcal{P}}))$ being characterized as the convex hull of the points $\boldsymbol{\sigma}''^{(n)}, n=0,\dots,N_\mathsf{T}$ defined as

$$\boldsymbol{\sigma}''^{(0)}=\mathbf{0}\tag{141}$$

$$\boldsymbol{\sigma}''^{(n)}=\mu_{\mathcal{P}}\cdot\mathcal{H}(r_1,\dots,r_n)\sum_{\ell=1}^{n}\boldsymbol{e}_\ell,\quad n=1,\dots,N_\mathsf{T}.\tag{142}$$

It can be readily verified that all points of this convex hull can be reached when setting $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}})$. When doing so, the eigenbasis of $\boldsymbol{P}''$ is precisely $\boldsymbol{U}_{\boldsymbol{P}''} = \mathbf{I}$. But remember that the choice $\boldsymbol{U}_{\boldsymbol{P}''} = \mathbf{I}$ was required in the previous paragraph for constructing a Pareto optimal point of $\partial^+\boldsymbol{\omega}(\boldsymbol{p}'', \mathbb{U}^{N_\mathsf{T}\times N_\mathsf{T}})$. Consequently, this Pareto optimal point is also contained in the subset $\boldsymbol{\omega}(\boldsymbol{p}'', \boldsymbol{U}_{\boldsymbol{P}''}(\boldsymbol{p}'')) \subseteq \boldsymbol{\omega}(\boldsymbol{p}'', \mathbb{U}^{N_\mathsf{T}\times N_\mathsf{T}})$, and is thus necessarily a Pareto optimal point of $\boldsymbol{\omega}(\boldsymbol{p}'', \boldsymbol{U}_{\boldsymbol{P}''}(\boldsymbol{p}''))$, too.

3) We now ask whether all points of the overall Pareto border $\partial^+\boldsymbol{\omega}(\mathcal{P}(\mu_\mathcal{P}))$ are attained by the construction method specified above, i.e., whether

$$\partial^+\boldsymbol{\omega}(\mathcal{P}(\mu_\mathcal{P})) \subseteq \bigcup_{\boldsymbol{p}''\in\partial^+\boldsymbol{p}''(\mathcal{P}(\mu_\mathcal{P}))} \boldsymbol{\omega}(\boldsymbol{p}'', \mathbf{I}). \tag{143}$$

Let us write out $\boldsymbol{\omega}(\boldsymbol{p}'', \mathbf{I})$ by means of (120) as

$$\boldsymbol{\Pi}\boldsymbol{\omega}(\boldsymbol{p}'', \mathbf{I}) = \boldsymbol{p}'' \odot \big((\mu_\mathcal{Q}\boldsymbol{r})^{-1}\boldsymbol{p}'' + \boldsymbol{\xi}\big)^{-1}, \tag{144}$$

where $\boldsymbol{\Pi} \in \mathbb{P}^{N_\mathsf{T}}$ is a sorting permutation, '$\odot$' denotes componentwise multiplication, $\boldsymbol{r}^{-1}$ denotes the vector of entries r_i^{-1} (i.e., componentwise reciprocal), and $\operatorname{diag}(\boldsymbol{\xi}) = \boldsymbol{\Xi} = \mathbf{I} + (\mu_\mathcal{Q}\boldsymbol{\Lambda}_{\boldsymbol{R}})^{-1}$. The mapping $\boldsymbol{p}'' \mapsto \boldsymbol{p}'' \odot \big((\mu_\mathcal{Q}\boldsymbol{r})^{-1}\boldsymbol{p}'' + \boldsymbol{\xi}\big)^{-1}$ is clearly injective, since $\boldsymbol{\xi} > \mathbf{0}$. Additionally, it has the property that for any real unit-norm vector $\boldsymbol{e} \geq \mathbf{0}$, there exists a scalar $\epsilon > 0$ and a single feasible vector $\boldsymbol{p}'' \in \partial^+\boldsymbol{p}''(\mathcal{P}(\mu_\mathcal{P}))$ such that

$$\boldsymbol{p}'' \odot \big((\mu_\mathcal{Q}\boldsymbol{r})^{-1}\boldsymbol{p}'' + \boldsymbol{\xi}\big)^{-1} = \epsilon\boldsymbol{e}. \tag{145}$$

To see this, we first rewrite Expression (145) as

$$\boldsymbol{p}'' = \epsilon\boldsymbol{e} \odot \boldsymbol{\xi} \odot \big(\mathbf{1} - \epsilon\boldsymbol{e} \odot (\mu_\mathcal{Q}\boldsymbol{r})^{-1}\big)^{-1}. \tag{146}$$

Since $\boldsymbol{p}'' \geq \mathbf{0}$, the scalar ϵ must lie in the semi-open interval $\epsilon \in [0; \min_i \mu_\mathcal{Q} r_i/e_i[$. From taking the Euclidian norm of Expression (146), we obtain a function $\epsilon \mapsto \|\boldsymbol{p}''\|_2$ which bijectively maps $[0; \min_i \mu_\mathcal{Q} r_i/e_i[$ onto $\mathbb{R}_+$. Since any $\boldsymbol{p}'' \in \partial^+\boldsymbol{p}''(\mathcal{P}(\mu_\mathcal{P}))$ has finite norm, there must necessarily exist one single value of ϵ fulfilling

$$\epsilon\boldsymbol{e} \odot \boldsymbol{\xi} \odot \big(\mathbf{1} - \epsilon\boldsymbol{e} \odot (\mu_\mathcal{Q}\boldsymbol{r})^{-1}\big)^{-1} \in \partial^+\boldsymbol{p}''(\mathcal{P}(\mu_\mathcal{P})). \tag{147}$$

Consequently, all Pareto optimal points $\partial^+\boldsymbol{\omega}(\mathcal{P}(\mu_\mathcal{P}))$ can be reached by the construction method from paragraphs 2) and 3), so that we may write

$$\partial^+\boldsymbol{\omega}(\mathcal{P}(\mu_\mathcal{P})) = \boldsymbol{\omega}\big(\partial^+\boldsymbol{p}''(\mathcal{P}(\mu_\mathcal{P})), \mathbf{I}\big). \tag{148}$$

4) Now that we have established that the Pareto border $\partial^+\boldsymbol{\omega}(\mathcal{P}(\mu_\mathcal{P}))$ can be reached by setting $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}})$, we have that $\tilde{\boldsymbol{R}} = (\boldsymbol{R}^{-1} + \boldsymbol{P})^{-1}$ and $\hat{\boldsymbol{R}} = \boldsymbol{R} - \tilde{\boldsymbol{R}}$ acquire the same eigenbasis, up to a column permutation. Specifically, we have that the alignment $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}})$ implies $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) = \mathcal{C}(\boldsymbol{U}_{\hat{\boldsymbol{R}}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\tilde{\boldsymbol{R}}})$. But as a consequence of Theorem V.2, the alignment $\mathcal{C}(\boldsymbol{U}_{\hat{\boldsymbol{R}}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\tilde{\boldsymbol{R}}})$ leads to [cf. (34)]

$$\mathcal{C}(\boldsymbol{U}_{\boldsymbol{Q}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\hat{\boldsymbol{R}}}). \tag{149}$$

Hence, we obtain $\mathcal{C}(\boldsymbol{U}_{\boldsymbol{Q}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}})$. Since we know from Section VII-B that $\operatorname{rank}(\boldsymbol{P}^\star) = \operatorname{rank}(\boldsymbol{Q}^\star)$ at any joint optimum $(\boldsymbol{P}^\star, \boldsymbol{Q}^\star)$, we get the desired alignment property

$$\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) = \mathcal{C}(\boldsymbol{U}_{\boldsymbol{Q}}) \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}}). \tag{150}$$

Obviously, in case (150) is a strict inclusion, the eigenbases of $\boldsymbol{P}$ and $\boldsymbol{Q}$ should contain the eigenvectors of $\boldsymbol{R}$ associated to the largest eigenvalues of $\boldsymbol{R}$, hence

$$\mathcal{C}(\boldsymbol{U}_{\boldsymbol{P}}) = \mathcal{C}(\boldsymbol{U}_{\boldsymbol{Q}}) = \{\boldsymbol{u}_{\boldsymbol{R},1}, \ldots, \boldsymbol{u}_{\boldsymbol{R},r^\star}\} \subseteq \mathcal{C}(\boldsymbol{U}_{\boldsymbol{R}}), \tag{151}$$

which concludes the proof of Theorem VII.1.

H. Proof of Lemma VII.1

For any set $\mathcal{A} \subseteq \mathbb{R}_+^n$, the Pareto border $\partial^+\mathcal{A}$ is a subset of the front border $\partial^\mathrm{f}\mathcal{A}$. In fact, if it were not so, then there would exist a Pareto optimal point, say $\boldsymbol{a}' \in \partial^+\mathcal{A}$, which would not be the solution to

$$\max_{\substack{\boldsymbol{a}\in\mathcal{A}\\ \boldsymbol{a}=\nu\boldsymbol{a}'}} \nu \tag{152}$$

in that another $\boldsymbol{a}'' \in \mathcal{A}$ colinear with $\boldsymbol{a}'$ would exist that would have larger norm, i.e., $\|\boldsymbol{a}''\| > \|\boldsymbol{a}'\|$. Yet this is impossible by the definition of $\partial^+\mathcal{A}$, because $\boldsymbol{a}''$ would dominate $\boldsymbol{a}'$ in the sense $\boldsymbol{a}'' \geq \boldsymbol{a}'$, hence the contradiction.

It thus suffices to prove that $\partial^\mathrm{f}\boldsymbol{s}(\Gamma) \subseteq \partial^+\boldsymbol{s}(\Gamma)$ in order to conclude on set equality $\partial^\mathrm{f}\boldsymbol{s}(\Gamma) = \partial^+\boldsymbol{s}(\Gamma)$. For this purpose, take $\boldsymbol{s}'$ to be some point of the front border $\partial^\mathrm{f}\boldsymbol{s}(\Gamma)$. Assume that there would exist another point $\boldsymbol{s}'' \in \boldsymbol{s}(\Gamma)$ different from $\boldsymbol{s}'$ that dominates $\boldsymbol{s}'$, that is, $\boldsymbol{s}'' \geq \boldsymbol{s}'$. For belonging to the set $\boldsymbol{s}(\Gamma)$, which is the union

$$\boldsymbol{s}(\Gamma) = \bigcup_{\substack{(\mu_\mathcal{P},\mu_\mathcal{Q})\\ \mu_\mathcal{P}+(T-T_\tau)\mu_\mathcal{Q}\leq T\mu}} \bigcup_{\boldsymbol{P}\in\mathcal{P}(\mu_\mathcal{P})} \boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_\mathcal{Q})), \tag{153}$$

the point $\boldsymbol{s}''$ would be contained in at least one of the sets $\boldsymbol{s}(\boldsymbol{P}, \mathcal{Q}(\mu_\mathcal{Q}))$. Call $\boldsymbol{P}'' \in \mathcal{P}(\mu_\mathcal{P})$ a pilot Gram of rank $r_{\boldsymbol{P}''}$ such that $\boldsymbol{s}''$ lies in $\boldsymbol{s}(\boldsymbol{P}'', \mathcal{Q}(\mu_\mathcal{Q}))$. According to Theorem V.2, the set $\boldsymbol{s}(\boldsymbol{P}'', \mathcal{Q}(\mu_\mathcal{Q}))$ is a simplex consisting of all convex combinations of $r_{\boldsymbol{P}''}+1$ points $\boldsymbol{\sigma}^{(n)}, n = 0, \ldots, r_{\boldsymbol{P}''}$, with $\boldsymbol{\sigma}^{(0)} = \mathbf{0}$ and [cf. (33)]

$$\boldsymbol{\sigma}^{(n)} = \mathcal{H}(\omega_1, \ldots, \omega_n) \sum_{j=1}^{n} \boldsymbol{e}_j, \quad n \in \{1, \ldots, r_{\boldsymbol{P}''}\}, \tag{154}$$

where ω_i are the non-increasingly ordered eigenvalues of the generalized eigenvalue problem [cf. (31)]

$$\hat{\boldsymbol{R}}''\boldsymbol{v}_i = \omega_i(\mu_\mathcal{Q}^{-1}\mathbf{I} + \tilde{\boldsymbol{R}}'')\boldsymbol{v}_i \tag{155}$$

with $\tilde{\boldsymbol{R}}'' = (\boldsymbol{R}^{-1} + \boldsymbol{P}'')^{-1}$ and $\hat{\boldsymbol{R}}'' = \boldsymbol{R} - \tilde{\boldsymbol{R}}''$. Notice that the linearly independent vectors $\boldsymbol{\sigma}^{(n)}, n = 1, \ldots, r_{\boldsymbol{P}''}$, when linearly combined with non-negative coefficients, span the linear subspace of $\mathbb{R}_+^{N_\mathsf{T}}$ of vectors having non-increasingly sorted entries on positions 1 through $r_{\boldsymbol{P}''}$, and zero entries on positions $r_{\boldsymbol{P}''}+1$ through N_T. Consequently, both $\boldsymbol{s}'$ and $\boldsymbol{s}''$, which by definition have non-increasing non-negative entries, can be written as linear combinations

$$\boldsymbol{s}' = \sum_{n=1}^{r_{\boldsymbol{P}''}} \nu_n'\boldsymbol{\sigma}^{(n)} \qquad \boldsymbol{s}'' = \sum_{n=1}^{r_{\boldsymbol{P}''}} \nu_n''\boldsymbol{\sigma}^{(n)} \tag{156}$$

with unique non-negative coefficients ν_n' and ν_n''. Since $\boldsymbol{s}'' \in \boldsymbol{s}(\boldsymbol{P}'', \mathcal{Q}(\mu_\mathcal{Q}))$, the coefficients ν_n'' sum up to $\sum_n \nu_n'' \leq 1$.

Now, since $\boldsymbol{s}'$ and $\boldsymbol{s}''$ are distinct, and $\boldsymbol{s}' \leq \boldsymbol{s}''$ by assumption, we must have

$$\sum_{n=1}^{r_{\boldsymbol{P}''}} \nu'_n < \sum_{n=1}^{r_{\boldsymbol{P}''}} \nu''_n \leq 1. \tag{157}$$

Therefore $\boldsymbol{s}'$ lies in the interior of $\boldsymbol{s}(\boldsymbol{P}'', \mathcal{Q}(\mu_{\mathcal{Q}}))$. Consequently, for a small enough $\epsilon > 0$, the point $(1+\epsilon)\boldsymbol{s}'$ is element of $\boldsymbol{s}(\boldsymbol{P}'', \mathcal{Q}(\mu_{\mathcal{Q}}))$, and thus of $\boldsymbol{s}(\Gamma)$, which contradicts the initial assumption that $\boldsymbol{s}' \in \partial^{\mathrm{f}}\boldsymbol{s}(\Gamma)$. Hence $\partial^{\mathrm{f}}\boldsymbol{s}(\Gamma) = \partial^{+}\boldsymbol{s}(\Gamma)$.

I. Proof of Lemma VII.2

Clearly, maximizing $\nu(\boldsymbol{p}, \boldsymbol{e})$ as defined in (66) is equivalent to minimizing the function

$$\breve{\nu}(\boldsymbol{p}) = \frac{1}{\nu(\boldsymbol{p}, \boldsymbol{e})} + 1 = \frac{1 + \boldsymbol{r}^{\mathrm{T}}\boldsymbol{q}(\boldsymbol{p}, \boldsymbol{e})}{\eta(\boldsymbol{p}, \boldsymbol{e})}, \tag{158}$$

where contrary to $\nu(\boldsymbol{p}, \boldsymbol{e})$, the direction vector $\boldsymbol{e}$ is omitted in the notation of the function $\breve{\nu}(\boldsymbol{p})$. Writing the latter function out in full with help of definitions (64) and (65) yields

$$\begin{aligned}
\breve{\nu}(\boldsymbol{p}) &= \frac{1 + \frac{T\mu - \sum_i p_i}{T - T_\tau}\left(\sum_i e_i \frac{1+r_i p_i}{r_i^2 p_i}\right)^{-1}\left(\sum_i e_i \frac{1+r_i p_i}{r_i p_i}\right)}{\frac{T\mu - \sum_i p_i}{T - T_\tau}\left(\sum_i e_i \frac{1+r_i p_i}{r_i^2 p_i}\right)^{-1}} \\
&= \frac{T - T_\tau}{T\mu - \sum_i p_i}\left(\sum_i e_i \frac{1 + r_i p_i}{r_i^2 p_i}\right) + \sum_i e_i \frac{1 + r_i p_i}{r_i p_i} \\
&\triangleq (T - T_\tau)\big(\breve{\nu}_1(\boldsymbol{p}) + \breve{\nu}_2(\boldsymbol{p})\big) + \breve{\nu}_3(\boldsymbol{p}),
\end{aligned} \tag{159}$$

the three functions $\breve{\nu}_1(\boldsymbol{p})$, $\breve{\nu}_2(\boldsymbol{p})$, $\breve{\nu}_3(\boldsymbol{p})$ in the last line being

$$\breve{\nu}_1(\boldsymbol{p}) = \sum_j e_j r_j^{-2} \frac{1}{p_j(T\mu - \sum_i p_i)} \tag{160a}$$

$$\breve{\nu}_2(\boldsymbol{p}) = \sum_j e_j r_j^{-1} \frac{1}{T\mu - \sum_i p_i} \tag{160b}$$

$$\breve{\nu}_3(\boldsymbol{p}) = \sum_i e_i \frac{1 + r_i p_i}{r_i p_i}. \tag{160c}$$

We will now show that the three functions $\breve{\nu}_1(\boldsymbol{p})$, $\breve{\nu}_2(\boldsymbol{p})$ and $\breve{\nu}_3(\boldsymbol{p})$ are all convex functions of $\boldsymbol{p}$ on the interior of $\mathcal{D}(T\mu) \subset (0;\infty)^{N_{\mathsf{T}}}$, which we shall denote as $\mathrm{int}(\mathcal{D}(T\mu))$. It is easy to see that $\breve{\nu}_3$ is essentially a linear combination (plus a constant) of functions $1/p_i$ that are convex on the entire open orthant $(0,\infty)^{N_{\mathsf{T}}}$, and thus on $\mathrm{int}(\mathcal{D}(T\mu)) \subset (0;\infty)^{N_{\mathsf{T}}}$. Similarly, $\breve{\nu}_2$ is convex on the open half-space $\sum_i p_i < T\mu$, and thus on the subset $\mathrm{int}(\mathcal{D}(T\mu))$ thereof. Finally, $\breve{\nu}_1$ is a linear combination of functions $\frac{1}{p_j}\frac{1}{T\mu - \sum_i p_i}$, each of which is convex in $\boldsymbol{p}$ on $\mathrm{int}(\mathcal{D}(T\mu))$. This can be shown as follows: take a pair of points $\big(\boldsymbol{p}^{(1)}, \boldsymbol{p}^{(2)}\big) \in \mathrm{int}(\mathcal{D}(T\mu))^2$, then for any $\theta \in [0;1]$,

$$\begin{aligned}
&\frac{1}{\theta p_j^{(1)} + (1-\theta)p_j^{(2)}} \frac{1}{T\mu - \sum_i\big(\theta p_i^{(1)} + (1-\theta)p_i^{(2)}\big)} \\
&\leq \theta \frac{1}{p_j^{(1)}} \frac{1}{T\mu - \sum_i p_i^{(1)}} + (1-\theta)\frac{1}{p_j^{(2)}} \frac{1}{T\mu - \sum_i p_i^{(2)}},
\end{aligned} \tag{161}$$

because the left-hand side of the latter inequality is convex in $\theta \in [0;1]$, since it is of the form

$$A \frac{1}{1 + B\theta} \frac{1}{1 + C\theta} \tag{162}$$

with constants $A = \frac{1}{p_j^{(2)}(T\mu - \sum_i p_i^{(2)})} \geq 0$, $B = \frac{p_j^{(1)} - p_j^{(2)}}{p_j^{(2)}}$, $C = \frac{\sum_i (p_i^{(2)} - p_i^{(1)})}{T\mu - \sum_i p_i^{(2)}}$, and $1 + B\theta \geq 0$ and $1 + C\theta \geq 0$ by construction. The convexity of (162) is best seen by differentiating twice:

$$\begin{aligned}
\frac{\mathrm{d}^2}{\mathrm{d}\theta^2}\left[\frac{1}{1+B\theta}\frac{1}{1+C\theta}\right] &= \frac{2}{(1+B\theta)(1+C\theta)} \\
&\times \left(\frac{B^2}{(1+B\theta)^2} + \frac{C^2}{(1+C\theta)^2} + \frac{BC}{(1+B\theta)(1+C\theta)}\right).
\end{aligned} \tag{163}$$

The above expression is obviously positive if $BC \geq 0$. Otherwise, if $BC \leq 0$, then the expression between square brackets on the right-hand side of the last equality is lower bounded by

$$\begin{aligned}
&\frac{B^2}{(1+B\theta)^2} + \frac{C^2}{(1+C\theta)^2} + \frac{2BC}{(1+B\theta)(1+C\theta)} \\
&\qquad = \left[\frac{B}{(1+B\theta)} + \frac{C}{(1+C\theta)}\right]^2 \geq 0.
\end{aligned} \tag{164}$$

Hence (161), and all the three functions $\breve{\nu}_1$, $\breve{\nu}_2$ and $\breve{\nu}_3$ are convex in $\boldsymbol{p}$ on the open set $\mathrm{int}(\mathcal{D}(T\mu))$. Thus, $\breve{\nu}(\boldsymbol{p})$ is convex on $\mathrm{int}(\mathcal{D}(T\mu))$. Therefore $\nu(\boldsymbol{p}, \boldsymbol{e}) = 1/(\breve{\nu}(\boldsymbol{p}) - 1)$, which is a decreasing function of $\breve{\nu}(\boldsymbol{p}) > 1$, is quasi-concave in $\boldsymbol{p}$ on $\mathrm{int}(\mathcal{D}(T\mu))$, according to Definition VII.1. Since $\nu(\boldsymbol{p}, \boldsymbol{e})$ vanishes on the boundary of $\mathcal{D}(T\mu)$ and is continuous in the vicinity of this boundary, we conclude that $\boldsymbol{p} \mapsto \nu(\boldsymbol{p}, \boldsymbol{e})$ is quasi-concave on the closure $\mathcal{D}(T\mu)$.

REFERENCES

[1] B. Hassibi and B. Hochwald, "How much training is needed in multiple-antenna wireless links?" *IEEE Transactions on Information Theory*, vol. 49, no. 4, pp. 951–963, Apr. 2003.

[2] A. Pastore, M. Joham, and J. Fonollosa, "Joint pilot and precoder design for optimal throughput," in *Proc. IEEE International Symposium on Information Theory*, Aug. 2011, pp. 371–375.

[3] A. Soysal and S. Ulukus, "Joint channel estimation and resource allocation for MIMO systems–Part I: Single-user analysis," *IEEE Transactions on Wireless Communications*, vol. 9, no. 2, pp. 624–631, Feb. 2010.

[4] A. Aubry, A. Tulino, and S. Venkatesan, "Multiple-access channel capacity region with incomplete channel state information," in *Proc. of IEEE International Symposium on Information Theory*, Jun. 2010, pp. 2313–2317.

[5] A. Aubry, I. Esnaola, A. Tulino, and S. Venkatesan, "Achievable rate region for gaussian MIMO MAC with partial CSI," *IEEE Transactions on Information Theory*, vol. 59, no. 7, pp. 4139–4170, 2013.

[6] L. Tong, B. M. Sadler, and M. Dong, "Pilot-assisted wireless transmissions: general model, design criteria, and signal processing," *IEEE Signal Processing Magazine*, vol. 21, no. 6, pp. 12–25, Nov. 2004.

[7] M. Biguesh and A. Gershman, "Training-based MIMO channel estimation: a study of estimator tradeoffs and optimal training signals," *IEEE Transactions on Signal Processing*, vol. 54, no. 3, pp. 884–893, Mar. 2006.

[8] M. Médard, "The effect upon channel capacity in wireless communications of perfect and imperfect knowledge of the channel," *IEEE Transactions on Information Theory*, vol. 46, no. 3, pp. 933–946, May 2000.

[9] J. Baltersee, G. Fock, and H. Meyr, "Achievable rate of MIMO channels with data-aided channel estimation and perfect interleaving," *IEEE Journal on Selected Areas in Communications*, vol. 19, no. 12, pp. 2358–2368, Dec. 2001.

[10] T. Yoo, E. Yoon, and A. Goldsmith, "MIMO capacity with channel uncertainty: Does feedback help?" in *Proc. IEEE Global Telecommunications Conference*, vol. 1, Dec. 2004, pp. 96–100.

[11] L. Musavian, M. Dohler, M. Nakhai, and A. Aghvami, "Transmitter design in partially coherent antenna systems," in *Proc. IEEE International Conference on Communications*, vol. 4, May 2005, pp. 2261–2265.

[12] A. Lozano, "Interplay of spectral efficiency, power and Doppler spectrum for reference-signal-assisted wireless communication," *IEEE Transactions on Wireless Communications*, vol. 7, no. 12, pp. 5020–5029, Dec. 2008.

[13] M. Ding and S. Blostein, "Maximum mutual information design for MIMO systems with imperfect channel knowledge," *IEEE Transactions on Information Theory*, vol. 56, no. 10, pp. 4793–4801, Oct. 2010.

[14] A. Lozano, A. M. Tulino, and S. Verdú, "Multiantenna capacity: myths and realities," in *Space-time Wireless Systems: From Array Processing to MIMO Communications*, H. Bölcskei, Ed. Cambridge: Cambridge University Press, 2006.

[15] E. Jorswieck and H. Boche, "Channel capacity and capacity-range of beamforming in MIMO wireless systems under correlated fading with covariance feedback," *IEEE Transactions on Wireless Communications*, vol. 3, no. 5, pp. 1543–1553, Sep. 2004.

[16] ——, "Optimal Transmission Strategies and Impact of Correlation in Multiantenna Systems with Different Types of Channel State Information," *IEEE Transactions on Signal Processing*, vol. 52, no. 12, pp. 3440–3453, Dec. 2004.

[17] S. Jafar and A. Goldsmith, "Transmitter optimization and optimality of beamforming for multiple antenna systems," *IEEE Transactions on Wireless Communications*, vol. 3, no. 4, pp. 1165–1175, Jul. 2004.

[18] T. Yoo and A. Goldsmith, "Capacity and power allocation for fading MIMO channels with channel estimation error," *IEEE Transactions on Information Theory*, vol. 52, no. 5, pp. 2203–2214, May 2006.

[19] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, Mar. 2004.

[20] A. Tulino, A. Lozano, and S. Verdú, "Capacity-achieving input covariance for single-user multi-antenna channels," *IEEE Transactions on Wireless Communications*, vol. 5, no. 3, pp. 662–671, Mar. 2006.

[21] E. Gauthier, A. Yongaçoglu, and J. Chouinard, "Capacity of multiple antenna systems in Rayleigh fading channels," in *Canadian Conference on Electrical and Computer Engineering*, vol. 1, 2000, pp. 275–279.

[22] A. Grant, "Rayleigh fading multi-antenna channels," *EURASIP Journal of Applied Signal Processing*, vol. 2002, no. 1, pp. 316–329, Jan. 2002.

[23] O. Oyman, R. Nabar, H. Bölcskei, and A. Paulraj, "Tight lower bounds on the ergodic capacity of Rayleigh fading MIMO channels," in *Proc. IEEE Global Telecommunications Conference*, vol. 2, Nov. 2002, pp. 1172–1176.

[24] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, Feb. 1990.