

Hardware-Limited Task-Based Quantization

Nir Shlezinger¹, Member, IEEE, Yonina C. Eldar², Fellow, IEEE,
and Miguel R. D. Rodrigues³, Senior Member, IEEE

Abstract—Quantization plays a critical role in digital signal processing systems. Quantizers are typically designed to obtain an accurate digital representation of the input signal, operating independently of the system task, and are commonly implemented using serial scalar analog-to-digital converters (ADCs). In this work, we study hardware-limited task-based quantization, where a system utilizing a serial scalar ADC is designed to provide a suitable representation in order to allow the recovery of a parameter vector underlying the input signal. We propose hardware-limited task-based quantization systems for a fixed and finite quantization resolution, and characterize their achievable distortion. We then apply the analysis to the practical setups of channel estimation and eigen-spectrum recovery from quantized measurements. Our results illustrate that properly designed hardware-limited systems can approach the optimal performance achievable with vector quantizers, and that by taking the underlying task into account, the quantization error can be made negligible with a relatively small number of bits.

Index Terms—Quantization, analog-to-digital conversion.

I. INTRODUCTION

QUANTIZATION refers to the representation of a continuous-amplitude signal using a finite dictionary, or equivalently, a finite number of bits [1]. Quantizers are implemented in digital signal processing systems using analog-to-digital converters (ADCs), which typically operate in a serial scalar manner due to hardware-limitations. In such systems, each incoming continuous-amplitude sample is represented in digital form using the same mechanism [2]. The quantized representation is commonly selected to accurately match the original signal, in the sense of minimizing some distortion measure, such that the signal can be recovered with minimal error from the quantized measurements [3, Ch. 10], [4].

Quantization design is typically performed regardless of the system task. However, in many signal processing applications, the goal is not to recover the actual signal, but to capture certain

parameters, such as an underlying model or unknown channel, from the quantized signal [5]. We refer to systems where one wishes to extract some information from the quantized signal, rather than recovering the signal itself, as *task-based quantization*, and to such systems operating with serial scalar ADCs as *hardware-limited task-based quantization* systems.

Hardware-limited quantization with low resolution is the focus of growing interest over recent years due to the increasing complexity and bitrate demands of modern signal processing and communications systems. Common tasks considered with low resolution hardware-limited quantization include multiple-input multiple-output (MIMO) communications [6]–[12], channel estimation [11]–[17], subspace estimation [18], time difference of arrival estimation [19], and direction of arrival (DOA) estimation [20], [21]. In these works it is assumed that quantization is carried out *separately from the system task*, typically using *fixed uniform* low-precision quantizers, e.g., one-bit quantization of a scalar value is implemented using the sign function [6]. These quantized measurements are then processed in the digital domain using different inference algorithms, such as linear estimators [11], [14], maximum-likelihood estimation [15], and approximate message passing based algorithms [10], [16], [17]. As these works focus only on the digital processing, they do not provide guidelines to designing quantization systems with a small and finite number of bits by acknowledging the task of the system.

When hardware-limitations are not present, task-based quantization systems can take advantage of joint vector quantization, which is known to be superior to serial scalar quantization [22, Ch. 22.2]. Previous works on task-based quantization without hardware limitations are divided according to whether the parameter vector is modeled as a random vector, namely, a Bayesian setup, or as a deterministic unknown parameter. When the signal parameter is random, task-based quantization can be viewed as an indirect lossy source coding problem¹ [1, Section V-G]. For this setup with a stationary source that is related to the observation vector via a stationary memoryless channel, Witsenhausen proved in [23] that the rate-distortion function, namely, the minimal number of bits required to obtain a given representation accuracy determined by the distortion measure, is asymptotically equivalent to the rate-distortion function for representing the observed signal – instead of the signal parameter – with a surrogate distortion measure. Under mean-squared error (MSE) distortion, Wolf and Ziv proved in [24] that this equivalence also holds for finite signal size, and the work [5] provided guidelines to the optimal joint quantization and estimation scheme. Recently, Kostina and Verdu characterized nonasymptotic bounds on the rate-distortion functions for indirect as well as direct lossy source coding with

¹Direct lossy source coding typically refers to the standard quantization setup where the task of the system is to recover the quantized signal, while indirect source coding refers to task-based quantization [25].

Manuscript received July 22, 2018; revised December 9, 2018 and July 28, 2019; accepted July 28, 2019. Date of publication August 19, 2019; date of current version September 12, 2019. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Youngchul Sung. This work was supported in part by the European Unions Horizon 2020 research and innovation program under Grant 646804-ERC-COG-BNYQ, in part by the Israel Science Foundation under Grant 0100101, and in part by the Royal Society International Exchange scheme IE 160348. This paper was presented in part at the 20th IEEE International Workshop on Signal Processing Advances in Wireless Communications, Cannes, France, July 2–5, 2019. (*Corresponding author: Nir Shlezinger.*)

N. Shlezinger and Y. C. Eldar are with the Faculty of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot 7610001, Israel (e-mail: nirshlezinger1@gmail.com; yonina@weizmann.ac.il).

M. R. D. Rodrigues is with the Department of Electrical Engineering, University College London, London WC1E 6BT, U.K. (e-mail: m.rodrigues@ucl.ac.uk).

Digital Object Identifier 10.1109/TSP.2019.2935864

arbitrary distortion measures [25], [26], by considering single-shot quantization, and specialized the bounds for i.i.d. signals with separable distortion. The indirect source coding framework was also used to study conversion of continuous-time signals into quantized discrete-time signals in [27], [28]. The focus in the works [23]–[28] is on the *optimal* tradeoff between quantization rate and achievable distortion. Consequently, their results cannot be applied to quantify the achievable performance of practical hardware-limited systems utilizing serial scalar ADCs.

For a signal parameter modeled as deterministic and unknown, [29] studied detection from quantized observations, i.e., recovering a scalar binary parameter, while [30] treated detection from quantized prior probabilities. Quantization for the recovery of a scalar parameter taking values on a discrete finite set was studied in [31]. The design of quantizers for the recovery of a vector parameter taking values on a continuous set was considered in [32], which proposed an adaptive algorithm for tuning the quantizer. In all the works above, i.e., [29]–[32], as well as in [5], the analysis assumes vector quantizers with high resolution, where the number of bits used for representing the quantized signal can be made arbitrarily large. They do not consider practical systems that utilize serial scalar quantizers with a fixed and finite number of bits.

A. Main Contributions

In this work we study quantization systems which are hardware-limited to utilize practical serial scalar ADCs operating with a fixed number of bits, for the task of acquiring a random parameter vector taking values on a continuous set. Our approach is to account for the task in design of the quantization system, thus mitigating the effect of the structure imposed by hardware limitations. This hardware-limited task-based framework fits a broad range of signal processing and communications systems. In particular, a task may be any data processing or inference objective, and various constraints encountered in practice, such as sampling rate limitations, delay restrictions, and imposed structures, can be represented as hardware limitations. Consequently, although our derivation here considers quantization for the task of estimating a desired parameter vector while restricted to using scalar ADCs, the underlying approach, i.e., to overcome hardware limitations by properly accounting for a task in the system design, is applicable to a much broader family of systems.

We first consider the case where the observations and the desired vector are related such that the minimum mean-squared error (MMSE) estimate is a linear function of the observations. Such relationships are commonly encountered in channel estimation and signal recovery problems, e.g., [5], [8]–[16]. We focus on structured systems implementing uniform quantization with linear processing, allowing analog combining prior to digital processing. This approach was previously studied in the context of MIMO communications [6], [33]–[35]. For this setup, we derive the hardware-limited task-based quantization system which minimizes the MSE, and characterize the achievable distortion. The proposed system accounts for the task by reducing the number of quantized samples via an appropriate linear transformation to be not larger than the size of the desired signal. It then rotates the quantized samples to have identical variance. Quantization is performed based on a waterfilling-type expression, accounting for the serial operation and the limited dynamic range of practical ADCs.

In addition, we characterize the minimal achievable distortion of two suboptimal approaches: We first discuss systems in which

processing is carried out only in the digital domain, as is the structure considered in the majority of the literature of tasks performed with low resolution quantization, e.g., [13]–[16], [18]–[20]. Then, we study systems which quantize the MMSE estimate, an approach which is known to be optimal when using vector quantizers [24], and was also proposed for compressed sensing with quantized measurements [36]. Surprisingly, we show that, unlike when vector quantizers are employed, in the presence of serial scalar ADCs, quantizing the MMSE estimate is generally suboptimal. We provide a necessary and sufficient condition for this approach to coincide with the optimal design.

Next, we extend the proposed system to scenarios where the observations and the desired vector are related via an arbitrary stochastic model. In particular, we identify the main design guidelines associated with the case where the MMSE estimate is a linear function of the observations, and discuss how they may be applied for arbitrary models. Then, we explicitly show how these guidelines can be used to construct a hardware-limited task-based quantization system for scenarios in which the desired vector can be recovered from the empirical covariance of the observations, as in [5], [18]–[21].

Finally, we apply our results to two practical setups: Channel estimation from quantized measurements [11]–[16] and eigen-spectrum estimation from quantized measurements [5]. We demonstrate that, by properly accounting for the presence of serial scalar ADCs, practical hardware-limited systems operating with a relatively small number of bits approach the optimal performance, achievable with vector quantizers, in practical and relevant scenarios. Furthermore, we show that hardware-limited quantizers designed accounting for the task of the system can substantially outperform task-ignorant systems utilizing vector quantizers. This gain is mainly achieved by applying task-based linear analog processing, in addition to the digital processing.

B. Organization and Notations

The rest of this paper is organized as follows: Section II briefly reviews some preliminaries in quantization theory, and formulates the hardware-limited task-based quantization setup. Section III discusses task-based quantization with vector quantizers. Section IV studies hardware-limited task-based quantization when the MMSE estimate is linear, and Section V extends the proposed design to arbitrary setups. Section VI presents the application of the results in a numerical study. Section VII provides some concluding remarks. Detailed proofs of the results are given in the appendix.

Throughout the paper, we use boldface lower-case letters for vectors, e.g., \mathbf{x} ; the i th element of \mathbf{x} is written as $(\mathbf{x})_i$. Matrices are denoted with boldface upper-case letters, e.g., \mathbf{M} , and $(\mathbf{M})_{i,j}$ is its (i,j) th element. Sets are denoted with calligraphic letters, e.g., \mathcal{X} , and \mathcal{X}^n is the n th order Cartesian power of \mathcal{X} . Transpose, Euclidean norm, trace, stochastic expectation, sign, and mutual information are written as $(\cdot)^T$, $\|\cdot\|$, $\text{Tr}(\cdot)$, $\mathbb{E}\{\cdot\}$, $\text{sign}(\cdot)$, and $I(\cdot;\cdot)$, respectively, and \mathcal{R} is the set of real numbers. We use a^+ to denote $\max(a, 0)$, and \mathbf{I}_n is the $n \times n$ identity matrix. All logarithms are taken to basis 2.

II. PRELIMINARIES AND PROBLEM STATEMENT

A. Preliminaries in Quantization Theory

To formulate the hardware-limited task-based quantization problem, we first review standard quantization notations, after




Fig. 1. Quantizer illustration.

which we discuss task-based quantization. To that aim, we recall the definition of a quantizer:

Definition 1 (Quantizer): A quantizer $Q_M^{n,k}(\cdot)$ with $\log M$ bits, input size n , input alphabet \mathcal{X} , output size k , and output alphabet $\hat{\mathcal{X}}$, consists of: 1) An encoding function $g_n^e: \mathcal{X}^n \mapsto \{1, 2, \dots, M\} \triangleq \mathcal{M}$ which maps the input into a discrete index $i \in \mathcal{M}$. 2) A decoding function $g_k^d: \mathcal{M} \mapsto \hat{\mathcal{X}}^k$ which maps each index $i \in \mathcal{M}$ into a codeword $q_i \in \hat{\mathcal{X}}^k$.

We write the output of the quantizer with input $x \in \mathcal{X}^n$ as $\hat{x} = g_k^d(g_n^e(x)) \triangleq Q_M^{n,k}(x)$. *Scalar quantizers* operate on a scalar input, i.e., $n = 1$ and \mathcal{X} is a scalar space, while *vector quantizers* have a multivariate input. The set of codewords $\{q_i\}_{i=1}^M$ is referred to as the *quantization codebook*. When the input size and output size are equal, namely, $n = k$, we write $Q_M^n(\cdot) \triangleq Q_M^{n,n}(\cdot)$. An illustration is given in Fig. 1.

1) **Standard Quantization:** In the standard quantization problem, a $Q_M^n(\cdot)$ quantizer is designed to minimize some distortion measure $d_n: \mathcal{X}^n \times \hat{\mathcal{X}}^n \mapsto \mathcal{R}^+$ between its input and its output. The performance of a quantizer is therefore characterized using two measures: The quantization rate, defined as $R \triangleq \frac{1}{n} \log M$, and the expected distortion $\mathbb{E}\{d_n(x, \hat{x})\}$. For a fixed input size n and codebook size M , the optimal quantizer is thus given by

$$Q_M^{n,\text{opt}}(\cdot) = \arg \min_{Q_M^n(\cdot)} \mathbb{E}\{d_n(x, Q_M^n(x))\}. \quad (1)$$

Characterizing the optimal quantizer via (1) and the optimal tradeoff between distortion and quantization rate is in general a very difficult task. Consequently, optimal quantizers are typically studied assuming either high quantization rate, i.e., $R \rightarrow \infty$, see, e.g., [37], or asymptotically large input size, namely, $n \rightarrow \infty$, typically with i.i.d. inputs,² via rate-distortion theory [3, Ch. 10]. For example, when the quantizer input consists of i.i.d. random variables (RVs) with probability measure f_x , and the distortion measure can be written as $d_n(x, \hat{x}) = \frac{1}{n} \sum_{i=1}^n d((x)_i, (\hat{x})_i)$ for some $d: \mathcal{X} \times \mathcal{X} \mapsto \mathcal{R}^+$, then the optimal distortion in the limit $n \rightarrow \infty$ for a given rate R is given by the distortion-rate function:

Definition 2 (Distortion-rate function): The distortion-rate function for input $x \in \mathcal{X}$ with respect to the distortion measure $d: \mathcal{X} \times \mathcal{X} \mapsto \mathcal{R}^+$ is defined as

$$D_x(R) = \min_{f_{\hat{x}|x}: I(\hat{x}; x) \leq R} \mathbb{E}\{d(\hat{x}, x)\}. \quad (2)$$

The conditional distribution which obtains the minima in (2), $f_{\hat{x}|x}^{\text{opt}}$, is referred to as the *optimal distortion-rate distribution*, and the marginal distribution $f_{\hat{x}} = \int f_{\hat{x}|x}^{\text{opt}} f_x$ is referred to henceforth as the *optimal marginal distortion-rate distribution*.

Comparing high quantization rate analysis for scalar quantizers and rate-distortion theory for vector quantizers demonstrates

²Rate-distortion theory can also be used for non i.i.d. signals, see, e.g., [38, Ch. 5]. However, the simple classical expression, as given by the distortion-rate function in Def. 2, requires the observed signal to have i.i.d. entries.




Fig. 2. Task-based quantization illustration.




Fig. 3. Hardware-limited task-based quantizer.

the sub-optimality of serial scalar quantization. For example, for large R , even for i.i.d. inputs, vector quantization outperforms serial scalar quantization, with a distortion gap of 4.35 dB for Gaussian inputs with the MSE distortion [22, Ch. 23.2].

2) **Task-Based Quantization:** In task-based quantization the design objective of the quantizer is some task other than minimizing the distortion between its input and output. In the following, we focus on the generic task of acquiring a zero-mean random vector $s \in \mathcal{R}^k$ from a measured zero-mean random vector $x \in \mathcal{R}^n$, where the entries of x and s have finite variance, and $n \geq k > 0$. This formulation accommodates a broad range of tasks, including channel estimation, covariance estimation, and source localization. A natural distortion measure for such setups is the MSE, which we consider throughout the paper. An illustration of a task-based quantization system is depicted in Fig. 2.

B. Problem Formulation

In this work we study task-based quantization with hardware limitations. As discussed in the introduction, practical digital signal processing systems typically obtain a digital representation of physical analog signals using serial scalar ADCs. We refer to task-based quantization with serial scalar ADCs as *hardware-limited task-based quantization*. Since in such systems, each continuous-amplitude sample is converted into a discrete representation using a single quantization rule, this operation can be modeled using *identical scalar quantizers*. Consequently, the system we consider is modeled using the setup depicted in Fig. 3. The observed signal $x \in \mathcal{R}^n$ is projected into \mathcal{R}^p , $p \leq n$, using some mapping $h_a(\cdot)$, which represents the pre-quantization processing carried out in the analog domain. Since general mappings may be difficult to implement in analog, we henceforth restrict $h_a(\cdot)$ to be a linear function, namely, we only allow *linear analog combining*, as in, e.g., [6], [33]. In this case, $h_a(x) = Ax$ for some $A \in \mathcal{R}^{p \times n}$.

Each entry of $h_a(x)$ is quantized using the same scalar quantizer with resolution $\tilde{M}_p \triangleq \lfloor M^{1/p} \rfloor$, denoted $Q_{\tilde{M}_p}^1(\cdot)$. The overall number of quantization levels is thus $(\tilde{M}_p)^p \leq M$. We note that M , which represents the memory requirement of the system, is also directly related to the ADC power consumption. However, for the same overall number of quantization levels M , different selections of p may result in different power consumptions, depending on the physical implementation of the ADC, see, e.g., [35]. In the following we keep the value of M fixed and finite, i.e., the memory requirement, which is independent

of the specific implementation of the ADC, is the same for all considered systems.

The representation of s , denoted \hat{s} , is obtained as the output of some post-quantization mapping $h_d : \mathcal{R}^p \mapsto \mathcal{R}^k$, applied to the output of the identical scalar quantizers. The mapping $h_d(\cdot)$ represents the joint-processing carried out in the digital domain. The quantized representation \hat{s} can be written as

$$\hat{s} = h_d \left(Q_{M_p}^1((h_a(\mathbf{x}))_1), \dots, Q_{M_p}^1((h_a(\mathbf{x}))_p) \right). \quad (3)$$

The novelty of the model in Fig. 3, compared to previous works on quantization for specific tasks with serial scalar ADCs, e.g., [8]–[21], is in the introduction of the additional linear processing carried out in the analog domain, represented by the mapping $h_a(\cdot)$. The concept of using analog combining prior to digital processing was previously studied in the context of MIMO communications in [6], [7], [33]–[35]. The motivation for introducing $h_a(\cdot)$ is to reduce the dimensionality of the input to the ADC, thus facilitating a more accurate quantization without increasing the overall number of bits, $\log M$. As shown in the following sections, by properly designing $h_a(\cdot)$, this approach can substantially improve the performance of task-based quantizers operating with serial scalar ADCs.

Our analysis of hardware-limited task-based quantization, focusing on the MSE distortion, consists of three parts:

- 1) As a preliminary step, in Section III, we discuss non hardware-limited task-based quantization systems, namely, systems implementing task-based quantization using optimal vector quantizers instead of serial scalar ADCs. The purpose of this analysis is to serve as a basis for comparing the performance of hardware-limited task-based quantizers to vector quantizers.
- 2) Next, in Section IV, we focus on the case where \mathbf{x} and s are related such that the MMSE estimate of s from \mathbf{x} is a linear function of \mathbf{x} . Such relationships arise in various channel estimation and signal recovery setups, e.g., [5], [8]–[16]. For this setting, we propose a hardware-limited quantization system design, and characterize its achievable distortion. We also characterize the minimal achievable distortion when no pre-quantization processing is carried out, as well as when the analog combiner is designed to recover the MMSE estimate.
- 3) Then, in Section V, we use the characterization of $h_a(\cdot)$ and $h_d(\cdot)$ given in Section IV for linear models, to provide guidelines for designing $h_a(\cdot)$ and $h_d(\cdot)$ under arbitrary relationships between \mathbf{x} and s . We suggest a concrete design for cases in which s can be estimated from the second-order statistical moments of \mathbf{x} , as in [5], [18]–[21].

Our analysis shows that, unlike when vector quantizers are applied, the optimal strategy for systems utilizing serial scalar ADCs is not to quantize the MMSE estimate. Instead, the input to the ADC is rotated to account for the identical quantization rule of serial scalar ADCs, and includes a waterfilling-type expression to account for the limited dynamic range. Furthermore, our numerical comparison presented in Section VI demonstrates that the proposed system, which uses simple hardware, can approach the performance of the optimal vector quantizer.

III. TASK-BASED VECTOR QUANTIZATION

As a preliminary step towards our study of hardware-limited task-based quantizers, we consider task-based quantization which utilizes vector quantizers without hardware limitations.




Fig. 4. Task-ignorant quantizer.

We focus on two approaches for task-based quantization: In the first, referred to as *optimal task-based quantization*, the quantizer $Q_{n,k}^M(\cdot)$ in Fig. 2 is designed to recover the desired vector s . In the second strategy, described in Fig. 4 and referred to as *task-ignorant quantization*, the quantizer is designed to recover the observed vector \mathbf{x} separately from the task, and s is estimated from the quantized representation. The optimal task-based quantizer obtains the minimal achievable distortion for a given quantization rate, while the task-ignorant quantizer represents the best system one can construct when the quantizer is designed separately from the task. The approaches we discuss below are based on joint vector quantization, and thus cannot be implemented using practical serial scalar ADCs.

1) *Optimal Task-Based Quantizer*: For the MSE distortion, the optimal quantizer is constructed by first obtaining the MMSE estimate of s from \mathbf{x} , $\tilde{s} = \mathbb{E}\{s|\mathbf{x}\}$, and then quantizing the estimate [24]. This leads to a minimal distortion given by

$$\begin{aligned} \min_{Q_M^{n,k}(\cdot)} \mathbb{E} \left\{ \left\| s - Q_M^{n,k}(\mathbf{x}) \right\|^2 \right\} \\ = \mathbb{E} \left\{ \left\| s - \tilde{s} \right\|^2 \right\} + \min_{Q_M^k(\cdot)} \mathbb{E} \left\{ \left\| \tilde{s} - Q_M^k(\tilde{s}) \right\|^2 \right\}. \quad (4) \end{aligned}$$

It follows from (4) that the minimal distortion is the sum of the minimal estimation error of s from \mathbf{x} , and the minimal distortion in quantizing the MMSE estimate \tilde{s} . The latter can be obtained explicitly under a high quantization rate assumption, i.e., $\frac{1}{k} \log M \rightarrow \infty$, using fine quantization analysis, as was done in [5], or alternatively, when \tilde{s} has i.i.d. entries and k tends to infinity, using rate-distortion theory. For finite k , n and M , the minimal distortion in quantizing the MMSE estimate may be bounded as stated in the following proposition:

Proposition 1: For any random vector $\tilde{c} \in \mathcal{R}^k$ with probability measure $f_{\tilde{c}}$ independent of \tilde{s} , the minimal MSE in quantizing the MMSE estimate \tilde{s} using a Q_M^k quantizer satisfies

$$\begin{aligned} D_{\tilde{s}}(\log M) &\leq \min_{Q_M^k(\cdot)} \mathbb{E} \left\{ \left\| \tilde{s} - Q_M^k(\tilde{s}) \right\|^2 \right\} \\ &\leq \mathbb{E} \left\{ \int_0^\infty \left[\Pr \left(\left\| \tilde{c} - \tilde{s} \right\|^2 > t \mid \tilde{s} \right) \right]^M dt \right\}. \quad (5) \end{aligned}$$

Proof: See Appendix A.

The bounds in (5) are used in the sequel for comparing the performance of hardware-limited task-based quantization to the optimal performance achievable using vector quantizers. These bounds are required since evaluating the minimal distortion of a vector quantizer, i.e., the middle term in (5), is a challenging task for finite signal size and quantization resolution. The upper bound in (5) is the exact performance of random coding, which is known to provide a relatively tight bound for fixed blocklengths [26], and to asymptotically achieve the distortion-rate curve [22, Ch. 23.2]. A reasonable assignment for $f_{\tilde{c}}$ in (5) is the optimal marginal distortion-rate distribution; with this distribution the distortion of quantizers with i.i.d. random codewords coincides

with the distortion-rate function for sources generating asymptotically large number of i.i.d. realizations of \tilde{s} [22, Ch. 24.2]. Alternative upper bounds, which may be tighter and simpler to compute compared to the one used in (5) yet are not analytically expressible, can be obtained by numerically evaluating the distortion of iterative and data-driven vector quantizer designs, such as the extension of Lloyd's algorithm to multivariate inputs [39]. In general, the distortion-rate function and the optimal marginal distribution can be obtained using iterative algorithms, e.g., the Blahut-Arimoto algorithm [3, Ch. 10.8] and its extensions to continuous-valued RVs [40].

2) *Task-Ignorant Quantizer*: When the quantizer operates independently of the task, the desired vector s must be estimated directly from the quantized observations. For the optimal quantizer and estimator for this setup, $Q_M^n(\cdot)$ minimizes the MSE between its output and x , and s is estimated from the output of the quantizer using the MMSE estimator. From the orthogonality principle, the resulting MSE in estimating s is

$$\begin{aligned} & \mathbb{E} \left\{ \left\| s - \mathbb{E} \{ s | Q_M^n(x) \} \right\|^2 \right\} \\ &= \mathbb{E} \left\{ \left\| s - \tilde{s} \right\|^2 \right\} + \mathbb{E} \left\{ \left\| \tilde{s} - \mathbb{E} \{ s | Q_M^n(x) \} \right\|^2 \right\} \\ &\stackrel{(a)}{=} \mathbb{E} \left\{ \left\| s - \tilde{s} \right\|^2 \right\} + \mathbb{E} \left\{ \left\| \tilde{s} - \mathbb{E} \{ \tilde{s} | Q_M^n(x) \} \right\|^2 \right\}, \quad (6) \end{aligned}$$

where (a) follows since $s \mapsto x \mapsto Q_M^n(x)$ form a Markov chain, thus, by [41, Prop. 4], $\mathbb{E} \{ s | Q_M^n(x) \} = \mathbb{E} \{ \tilde{s} | Q_M^n(x) \}$. The relation in (6) shows that the distortion of the task-ignorant quantizer is given by the sum of the estimation error of the MMSE estimate \tilde{s} and the estimation error of the MMSE estimate of \tilde{s} from the quantizer output $Q_M^n(x)$. The main difference between (6) and the optimal estimation error in (4) is that in (6) the quantizer is fixed, while in (4) it can be set to minimize the estimation error.

In order to compute (6), the distribution of \tilde{s} given $Q_M^n(x)$ is required, which may be difficult to characterize. One scenario in which this requirement can be relaxed is when \tilde{s} is a linear function of x , i.e., $\tilde{s} = \Gamma x$ for some $\Gamma \in \mathcal{R}^{k \times n}$. To formulate the resulting distortion, let Σ_x and $\Sigma_{Q_M^n(x)}$ be the covariance matrices of x and of $Q_M^n(x)$, respectively. The MSE in this case is stated in the following proposition.

Proposition 2: When $\tilde{s} = \Gamma x$, x is zero-mean, and $Q_M^n(x)$ is the optimal quantizer of x , then

$$\mathbb{E} \left\{ \left\| \tilde{s} - \mathbb{E} \{ \tilde{s} | Q_M^n(x) \} \right\|^2 \right\} = \text{Tr}(\Gamma^T \Gamma (\Sigma_x - \Sigma_{Q_M^n(x)})). \quad (7)$$

Proof: When \tilde{s} is the linear MMSE estimator, the second summand in (6) can be written as $\mathbb{E} \{ \left\| \Gamma(x - \mathbb{E}\{x | Q_M^n(x)\}) \right\|^2 \}$. Therefore,

$$\begin{aligned} & \mathbb{E} \left\{ \left\| \tilde{s} - \mathbb{E} \{ \tilde{s} | Q_M^n(x) \} \right\|^2 \right\} \stackrel{(a)}{=} \mathbb{E} \left\{ \left\| \Gamma(x - Q_M^n(x)) \right\|^2 \right\} \\ &= \text{Tr} \left(\Gamma^T \Gamma \mathbb{E} \left\{ (x - Q_M^n(x))(x - Q_M^n(x))^T \right\} \right) \\ &\stackrel{(b)}{=} \text{Tr} \left(\Gamma^T \Gamma (\Sigma_x - \Sigma_{Q_M^n(x)}) \right), \quad (8) \end{aligned}$$

where (a) follows since $Q_M^n(x)$ is the optimal quantizer of x in the MSE sense, hence $Q_M^n(x) = \mathbb{E}\{x | Q_M^n(x)\}$; and (b) is a result of the fact that the optimal quantizer is uncorrelated with the quantization error [1, Sec. III]. ■

Proposition 2 suggests that, when \tilde{s} is a linear function of x , the distortion can be evaluated using only the covariance matrix of the task-ignorant quantizer $Q_M^n(x)$. Nonetheless, the covariance of the quantizer which minimizes the distortion with respect to x is typically difficult to compute for finite M . Since $I(x; Q_M^n(x)) \leq \log M$ [22, Thm. 23.2], a possible approach to approximate the distortion is to evaluate Proposition 2 with the covariance matrix of the output distribution which obtains the distortion-rate function $D_x(\log M)$, instead of $\Sigma_{Q_M^n(x)}$. This replacement provides a reliable characterization of the performance of random codes distributed via the optimal marginal distortion-rate distribution for large M . In the numerical study in Section VI we illustrate that (7) approaches the performance of the optimal quantizer designed to recover x .

IV. HARDWARE-LIMITED TASK-BASED QUANTIZATION SYSTEMS DESIGN

A. Model Assumptions

We now study the design of hardware-limited task-based quantization systems illustrated in Fig. 3. As stated in the problem formulation, we consider the case where n , k and $\log M$ are fixed and finite, namely, we do not assume high quantization rate or arbitrarily large inputs. In such cases, explicitly characterizing the optimal quantization system and the minimal achievable distortion is a very difficult task, just as characterizing the minimal achievable distortion in lossy source coding with fixed block-lengths is difficult [25], [26]. Consequently, in the following section we focus on scenarios in which the stochastic relationship between the vector of interest s and the observation vector x are such that the MMSE estimate of s from x , $\tilde{s} = \mathbb{E}\{s|x\}$, is a linear function of x . Additionally, we restrict the digital mapping $h_d(\cdot)$ to be linear, namely, $h_d(u) = Bu$, $B \in \mathcal{R}^{k \times p}$. Since the MMSE estimate is linear here, this constraint is not expected to have a notable effect on the overall performance, especially when the error due to quantization is small.

By focusing on these setups, we are able to explicitly derive the achievable distortion and to characterize the system which achieves minimal distortion. This derivation reveals some non-trivial insights. For example, we show that the optimal approach when using vector quantizers, namely, to quantize the MMSE estimate, is no longer optimal in the presence of serial scalar ADCs. Furthermore, as detailed in Section V, this analysis provides general guidelines for designing hardware-limited task-based quantization systems, which can be used for any relationship between s and x .

To design a system which operates with simple scalar uniform quantizers, we carry out our analysis assuming dithered quantization [42]. Using dithered quantizers results in some favorable properties of the quantized signal, elaborated on in the sequel, which facilitate the analysis. These properties are also approximately satisfied without dithering for many input distributions [43]. Therefore, by considering dithered quantization, we are able to rigorously derive the optimal system, where in practice the resulting system can approach the optimal performance using standard uniform quantizers without dithering.

More specifically, we assume the identical scalar quantizers $Q_{M_p}^1(\cdot)$ implement non-subtractive uniform dithered quantization [42]. Unlike subtractive dithered quantization, considered in, e.g., [46], non-subtractive quantizers do not require the realization of the dithered signal to be subtracted from the quantizer




Fig. 5. Dithered uniform quantization illustration.

output in the digital domain, resulting in a practical structure [42]. An illustration is depicted in Fig. 5.

To formulate the input-output relationship of the serial ADC, let γ denote the dynamic range of the quantizer, and define $\Delta_p \triangleq \frac{2\gamma}{\tilde{M}_p}$ as the quantization spacing. The uniform quantizer is designed to operate within the dynamic range, namely, the amplitude of the input is not larger than γ with sufficiently large probability. To guarantee this, we fix γ to be some multiple η of the maximal standard deviation of the input. By Chebyshev's inequality [3, Pg. 64], for $\eta \geq 3$ the amplitude of the input is smaller than the dynamic range with probability over 89% for any input distribution. We assume that $\eta < \sqrt{3}\tilde{M}_p$, such that the variable $\kappa_p \triangleq \eta^2(1 - \frac{\eta^2}{3\tilde{M}_p^2})^{-1}$ is strictly positive. Note that $\eta = 3$ satisfies this requirement for any $\tilde{M}_p \geq 2$, i.e., the ADC is implemented using scalar quantizers with at least one bit. The output of the serial scalar ADC with input sequence y_1, y_2, \dots, y_p can be written as $Q_{\tilde{M}_p}^l(y_i) = q_p(y_i + z_i)$, where z_1, z_2, \dots, z_p are i.i.d. RVs uniformly distributed over $[-\frac{\Delta_p}{2}, \frac{\Delta_p}{2}]$, mutually independent of the input, representing the dither signal. The function $q_p(\cdot)$, which implements uniform quantization, is given by

$$q_p(y) = \begin{cases} -\gamma + \Delta_p(l + \frac{1}{2}) & y - l \cdot \Delta_p + \gamma \in [0, \Delta_p] \\ \text{sign}(y) \left(\gamma - \frac{\Delta_p}{2} \right) & |y| > \gamma. \end{cases} \quad l \in \{0, 1, \dots, \tilde{M}_p - 1\}$$

Note that when $\tilde{M}_p = 2$, the resulting quantizer is a standard one-bit sign quantizer of the form $q_p(y) = c \cdot \text{sign}(y)$, where the constant $c > 0$ is determined by the dynamic range γ .

Dithered quantizers significantly facilitate the analysis, due to the following favorable properties: When the input is inside the dynamic range of the quantizer, the output can be written as the sum of the input and an additive zero-mean white quantization noise signal, which is uncorrelated with the input. This model allows us to accurately characterize the quantization system which minimizes the MSE in Theorem 1. The favorable properties of dithered quantization are also satisfied in uniform quantization *without dithering* for inputs with bandlimited characteristic function, and are approximately satisfied for various families of input distributions, including the Gaussian distribution [43]. Furthermore, for those input distributions under which the favorable additive uncorrelated quantization noise model holds without dithering, the energy of the quantization noise is smaller without dithering than it is with dithering. Consequently, while in the following analysis we assume dithered quantization, exploiting the fact that the resulting quantization noise is white and uncorrelated with the input, the proposed system can also be applied without dithering. Under input distributions for which

the uncorrelated quantization noise model approximately holds, we expect lower distortion values to be achievable compared to the dithered case. This behavior is demonstrated in the simulations study in Section VI, where we show that applying the proposed system without dithering yields improved performance, due to the reduced energy of the quantization noise.

Since the dithered quantization operation can be modeled as adding uncorrelated noise, designing hardware-limited task-based quantization systems bears some similarity to linear transceiver design in MIMO communications. However, there are several fundamental differences between the two models. In hardware-limited quantization the additive quantization noise depends on the number of scalar quantizers due to the constraint on the overall number of bits, while in MIMO communications the channel noise does not depend on how many antennas are used. Furthermore, this additive noise is uncorrelated with the input but not independent, as commonly assumed in MIMO communications, and this model holds only when the quantized input is within the dynamic range. These differences between MIMO communications and hardware-limited task-based quantization result in a different system design.

B. Hardware-Limited Task-Based Quantizer Design

We now characterize the hardware-limited task-based quantizer which minimizes the MSE under the system model detailed in the previous subsection. Recall that here, the linear operators $h_a(\cdot)$ and $h_d(\cdot)$, denoting the analog combining and digital processing in Fig. 3, respectively, are represented using the matrices \mathbf{A} and \mathbf{B} , respectively. Our characterization yields the analog combining matrix and digital processing matrix, denoted \mathbf{A}° and \mathbf{B}° , respectively, and the corresponding dynamic range γ . Since for any quantized representation $\hat{\mathbf{s}}$, it follows from the orthogonality principle that the MSE, $\mathbb{E}\{\|\mathbf{s} - \hat{\mathbf{s}}\|^2\}$, equals the sum of the estimation error of the MMSE estimate, $\mathbb{E}\{\|\mathbf{s} - \tilde{\mathbf{s}}\|^2\}$, and the distortion with respect to the MMSE estimate, $\mathbb{E}\{\|\tilde{\mathbf{s}} - \hat{\mathbf{s}}\|^2\}$, in the following we characterize the performance of the proposed systems via the distortion with respect to $\tilde{\mathbf{s}}$.

Let $\mathbf{\Gamma}$ be the MSE optimal transformation of \mathbf{x} , namely, $\tilde{\mathbf{s}} = \mathbf{\Gamma}\mathbf{x}$, and let $\mathbf{\Sigma}_x$ be the covariance matrix of \mathbf{x} , assumed to be non-singular. Before we study the overall hardware-limited task-based quantization system, we first derive the digital processing matrix which minimizes the MSE for a given analog combining matrix \mathbf{A} and the resulting MSE, as stated in the following lemma:

Lemma 1: For any analog combining matrix \mathbf{A} and dynamic range γ such that $\Pr(|(\mathbf{A}\mathbf{x})_l + z_l| > \gamma) = 0$, namely, the quantizers operate within their dynamic range with probability one, the digital processing matrix which minimizes the MSE is given by

$$\mathbf{B}^\circ(\mathbf{A}) = \mathbf{\Gamma}\mathbf{\Sigma}_x\mathbf{A}^T \left(\mathbf{A}\mathbf{\Sigma}_x\mathbf{A}^T + \frac{2\gamma^2}{3\tilde{M}_p^2}\mathbf{I}_p \right)^{-1},$$

and the minimal achievable MSE is

$$\begin{aligned} \text{MSE}(\mathbf{A}) &= \min_{\mathbf{B}} \mathbb{E}\{\|\tilde{\mathbf{s}} - \hat{\mathbf{s}}\|^2\} \\ &= \text{Tr} \left(\mathbf{\Gamma}\mathbf{\Sigma}_x\mathbf{\Gamma}^T - \mathbf{\Gamma}\mathbf{\Sigma}_x\mathbf{A}^T \left(\mathbf{A}\mathbf{\Sigma}_x\mathbf{A}^T + \frac{2\gamma^2}{3\tilde{M}_p^2}\mathbf{I}_p \right)^{-1} \mathbf{A}\mathbf{\Sigma}_x\mathbf{\Gamma}^T \right). \end{aligned}$$

Proof: See Appendix B.

The digital processing matrix in Lemma 1 is the linear MMSE estimator of \mathbf{s} from the vector $\mathbf{A}\mathbf{x} + \mathbf{e}$, where \mathbf{e} represents the quantization noise, which is white and uncorrelated with $\mathbf{A}\mathbf{x}$. This stochastic representation is a result of the usage of dithered quantizers with the amplitude of the input to the quantizers being within the dynamic range. The resulting requirement, namely, to set the dynamic range γ such that $\Pr(|(\mathbf{A}\mathbf{x})_l + z_l| > \gamma) = 0$, is quite restrictive, as it holds only when the entries of \mathbf{x} have a finite support. This model can be satisfied along with the requirement for the MMSE estimator to be linear in some setups, for example, when $\mathbf{s} = \mathbf{\Gamma}\mathbf{x} + \mathbf{w}$ where \mathbf{w} is a zero-mean random vector independent of \mathbf{x} . However, in many scenarios of interest the entries of \mathbf{x} have an infinite support. In this case, there is always some probability that the input amplitude will exceed the dynamic range, resulting in some level of correlation between $\mathbf{A}\mathbf{x}$ and \mathbf{e} .

Nonetheless, in the following we use the model on which Lemma 1 is based, namely, that the output of the dithered quantizer can be written as its input corrupted by additive uncorrelated white noise, to design hardware-limited task-based quantizers also when the entries of \mathbf{x} have infinite support. Specifically, when the dynamic range is set such that the probability of overloading the quantizer is sufficiently small, namely, $\Pr(|(\mathbf{A}\mathbf{x})_l + z_l| > \gamma) \approx 0$ for each l , then modeling $\mathbf{A}\mathbf{x}$ and \mathbf{e} as uncorrelated becomes a reliable approximation. Therefore, in order to use Lemma 1 to design hardware-limited task-based quantizers, we explicitly require the uniform quantizers to operate within their dynamic range with high probability, as explained in Subsection IV-A, and set the value of γ accordingly. In Section VI we numerically demonstrate that this model leads to hardware-limited task-based quantizers which are capable of approaching the performance achievable using vector quantizers for inputs with infinite support.

We now use Lemma 1 to obtain the analog combining matrix \mathbf{A}° which minimizes MSE and the resulting system. Define the matrix $\tilde{\mathbf{\Gamma}} \triangleq \mathbf{\Gamma}\Sigma_{\mathbf{x}}^{1/2}$, and let $\{\lambda_{\tilde{\mathbf{\Gamma}},i}\}$ be its singular values arranged in a descending order. Note that for $i > \text{rank}(\tilde{\mathbf{\Gamma}})$, $\lambda_{\tilde{\mathbf{\Gamma}},i} = 0$.

The hardware-limited task-based quantization system based on the structure detailed in Subsection II-B which minimizes the MSE under the model assumptions of Subsection IV-A is given in the following theorem:

Theorem 1: For the hardware-limited quantization system based on the model detailed in Subsection IV-A, the analog combining matrix \mathbf{A}° is given by $\mathbf{A}^\circ = \mathbf{U}_A \mathbf{\Lambda}_A \mathbf{V}_A^T \Sigma_{\mathbf{x}}^{-1/2}$, where

- $\mathbf{V}_A \in \mathcal{R}^{n \times n}$ is the right singular vectors matrix of $\tilde{\mathbf{\Gamma}}$.
- $\mathbf{\Lambda}_A \in \mathcal{R}^{p \times n}$ is a diagonal matrix with diagonal entries

$$(\mathbf{\Lambda}_A)_{i,i}^2 = \frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p} \left(\zeta \cdot \lambda_{\tilde{\mathbf{\Gamma}},i} - 1 \right)^+, \quad (9a)$$

where ζ is set such that $\frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p} \sum_{i=1}^p \left(\zeta \cdot \lambda_{\tilde{\mathbf{\Gamma}},i} - 1 \right)^+ = 1$.

- $\mathbf{U}_A \in \mathcal{R}^{p \times p}$ is a unitary matrix which guarantees that $\mathbf{U}_A \mathbf{\Lambda}_A \mathbf{\Lambda}_A^T \mathbf{U}_A^T$ has identical diagonal entries, namely, $\mathbf{U}_A \mathbf{\Lambda}_A \mathbf{\Lambda}_A^T \mathbf{U}_A^T$ is weakly majorized by all possible rotations of $\mathbf{\Lambda}_A \mathbf{\Lambda}_A^T$ [47, Cor. 2.1]. The matrix \mathbf{U}_A can be obtained³ via [47, Alg. 2.2].

³The existence of the unitary matrix \mathbf{U}_A is guaranteed by [47, Cor. 2.1]. However, this matrix is not unique as, e.g., both \mathbf{U}_A and $-\mathbf{U}_A$ result in a rotation of $\mathbf{\Lambda}_A \mathbf{\Lambda}_A^T$ having identical diagonal entries.

The dynamic range of the ADC is given by

$$\gamma^2 = \frac{\kappa_p}{p} = \frac{\eta^2}{p} \left(1 - \frac{\eta^2}{3\tilde{M}_p^2} \right)^{-1}, \quad (9b)$$

and the digital processing matrix is equal to

$$\mathbf{B}^\circ(\mathbf{A}^\circ) = \tilde{\mathbf{\Gamma}} \mathbf{V}_A \mathbf{\Lambda}_A^T \left(\mathbf{\Lambda}_A \mathbf{\Lambda}_A^T + \frac{2\gamma^2}{3\tilde{M}_p^2} \mathbf{I}_p \right)^{-1} \mathbf{U}_A^T. \quad (9c)$$

The resulting minimal achievable distortion is

$$\mathbb{E} \left\{ \|\tilde{\mathbf{s}} - \hat{\mathbf{s}}\|^2 \right\} = \begin{cases} \sum_{i=1}^k \frac{\lambda_{\tilde{\mathbf{\Gamma}},i}^2}{(\zeta \cdot \lambda_{\tilde{\mathbf{\Gamma}},i-1})^{+1}}, & p \geq k \\ \sum_{i=1}^p \frac{\lambda_{\tilde{\mathbf{\Gamma}},i}^2}{(\zeta \cdot \lambda_{\tilde{\mathbf{\Gamma}},i-1})^{+1}} + \sum_{i=p+1}^k \lambda_{\tilde{\mathbf{\Gamma}},i}^2, & p < k. \end{cases} \quad (9d)$$

Proof: See Appendix C.

We note that, unlike task-based vector quantizers, for hardware-limited systems detailed in Subsection II-B, recovering the MMSE estimate $\tilde{\mathbf{s}}$ in the analog domain is sub-optimal. Since the quantization is carried out using a serial scalar ADC, the proposed analog combining rotates the input to the ADC such that each entry has *identical variance*, accounting for the fact that the same quantization rule is applied to each entry. Furthermore, the analog combiner includes a waterfilling-type expression over its singular values, which accounts for the finite dynamic range of the ADC. In particular, the waterfilling allows the resulting system to balance the estimation and quantization errors. To see this, we note from Appendix C that the matrix $\mathbf{\Lambda}_A$ determines the dynamic range γ . Consequently, by potentially nulling the diagonal entries corresponding to the less dominant singular values $\{\lambda_{\tilde{\mathbf{\Gamma}},i}\}$, the quantization system reduces the dynamic range. This yields more precise quantization and reduces the quantization error, at the cost of a small estimation error.

The quantization system in Theorem 1 minimizes the MSE under the model detailed in Subsection IV-A. This model is restricted to uniform quantization mappings, as these quantizers faithfully represent typical serial ADCs. The performance can be further improved by allowing for non-uniform quantizers, as shown in [44], which used data-driven machine learning methods to optimize the overall system. The model-based analysis of task-based quantization with non-uniform quantizers is left for future investigation.

Theorem 1 also provides guidelines to selecting the dimensions of the output of the analog combiner, as stated in the following corollary:

Corollary 1: In order to minimize the MSE, p must not be larger than the rank of the covariance matrix of $\tilde{\mathbf{s}}$.

Proof: See Appendix D.

Corollary 1 indicates that analog combining should project the observed vector such that the signal which undergoes the serial scalar quantization has reduced dimensionality, not larger than the rank of the covariance of $\tilde{\mathbf{s}}$. This follows since, by reducing the dimensionality of the inputs to the ADC while keeping the overall number of quantization levels M fixed, the quantization error induced by the scalar quantization is reduced. The exact optimal value of p is determined by the values of the non-zero singular values $\{\lambda_{\tilde{\mathbf{\Gamma}},i}\}$. In particular, the MSE expression in Theorem 1 implies that decreasing p below the number of non-zero singular values results in a tradeoff between improving quantization precision and increasing the estimation error. In the numerical analysis in Section VI we demonstrate

that using the proposed hardware-limited task-based system, the quantization error is made negligible for relatively small M , and the performance approaches that of the MMSE estimator.

Finally, we show that when the quantization resolution is sufficiently large, the proposed system produces the MMSE estimate \tilde{s} . To that aim, we assume that the covariance matrix of \tilde{s} is non-singular, thus we set $p = k$. When the quantization resolution is such that \tilde{M}_k is sufficiently large, the quantization noise introduced by the ADC becomes negligible, and the output of the system can be written as

$$\begin{aligned}\hat{s} &\approx B^\circ \mathbf{A}^\circ \mathbf{x} \\ &\approx \tilde{\Gamma} \mathbf{V}_A \Lambda_A^T (\Lambda_A \Lambda_A^T)^{-1} \Lambda_A \mathbf{V}_A^T \Sigma_x^{-1/2} \mathbf{x}.\end{aligned}\quad (10)$$

Furthermore, for large \tilde{M}_k , the parameter ζ becomes $\zeta \approx 3\tilde{M}_k^2 \cdot k / (2\kappa_k \sum_{i=1}^k \lambda_{\tilde{\Gamma},i})$. Thus, the diagonal entries in (9a) become $(\Lambda_A)_{i,i}^2 \approx \lambda_{\tilde{\Gamma},i} / (\sum_{i=1}^k \lambda_{\tilde{\Gamma},i})$. By writing the singular value decomposition (SVD) $\tilde{\Gamma} = \mathbf{U}_{\tilde{\Gamma}} \Lambda_{\tilde{\Gamma}} \mathbf{V}_{\tilde{\Gamma}}^T$ in (10), and recalling that $\mathbf{V}_{\tilde{\Gamma}} = \mathbf{V}_A$, we have

$$\begin{aligned}\hat{s} &\approx \mathbf{U}_{\tilde{\Gamma}} \Lambda_{\tilde{\Gamma}} \Lambda_A^T (\Lambda_A \Lambda_A^T)^{-1} \Lambda_A \mathbf{V}_{\tilde{\Gamma}}^T \Sigma_x^{-1/2} \mathbf{x} \\ &\stackrel{(a)}{=} \mathbf{U}_{\tilde{\Gamma}} \Lambda_{\tilde{\Gamma}} \mathbf{V}_{\tilde{\Gamma}}^T \Sigma_x^{-1/2} \mathbf{x} \stackrel{(b)}{=} \Gamma \mathbf{x} = \tilde{s},\end{aligned}$$

where (a) follows since for this setting of Λ_A , $\Lambda_{\tilde{\Gamma}} \Lambda_A^T (\Lambda_A \Lambda_A^T)^{-1} \Lambda_A = \Lambda_{\tilde{\Gamma}}$, and (b) follows since $\tilde{\Gamma} = \Gamma \Sigma_x^{1/2}$. Consequently, for sufficiently large quantization resolution, \hat{s} approaches the MMSE estimate \tilde{s} .

C. Suboptimal Quantization Systems

In the previous subsection we characterized the hardware-limited task-based quantization system which minimizes the MSE under the model assumptions of Subsection IV-A. In the following we study two suboptimal systems of interest: a system which does not carry out any processing in the analog domain, and a system which mimics the optimal vector task-based quantizer by quantizing the MMSE estimate. Our results in the following are based on the characterization of the achievable MSE for a fixed analog combining matrix \mathbf{A} in Lemma 1.

We begin with the suboptimal case where processing is carried out only in the digital domain. Here, $p = n$, and the analog combiner is given by $\mathbf{A} = \mathbf{I}_n$. This structure accommodates the majority of systems studied in the literature in the context of tasks performed with low precision ADCs, e.g., [13]–[16], [18]–[20]. The digital processing matrix which minimizes the MSE for this case and the resulting MSE are stated in the following corollary:

Corollary 2: When the analog combiner is $\mathbf{A} = \mathbf{I}_n$, the minimal achievable MSE is given by

$$\mathbb{E} \left\{ \|\tilde{s} - \hat{s}\|^2 \right\} = \text{Tr} \left(\tilde{\Gamma}^T \tilde{\Gamma} \left(\mathbf{I}_n + \frac{3\tilde{M}_n^2}{2\kappa_n \sigma_{\mathbf{x},\max}^2} \Sigma_{\mathbf{x}} \right)^{-1} \right), \quad (11a)$$

and the corresponding digital matrix is

$$B^\circ(\mathbf{I}_n) = \Gamma \Sigma_{\mathbf{x}} \left(\Sigma_{\mathbf{x}} + \frac{2\kappa_n \sigma_{\mathbf{x},\max}^2}{3\tilde{M}_n^2} \mathbf{I}_n \right)^{-1}, \quad (11b)$$

where $\sigma_{\mathbf{x},\max}^2 \triangleq \max_{i=1,\dots,n} ((\Sigma_{\mathbf{x}})_{i,i})$.

Proof: The corollary follows directly from Lemma 1. In particular, (11b) is obtained from the digital processing matrix in Lemma 1 by setting $\mathbf{A} = \mathbf{I}_n$, and (11a) is obtained from the resulting MSE via the matrix inversion lemma. ■

The resulting suboptimal system bears some similarity to the task-ignorant system discussed in Section III in the sense that quantization is carried out independently of the task. However, the system discussed in Section III performs joint vector quantization, while (11a) is achievable with a serial ADC. As a result, the system considered here can operate only when $\log M \geq n$, otherwise the scalar quantizers are assigned zero bits, while the task-ignorant system of Section III can operate with any positive value of $\log M$.

Next, we consider a system in which the analog combining is designed to recover the MMSE estimate \tilde{s} . Here, $p = k$, and $\mathbf{A} = \tilde{\Gamma}$. As noted in the discussion following Theorem 1, this approach is suboptimal when working with serial scalar ADCs, unlike the case with vector quantizers discussed in Section III. The digital processing matrix which minimizes the MSE for this setup and the resulting MSE are stated in the following corollary:

Corollary 3: When the analog combiner is $\mathbf{A} = \tilde{\Gamma}$, the minimal achievable MSE is given by

$$\mathbb{E} \left\{ \|\tilde{s} - \hat{s}\|^2 \right\} = \text{Tr} \left(\tilde{\Gamma}^T \tilde{\Gamma} \left(\mathbf{I}_n + \frac{3\tilde{M}_k^2}{2\kappa_k \sigma_{\tilde{s},\max}^2} \tilde{\Gamma}^T \tilde{\Gamma} \right)^{-1} \right), \quad (12a)$$

and the corresponding digital matrix is

$$B^\circ(\tilde{\Gamma}) = \tilde{\Gamma} \tilde{\Gamma}^T \left(\tilde{\Gamma} \tilde{\Gamma}^T + \frac{2\kappa_k \sigma_{\tilde{s},\max}^2}{3\tilde{M}_k^2} \mathbf{I}_k \right)^{-1}, \quad (12b)$$

where $\sigma_{\tilde{s},\max}^2 \triangleq \max_{i=1,\dots,k} (\mathbb{E}\{(\tilde{s})_i^2\})$.

Proof: The corollary follows directly from Lemma 1 using the same arguments as in the proof of Corollary 2. ■

The approach of quantizing the MMSE estimate is in general suboptimal. When the entries of \tilde{s} are not linearly dependent, namely, the covariance matrix of \tilde{s} is non-singular [49, Ch. 8.1], designing the analog combiner to recover the MMSE estimate minimizes the MSE if and only if the conditions stated in the following corollary is satisfied:

Corollary 4: When the covariance matrix of \tilde{s} is non-singular, quantizing the MMSE estimate minimizes the MSE if and only if the covariance matrix of \tilde{s} is $\frac{1}{k} \mathbf{I}_k$.

Proof: See Appendix E.

Corollary 4 indicates that, except for very specific statistical relationships between \mathbf{x} and \mathbf{s} , quantizing the entries of the MMSE estimate vector is purely suboptimal. In the numerical study presented in Section VI we numerically evaluate the achievable MSE of the considered systems, and illustrate that both the system proposed in Theorem 1 and the suboptimal system discussed in Corollary 4 are able to approach the performance of the optimal vector quantizer for large number of quantization levels M , and that the system of Theorem 1 outperforms the suboptimal system in Corollary 4 for all considered values of M . Additionally, we illustrate that for large k and relatively small M , a notable gap in MSE is observed between the hardware-limited task-based quantizer of Theorem 1 and the suboptimal system of Corollary 4.

V. GUIDELINES FOR HARDWARE-LIMITED TASK-BASED QUANTIZATION FOR ARBITRARY MODELS

A. Design Guidelines

In the previous section we characterized the distortion of hardware-limited task-based quantizers when the digital mapping $h_d(\cdot)$ is a linear function. To that aim, we designed the analog combining $h_a(\cdot)$ and the digital mapping $h_d(\cdot)$ such that, if the quantization error induced by the serial scalar ADC is negligible, then the resulting output approaches the MMSE estimate. In scenarios where the MMSE estimate \tilde{s} is linear, one can design linear $h_a(\cdot)$ and $h_d(\cdot)$ such that the resulting quantized representation approaches \tilde{s} as M increases for any value of $p \in [k, n]$, which denotes the dimensions of the output of the analog linear mapping. In particular, it was noted that when p decreases to the rank of the covariance matrix of \tilde{s} , the performance of the quantization system improves. This improvement follows as more bits can be assigned to the ADC, thus reducing the error induced by scalar quantization without modifying the overall number of bits used by the system, $\log M$.

The principles used for designing the linear analog and digital mappings $h_a(\cdot)$ $h_d(\cdot)$ for the relationship between s and x considered in Section IV also suggest guidelines for designing hardware-limited task-based quantization systems for arbitrary joint distributions of s and x with finite-variance entries. In particular, we propose to set $h_a(\cdot)$ and $h_d(\cdot)$ according to the following guidelines:

- 1) The mappings $h_a(\cdot)$, $h_d(\cdot)$ are such that when \tilde{M}_p is large enough, $h_d(h_a(x))$ approaches the MMSE estimate \tilde{s} .
- 2) The size of the output of the analog linear mapping, p , is as small as possible.

The first guideline implies that when the quantization error induced by the ADC is sufficiently small, the output of the system approaches the MMSE estimate, thus reducing the estimation error. The second guideline guarantees that more bits are assigned to the serial scalar ADC, thus decreasing the quantization error. Generally speaking, these guidelines provide the ability to balance quantization and estimation errors.

We note that in some scenarios, it may not be possible to obtain or approximate the MMSE estimate, \tilde{s} , from a linear function of x of reduced dimensions. For example, when s is estimated from the second-order statistical moments of x , as in eigen-spectrum estimation [5], subspace learning [18], DOA estimation [20], [21], and source localization [19]. In such cases, \tilde{s} generally cannot be obtained from a linear function of x of reduced dimensions. Nonetheless, the proposed guidelines can still be applied to design the quantization system. As an illustrative example, in the following subsection we explicitly show how these guidelines can be used for recovering the empirical covariance of an input signal.

B. Example: Recovery From Empirical Covariance

We next demonstrate how the guidelines for designing hardware-limited task-based quantization systems discussed in Subsection V-A can be applied for recovering the empirical covariance of the input. Unlike the results presented in Section IV, here we will not be able to explicitly characterize the resulting distortion. However, in the numerical study carried out in Section VI we empirically illustrate the benefits of the proposed design, and show that it outperforms processing the observations

only in the digital domain, which is the more popular approach in the literature.

In particular, consider the case where the observed vector x consists of n_x zero-mean i.i.d. $m_x \times 1$ vectors $\{x_i\}_{i=1}^{n_x}$, i.e., $x = [x_1^T, x_2^T, \dots, x_{n_x}^T]^T$ and $n = n_x \cdot m_x$. The desired vector s (or its MMSE estimate \tilde{s}) can be recovered from the empirical covariance of $\{x_i\}_{i=1}^{n_x}$, namely, from

$$\mathbf{R}_x \triangleq \frac{1}{n_x} \sum_{i=1}^{n_x} x_i x_i^T.$$

For example, when s is the eigenspectrum of x , the MMSE estimate \tilde{s} is obtained from \mathbf{R}_x via [5, Eq. (22)]. At first glance, the proposed guidelines cannot be used here, as, in general, for $p < n$ there exists no linear transformation $h_a: \mathcal{R}^n \mapsto \mathcal{R}^p$ such that \mathbf{R}_x can be recovered from $h_a(x)$. However, an approximation of \mathbf{R}_x can be obtained via the following steps:

- Divide the set $\{x_i\}_{i=1}^{n_x}$ into n_s distinct sets, each consisting of $m_s = \frac{n_x}{n_s}$ vectors, namely, the l th set is given by $\{x_i\}_{i=(l-1)m_s+1}^{l m_s}$, $l \in \{1, 2, \dots, n_s\}$.
- Fix the analog combining such that the input to the serial scalar ADC consists of n_s vectors $\{z_l\}_{l=1}^{n_s}$, where $z_l = \sum_{i=(l-1)m_s+1}^{l m_s} x_i$. This is achieved by setting $h_a(x) = \mathbf{A}x$, where the entries of $\mathbf{A} \in \mathcal{R}^{n_s \cdot m_x \times n_x \cdot m_x}$ are given by

$$\begin{aligned} (\mathbf{A})_{(p_1-1)m_x+q_1, (p_2-1)m_x+q_2} \\ = \delta_{q_1, q_2} \sum_{l=1}^{m_s} \delta_{(p_1-1)m_s+l, p_2}, \end{aligned}$$

for $p_1 \in \{1, 2, \dots, n_s\}$, $p_2 \in \{1, 2, \dots, n_x\}$, $q_1, q_2 \in \{1, 2, \dots, m_s\}$.

- In the digital domain, we approximate \mathbf{R}_x from the quantized vectors $\tilde{z}_l \triangleq [Q_{M_p}^1((z_l)_1), \dots, Q_{M_p}^1((z_l)_{m_s})]^T$ via

$$\hat{\mathbf{R}}_x \triangleq \frac{1}{n_x} \sum_{l=1}^{n_s} \tilde{z}_l \tilde{z}_l^T.$$

The rationale behind these steps is that, as discussed in the previous subsection, it allows to balance quantization and estimation errors. To see this, we note that when the quantization error is negligible, and the number of sets n_s is sufficiently large, $\hat{\mathbf{R}}_x$ approaches the true covariance matrix by the law of large numbers. Furthermore, by decreasing the number of sets n_s , fewer scalar quantizers are needed, thus the quantization error induced by the serial scalar ADCs is reduced. As n_s increases to n_x , $\hat{\mathbf{R}}_x$ approaches the true covariance of \tilde{z}_l (up to a constant factor). The proposed guidelines thus provide the ability to trade quantization and estimation errors, which is expected to be most beneficial for small values of M , i.e., low quantization rates, where the quantization error becomes dominant. We expect to have an optimal value of n_s in the range $[1, n_x]$ for each value of M . This behavior, as well as the benefits of the proposed approach for finite values of M , are illustrated in the empirical study in Subsection VI-B. In particular, in Subsection VI-B it is illustrated that for the problem of eigenspectrum recovery, a quantization system designed according to the above guidelines outperforms a system which performs no analog combining prior to quantization, and that the performance gap depends on the overall quantization levels M and on the number of sets n_s .

VI. APPLICATIONS AND NUMERICAL STUDY

In this section we study the application of the hardware-limited task-based quantization systems proposed in Sections IV–V, in two scenarios involving parameter acquisition from quantized measurements: First, in Subsection VI-A, we study the achievable MSE in estimating a scalar channel with finite intersymbol interference (ISI) from a fixed number of quantized measurements, as in, e.g., [13]–[15], using the hardware-limited task-based quantizer proposed in Section IV. Then, in Subsection VI-B, we consider the problem of estimating the eigen-spectrum from a set of i.i.d. measurements, see, e.g., [5], and evaluate the achievable distortion of the quantization system design detailed in Section V.

A. ISI Channel Estimation

We first consider the estimation of a scalar ISI channel from quantized observations, as in [13]–[15]. In this scenario, the parameter vector \mathbf{s} represents the coefficients of a multipath channel with k taps. The channel is estimated from a set of $n = 120$ noisy observations \mathbf{x} , given by [13, Eq. (1)]

$$(\mathbf{x})_i = \sum_{l=1}^k (\mathbf{s})_l a_{i-l+1} + w_i, \quad i \in \{1, 2, \dots, n\}, \quad (13)$$

where a_i is a deterministic known training sequence, and $\{w_i\}_{i=1}^n$ are samples from an i.i.d. zero-mean unit variance Gaussian noise process independent of \mathbf{s} . In particular, the channel \mathbf{s} is modeled as a zero-mean Gaussian vector with covariance matrix $\Sigma_{\mathbf{s}}$, given by $(\Sigma_{\mathbf{s}})_{i,j} = e^{-|i-j|}$, $i, j \in \{1, 2, \dots, k\} \triangleq \mathcal{K}$, and the training sequence is given by $a_i = \cos(\frac{2\pi i}{n})$ for $i > 0$ and $a_i = 0$ otherwise. Note that \mathbf{s} and \mathbf{x} are jointly Gaussian, and thus the MMSE estimator $\tilde{\mathbf{s}}$ is a linear function of \mathbf{x} .

In the following we evaluate the achievable distortion of the resulting hardware-limited task-based quantization for this setup, and compare the achievable distortion to that of the optimal task-based quantizer and of the task-ignorant quantizer discussed in Section III. To that aim, we consider two channels: one with $k = 2$ taps and one with $k = 8$ taps, and let the overall number of quantization bits be $\log M \in [2 \cdot k, 10 \cdot k]$. As $\log M$ is strictly smaller than n , any quantization system which is based on applying serial scalar quantization to the observation \mathbf{x} without any processing in the analog domain, such as the system discussed in Corollary 2, as well as the quantization systems considered in [13], [14], cannot be implemented here.

In the numerical study we evaluate the following quantities:

- The MMSE $\mathbb{E}\{\|\mathbf{s} - \tilde{\mathbf{s}}\|^2\}$, which is the optimal distortion of a system with no quantization.
- For hardware-limited task-based quantization, we compute the achievable distortion of the system derived in Theorem 1. Since the covariance matrix of $\tilde{\mathbf{s}}$ is non-singular for the considered setup, we set $p = k$ following Corollary 1. Additionally, we compute the MSE of a system which recovers the MMSE estimate $\tilde{\mathbf{s}}$ in the analog domain, based on Corollary 3. Furthermore, since dithering increases the energy of the quantization noise, we also compute the achievable MSE of the proposed systems when the ADCs implement uniform quantization without dithering. The MSE of all these systems is computed by empirical averaging over 10000 Monte Carlo




Fig. 6. Distortion comparison, channel estimation, $k = 2$.

simulations. In order to avoid cluttering, we do not depict the theoretical performance of the systems in Theorem 1 and Corollary 3 with dithering, computed via (9d) and (12a), respectively. However, we note that the empirical performance depicted here coincides with the theoretical MSE.

- For the optimal task-based system, we evaluate the bounds in Proposition 1, where the lower bound is computed using the reverse waterfilling algorithm for multivariate Gaussian sources [3, Ch. 10.3], and the upper bound is computed via (5), by setting $f_{\bar{c}}$ to be the optimal marginal distortion-rate distribution, which is a zero-mean Gaussian distribution whose covariance is obtained as in [3, Ch. 10.3].
- For the task-ignorant system, we numerically evaluate the distortion in (6) by letting Q_M^n be the quantizer in which the codewords are generated i.i.d. from the optimal marginal distortion-rate probability measure with respect to \mathbf{x} , averaging the performance over 20000 Monte Carlo simulations. As this computation becomes prohibitive for large values of M , we evaluate (6) only for $\log M \leq 16$. We also compute the approximate achievable distortion of Proposition 2, where the quantized output distribution is set to the optimal marginal distortion-rate distribution for quantizing \mathbf{x} , as proposed in the discussion following Proposition 2. This quantity is computed for all considered values of M , and is shown to provide a good approximation of (6) for large values of M .

Figs. 6–7 depict the distortions for $k = 2$ and for $k = 8$, respectively. Observing Figs. 6–7, we note that hardware-limited task-based quantizers substantially outperform task-ignorant vector quantization, and approach the optimal performance as M increases. In particular, when each scalar quantizer uses at least five bits, i.e., $\log M \geq 5k$, the quantization error becomes negligible and the overall distortion is effectively the minimum achievable estimation error, i.e., the MMSE.

Furthermore, we note that the gap between the system of Theorem 1 and the system which quantizes the MMSE estimate of Corollary 3 is small for $k = 2$, and becomes notable for $k = 8$ at small values of M . However, it is emphasized that, as expected, the system of Theorem 1 outperforms the approach of quantizing the MMSE estimate, which is known to be optimal when using vector quantizers, for all considered values of M . It is also noted that the proposed hardware-limited task-based




Fig. 7. Distortion comparison, channel estimation, $k = 8$.

quantizers, designed assuming dithered uniform quantizers, obtain improved performance without dithering. This follows since the favorable properties of dithered quantization discussed in Subsection IV-A, which are accounted for in the design of the systems in Section IV, are approximately satisfied also for non-dithered standard quantization, as noted in [43], without the excess distortion induced by dithering. This illustrates that our proposed design can be applied also without dithering, and that the resulting performance is improved compared to systems implementing dithered quantization.

Regarding the achievable distortion of the non hardware-limited system discussed in Section III, we note that the gap between the lower and the upper bounds on the optimal achievable distortion in Proposition 1, which is represented by the shaded region in the figure, narrows down from an MSE gap of 0.16 at $\log_2 M = 4$ to 0.0158 at $\log_2 M \geq 10$ for $k = 2$, and from 0.095 at $\log_2 M = 16$ to 0.06 at $\log_2 M \geq 36$ for $k = 8$. Since the systems proposed in Section IV can be obtained as a special case of the task-based vector quantizer depicted in Fig. 2, combining Theorem 1 and Proposition 1 provides relatively tight bounds on the optimal distortion for finite quantization resolution. We also note that the empirical distortion of the task-ignorant system is higher than the upper bound on the optimal distortion for all values of M , and from Fig. 6 we note that the approximated distortion of the task-ignorant system, computed via Proposition 2, provides an excellent match with the empirical distortion for $\log_2 M > 10$, without the need to empirically average over multiple codes.

B. Eigen-Spectrum Estimation

We next consider the problem of estimating the eigen-spectrum of a multivariate Gaussian source from quantized measurements, as in [5]. Here, the desired vector \mathbf{s} represents the eigenvalues of a $k \times k$ covariance matrix, $\Sigma_{\mathbf{x}} \triangleq \mathbf{U}\mathbf{\Lambda}\mathbf{U}^H$, where \mathbf{U} is a deterministic known unitary matrix, and $\mathbf{\Lambda}$ is a diagonal matrix with diagonal entries $(\mathbf{\Lambda})_{i,i} = (\mathbf{s})_i$, $i \in \mathcal{K}$. Following [5], the entries of \mathbf{s} are mutually independent, and each entry $(\mathbf{s})_i$ obeys an inverse gamma distribution with shape parameter $\alpha_i > 2$ and scale parameter $\beta_i > 0$, $i \in \mathcal{K}$. The vector \mathbf{x} consists of n_x random vectors $\{\mathbf{x}_i\}_{i=1}^{n_x}$, which, given \mathbf{s} , are i.i.d. zero-mean Gaussian with covariance $\Sigma_{\mathbf{x}}$.

Note that for the considered scenario, \mathbf{x} and \mathbf{s} are uncorrelated, as, by the law of total expectation, $\mathbb{E}\{\mathbf{x}\mathbf{s}^T\} = \mathbb{E}\{\mathbb{E}\{\mathbf{x}|\mathbf{s}\}\mathbf{s}^T\} = \mathbf{0}$. Hence, the linear MMSE estimator of \mathbf{s} from \mathbf{x} is the expected value $\mathbb{E}\{\mathbf{s}\}$, and it thus makes little sense to design the quantization system to approach the linear MMSE estimator, which is main design principle of the systems proposed in Section IV. On the other hand, the (non-linear) MMSE estimator for this scenario is given by [5, Thm. 4]

$$\tilde{s}_i = \frac{1}{\alpha_i - \frac{1}{2}n_x - 1} \left(\beta_i + \frac{1}{2} \left(\mathbf{U}^H \left(\sum_{l=1}^{n_x} \mathbf{x}_l \mathbf{x}_l^T \right) \mathbf{U} \right)_{i,i} \right),$$

$i \in \mathcal{K}$. Consequently, the eigen-spectrum \mathbf{s} can be estimated from the empirical covariance of $\{\mathbf{x}_i\}_{i=1}^{n_x}$, and we thus apply the quantizer design proposed in Section V. To implement the hardware-limited task-based quantizer, we divide the n_x realizations $\{\mathbf{x}_i\}_{i=1}^{n_x}$ into n_s distinct sets of size $m_s = \frac{n_x}{n_s}$. The quantization system of Section V is then used to produce a quantized estimation of $\frac{1}{n_x} \sum_{i=1}^{n_x} \mathbf{x}_i \mathbf{x}_i^T$, denoted $\hat{\mathbf{R}}_{\mathbf{x}}$, using $p = k \cdot n_x$ identical scalar quantizers. The eigen-spectrum is then estimated via

$$\hat{\tilde{s}}_i = \frac{1}{\alpha_i - \frac{1}{2}n_x - 1} \left(\beta_i + \frac{1}{2} \left(\mathbf{U}^H \hat{\mathbf{R}}_{\mathbf{x}} \mathbf{U} \right)_{i,i} \right), \quad (14)$$

$i \in \mathcal{K}$. Note that when $n_s = n_x$, no analog combining is performed, and the quantization system results the standard approach of estimating from uniformly quantized measurements.

In the numerical study we consider two setups:

- In the first setup, we fix $k = 2$, and set the inverse gamma distribution parameters to $\{\alpha_i\}_{i=1}^k = \{5.5, 6.5\}$ and $\{\beta_i\}_{i=1}^k = \{8.4, 11.6\}$, thus, each entry of \mathbf{s} has approximately unit variance [45]. The observed vector \mathbf{x} consists of $n_x = 20$ realizations, thus $n = n_x \cdot k = 40$, and the unitary matrix \mathbf{U} is set to the $k \times k$ discrete Fourier transform (DFT) matrix.
- In the second setup, we fix $k = 4$, and set the inverse gamma distribution parameters to $\{\alpha_i\}_{i=1}^k = \{4, 5, 6, 7\}$ and $\{\beta_i\}_{i=1}^k = \{4.2, 6.9, 10, 13.4\}$, which again results each entry of \mathbf{s} having approximately unit variance. The observed vector \mathbf{x} consists of $n_x = 60$ realizations, thus $n = 240$, and \mathbf{U} is set to the identity matrix.

For both setups we numerically evaluate the achievable distortion of the quantization system proposed in Section V, where the scalar quantization is carried out using standard (non-dithered) uniform mapping. These distortions are compared to the achievable distortion of the optimal task-based system in Proposition 1, which is computed with $f_{\hat{\mathbf{e}}}$ set to the probability measure of \mathbf{s} . The actual minimal achievable distortion is thus upper bounded by this achievable rate, and lower bounded by the MMSE (dashed black curve in Figs. 8–9). The distortions are also compared to the distortion of the optimal linear estimator, which, in this case, is the mean value $\mathbb{E}\{\mathbf{s}\}$ (dashed blue curve in Figs. 8–9), and is thus independent of the observed signal and the quantization system. The distortions are computed for quantization rate $R \in [0.1, 3]$, resulting in $\log M \in [4, 120]$ for the first setup and $\log M \in [24, 720]$ for the second setup. Note that unlike the numerical study presented in Subsection VI-A, here we do not evaluate the performance of the task-ignorant vector quantizer, since the MMSE estimator is not linear, thus the approximation in Proposition 2 does not hold, and explicitly




Fig. 8. Distortion comparison, eigen-spectrum estimation, $k = 2$.




Fig. 9. Distortion comparison, eigen-spectrum estimation, $k = 4$.

computing the distortion in (6) by Monte-Carlo simulations is not feasible for large values of M .

In Fig. 8 we depict the achievable distortion of the proposed hardware-limited task-based quantization systems for the first scenario, where we used $n_s \in \{2, 4, 5, 10, 20\}$, while Fig. 9 depicts the corresponding distortions for the second scenario with $n_s \in \{12, 15, 20, 30, 60\}$. Observing Figs. 8–9, we first note that the performance of the optimal task-based quantizer can still be approached within a small gap by the proposed hardware-limited task-based quantization by properly selecting the number of sets n_s . In this context, we note that the best selection of n_s depends on the overall number of bits $\log M$. When $\log M$ is small, the distortion is dominated by the quantization error induced by the scalar quantizers, and thus using less partitions, which allows to assign more bits to the scalar quantizers, is beneficial. However, as $\log M$ increases, the error in estimating the empirical covariance induced by averaging over each set becomes dominant, and thus using more sets of smaller size achieves better performance. In particular, for the first setup, it is observed in Fig. 8 that using $n_s = 10$ results in the smallest distortion for most considered values of $\log M$, and in fact, achieves an MSE which is less than 0.6 from the MMSE for $\log M = 40$, i.e., for quantization rate as small as $R = \frac{1}{3}$. For the second setup, it is observed in Fig. 8 that using $n_s = 20$ is the best selection for $\log M \in [160, 480]$, i.e., $R \in [\frac{2}{3}, 2]$,

while $n_s = 30$ achieves the best performance for $\log M \geq 480$, namely $R \geq 2$.

The approach of estimating from uniformly quantized measurements without any analog combining, namely, $n_s = n_x$, achieves poor performance for $\log M \leq 60$ when $k = 2$, and for all considered values of $\log M$ when $k = 4$. In fact, for the scenario in Fig. 9, the approach of estimating from uniformly quantized measurements is outperformed by the data ignorant mean estimate for most considered values of $\log M$. The results presented in this section thus demonstrate the gain of the proposed hardware-limited task-based quantization system design for systems operating with finite and relatively small number of bits $\log M$, in scenarios where the MMSE estimate is not a linear function of the observations.

VII. CONCLUSION

In this work we studied hardware-limited task-based quantization systems, operating with practical serial scalar ADCs, for finite-size signals with finite-resolution quantization. We characterized the hardware-limited task-based quantizer which minimizes the MSE when the MMSE estimate of the desired signal is a linear function of the observed signal, and demonstrated that the analysis leads to design principles applicable to a much wider range of settings. We showed that, unlike when vector quantizers are used, quantizing the MMSE estimate is generally not optimal. Finally, we applied our results to the relevant problems of channel estimation in ISI channels and eigen-spectrum estimation. For these scenarios, we showed that the performance of the optimal task-based vector quantizer can be approached with a practical system operating with a serial scalar ADC. Furthermore, we demonstrated that by properly accounting for the task in the design of the quantizer, hardware-limited systems can substantially outperform task-ignorant vector quantizers. The proposed hardware-limited task-based framework and the gains observed in this work for the task of signal recovery give rise to a multitude of future research directions by considering different tasks and hardware-limitations, including communications with low-resolution ADCs and hardware-limited task-based quantization for classification.

APPENDIX

A. Proof of Proposition 1

The lower bound in (5) follows from Shannon's converse [48], which states that for any quantizer $Q_M^k, I(\tilde{\mathbf{s}}; Q_M^k(\tilde{\mathbf{s}})) \leq \log M$, and thus

$$\begin{aligned} & \min_{Q_M^k(\cdot)} \mathbb{E} \left\{ \|\tilde{\mathbf{s}} - Q_M^k(\tilde{\mathbf{s}})\|^2 \right\} \\ & \geq \min_{Q_M^k(\cdot): I(\mathbf{s}; Q_M^k(\mathbf{s})) \leq \log M} \mathbb{E} \left\{ \|\tilde{\mathbf{s}} - Q_M^k(\tilde{\mathbf{s}})\|^2 \right\} \\ & \stackrel{(a)}{\geq} D_{\tilde{\mathbf{s}}}(\log M), \end{aligned} \quad (\text{A.1})$$

where (a) follows from Def. 2. The upper bound in (5) is obtained by considering the quantizer in which the codewords $\{\mathbf{q}_i\}_{i=1}^M$ are generated i.i.d. from the probability measure $f_{\tilde{\mathbf{e}}}$, and the quantizer output is set to

$$Q_M^k(\mathbf{y}) = \arg \min_{\{\mathbf{q}_i\}} \|\mathbf{y} - \mathbf{q}_i\|^2, \forall \mathbf{y} \in \mathcal{R}^k.$$

The MSE of this quantizer is given by

$$\begin{aligned}
\mathbb{E} \left\{ \left\| Q_M^k(\tilde{\mathbf{s}}) - \tilde{\mathbf{s}} \right\|^2 \right\} &= \mathbb{E} \left\{ \min_i \left(\|\mathbf{q}_i - \tilde{\mathbf{s}}\|^2 \right) \right\} \\
&= \mathbb{E} \left\{ \mathbb{E} \left\{ \min_i \left(\|\mathbf{q}_i - \tilde{\mathbf{s}}\|^2 \right) \mid \tilde{\mathbf{s}} \right\} \right\} \\
&\stackrel{(a)}{=} \mathbb{E} \left\{ \int_0^\infty \Pr \left\{ \min_i \left(\|\mathbf{q}_i - \tilde{\mathbf{s}}\|^2 > t \right) \mid \tilde{\mathbf{s}} \right\} dt \right\} \\
&\stackrel{(b)}{=} \mathbb{E} \left\{ \int_0^\infty \left[\Pr \left(\|\tilde{\mathbf{c}} - \tilde{\mathbf{s}}\|^2 > t \mid \tilde{\mathbf{s}} \right) \right]^M dt \right\}, \quad (\text{A.2})
\end{aligned}$$

where (a) follows from [49, Ch. 5.3] as $\|\mathbf{q}_i - \tilde{\mathbf{s}}\|^2$ is non-negative for any realization of $\tilde{\mathbf{s}}$, and (b) follows from the code construction. Eqns. (A.1)–(A.2) prove (5). ■

B. Proof of Lemma 1

To prove the lemma, we first show that the output of the ADC can be written as the sum of its input and uncorrelated noise, and then we derive the optimal digital processing matrix.

Note that when $\Pr(|(\mathbf{A}\mathbf{x})_l + z_l| > \gamma) = 0$, it follows from [42, Thm. 2] that, since the dither signal z_l is uniformly distributed over $(-\Delta_p/2, \Delta_p/2]$, then the output of the scalar ADC can be written as $\mathbf{A}\mathbf{x} + e$, where the quantization noise e has uncorrelated zero-mean entries with variance $\frac{\Delta_p^2}{6}$. Furthermore, by [42, Thm. 2] e satisfies $\mathbb{E}\{e|\mathbf{A}\mathbf{x}\} = \mathbb{E}\{e\} = \mathbf{0}$, and thus

$$\begin{aligned}
\mathbb{E}\{\mathbf{x}e^T\} &= \mathbb{E}\{\mathbf{x}\mathbb{E}\{e^T|\mathbf{x}\}\} \\
&\stackrel{(a)}{=} \mathbb{E}\{\mathbf{x}\mathbb{E}\{e^T|\mathbf{A}\mathbf{x}\}\} = \mathbf{0},
\end{aligned}$$

i.e., e is uncorrelated with \mathbf{x} . Here, (a) follows since the quantization noise e depends on the observed vector \mathbf{x} only via the input to the serial scalar quantizer $\mathbf{A}\mathbf{x}$, and thus the conditional distribution of e given some realization $\mathbf{x} = \mathbf{x}'$ is equal to the conditional distribution of e given $\mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x}'$. Thus, the optimal digital processing matrix, which results in $\hat{\mathbf{s}}$ being the linear MMSE estimator of $\tilde{\mathbf{s}}$ from $\mathbf{A}\mathbf{x} + e$, is

$$\begin{aligned}
\mathbf{B}^\circ(\mathbf{A}) &= \mathbb{E}\{\tilde{\mathbf{s}}(\mathbf{A}\mathbf{x} + e)^T\} \left(\mathbb{E}\{(\mathbf{A}\mathbf{x} + e)(\mathbf{A}\mathbf{x} + e)^T\} \right)^{-1} \\
&= \mathbb{E}\{\Gamma\mathbf{x}(\mathbf{A}\mathbf{x})^T\} \left(\mathbb{E}\{(\mathbf{A}\mathbf{x})(\mathbf{A}\mathbf{x})^T\} + \mathbb{E}\{ee^T\} \right)^{-1} \\
&\stackrel{(a)}{=} \Gamma\Sigma_{\mathbf{x}}\mathbf{A}^T \left(\mathbf{A}\Sigma_{\mathbf{x}}\mathbf{A}^T + \frac{\Delta_p^2}{6}\mathbf{I}_p \right)^{-1}, \quad (\text{B.1})
\end{aligned}$$

where (a) is a result of [42, Thm. 2]. The MSE of the linear MMSE estimate is thus given by

$$\begin{aligned}
\text{MSE}(\mathbf{A}) &= \text{Tr}(\Gamma\Sigma_{\mathbf{x}}\Gamma^T) \\
&\quad - \text{Tr} \left(\Gamma\Sigma_{\mathbf{x}}\mathbf{A}^T \left(\mathbf{A}\Sigma_{\mathbf{x}}\mathbf{A}^T + \frac{\Delta_p^2}{6}\mathbf{I}_p \right)^{-1} \mathbf{A}\Sigma_{\mathbf{x}}\Gamma^T \right). \quad (\text{B.2})
\end{aligned}$$

Plugging the quantization spacing $\Delta_p = \frac{2\gamma}{M_p}$ into (B.1)–(B.2) proves the lemma. ■

C. Proof of Theorem 1

To prove the theorem, we first derive the optimal unitary rotation for a given \mathbf{A} . Then, we characterize the optimal analog combining and the resulting MSE. We use the fact that for a fixed matrix \mathbf{A} , the resulting MSE is given in Lemma 1.

Recall that the dynamic threshold is set to a multiple η of the maximal standard deviation of the quantizer input. Therefore,

$$\begin{aligned}
\gamma^2 &= \eta^2 \max_{l=1,\dots,p} \mathbb{E} \left\{ ((\mathbf{A}\mathbf{x})_l + z_l)^2 \right\} \\
&\stackrel{(a)}{=} \eta^2 \max_{l=1,\dots,p} \mathbb{E} \left\{ (\mathbf{A}\mathbf{x})_l^2 \right\} + \eta^2 \frac{\gamma^2}{3M_p^2}, \quad (\text{C.1})
\end{aligned}$$

where (a) follows since z_l is independent of \mathbf{x} and its variance equals $\frac{\Delta_p^2}{12} = \frac{\gamma^2}{3M_p^2}$. From (C.1),

$$\begin{aligned}
\gamma^2 &= \eta^2 \left(1 - \frac{\eta^2}{3M_p^2} \right)^{-1} \max_{l=1,\dots,p} \mathbb{E} \left\{ (\mathbf{A}\mathbf{x})_l^2 \right\} \\
&= \kappa_p \max_{l=1,\dots,p} \mathbb{E} \left\{ (\mathbf{A}\mathbf{x})_l^2 \right\}. \quad (\text{C.2})
\end{aligned}$$

Substituting (C.2) to the expression for $\text{MSE}(\mathbf{A})$ in Lemma 1 yields

$$\begin{aligned}
\text{MSE}(\mathbf{A}) &= \text{Tr} \left(\Gamma\Sigma_{\mathbf{x}}\Gamma^T - \Gamma\Sigma_{\mathbf{x}}\mathbf{A}^T \left(\mathbf{A}\Sigma_{\mathbf{x}}\mathbf{A}^T + \frac{2\kappa_p}{3M_p^2} \max_{l=1,\dots,p} \mathbb{E} \left\{ (\mathbf{A}\mathbf{x})_l^2 \right\} \mathbf{I}_p \right)^{-1} \mathbf{A}\Sigma_{\mathbf{x}}\Gamma^T \right). \quad (\text{C.3})
\end{aligned}$$

Using (C.3), we can find for each analog combining matrix \mathbf{A} an optimal unitary rotation, which minimizes the MSE, as stated in the following lemma:

Lemma C.1: For every matrix $\mathbf{A} \in \mathcal{R}^{p \times n}$ there exists a unitary matrix $\mathbf{U}_A \in \mathcal{R}^{p \times p}$ such that

$$\begin{aligned}
\text{MSE}(\mathbf{A}) &\geq \text{MSE}(\mathbf{U}_A\mathbf{A}) \\
&= \text{Tr} \left(\Gamma\Sigma_{\mathbf{x}}\Gamma^T - \Gamma\Sigma_{\mathbf{x}}\mathbf{A}^T \left(\mathbf{A}\Sigma_{\mathbf{x}}\mathbf{A}^T + \frac{2\kappa_p}{3M_p^2 \cdot p} \text{Tr}(\mathbf{A}\Sigma_{\mathbf{x}}\mathbf{A}^T) \mathbf{I}_p \right)^{-1} \mathbf{A}\Sigma_{\mathbf{x}}\Gamma^T \right). \quad (\text{C.4})
\end{aligned}$$

Proof: Note that for any unitary matrix \mathbf{U}_A , it follows from (C.3) that

$$\begin{aligned}
\text{MSE}(\mathbf{U}_A\mathbf{A}) &= \text{Tr}(\Gamma\Sigma_{\mathbf{x}}\Gamma^T) \\
&\quad - \text{Tr} \left(\Gamma\Sigma_{\mathbf{x}}\mathbf{A}^T \left(\mathbf{A}\Sigma_{\mathbf{x}}\mathbf{A}^T + \frac{2\kappa_p}{3M_p^2} \max_{l=1,\dots,p} \mathbb{E} \left\{ (\mathbf{U}_A\mathbf{A}\mathbf{x})_l^2 \right\} \mathbf{I}_p \right)^{-1} \mathbf{A}\Sigma_{\mathbf{x}}\Gamma^T \right). \quad (\text{C.5})
\end{aligned}$$

For each pair of positive semi-definite symmetric matrices $\mathbf{M}_1, \mathbf{M}_2$, the scalar function $h(\alpha) = \text{Tr}(\mathbf{M}_1(\mathbf{M}_2 + \alpha\mathbf{I})^{-1})$ is

monotonically decreasing for $\alpha > 0$. Thus, by (C.5), the unitary U_A which minimizes the MSE is given by

$$\begin{aligned} U_A &= \arg \min_U \max_{l=1, \dots, p} \mathbb{E} \left\{ (U A \mathbf{x})_l^2 \right\} \\ &= \arg \min_U \max_{l=1, \dots, p} (U A \Sigma_{\mathbf{x}} A^T U^T)_{l,l}. \end{aligned} \quad (\text{C.6})$$

From majorization theory [47, Cor. 2.4] it follows that $\min_U \max_{l=1, \dots, p} (U A \Sigma_{\mathbf{x}} A^T U^T)_{l,l} = \frac{1}{p} \text{Tr}(A \Sigma_{\mathbf{x}} A^T)$. Furthermore, a unitary matrix which solves (C.6) can be obtained using the iterative algorithm in [47, Alg. 2.2]. Plugging this into (C.5) proves the lemma. ■

Finally, we characterize the matrix A which minimizes (C.4). To that aim, define $\tilde{A} \triangleq A \Sigma_{\mathbf{x}}^{1/2}$, and recall the notation $\tilde{\Gamma} = \Gamma \Sigma_{\mathbf{x}}^{1/2}$. It follows from (C.4) that \tilde{A} must be set to

$$\begin{aligned} \tilde{A}^\circ &= \arg \max_{\tilde{A}} \text{Tr} \left(\tilde{\Gamma} \tilde{A}^T \left(\tilde{A} \tilde{A}^T + \frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p} \right. \right. \\ &\quad \left. \left. \times \text{Tr} \left(\tilde{A} \tilde{A}^T \right) I_p \right)^{-1} \tilde{A} \tilde{\Gamma}^T \right). \end{aligned} \quad (\text{C.7})$$

Note that the right hand side of (C.7) is invariant to replacing \tilde{A} with $\alpha \cdot U \tilde{A}$ for any $\alpha > 0$ and for any unitary U . Consequently, we can fix $\text{Tr}(\tilde{A} \tilde{A}^T) = 1$, and write $\tilde{A} = \Lambda V^T$, where $\Lambda \in \mathcal{R}^{p \times n}$ is a diagonal matrix whose diagonal entries are arranged in a descending order, and $V \in \mathcal{R}^{n \times n}$ is unitary. Under this setting, solving (C.7) reduces to solving

$$\begin{aligned} \arg \max_{\Lambda, V} \text{Tr} \left(\tilde{\Gamma}^T \tilde{\Gamma} V \Lambda^T \left(\Lambda \Lambda^T + \frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p} I_p \right)^{-1} \Lambda V^T \right), \\ \text{subject to } \text{Tr}(\Lambda \Lambda^T) = 1. \end{aligned} \quad (\text{C.8})$$

We use Λ_A, V_A to denote the optimizing matrices of (C.8).

Let $\tilde{\Lambda} \triangleq \Lambda^T (\Lambda \Lambda^T + \frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p} I_p)^{-1} \Lambda$. Clearly, $\tilde{\Lambda}$ is a diagonal matrix with diagonal entries $(\tilde{\Lambda})_{l,l} = \frac{(\Lambda)_{l,l}^2}{(\Lambda)_{l,l}^2 + \frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p}}$, for all $l \in$

$\{1, 2, \dots, p\}$, thus, the diagonal entries of $\tilde{\Lambda}$ are arranged in descending order. By [50, Thm. II.1], the optimal unitary matrix V_A is the right singular vectors matrix of $\tilde{\Gamma}$. The fact that the diagonal entries of $\tilde{\Lambda}$ are arranged in descending order results in the optimal V_A being the right singular vectors matrix of $\tilde{\Gamma}$ instead of a permutation of this matrix [50, Thm. II.1].

With this setting, (C.8) becomes

$$\begin{aligned} \arg \max_{\Lambda} \sum_{i=1}^{\min(k,p)} \lambda_{\tilde{\Gamma},i}^2 \cdot \frac{(\Lambda)_{i,i}^2}{(\Lambda)_{i,i}^2 + \frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p}} \\ \text{subject to } \sum_{i=1}^p (\Lambda)_{i,i}^2 = 1. \end{aligned} \quad (\text{C.9})$$

To solve (C.9), we write $\alpha_i \triangleq (\Lambda)_{i,i}^2$. With this setting, (C.9) is concave with respect to $\{\alpha_i\}_{i=1}^p$. Thus, the KKT conditions are necessary and sufficient for optimality [51, Ch. 5.5.3]. The

resulting $\{\alpha_i\}_{i=1}^p$ satisfy $\alpha_i = 0$ for $i > k$ and for $i \leq k$,

$$\begin{aligned} \alpha_i &= \left(\sqrt{\frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p \cdot \beta}} \lambda_{\tilde{\Gamma},i} - \frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p} \right)^+ \\ &= \frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p} \left(\sqrt{\frac{3\tilde{M}_p^2 \cdot p}{2\kappa_p \cdot \beta}} \lambda_{\tilde{\Gamma},i} - 1 \right)^+, \end{aligned} \quad (\text{C.10})$$

where β is set such that $\sum_{i=1}^p \alpha_i = 1$.

By defining $\zeta \triangleq (\frac{3\tilde{M}_p^2 \cdot p}{2\kappa_p \cdot \beta})^{1/2}$, it follows that the diagonal entries of the optimal diagonal matrix Λ_A satisfy

$$(\Lambda_A)_{i,i}^2 = \begin{cases} \frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p} (\zeta \cdot \lambda_{\tilde{\Gamma},i} - 1)^+ & i \leq \min(k, p) \\ 0 & i > \min(k, p), \end{cases}$$

where $\zeta > 0$ is set such that $\frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p} \sum_{i=1}^p (\zeta \cdot \lambda_{\tilde{\Gamma},i} - 1)^+ = 1$.

The optimal analog combining is thus given by $A^\circ = U_A \Lambda_A V_A^T \Sigma_{\mathbf{x}}^{-1/2}$. Note that under this setting, the dynamic range in (C.2) is given by $\gamma^2 = \frac{\kappa_p}{p} \text{Tr}(\Lambda_A \Lambda_A^T) = \frac{\kappa_p}{p}$, proving (9b). The resulting optimal MSE can be written as

$$\begin{aligned} \text{MSE}(A^\circ) &= \text{MSE}(\tilde{A}^\circ) \\ &= \text{Tr}(\Gamma \Sigma_{\mathbf{x}} \Gamma^T) - \sum_{i=1}^k \lambda_{\tilde{\Gamma},i}^2 \cdot \frac{(\Lambda_A)_{i,i}^2}{(\Lambda_A)_{i,i}^2 + \frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p}} \\ &= \sum_{i=1}^k \lambda_{\tilde{\Gamma},i}^2 - \sum_{i=1}^{\min(k,p)} \lambda_{\tilde{\Gamma},i}^2 \frac{(\zeta \cdot \lambda_{\tilde{\Gamma},i} - 1)^+}{(\zeta \cdot \lambda_{\tilde{\Gamma},i} - 1)^+ + 1}. \end{aligned} \quad (\text{C.11})$$

Since $\lambda_{\tilde{\Gamma},i}^2 \cdot \left(1 - \frac{(\zeta \cdot \lambda_{\tilde{\Gamma},i} - 1)^+}{(\zeta \cdot \lambda_{\tilde{\Gamma},i} - 1)^+ + 1} \right) = \frac{\lambda_{\tilde{\Gamma},i}^2}{(\zeta \cdot \lambda_{\tilde{\Gamma},i} - 1)^+ + 1}$, it follows that (C.11) results in

$$\text{MSE}(A^\circ) = \begin{cases} \sum_{i=1}^k \frac{\lambda_{\tilde{\Gamma},i}^2}{(\zeta \cdot \lambda_{\tilde{\Gamma},i} - 1)^+ + 1}, & p \geq k \\ \sum_{i=1}^p \frac{\lambda_{\tilde{\Gamma},i}^2}{(\zeta \cdot \lambda_{\tilde{\Gamma},i} - 1)^+ + 1} + \sum_{i=p+1}^k \lambda_{\tilde{\Gamma},i}^2, & p < k. \end{cases}$$

Combining this with the design of A° proves the theorem. ■

D. Proof of Corollary 1

We note that the MSE in Theorem 1 decreases as ζ increases. Therefore, p must be set such that ζ is maximized. Let r denote the number of non-zero singular values $\{\lambda_{\tilde{\Gamma},i}\}$. From the definition of $\tilde{\Gamma}$ it follows that r is also the rank of the covariance matrix of $\tilde{\mathbf{s}}$. When $p \geq r$, ζ is set such that

$$\frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p} \sum_{i=1}^r (\zeta \cdot \lambda_{\tilde{\Gamma},i} - 1)^+ = 1. \quad (\text{D.1})$$

Since \tilde{M}_p grows exponentially with $1/p$, it follows that as p decreases, the value of $\frac{2\kappa_p}{3\tilde{M}_p^2 \cdot p}$, also decreases, and thus ζ is maximized for $p \geq r$ when $p = r$, proving the corollary. ■

E. Proof of Corollary 4

Since the optimal \mathbf{A}° for any p , including $p = k$, is given in Theorem 1, in the following we find a necessary and sufficient condition for which $\mathbf{A}^\circ = \mathbf{\Gamma}$, or equivalently, $\mathbf{U}_A \mathbf{\Lambda}_A \mathbf{V}_A^T = \tilde{\mathbf{\Gamma}}$. Note that \mathbf{V}_A already equals the right-singular matrix of $\tilde{\mathbf{\Gamma}}$, and thus $\mathbf{A} = \mathbf{\Gamma}$ is optimal if and only if $\mathbf{U}_A \mathbf{\Lambda}_A \mathbf{\Lambda}_A^T \mathbf{U}_A = \tilde{\mathbf{\Gamma}} \tilde{\mathbf{\Gamma}}^T$. This condition is satisfied only if

$$(\mathbf{\Lambda}_A)_{i,i}^2 = \lambda_{\tilde{\mathbf{\Gamma}},i}^2, \forall i \in \mathcal{K}. \quad (\text{E.1})$$

Since $\tilde{\mathbf{\Gamma}} \tilde{\mathbf{\Gamma}}^T$ is the covariance matrix of the MMSE estimate $\tilde{\mathbf{s}}$, the fact that the covariance matrix is non-singular implies that $\lambda_{\tilde{\mathbf{\Gamma}},i} \neq 0$ for all $i \in \mathcal{K}$. Combining this with (9a) results in

$$\lambda_{\tilde{\mathbf{\Gamma}},i}^2 = \frac{2\kappa_k}{3\tilde{M}_k^2 \cdot k} \left(\zeta \cdot \lambda_{\tilde{\mathbf{\Gamma}},i} - 1 \right)^+ = \frac{2\kappa_k}{3\tilde{M}_k^2 \cdot k} \left(\zeta \cdot \lambda_{\tilde{\mathbf{\Gamma}},i} - 1 \right). \quad (\text{E.2})$$

Consequently, $\zeta \cdot \lambda_{\tilde{\mathbf{\Gamma}},i} - 1 > 0$ for all $i \in \mathcal{K}$, and thus, from the condition $\frac{2\kappa_p}{3\tilde{M}_p^2} \sum_{i=1}^p (\zeta \cdot \lambda_{\tilde{\mathbf{\Gamma}},i} - 1)^+ = 1$, we have that $\zeta = k \left(\frac{3\tilde{M}_k^2}{2\kappa_k} + 1 \right) \left(\sum_{i=1}^k \lambda_{\tilde{\mathbf{\Gamma}},i} \right)^{-1}$. Plugging this into (E.2) results in

$$\lambda_{\tilde{\mathbf{\Gamma}},i}^2 = \lambda_{\tilde{\mathbf{\Gamma}},i} \left(\sum_{i=1}^k \lambda_{\tilde{\mathbf{\Gamma}},i} \right)^{-1} \quad \forall i \in \mathcal{K}. \quad (\text{E.3})$$

Note that (E.3) is satisfied if and only if $\lambda_{\tilde{\mathbf{\Gamma}},i} = \frac{1}{\sqrt{k}}$. As a result, $\mathbf{\Lambda}_A$ in (E.1) has identical diagonal entries, \mathbf{U}_A can be any permutation matrix, and both the condition on \mathbf{U}_A in Theorem 1 as well as the condition $\mathbf{U}_A \mathbf{\Lambda}_A \mathbf{\Lambda}_A^T \mathbf{U}_A = \tilde{\mathbf{\Gamma}} \tilde{\mathbf{\Gamma}}^T$ are satisfied here. Consequently, the condition $\lambda_{\tilde{\mathbf{\Gamma}},i} = \frac{1}{\sqrt{k}}$ is not only necessary for $\mathbf{A} = \mathbf{\Gamma}$ to be optimal, but it is also sufficient. Noting that in this case, the covariance matrix of $\tilde{\mathbf{s}}$ is $\frac{1}{k} \mathbf{I}_k$, concludes the proof of the corollary. ■

REFERENCES

- [1] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2325–2383, Oct. 1998.
- [2] Y. C. Eldar, *Sampling Theory: Beyond Bandlimited Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2015.
- [3] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 2006.
- [4] T. Berger and J. D. Gibson, "Lossy source coding," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2693–2723, Oct. 1998.
- [5] M. R. D. Rodrigues, N. Deligiannis, L. Lai, and Y. C. Eldar, "Rate-distortion trade-offs in acquisition of signal parameters," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, New Orleans, LA, USA, Mar. 2017, pp. 6105–6109.
- [6] S. Rini, L. Barlett, E. Erkip, and Y. C. Eldar, "A general framework for MIMO receivers with low-resolution quantization," in *Proc. IEEE Inf. Theory Workshop*, Kaohsiung, Taiwan, Nov. 2017, pp. 599–603.
- [7] J. Choi, J. Sung, B. L. Evans, and A. Gatherer, "Antenna selection for large-scale MIMO systems with low-resolution ADCs," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Calgary, AB, Canada, Apr. 2018, pp. 3594–3598.
- [8] J. Choi, B. L. Evans, and A. Gatherer, "Resolution-adaptive hybrid MIMO architectures for millimeter wave communications," *IEEE Trans. Signal Process.*, vol. 65, no. 23, pp. 6201–6216, Dec. 2017.
- [9] J. Mo, A. Alkhateeb, S. Abu-Surra, and R. W. Heath, "Hybrid architectures with few-bit ADC receivers: Achievable rates and energy-rate tradeoffs," *IEEE Trans. Wireless Commun.*, vol. 16, no. 4, pp. 2274–2287, Apr. 2017.
- [10] T. C. Zhang, C. K. Wen, S. Jin, and T. Jiang, "Mixed-ADC massive MIMO detectors: Performance analysis and design optimization," *IEEE Trans. Wireless Commun.*, vol. 15, no. 11, pp. 7738–7752, Nov. 2017.
- [11] Y. Li, C. Tao, G. Seco-Granados, A. Mezghani, A. L. Swindlehurst, and L. Liu, "Channel estimation and performance analysis of one-bit massive MIMO systems," *IEEE Trans. Signal Process.*, vol. 65, no. 15, pp. 4075–4089, Aug. 2017.
- [12] J. Choi, J. Mo, and R. W. Heath, "Near maximum-likelihood detector and channel estimator for uplink multiuser massive MIMO systems with one-bit ADCs," *IEEE Trans. Commun.*, vol. 64, no. 5, pp. 2005–2018, May 2016.
- [13] G. Zeitler, G. Kramer, and A. C. Singer, "Bayesian parameter estimation using single-bit dithered quantization," *IEEE Trans. Signal Process.*, vol. 60, no. 6, pp. 2713–2726, Jun. 2012.
- [14] O. Dabeer and U. Madhow, "Channel estimation with low-precision analog-to-digital conversion," in *Proc. IEEE Int. Conf. Commun.*, Cape Town, South Africa, May 2010, pp. 1–6.
- [15] M. S. Stein, S. Bar, J. A. Nossek, and J. Tabrikian, "Performance analysis for channel estimation with 1-bit ADC and unknown quantization threshold," *IEEE Trans. Signal Process.*, vol. 66, no. 10, pp. 2557–2571, May 2018.
- [16] J. Mo, P. Schniter, and R. W. Heath, "Channel estimation in broadband millimeter wave MIMO systems with few-bit ADCs," *IEEE Trans. Signal Process.*, vol. 66, no. 5, pp. 1141–1154, Mar. 2018.
- [17] J. Sung, J. Choi, and B. L. Evans, "Narrowband channel estimation for hybrid beamforming millimeter wave communication systems with one-bit quantization," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Calgary, AB, Canada, Apr. 2018, pp. 3914–3918.
- [18] Y. Chi and H. Fu, "Subspace learning from bits," *IEEE Trans. Signal Process.*, vol. 65, no. 17, pp. 4429–4442, Sep. 2017.
- [19] R. M. Corey and A. C. Singer, "Wideband source localization using one-bit quantized arrays," in *Proc. IEEE 7th Int. Workshop Comput. Adv. Multi-Sensor Adaptive Process.*, Curacao, Dutch Antilles, Dec. 2017, pp. 1–5.
- [20] K. Yu, Y. D. Zhang, M. Bao, Y. Hu, and Z. Wang, "DOA estimation from one-bit compressed array data via joint sparse representation," *IEEE Signal Process. Lett.*, vol. 23, no. 8, pp. 1279–1283, Sep. 2016.
- [21] C. L. Liu and P. P. Vaidyanathan, "One-bit sparse array DOA estimation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, New Orleans, LA, USA, Mar. 2017, pp. 3126–3130.
- [22] Y. Polyanskiy and Y. Wu, *Lecture Notes on Information Theory*. Cambridge, MA, USA: MIT, 2015.
- [23] H. Witsenhausen, "Indirect rate distortion problems," *IEEE Trans. Inf. Theory*, vol. IT-26, no. 5, pp. 518–521, Sep. 1980.
- [24] J. K. Wolf and J. Ziv, "Transmission of noisy information to a noisy receiver with minimum distortion," *IEEE Trans. Inf. Theory*, vol. IT-16, no. 4, pp. 406–411, Jul. 1970.
- [25] V. Kostina and S. Verdú, "Nonasymptotic noisy lossy source coding," *IEEE Trans. Inf. Theory*, vol. 62, no. 11, pp. 6111–6123, Nov. 2016.
- [26] V. Kostina and S. Verdú, "Fixed-length lossy compression in the finite blocklength regime," *IEEE Trans. Inf. Theory*, vol. 58, no. 3, pp. 3309–3338, Jun. 2012.
- [27] A. Kipnis, A. J. Goldsmith, and Y. C. Eldar, "Fundamental distortion limits of analog-to-digital compression," *IEEE Trans. Inf. Theory*, vol. 64, no. 9, pp. 6013–6033, Sep. 2018.
- [28] A. Kipnis, Y. C. Eldar, and A. J. Goldsmith, "Analog-to-digital compression: A new paradigm for converting signals to bits," *IEEE Signal Process. Mag.*, vol. 35, no. 3, pp. 16–39, Mar. 2018.
- [29] R. Gupta and A. O. Hero, "High-rate vector quantization for detection," *IEEE Trans. Inf. Theory*, vol. 49, no. 8, pp. 1951–1969, Aug. 2003.
- [30] K. R. Varshney and L. R. Varshney, "Quantization of prior probabilities for hypothesis testing," *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 4553–4562, Oct. 2008.
- [31] B. M. Dogahe and M. N. Murthi, "Quantization for classification accuracy in high-rate quantizers," in *Proc. IEEE Digit. Signal Process. Signal Process. Educ. Meeting*, Sedona, AZ, USA, Jan. 2011, pp. 277–282.
- [32] P. Bianchi and J. Jakubowicz, "Adaptive learning vector quantization for online parametric estimation," *IEEE Trans. Signal Process.*, vol. 61, no. 12, pp. 3119–3128, Jun. 2013.
- [33] S. S. Ioushua and Y. C. Eldar, "A family of hybrid analog digital beamforming methods for massive MIMO systems," *IEEE Trans. Signal Process.*, vol. 16, no. 12, pp. 3243–3257, Jun. 2019.
- [34] A. Alkhateeb, J. Mo, N. Gonzalez-Prelcic, and R. W. Heath, "MIMO precoding and combining solutions for millimeter-wave systems," *IEEE Commun. Mag.*, vol. 52, no. 12, pp. 122–131, Dec. 2014.

- [35] W. B. Abbas, F. Gomez-Cuba, and M. Zorzi, "Millimeter wave receiver efficiency: A comprehensive comparison of beamforming schemes with low resolution ADCs," *IEEE Trans. Wireless Commun.* vol. 16, no. 12, pp. 8131–8146, Dec. 2017.
- [36] A. Kipnis, G. Reeves, Y. C. Eldar, and A. Goldsmith, "Compressed sensing under optimal quantization," in *Proc. IEEE Int. Symp. Inf. Theory*, Aachen, Germany, Jun. 2017, pp. 2148–2152.
- [37] J. Li, N. Chaddha, and R. M. Gray, "Asymptotic performance of vector quantizers with a perceptual distortion measure," *IEEE Trans. Inf. Theory*, vol. 45, no. 4, pp. 1082–1091, May 1999.
- [38] T. S. Han, *Information-Spectrum Methods in Information Theory*. Berlin, Germany: Springer, 2003.
- [39] Y. Linde, A. Buzo, and R. Gray, "An algorithm for vector quantizer design," *IEEE Trans. Commun.*, vol. COM-28, no. 1, pp. 84–95, Jan. 1980.
- [40] J. Dauwels, "Numerical computation of the capacity of continuous memoryless channels," in *Proc. 26th Symp. Inf. Theory BENELUX*, 2005, pp. 221–228.
- [41] O. Rioul, "Information theoretic proofs of entropy power inequalities," *IEEE Trans. Inf. Theory*, vol. 57, no. 1, pp. 33–55, Jan. 2011.
- [42] R. M. Gray and T. G. Stockholm, "Dithered quantizers," *IEEE Trans. Inf. Theory*, vol. 39, no. 3, pp. 805–812, Mar. 1993.
- [43] B. Widrow, I. Kollar, and M. C. Liu, "Statistical theory of quantization," *IEEE Trans. Instrum. Meas.*, vol. 45, no. 2, pp. 353–361, Apr. 1996.
- [44] M. Shohat, G. Tsintsadze, N. Shlezinger, and Y. C. Eldar, "Deep quantization for MIMO channel estimation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Brighton, U.K., May 2019, pp. 3912–3916.
- [45] A. Llera and C. F. Beckmann, "Estimating an inverse gamma distribution," 2016, *arXiv:1605.01019*.
- [46] R. Zamir and M. Feder, "On universal quantization by randomized uniform lattice quantizers," *IEEE Trans. Inf. Theory*, vol. 88, no. 2, pp. 428–436, May 1992.
- [47] D. P. Palomar and Y. Jiang, *MIMO Transceiver Design via Majorization Theory*. Delft, The Netherlands: Now Publishers, 2007.
- [48] C. E. Shannon, "Coding theorems for a discrete source with a fidelity criterion," *IRE Int. Conv. Rec.*, vol. 7, pp. 142–163, Mar. 1959.
- [49] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 3rd ed. New York, NY, USA: McGraw-Hill, 1991.
- [50] J. B. Lassaré, "A trace inequality for matrix product," *IEEE Trans. Autom. Control*, vol. 40, no. 8, pp. 1500–1501, Aug. 1995.
- [51] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.



Nir Shlezinger (M'17) received the B.Sc., M.Sc., and Ph.D. degrees in 2011, 2013, and 2017, respectively, from Ben-Gurion University, Beersheba, Israel, all in electrical and computer engineering. From 2017 to 2019, he was a Postdoctoral Researcher with the Technion, Israel Institute of Technology. From 2009 to 2013, he was an Engineer with Yitran Communications. He is currently a Postdoctoral Researcher with the Signal Acquisition Modeling and Processing Laboratory, Weizmann Institute of Science, Rehovot, Israel. His research interests include information theory and signal processing for communications.



Yonina C. Eldar (S'98–M'02–SM'07–F'12) received the B.Sc. degrees in physics and electrical engineering in 1995 and 1996, respectively, both from Tel Aviv University, Tel Aviv, Israel, and the Ph.D. degree in electrical engineering and computer science in 2002 from the Massachusetts Institute of Technology, Cambridge, MA, USA.

She is currently a Professor with the Department of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, Israel. She was previously a Professor with the Department of Electrical Engineering, Technion, where she held the Edwards Chair in Engineering. She is also a Visiting Professor at MIT, a Visiting Scientist at the Broad Institute, and an Adjunct Professor at Duke University, and was a Visiting Professor at Stanford. Her research interests are in the broad areas of statistical signal processing, sampling theory and compressed sensing, learning and optimization methods, and their applications to biology and optics.

Dr. Eldar is a member of the Israel Academy of Sciences and Humanities (elected 2017) and a EURASIP Fellow. She was a member of the Young Israel Academy of Science and Humanities and the Israel Committee for Higher Education. She is the Editor-in-Chief for the Foundations and Trends in Signal Processing, a member of the IEEE Sensor Array and Multichannel Technical Committee, and serves on several other IEEE committees. In the past, she was a Signal Processing Society Distinguished Lecturer, member of the IEEE Signal Processing Theory and Methods and Bio Imaging Signal Processing technical committees, and served as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the *EURASIP Journal of Signal Processing*, the *SIAM Journal on Matrix Analysis and Applications*, and the *SIAM Journal on Imaging Sciences*. She was a Co-Chair and a Technical Co-Chair of several international conferences and workshops. She was the recipient of the numerous awards for excellence in research and teaching, including the IEEE Signal Processing Society Technical Achievement Award (2013), the IEEE/AESS Fred Nathanson Memorial Radar Award (2014), and the IEEE Kiyoo Tomiyasu Award (2016). She was a Horev Fellow of the Leaders in Science and Technology program at the Technion and an Alon Fellow. She received the Michael Bruno Memorial Award from the Rothschild Foundation, the Weizmann Prize for Exact Sciences, the Wolf Foundation Krill Prize for Excellence in Scientific Research, the Henry Taub Prize for Excellence in Research (twice), the Hershel Rich Innovation Award (three times), the Award for Women with Distinguished Contributions, the Andre and Bella Meyer Lectureship, the Career Development Chair at the Technion, the Muriel and David Jacknow Award for Excellence in Teaching, and the Technions Award for Excellence in Teaching (twice). She was also the recipient of the several best paper awards and best demo awards together with her research students and colleagues including the SIAM outstanding Paper Prize and the IET Circuits, Devices and Systems Premium Award, and was selected as one of the 50 most influential women in Israel.



Miguel R. D. Rodrigues (SM'15) is currently a Reader in information theory and processing with the Department of Electronic and Electrical Engineering, University College London, London, U.K., and a Faculty Fellow with the Turing Institute, London, U.K. He was with the Department of Computer Science, University of Porto, Portugal, rising through the ranks from Assistant to Associate Professor. He also held research positions at Princeton University, Cambridge University, and Duke University. His research interests which lie in the general areas of information theory and processing have led to nearly 200 publications in leading journals and conferences in the field, including the prestigious IEEE Communications and Information Theory Societies Joint Paper Award 2011. He has coauthored a book *Information-Theoretic Methods in Data Science* (Cambridge Univ. Press, 2020).