

Dimensionally-Efficient Transmission and Storage of Unitary Matrices

Juan Vidal Alegría

Abstract—Unitary matrices are the basis of a large number of signal processing applications. In many of these applications, finding ways to efficiently store, and even transmit these matrices, can significantly reduce memory and throughput requirements. In this work, we study the problem of efficient transmission and storage of unitary matrices. Specifically, we explicitly derive a dimensionally-efficient parametrization (DEP) for unitary matrices that allows identifying them with sequences of real numbers, where the dimension coincides with the dimension of the unitary group where they lie. We also characterize its inverse map that allows retrieving the original unitary matrices from their DEP. The proposed approach effectively allows halving the dimension with respect to naively considering all the entries of each unitary matrix, thus reducing the resources required to store and transmit these matrices. Furthermore, we show that the sequence of real numbers associated to the proposed DEP is bounded, and we delimit the interval where these numbers are contained, facilitating the implementation of quantization approaches with limited distortion. On the other hand, we outline ways to further reduce the dimension of the DEP when considering more restrictive constraints for matrices that show up in certain applications. The numerical results showcase the potential of the proposed approach in general settings, as well as in three specific applications of current interest for wireless communications research.

Index Terms—Unitary matrices, dimensionally-efficient parametrization (DEP), multiple-input multiple-output (MIMO), reconfigurable surfaces, distributed-MIMO (D-MIMO).

I. INTRODUCTION

UNITARY matrices have a crucial role in signal processing, finding application in a long list of diverse fields that include wireless communications [1], biomedical engineering [2], microwave theory [3], and quantum computation [4]. Unitary transformations have unique properties which allow them to preserve norms and inner products, while they can be inverted at low complexity by simple conjugate transposition. These and other features make them one of the basic building blocks for spectral analysis, quantum processing, independent component analysis (ICA), and other fundamental signal processing techniques.

Many of the applications that employ unitary matrices require storing these matrices so that they can be accessible at different instances. In the context of quantum computing, finding ways to physically store unitary matrices in quantum systems is a research area in itself [5], [6]. However, the

current work focuses on the storage of unitary matrices using classical computing systems, specifically in the context of the widely employed digital technology.¹ In this sense, the goal is to store unitary matrices effectively while using the least possible memory resources, e.g., mapping unitary matrices with the shortest binary sequences that allow recovering them under given error. To the best of our knowledge, this problem has not been explicitly studied in previous literature.

The reduced dimensionality of the unitary group allows parameterizing unitary matrices through a sequence of real numbers half as long as the one required to parameterize arbitrary complex matrices. Hence, a straightforward way to store unitary matrices efficiently is to define a suitable *dimensionally-efficient parametrization* (DEP), i.e., a parametrization having minimum number of dimensions, and to store instead the resulting sequence of minimum dimension that allows retrieving each unitary matrix without loss. In [7] a DEP is defined to represent qudits as a sequence of $N^2 - 1$ real numbers, which coincides to the real dimension of the special unitary group (a restricted version of the unitary group) where the qudits lie. In [8], a similar DEP is considered for towards the goal of training recurrent neural networks (RNNs) with unitary transition matrices. In this work, we show how to practically exploit such parametrizations to efficiently store unitary matrices.

In some applications, specially when dealing with wireless communication systems, it may even be necessary to send unitary matrices to other remote locations so that different systems can employ them to perform their tasks. For example, in order to achieve full channel capacity in multiple-input multiple-output (MIMO) systems, the transmitter has to precode its symbols using the right unitary matrix from the singular-value decomposition (SVD) of the MIMO channel matrix. This requires to consider channel state information (CSI) feedback strategies to transmit the respective unitary matrix, which is typically only known at the receiver where the channel is estimated. This has been a topic widely studied in the literature with numerous solutions which may constitute the most relevant state-of-the-art for our study [9]–[12]. In this context, [9] considers a DEP based on Givens rotations for unitary matrices up to a diagonal phase uncertainty. This DEP is then used to reduce the CSI feedback rate by performing optimized quantization upon it. A similar approach based on

This work was supported by the Swedish Foundation of Strategic Research (CHI19-0001).

J. Vidal Alegría is with Lund University, SE-22100 Lund, Sweden (corresponding email: juan.vidal_alegría@eit.lth.se).

¹Classical computing is still the main technology used to perform signal processing tasks, while there seems to be a necessity for their coexistence with quantum computing systems towards future technological advances [4].

Givens rotations is considered in [10], which further proposes differential schemes to update the respective CSI. Other differential techniques for transmission of unitary matrices have been considered in [11], [12].

Due to the intrinsic equivalence between the rate required to transmit a unitary matrix, and the memory required to store it [13], many of the methods for transmitting unitary matrices can be directly reused for the storage counterpart. However, the main limitation of the approaches in [9]–[12] lie in the fact that they are tailored to the CSI feedback application, e.g., by assuming specific channel models with given spatial distributions or time-correlation structures, or by ignoring the diagonal phase uncertainty of the transmitted unitary matrices. Moreover, the differential techniques [10]–[12] have the extra limitation that they still require an initial full-estimate of the respective unitary matrix. Thus, finding techniques to store and transmit unitary matrices in general settings remains an open problem.

In the context of modern wireless communications research, several promising enabling technologies have been proposed towards future generations of wireless systems. Some of these enabling technologies depend strongly on the possibility to effectively store and transmit unitary matrices. For example, the recently proposed beyond-diagonal reconfigurable intelligent surface (BD-RIS) [14]–[16], which extends the widespread reconfigurable intelligent surface (RIS) [17], [18] by including inter-element reconfigurable connections to further enhance performance, leads to reflection matrices with unitary constraints. Since the optimal reflection matrices are typically computed at a transceiver node with channel estimation capabilities, forwarding the computed matrices to the BD-RIS, where they are physically applied, may directly benefit from efficient methods to send and store unitary matrices. Another recent proposal is to use lossless microwave analog computing networks for MIMO processing [19]–[21], which are characterized by scattering matrices with unitary constraints. In this case, sharing these matrices between the nodes where they are computed and the nodes where they are applied may also benefit from efficient storage and transmission of unitary matrices. To the best of our knowledge, this issue has been largely overlooked in existing literature.

The main contributions of this paper are summarized next:

- We derive through explicit transformations a DEP for unitary matrices to map arbitrary unitary matrices with sequences of real numbers with minimum dimension.
- We study the time-complexity associated to computing the proposed DEP, as well as to retrieving the original unitary matrix from it, and conclude that the order does not increase compared with performing SVD—which is often a pre-requisite for the considered applications.
- We show that the proposed DEP maintains the property of boundedness, and derive a range within which the resulting real numbers should be contained.
- We outline how to adapt the proposed DEP to incorporate extra constraints such as symmetry, real-valuedness (i.e., for orthogonal matrices), or phase uncertainty.

- We study three wireless communications applications, specifically in the context of modern MIMO technologies as mentioned above, and we show how they may benefit from the considered approach.
- We numerically analyze the potential of the proposed DEP in the task of storage and transmission of unitary matrices in general tasks, as well as in the specific wireless communication applications previously mentioned.

Some of the methods hereby exploited may appear trivial for readers familiar with Lie theory [22]. However, an important contribution of this work is to facilitate their practical exploitation, and to point out their significance in the context of storage and transmission of unitary matrices—specially in the context of wireless communications. On the other hand, the proof of boundedness of the proposed DEP for unitary matrices poses, to the best of our knowledge, a novel theoretical contribution, further providing a useful explicit range for the values of the entries of this DEP. The previous result may be exploited to implement simple quantization methods with bounded error.

II. PROBLEM FORMULATION

Consider the transmission/storage² of an $N \times N$ unitary matrix, $\mathbf{U} \in \mathcal{U}(N)$, where $\mathcal{U}(N)$ denotes the unitary group. The straightforward approach to parameterize $\mathbf{U}$ is through the N^2 complex numbers composing it. Since each complex number may be characterized by two real numbers (its real and imaginary part), this would correspond to characterizing $\mathbf{U}$ by $2N^2$ real numbers. Thus, in order to transmit $\mathbf{U}$ we could simply send the sequence of $2N^2$ real numbers composing it, e.g., through $2N^2$ uses of a scalar channel. This naive approach would still be the most dimensionally-efficient way to transmit an arbitrary complex matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$, where we assume no specific structure on $\mathbf{A}$. The reason is that the space $\mathbb{C}^{N \times N}$ of $N \times N$ complex matrices is composed of $2N^2$ real dimensions, so in order to differentiate $\mathbf{A}$ from any other arbitrary point in $\mathbb{C}^{N \times N}$ we would need to parameterize it with at least $2N^2$ real numbers. However, it is well known that $\mathcal{U}(N)$ corresponds to a topological (Riemannian)³ manifold of N^2 real dimensions [23]. Hence, it should be possible to parameterize without loss any unitary matrix as a sequence of N^2 real numbers, hereby referred to as a DEP for $\mathcal{U}(N)$.

The notion of $\mathcal{U}(N)$ being considered a topological (Riemannian) manifold of (real) dimension N^2 inherently means that there exist local homeomorphic (bijective) maps between $\mathcal{U}(N)$ and $\mathbb{R}^{N^2}$ [23]. However, even if these maps (also termed local charts/coordinates) can be explicitly defined, the term “local” means that each of these may only apply to neighborhoods around specific elements of $\mathcal{U}(N)$ (and their respective image within $\mathbb{R}^{N^2}$). Due to the nature of

²In this work we can interchange the word transmission (transmit) with storage (store), since both information rate and level of compression are similarly influenced by the parametrization of the input [13].

³In order to consider $\mathcal{U}(N)$ as a Riemannian manifold it should further be equipped with a Riemannian metric. However, this work does not directly require such tool.

the problem under consideration, it is in fact desirable to have a more general construction: a unique map that can be applied to any element of $\mathcal{U}(N)$ and leads to distinguishable elements within $\mathbb{R}^{N^2}$. This way, if we have one transmitter and one receiver (conversely one system storing a unitary matrix and another one accessing it) they should both have a common function to translate from unitary matrices to their DEPs without having to know further details about the neighborhoods around which these translations would work. An initial goal of this work is thus to explicitly define a DEP for $\mathcal{U}(N)$, which consists of an injective function

$$\mathbf{f} : \mathcal{U}(N) \rightarrow \mathbb{R}^{N^2}, \quad (1)$$

which allows to unequivocally identify an arbitrary unitary matrix $\mathbf{U} \in \mathcal{U}(N)$ with a real vector $\mathbf{v} = \mathbf{f}(\mathbf{U}) \in \mathbb{R}^{N^2}$.⁴

Another valuable property of the unitary group is its compactness [22], [23]. This implies that $\mathcal{U}(N)$ is a closed and bounded (under the standard metric) subspace of $\mathbb{C}^{N \times N}$, i.e., in the same way as the interval $[0, 1]$ is a closed and bounded subspace of $\mathbb{R}$. The reason why this property is important is that, if we consider digital systems, in order to transmit/store a real number $r \in \mathbb{R}$ we first need to be able to translate it to a finite binary sequence. This is usually done by quantization [24], which corresponds to approximating r with the closest value $\hat{r} \in \mathcal{Q}$ within a finite subset $\mathcal{Q} \subset \mathbb{R}$ ($|\mathcal{Q}| = M$). From the finite property, the quantized subset for any quantization process is inevitably contained within a bounded (compact) interval, i.e., $\mathcal{Q} \subset [a, b]$. Thus, if the real number r takes values much larger than b (or much smaller than a), the quantization process would incur significant distortion and information loss due to clipping effects [25]. On the other hand, if we know that r is contained within a bounded subset of $\mathbb{R}$, e.g., $r \in [a', b']$, we can avoid clipping altogether by adjusting accordingly the quantization interval. Hence, it is desirable that our DEP $\mathbf{f}(\mathbf{U}) \in \mathbb{R}^{N^2}$ maintains the boundedness property inherent to the unitary group, i.e., such that $\mathbf{f}(\mathbf{U}) \in [\tau_{\min}, \tau_{\max}]^{N^2}$. This would allow us to limit the distortion associated to quantizing $\mathbf{f}(\mathbf{U})$ since we could adjust the quantization intervals to avoid/limit clipping effects. Moreover, we would like to characterize the tightest interval $[\tau_{\min}, \tau_{\max}]$ such that the image of $\mathcal{U}(N)$ under $\mathbf{f}$ approximately coincides with $[\tau_{\min}, \tau_{\max}]^{N^2}$. This would allow us to efficiently quantize the entries of $\mathbf{f}(\mathbf{U}) \in [\tau_{\min}, \tau_{\max}]^{N^2}$ by tightly adjusting the quantization intervals.

Altogether, this work aims at defining an injective function

$$\mathbf{f} : \mathcal{U}(N) \rightarrow [a, b]^{N^2}, \quad (2)$$

with a codomain $[a, b]^{N^2}$ as tight as possible to its image $\mathbf{f}(\mathcal{U}(N))$. Ideally, we would like $\mathbf{f}$ to be also surjective (or to be able to precisely delimit its image) so that its codomain can be made equal to its image, allowing for perfect adjustment of the dynamic range of the quantizers potentially employed

⁴Note that surjectivity can always be enforced by restricting the output domain since we may assume that we will always start with a given unitary matrix, leading to a well-defined inverse function $\mathbf{f}^{-1} : \mathbf{f}(\mathcal{U}(N)) \rightarrow \mathcal{U}(N)$.

to digitalize the outputs of $\mathbf{f}$. However, ensuring surjectivity is a complex task which only offers a limited gain in terms of quantization resolution. Thus, we consider that surjectivity is not a formal requirement.

III. UNITARY MATRIX DIMENSIONALLY-EFFICIENT PARAMETRIZATION

A. The exponential map for $\mathcal{U}(N)$

Let us start by introducing the exponential map. For general Lie groups, the exponential map is defined as a map between the Lie algebra associated to the Lie group and the group itself. Alternatively, we can say that the exponential map links the tangent plane at the identity element with the Lie group, which gives a more geometric interpretation (one may think of a sphere being unfolded into its tangent plane). In the case of the unitary group, we can write

$$\exp : \mathfrak{u}(N) \rightarrow \mathcal{U}(N), \quad (3)$$

where $\mathfrak{u}(N)$ corresponds to the unitary Lie algebra, which is defined as [22], [23], [26].

$$\mathfrak{u}(N) \triangleq T_{\mathbf{I}_N} \mathcal{U}(N) = \{\mathbf{X} \in \mathbb{C}^{N \times N} : \mathbf{X}^H + \mathbf{X} = \mathbf{0}\}, \quad (4)$$

with $T_{\mathbf{I}_N} \mathcal{U}(N)$ denoting the tangent plane to $\mathcal{U}(N)$ at the identity element, $\mathbf{I}_N$. From (4), we can note the correspondence between $\mathfrak{u}(N)$ and the space of skew-Hermitian matrices [22], [23].

The exponential map is a useful tool to link a Lie algebra with a Lie group. However, the premise of this work is that, starting with a unitary matrix, we would like to link it to a DEP. To this end, we should also consider the inverse exponential map, which allows mapping a Lie group to its Lie algebra. For the unitary group, we can write

$$\exp^{-1} : \mathcal{U}(N) \rightarrow \mathfrak{u}(N). \quad (5)$$

A Lie algebra corresponds to a real vector space of the same dimension as its Lie group, endorsing it with nice properties to operate with its elements by exploiting standard notions from linear algebra [22], [23]. On the other hand, the exponential map for any connected, compact matrix Lie group (as the unitary group) is always surjective. This means that the inverse exponential map, defined in (5) for the unitary group, is injective. Thus, this map fulfills our initial requirement for unequivocally identifying unitary matrices with elements of their Lie algebra. We next focus on providing explicit expressions for computing the exponential map, and the inverse exponential map, in the context of $\mathcal{U}(N)$.

For any matrix Lie group (as $\mathcal{U}(N)$), the exponential map coincides with the matrix exponential [22], [23]. Conversely, we may identify the inverse exponential map with the matrix logarithm (we can substitute $\exp^{-1}$ in (5) with $\log$). Note that, within the convergence region of the matrix logarithm, we have that $\exp(\log(\mathbf{A})) = \mathbf{A}$ [22, Theorem 2.8]. Since both unitary matrices and skew-Hermitian matrices fall within the class of normal matrices [27] (and are thus unitarily diagonalizable), we can give an explicit expression for the

matrix exponential of a skew-Hermitian matrix, and for the matrix logarithm of a unitary matrix, based on their eigenvalue decompositions (EVDs). Specifically, for $\mathbf{X} \in \mathfrak{u}(N)$ we have

$$\exp(\mathbf{X}) = \mathbf{U}_\mathbf{X} \exp(\boldsymbol{\Lambda}_\mathbf{X}) \mathbf{U}_\mathbf{X}^\mathbf{H}, \quad (6)$$

and for $\mathbf{U} \in \mathcal{U}(N)$ we have

$$\log(\mathbf{U}) = \mathbf{U}_\mathbf{U} \log(\boldsymbol{\Lambda}_\mathbf{U}) \mathbf{U}_\mathbf{U}^\mathbf{H}, \quad (7)$$

where we have considered the EVDs $\mathbf{X} = \mathbf{U}_\mathbf{X} \boldsymbol{\Lambda}_\mathbf{X} \mathbf{U}_\mathbf{X}^\mathbf{H}$ and $\mathbf{U} = \mathbf{U}_\mathbf{U} \boldsymbol{\Lambda}_\mathbf{U} \mathbf{U}_\mathbf{U}^\mathbf{H}$, with $\boldsymbol{\Lambda}_\mathbf{X}$ and $\boldsymbol{\Lambda}_\mathbf{U}$ corresponding to diagonal matrices, and $\mathbf{U}_\mathbf{X}$ and $\mathbf{U}_\mathbf{U}$ corresponding to unitary matrices. Note that, for diagonal matrices, the exponential/logarithm is defined as the traditional complex exponential/logarithm applied to the diagonal entries. On the other hand, for unitary matrices that do not fulfill the explicit convergence criterion of the matrix logarithm, given by $\|\mathbf{U} - \mathbf{I}_N\| < 1$ [22, Theorem 2.8], we can still have a well defined inverse exponential map by considering the complex logarithm [28] in (7), which allows having real negative diagonal entries in $\boldsymbol{\Lambda}_\mathbf{U}$ while ensuring $\exp(\log(\mathbf{U})) = \mathbf{U}$. The only remaining cases leading to an ill-defined logarithm in (7) would then correspond to when $\boldsymbol{\Lambda}_\mathbf{U}$ has diagonal entries equal to 0. However, this can never happen when considering unitary matrices since their eigenvalues always lie on the complex unit circle [29, Section 4.5], [30], i.e., the diagonal elements of $\boldsymbol{\Lambda}_\mathbf{U}$ are always unimodular (thus non-zero).

Remark 1: The relation between unitary matrices and Skew-Hermitian matrices may be understood from the standard scalar/element-wise complex exponential and logarithm. Unitary matrices have eigenvalues which are modulus-1 complex numbers, while Skew-Hermitian matrices (which may be seen as Hermitian matrices scaled by the imaginary unit) have eigenvalues which are purely imaginary [30]. Hence, (7) and (17) map the elements from $\mathfrak{u}(N)$ with the elements from $\mathcal{U}(N)$ through their eigenvalues, in the same way that the standard complex logarithm and exponential map modulus-1 complex numbers with purely imaginary numbers.

Remark 2: An alternative injective map between Lie algebra $\mathfrak{u}(N)$ and the Lie Group $\mathcal{U}(N)$ could be the Cayley transform, as considered in [31] for the alternative problem of generating unitary space-time codes. However, this transform would not work for cases where the input unitary matrix has any eigenvalues at -1, which limits its applicability. Moreover, the resulting image would not be bounded since when inputting a unitary matrix with eigenvalues close to -1 the resulting Skew-Hermitian matrix would have the respective eigenvalues approaching infinity, reducing practicality since having a compact image in (2) would not be possible.

We have now defined the maps that allow unequivocally translating unitary matrices to elements of the real vector space $\mathfrak{u}(N)$, and vice-versa. Our next goal is to exploit the real vector space property of $\mathfrak{u}(N)$ to parameterize its elements as dimensionally-efficient sequences of real numbers within a well-defined basis.

B. A basis for $\mathfrak{u}(N)$

Let us consider the standard Hilbert-Schmidt inner product $\langle \mathbf{A}, \mathbf{B} \rangle = \text{trace}(\mathbf{A}^\mathbf{H} \mathbf{B})$, which induces the squared Frobenius norm for $\mathbf{A} = \mathbf{B}$.⁵ In this section, we will propose an orthonormal basis for $\mathfrak{u}(N)$ under the previous considerations. This basis will allow us to extract the N^2 real coordinates associated to each element $\mathbf{X} \in \mathfrak{u}(N)$, by simply applying a unitary transformation on the vectorized $\text{vec}(\mathbf{X})$.

As noted in Section III-A, $\mathfrak{u}(N)$ corresponds to a real vector space which may be identified with the space of skew-Hermitian matrices (4). The generalized Gell-Mann matrices [7], [32], offer an orthogonal basis for the space of $N \times N$ Hermitian traceless matrices. By further including the identity matrix $\mathbf{I}_N$ as an extra basis element, the complete basis would actually span the whole space of $N \times N$ Hermitian matrices (traceless or not) [32]. On the other hand, any skew-Hermitian matrix can be generated as a Hermitian matrix scaled by the complex unit j , since this exchanges the properties of the real and imaginary parts. Note that the real part of a Hermitian matrix is symmetric and its imaginary part is antisymmetric, while the converse is true for skew-Hermitian matrices. Hence, we can obtain an orthogonal basis for $\mathfrak{u}(N)$ by including the generalized Gell-Mann matrices, together with the identity matrix, each scaled by j . We can further divide each basis element by its Frobenius norm to get a normalized basis within the considered framework. This gives the basis

$$\mathcal{B} = \{\mathbf{B}_1, \dots, \mathbf{B}_{N^2}\}, \quad (8)$$

whose elements are defined in (10). The last three lines of (10) correspond to the scaled generalized Gell-Mann matrices consisting of $N-1$ imaginary diagonal matrices, $N(N-1)/2$ imaginary symmetric matrices, and $N(N-1)/2$ real anti-symmetric matrices. Note that the arrangement of the basis elements is irrelevant as long as the basis contains all the N^2 mutually-orthogonal elements. However, a convention should be adopted so that different systems, e.g., a transmitter and a receiver, may coherently retrieve unitary matrices from their DEPs. For the symmetric and anti-symmetric elements in the last two rows of (10), one possible arrangement is obtained by row traversing the upper triangular part, which gives

$$k(n) = \left\lfloor \frac{2N-1 - \sqrt{(2N-1)^2 - 8n}}{2} \right\rfloor + 1, \quad (9a)$$

$$l(n) = k(n) - (n \bmod (N - k_n)). \quad (9b)$$

Note that $k(n)$ in (9a) simply corresponds to a sequence of $N-1$ 1s, followed by $N-2$ 2s, and so forth until reaching the last element equaling $N-1$.

Now that we have defined a basis for $\mathfrak{u}(N)$ it only remains to define the transformation that allows extracting the N^2 real coordinates associated to each element $\mathbf{X} \in \mathfrak{u}(N)$.

⁵The Frobenius norm is specially interesting here since it ensures coherence between the norm of a matrix and the norm of its vectorized form, i.e., $\|\mathbf{A}\|_F = \|\text{vec}(\mathbf{A})\|$.

$$B_n = \begin{cases} \frac{j}{\sqrt{N}} \mathbf{I}_N, & n = 1 \\ \frac{j}{\sqrt{n(n-1)}} \left(\sum_{i=1}^{n-1} \mathbf{E}_{i,i} - n \mathbf{E}_{n,n} \right), & n = 2, \dots, N \\ \frac{j}{\sqrt{2}} (\mathbf{E}_{k(n'), \ell(n')} + \mathbf{E}_{k(n''), \ell(n'')}), & n = N+1, \dots, \frac{(N+1)N}{2}, \quad n' = n - N \\ \frac{j}{\sqrt{2}} (\mathbf{E}_{k(n''), \ell(n'')} - \mathbf{E}_{k(n'), \ell(n')}), & n = \frac{(N+1)N}{2} + 1, \dots, N^2, \quad n'' = n - \frac{(N+1)N}{2} \end{cases} \quad (10)$$

C. Unitary change of basis

In Section III-A, we argued that $\mathfrak{u}(N)$ corresponds to a real vector space from its definition as a Lie algebra. As for any vector space, a change of basis can then be identified with a linear transformation, i.e., a matrix multiplication. Given the basis $\mathcal{B}$ defined in Section III-B, whose elements are given in (10), we can write for each $\mathbf{X} \in \mathfrak{u}(N)$

$$\mathbf{X} = \alpha_1 \mathbf{B}_1 + \dots + \alpha_{N^2} \mathbf{B}_{N^2}, \quad (11)$$

where $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_{N^2}]^T$ is the vector of real coordinates for $\mathbf{X}$. By vectorizing (11) we can rewrite it as

$$\text{vec}(\mathbf{X}) = \mathbf{U}_B \boldsymbol{\alpha}, \quad (12)$$

where $\mathbf{U}_B = [\text{vec}(\mathbf{B}_1), \dots, \text{vec}(\mathbf{B}_{N^2})]$. Since the Hilbert-Schmidt inner product coincides with the dot product of the vectorized input matrices, i.e., $\text{trace}(\mathbf{A}^H \mathbf{B}) = \text{vec}(\mathbf{A})^H \text{vec}(\mathbf{B})$, the columns $\{\text{vec}(\mathbf{B}_i)\}_{1 \leq i \leq N^2}$ of $\mathbf{U}_B$ are also orthonormal in the usual sense.⁶ Hence, the $N^2 \times N^2$ matrix $\mathbf{U}_B$ corresponds itself to a unitary transformation, i.e., $\mathbf{U}_B^H \mathbf{U}_B = \mathbf{U}_B \mathbf{U}_B^H = \mathbf{I}_{N^2}$. The N^2 real coordinates associated to a given $\mathbf{X} \in \mathfrak{u}(N)$ can be thus obtained as

$$\boldsymbol{\alpha} = \mathbf{U}_B^H \text{vec}(\mathbf{X}). \quad (13)$$

The unitary property of the considered change of basis will be helpful in the forthcoming analysis. However, if the basis $\mathcal{B}$ was defined in any other way such that orthonormality is not achieved, the change of basis in (13) would be applied by multiplying the inverse of $\mathbf{U}_B$ instead of its conjugate transpose. This may lead to higher computational complexity due to the matrix inversion, which is not required if $\mathbf{U}_B$ is unitary. However, since the same basis should be considered in any system employing this parametrization, said matrix inversion could be precomputed and reused indefinitely, incurring negligible complexity increase in the long run.

Remark 3: In some practical applications, e.g., in MIMO CSI feedback, or in quantum state characterization, a uni-modular scaling of the employed unitary matrix $\mathbf{U}$ may have no major impact on performance. For such cases, the first basis element in (10) may be ignored, and the resulting DEP would be reduced to $N^2 - 1$ dimensions. This is equivalent to parameterize instead the projection of $\mathbf{U}$ to the special unitary group $\mathcal{SU}(N)$, for which DEPs have been studied in [7]. Nevertheless, general applications, as the ones in Section V-B and V-C, may still require to resolve this uncertainty, preventing any further dimensionality reduction. Moreover, for

large enough N , the resulting dimensionality reduction from projecting onto $\mathcal{SU}(N)$ would be negligible.

Remark 4: In some applications we may need to share orthogonal matrices instead of unitary matrices, e.g., when sharing graph Fourier transform matrices as in [33]. If we consider only orthogonal matrices corresponding to rotation matrices (having determinant 1), the DEP may be achieved in the same way as the proposed unitary matrix DEP, but removing the first $N(N+1)/2$ basis elements in (10), associated to purely imaginary entries. This would reduce the resulting dimensions to $N(N-1)/2$ real dimensions. On the other hand, general orthogonal matrices (with determinant ± 1) may be projected to rotation matrices by scaling the last row with the determinant. This would only require sending and extra bit associated to the respective determinant apart from the rotation matrix DEP.

D. Full parametrization and retrieval of unitary matrices

If we combine the injective inverse exponential map, and the orthonormal basis inversion previously mentioned, we can define the DEP described in (1), which allows parameterizing an arbitrary unitary matrix through a set of N^2 real numbers. Specifically, we can define $\mathbf{f}$ for each $\mathbf{U} \in \mathcal{U}(N)$ as

$$\mathbf{f}(\mathbf{U}) = \mathbf{U}_B^H \text{vec}(\log(\mathbf{U})). \quad (14)$$

On the other hand, the injectivity of $\mathbf{f}$ is ensured by the fact that $\log$ is injective, as previously mentioned, while the linear map associated to the change of basis $\mathbf{U}_B^H$ is bijective since it corresponds to a full-rank linear transformation.

Assume that we have an arbitrary unitary matrix $\mathbf{U} \in \mathcal{U}(N)$ which was perfectly (without errors) stored/received using its DEP $\boldsymbol{\alpha} = \mathbf{f}(\mathbf{U})$. We can then retrieve the original matrix $\mathbf{U}$ from its inverse map $\mathbf{U} = \mathbf{f}^{-1}(\boldsymbol{\alpha})$. Considering (14), we may define the inverse map as

$$\mathbf{f}^{-1}(\boldsymbol{\alpha}) = \exp(\text{vec}^{-1}(\mathbf{U}_B \boldsymbol{\alpha})), \quad (15)$$

where $\text{vec}^{-1}(\cdot)$ corresponds to inverse vectorization, which reorganizes the vector input into an $N \times N$ matrix. By restricting the domain of $\mathbf{f}^{-1}$ to the image $\mathbf{f}(\mathcal{U}(N))$ of the whole space of unitary matrices under $\mathbf{f}$, we may further ensure bijectivity so that the inverse corresponds to a well-defined function. However, the explicit definitions of the complex matrix exponential and logarithm from (6) and (7), respectively, allow us to relax this formality. The reason is that (15) will always retrieve the original $\mathbf{U}$ regardless of the considered image of the complex logarithm (for a similar reason that $\exp(j(x + \ell 2\pi)) = \exp(jx)$, with $\ell \in \mathbb{Z}$, in

⁶Note that the dot product corresponds to the standard inner product for complex vector spaces.

the scalar case). This fact may be further understood after discussing boundedness in Section IV.

E. Time-complexity analysis

Getting timely access to a unitary matrix from its DEP, and viceversa, is of utmost importance for delay sensitive applications, as the ones considered in Section V. The following proposition determines the time-complexity scaling of the proposed DEP.

Proposition 1: Assuming unlimited hardware resources, the time-complexity for computing the proposed DEP given in (14), as well as its inverse map given in (15), is $\mathcal{O}(N^3)$.

Proof: Let us start considering (14). The required operations consist of a matrix multiplication between an $N^2 \times N^2$ matrix and an N^2 vector, as well as a matrix logarithm of the form (7). The total number of operations to perform the matrix multiplication $U_B^H \text{vec}(\mathbf{X})$ is $\mathcal{O}(N^4)$. However, if we have enough hardware resources, we may parallelize it by performing simultaneously (e.g., in different hardware modules) the multiplication of each element of the vector $\text{vec}(\mathbf{X})$ with all the elements of the respective column of U_B , leading to time-complexity $\mathcal{O}(1)$. The N^2 sums of the resulting N^2 -sized vectors can also be parallelized by summing simultaneously different entries of each vector, leading to time-complexity $\mathcal{O}(N^2)$. Thus, the time complexity for the matrix multiplication is $\mathcal{O}(N^2)$. For the matrix logarithm, the operations required are an EVD with $\mathcal{O}(N^3)$, a logarithm of the N eigenvalues with $\mathcal{O}(N)$, and two matrix multiplications: one $\mathcal{O}(1)$ (assuming parallelization) of an $N \times N$ diagonal matrix with a $N \times N$ matrix, and one $\mathcal{O}(N^3)$ of two $N \times N$ matrices. Hence, the overall time-complexity is $\mathcal{O}(N^3)$, which is dominated by the EVD (and matrix multiplication) to perform (7).

For the inverse map (15), the required operations are essentially the same as for (14), but taken in reverse order. The only difference is that the logarithm of the N diagonal elements now corresponds to an exponential with the same complexity. Hence, the overall time-complexity is $\mathcal{O}(N^3)$, which is dominated by the EVD (and matrix multiplication) to perform (6). $\square$

It should be noted that most of the applications of interest for this work require at least pre-computing the SVD of an $L \times N$ matrix, which has complexity $\mathcal{O}(\max(L, N)LN)$. Hence, for $L \geq N$, computing the proposed DEP (or its inverse) would not increase the overall time-complexity order. On the other hand, if we compare the complexity of the proposed DEP with approaches based on the Cayley transform, as proposed in [31], the time-complexity order would also be $\mathcal{O}(N^3)$ since this transform requires inverting an $N \times N$ matrix.

IV. BOUNDEDNESS OF THE DIMENSIONALLY-EFFICIENT PARAMETRIZATION

In the previous section, we explicitly defined an injective function $\mathbf{f}$ that maps an arbitrary unitary matrix $\mathbf{U} \in \mathcal{U}(N)$

to its DEP $\mathbf{f}(\mathbf{U}) \in \mathbb{R}^{N^2}$. Next, we will show that $\mathbf{f}$ may be fixed such that the DEP maintains the boundedness property, inherent to the unitary matrix space $\mathcal{U}(N)$. Specifically, we will show that the image of $\mathcal{U}(N)$ under $\mathbf{f}$ may be contained within a compact (i.e., closed and bounded) subspace of the form $[a, b]^{N^2} \subset \mathbb{R}^{N^2}$, so that we achieve (2).

Let us consider the inverse exponential map defined in (7) for an arbitrary $\mathbf{U} \in \mathcal{U}(N)$. As previously noted, the eigenvalues of a unitary matrix always lie on the complex unit circle [29, Section 4.5], [30]. We may thus express

$$\Lambda_{\mathbf{U}} = \text{diag}(\exp(j\varphi_1), \dots, \exp(j\varphi_N)), \quad (16)$$

where $\varphi_n \in \mathbb{R}$, $\forall n$. From the periodicity of the complex exponential, we may further consider the common restriction $\varphi_n \in (-\pi, \pi]$, $\forall n$. If we then apply the complex logarithm to the diagonal elements, we get

$$\log(\Lambda_{\mathbf{U}}) = \text{diag}(j(\varphi_1 + \ell_1 2\pi), \dots, j(\varphi_N + \ell_N 2\pi)), \quad (17)$$

where $\ell_n \in \mathbb{Z}$, $\forall n$. Note that the space of Skew-Hermitian matrices (correspondingly $\mathfrak{u}(N)$) coincides with the space of normal matrices with purely imaginary eigenvalues since, for $\mathbf{X} \in \mathfrak{u}(N)$, we have

$$\mathbf{X} + \mathbf{X}^H = \mathbf{U}_{\mathbf{X}}(\Lambda_{\mathbf{X}} + \Lambda_{\mathbf{X}}^*)\mathbf{U}_{\mathbf{X}}^H, \quad (18)$$

which equals $\mathbf{0}$ if and only if $\Lambda_{\mathbf{X}} = -\Lambda_{\mathbf{X}}^*$, i.e., if and only if the eigenvalues are purely imaginary. Thus, we could in principle span the whole $\mathfrak{u}(N)$ by considering the complete logarithm definition from (17) in (7) for each $\mathbf{U} \in \mathcal{U}(N)$. However, given that the idea is to eventually go back to the unitary group, we may restrict without loss the definition of the complex logarithm to its principal branch, such that the outputs lie always in the interval $(-\pi, \pi]$ of the imaginary axis. This corresponds to considering a logarithm definition as (17), but where we fix $\ell_n = 0$, $\forall n$. In the remainder, we assume that any system employing the proposed unitary matrix DEP considers the latter definition such that, $\forall \mathbf{U} \in \mathcal{U}(N)$, the eigenvalues of $\log(\mathbf{U})$ may be expressed as $\lambda_n(\log(\mathbf{U})) = j\varphi_n$, with $\varphi_n \in (-\pi, \pi]$, $\forall n$. Note that this definition also allows to have a well defined logarithm for unitary matrices $\mathbf{U}$ having some eigenvalues equal to -1 , which would give $\varphi_n = \pi$.

The following proposition stresses that the proposed unitary matrix DEP can ensure boundedness, while it provides a specific bound on the resulting dynamic range.

Proposition 2: Assume $\mathbf{f}$ is defined by (14), where $\log$ is defined by (7) and (17) with $\ell_n = 0$, $\forall n$. The image of $\mathcal{U}(N)$ under $\mathbf{f}$ is contained within the compact (bounded) space $[-\sqrt{N}\pi, \sqrt{N}\pi]^{N^2} \subset \mathbb{R}^{N^2}$. In other words, $\forall \mathbf{U} \in \mathcal{U}(N)$ we have that the entries of $\mathbf{f}(\mathbf{U})$ fulfill

$$f_i(\mathbf{U}) \in [-\sqrt{N}\pi, \sqrt{N}\pi], \quad \forall i \in \{1, \dots, N^2\}. \quad (19)$$

Proof: Using the logarithm definition from (17) with $\ell_n = 0$, $\forall n$, we may express (14) for any $\mathbf{U} \in \mathcal{U}(N)$ as

$$\begin{aligned} \mathbf{f}(\mathbf{U}) &= U_B^H \text{vec}(U_{\mathbf{U}} \text{diag}(j\varphi_1, \dots, j\varphi_N) U_{\mathbf{U}}^H) \\ &= U_B^H (U_{\mathbf{U}}^* \otimes U_{\mathbf{U}}) \text{vec}(\text{diag}(j\varphi_1, \dots, j\varphi_N)). \end{aligned} \quad (20)$$

Given the orthonormal basis defined in Section III-C, the product $\mathbf{U}_B^H(\mathbf{U}_U^* \otimes \mathbf{U}_U)$ corresponds to a product of unitary matrices, giving itself an $N^2 \times N^2$ unitary matrix. Hence, the squared norm of (20) is given by

$$\|\mathbf{f}(\mathbf{U})\|^2 \triangleq \mathbf{f}^H(\mathbf{U})\mathbf{f}(\mathbf{U}) = \sum_{n=1}^N \varphi_n^2. \quad (21)$$

Note that, although not explicitly seen from (20), the selected basis derived in Section III-C ensures that each entry, $f_i(\mathbf{U})$, of $\mathbf{f}(\mathbf{U})$ is real, $\forall i \in \{1, \dots, N^2\}$. Since the considered logarithm definition further ensures that $\varphi_n \in (-\pi, \pi]$, we can upper-bound (21) as

$$\|\mathbf{f}(\mathbf{U})\|^2 \triangleq \sum_{i=1}^{N^2} f_i^2(\mathbf{U}) \leq N\pi^2, \quad (22)$$

which trivially leads to

$$|f_i(\mathbf{U})| \leq \sqrt{N}\pi, \quad \forall i \in \{1, \dots, N^2\}. \quad (23)$$

The proof finalizes by noting that (23) is equivalent to (19). $\square$

With the results presented in this section, we can now proclaim that we have explicitly specified an injective function

$$\mathbf{f} : \mathcal{U}(N) \rightarrow \left[-\sqrt{N}\pi, \sqrt{N}\pi\right]^{N^2}, \quad (24)$$

as well as its inverse map. Thus, we can use these maps to unequivocally parameterize unitary matrices as a set of N^2 real bounded numbers. This may allow reducing the amount of information required to identify an arbitrary unitary matrix with respect to naively identifying it through its N^2 complex entries (i.e., $2N^2$ real entries). Moreover, the boundedness property also allows for the use of effective quantization with bounded error. Note, however, that the real and imaginary parts of the entries of an $N \times N$ unitary matrix are also bounded since they are contained within the interval $[-1, 1]$. Next, we emphasize the value of the proposed DEP by presenting several applications from the field of wireless communications which may benefit from it.

V. WIRELESS COMMUNICATIONS APPLICATIONS

One of the most relevant modern advancements in the field of wireless communications is the use of MIMO [34], which can significantly improve spectral efficiency by exploiting the degrees of freedom in the spatial domain. Much of the MIMO processing depends on performing operations involving unitary matrices, e.g., diagonalization of the channel matrix. These unitary matrices should be stored, and sometimes even shared among different modules. Next, we specify three potential applications related to the field of MIMO communications that showcase the benefit of the unitary matrix DEP proposed in this work. However, as previously outlined, the scope of this work may go beyond these applications, finding potential utility within diverse fields as image processing, quantum computation, computer science, etc.

A. CSI-Feedback to Achieve MIMO Capacity

Perhaps the most direct application of the proposed DEP is within MIMO CSI feedback, which is a topic that has been widely studied in the literature [9]–[12]. Let us consider an uplink transmission of a N -antenna user equipment (UE) towards an M -antenna base station (BS) ($M \geq N$) under a narrowband channel. The received vector at the BS is then governed by the renowned MIMO equation [34]

$$\mathbf{y} = \mathbf{H}\mathbf{s} + \mathbf{n}, \quad (25)$$

where $\mathbf{H}$ is the $M \times N$ complex channel matrix, $\mathbf{s}$ is the vector of baseband symbols transmitted by each of the UE's antennas, and $\mathbf{n} \sim \mathcal{N}_{\mathbb{C}}(\mathbf{0}, \mathbf{I}_M)$ is the noise vector. In order to achieve the channel capacity, the UE should access the orthogonal spatial streams by precoding its data according to the right unitary matrix of the SVD of the channel, which allows subsequently allocating power to each spatial stream according to the waterfilling algorithm [35]. However, in modern MIMO-based systems, the channel is typically assumed to be estimated via uplink pilots [36], which means that only the BS has access to the estimate of $\mathbf{H}$. Thus, a CSI-feedback strategy should be employed so that the UE can gain access to the relevant CSI required to achieve capacity.

Assuming that the BS performs the SVD of the channel matrix, given by $\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H$,⁷ the proposed DEP could be employed to reduce the information rate required to feedback the matrix $\mathbf{V}$ to the UE. In order to achieve capacity, we should also send the N real numbers associated to the power allocation, so that the total number of real dimensions to be feedbacked to the UE would be $N^2 + N$. Since the power allocation values may be given as ratios with respect to the total available power, the whole $N^2 + N$ real dimensions are associated to bounded quantities, which can be effectively quantized with limited reconstruction error. Thus, for large enough N , we could essentially halve the dimension of the information required to feedback these quantities with respect to sending the complete $\mathbf{V}$ matrix together with the power allocation values, giving instead $2N^2 + N$ bounded real dimensions. Another feedback option with comparable dimensional-efficiency could be to send the Gramian of the channel, upon which the UE would have to perform SVD independently, leading to the same time-complexity order as for retrieving the respective unitary matrix from the considered DEP. However, the elements of the Gramian matrix are not bounded in general, which may lead to greater quantization errors. Moreover, the UE would also have to find the optimum power allocations from the eigenvalues of the received Gramian, leading to higher power consumption at the device, which may compromise its battery life.

B. Fully-Connected RIS Configurations

RIS is a wireless communication technology which has gained much attention in the context 6G and beyond research.

⁷Note that performing this SVD at the BS is needed to be able to apply the linear equalizer $\mathbf{U}^H$ that allows achieving capacity [34].

Typical RIS implementations consist of a surface equipped with a large number of passive reconfigurable elements, which may change their reflection properties to improve wireless communication links. We may consider a multi-user MIMO (MU-MIMO) scenario where an M -antenna BS serves K single-antenna UEs. The uplink system model may be described through an input-output relation as (25), but where the channel matrix is given by

$$\mathbf{H} = \mathbf{H}_0 + \mathbf{H}_1 \mathbf{\Theta} \mathbf{H}_2, \quad (26)$$

where $\mathbf{H}_0$ is the $M \times K$ direct channel matrix between the UEs and the BS, $\mathbf{H}_1$ is the $M \times N$ channel matrix between the BS and the RIS, $\mathbf{\Theta}$ is the $N \times N$ reconfigurable reflection matrix, and $\mathbf{H}_2$ is the $N \times K$ channel matrix between the UEs and the RIS.

Recent works have considered extended RIS architectures by assuming that the RIS elements can be interconnected, offering a greater degree of reconfigurability [14], [15]. These architectures fall within the umbrella of BD-RIS since the associated reflection matrix $\mathbf{\Theta}$ may depart from the common diagonal reflection matrix with unimodular entries considered in the conventional RIS. In fact, if we consider fully-connected BD-RIS implementations with lossless passive impedance networks, the reflection matrix becomes unitary [14], [15]. Due to the passive nature of these surfaces, when they are employed in MU-MIMO scenarios, the BS is typically in charge of estimating the channel and finding the appropriate BD-RIS configuration [37], [38], which can be unequivocally identified with its reflection matrix. Thus, the BS should then employ a side-link to forward the respective configuration to the BD-RIS so that it can apply it. Assuming a BD-RIS with full-reconfigurability, and exploiting the unitary constraint on the reflection matrix, the proposed unitary matrix DEP may be employed to reduce the overhead required to forward the respective configuration from the BS (or generic receiver) to the BD-RIS.

Remark 5: The proposed DEP is especially useful when considering a fully-connected lossless non-reciprocal BD-RIS [16], also found under the term fully-reconfigurable intelligent surface (FRIS) [37]. The reason is that such system would pose no further constraints on the reflection matrix other than the unitary constraint. However, the proposed DEP could be adapted to take into account symmetry constraints (associated to reciprocal BD-RIS) to further reduce dimensionality. This may be simply done by eliminating some of the basis elements in (10). Specifically, it is enough to remove the $N(N-1)/2$ real anti-symmetric elements, since the matrix logarithm of a symmetric unitary matrix only has purely imaginary entries.⁸ The resulting DEP would then be reduced to $N(N+1)/2$ bounded real numbers. On the other hand, in case of further restrictions on the interconnection structure of the BD-RIS, e.g., leading to a block diagonal reflection matrix, the dimen-

sionality of the proposed DEP may be similarly reduced by removing other unnecessary basis elements from (10).

An alternative to sending the unitary matrix associated to the BD-RIS reflection matrix is to send the impedance values associated to the configuration leading to such reflection matrix. If we consider lossless non-reciprocal BD-RIS implementations based on circulators, as illustrated in [37] for FRIS, the respective impedance matrix may be linked to $2N^2 + N$ real numbers (scaled by the imaginary unit), which would thus require more than twice the dimensions of the proposed DEP. For the case of reciprocal BD-RIS, the impedance matrix would correspond to a symmetric purely imaginary matrix [15], giving $N(N+1)/2$ real variables as in the proposed DEP after removing the unnecessary basis elements from (10). However, due to the relation between the reflection matrix and the impedance matrix, which is given through the Cayley transform [3], the resulting impedance values would not fulfill the boundedness property. This would make them less suitable for transmission since they could lead to unbounded quantization errors, as previously noted.

C. Sharing of Decentralized Analog Beamformers

As happens with the reflection matrix of the lossless impedance networks used in BD-RIS architectures, the narrowband response of a passive and lossless analog beamformer can be identified with a unitary matrix associated to the corresponding impedance network [3]. An example of such is the Butler matrix, proposed as an analog beamformer based on phase shifters and power dividers [39]. In [21], a generalized hybrid beamforming architecture is considered where the analog beamforming may be performed in decentralized modules. The resulting processing is described through the WAX framework [40], [41], which allows exploiting the trade-off between decentralized processing complexity and level of decentralization. The post-processed vector is then given by

$$\mathbf{z} = \mathbf{X}^H \mathbf{A}^H \mathbf{W}^H \mathbf{y}, \quad (27)$$

where $\mathbf{y}$ is the received vector (as in (25)), $\mathbf{X}$ is the baseband processing, $\mathbf{A}$ is a fixed analog combining module, and $\mathbf{W} = \text{diag}(\mathbf{W}_1, \dots, \mathbf{W}_{M_p})$ is the analog beamforming matrix, where each $\mathbf{W}_m$ corresponds to an $N \times N$ unitary matrix associated to the analog beamforming applied at decentralized module m . This general scheme, and other related ones, potentially rely on finding the analog beamformers at a central node, associated to a baseband unit (BBU), which should then be forwarded to the decentralized modules where they are physically applied. Alternatively, these analog beamformers may be found iteratively by updating and sharing them among the different decentralized modules. The interconnection bandwidth required to forward the unitary matrices (associated to the analog beamformers) from the central node to the decentralized nodes, or to iteratively share them among the different decentralized modules, may be effectively reduced by employing the proposed unitary matrix DEP.

⁸Note that if, $\mathbf{U} = \mathbf{U}^T$, the eigenvalue decomposition in (7) has real eigenvector matrix since $\mathbf{U}_U = \mathbf{U}_U^*$, while the resulting eigenvalues are purely imaginary.

VI. NUMERICAL RESULTS

A. Analysis of Reconstruction Error

To understand the potential of the proposed DEP for efficient transmission/storage of unitary matrices, we begin by analyzing the error incurred when retrieving the unitary matrices from noisy observations of the DEP. A common metric for measuring the error in a reconstructed matrix is the mean squared error (MSE), which may be computed using the matrix Frobenius norm as

$$\text{MSE} = \frac{1}{N^2} \mathbb{E}\{\|\mathbf{U} - \hat{\mathbf{U}}\|_F^2\}. \quad (28)$$

where the expectation averages throughout different realizations of the original matrix $\mathbf{U}$ and its reconstructed version $\hat{\mathbf{U}}$, while, for each realization, the error is averaged throughout the N^2 matrix entries. Note that, although MSE is not directly tailored to unitary matrices, the squared distance induced by the standard Hilbert-Schmidt inner product for unitary matrices (or trace distance) has direct correspondence with the MSE defined in (28) through a scaling factor of $N/2$. Another way to analyze the reconstruction error is to look at the fidelity between unitary operators [42], which is widely used in the context of quantum gates [43]. We consider the following normalized definition

$$F(\mathbf{U}, \hat{\mathbf{U}}) = \frac{1}{N} |\text{Tr}(\mathbf{U}\hat{\mathbf{U}})|, \quad (29)$$

which gives a value between 0 and 1, where 1 corresponds to perfect reconstruction. Note that this definition of fidelity ignores the global phase uncertainty, further focusing on orthogonal space retrieval. For performance comparison, aside from the proposed DEP, we will consider a naive approach, consisting of transmitting the full $2N^2$ real dimensions associated to the matrix entries, as well as an improved version where the resulting matrix is projected back to the space of unitary matrices.⁹ We will also consider a more advanced DEP based on Givens rotations, similar to what is considered in [9], [10], [33], which may also be expressed in terms of N^2 bounded real dimensions: $N^2 - N$ for the normalized Givens rotations, given by phases and amplitudes smaller than 1, and N phase values to resolve the uncertainty associated to multiplying by a diagonal matrix of unimodular entries.

In Fig. 1 we plot the MSE (left) and average fidelity (right) when retrieving a unitary matrix from its noisy DEP under Gaussian noise. We have averaged the results over 10^4 realizations of an isotropic (uniformly distributed) $\mathbf{U} \in \mathcal{U}(N)$. The results assume that each DEP entry suffers from additive white Gaussian noise (AWGN), where the capacity per channel use has full correspondence with the reconstruction signal-to-noise ratio (SNR) through the renowned $\log_2(1 + \text{SNR})$ equation [13]. For the DEPs, we consider N^2 channel uses (one per dimension), while for the naive approaches we consider $2N^2$ channel uses (one per real/imaginary matrix entry) with half

the reported capacity, which corresponds to equally dividing the capacity between the real and imaginary parts. It can be seen that the proposed DEP offers the best reconstruction performance in terms of both MSE and average fidelity. The DEP based on Givens rotations gets the same MSE slope as the proposed DEP since they both achieve the same number of dimensions. However, the proposed DEP attains a lower MSE offset, especially as the number of dimensions increase. As for the naive approaches, including the projection back to $\mathcal{U}(N)$ reduces the reconstruction MSE as expected, since the extra noise falling outside the unitary group is removed. However, the same slope is maintained with respect to the naive approach without projection due to the unchanged number of dimensions. In terms of Fidelity, the proposed DEP converges quickly to full-fidelity, while the naive approach including projection attains a higher fidelity than the DEP based on Givens rotations. Note that the definition of fidelity only applies to unitary matrices, while the naive approach without projection may give a non-unitary reconstructed matrix.

In practical scenarios, the DEP would be stored/transmitted in the digital domain using simple quantizers. The most basic and widespread way to perform quantization is uniform scalar quantization [24], which consists of dividing the range of input values into a discrete set of equally-sized intervals. Fig. 2 shows the MSE (left) and average fidelity when retrieving 10^4 isotropic unitary matrices from their uniformly quantized DEP. The quantization resolution corresponds to the bits used to quantize each of the N^2 entries of the DEPs, while these bits are equally divided between the real and imaginary entries in the naive approaches. In all cases, the quantization ranges have been fixed to the known bounds of the input values, which were derived in Proposition 2 for the proposed DEP, for the Givens rotations based DEP some lie in the $[0, 1]$ interval and some in the $(-\pi, \pi]$ interval [9], and for the naive approaches all values are in the $[0, 1]$ interval. We have also included results under *quantization overrange*, which consist of allowing the quantizer inputs to have a dynamic range beyond the quantization range by a certain margin (in percentage or scaling). The input range for the proposed DEP is assumed to coincide with (19), and symmetry around the midpoint is always assumed.

As we can see from Fig. 2 (left), the proposed DEP and the Givens rotations based one both attain the same MSE slope, which is greater than for the naive approaches due to the reduced dimensions. However, the proposed DEP has slightly worse performance than the Givens rotations based one, with greater loss as the matrix size increases. However, as we allow for some overrange, the proposed DEP can attain similar performance to the Givens rotations based one, with more ground for improvement as the matrix size increases. Note that the interval derived in Proposition 2 for bounding the proposed DEP is not necessarily tight, and more research may be needed on how to further reduce it. However, a reasonable explanation is that this may come from the fact that the proposed DEP generates outputs that are more concentrated around the mean, while in the Givens based DEP the entries may have a

⁹Projecting an arbitrary matrix to $\mathcal{U}(N)$ requires computing the SVD, leading to the same time-complexity order required to retrieve a unitary matrix from the proposed DEP.




Fig. 1: MSE (left) and average fidelity (right) versus capacity per channel use for AWGN channel.

probability distribution closer to uniform. Specifically, from the way the boundedness interval is derived, if one value is close to the interval boundary, the rest of the entries will have to be very far from it. The same happens with the naive approaches, where, if one entry is close to 1, the rest of the entries associated to that column/row are likely to be close to 0 due to the norm constraint. Hence, as the number of entries grow, it is less likely to have entries close to the boundaries of the input range, allowing for some gains from quantization overrange. In the case of the Givens rotations based DEP, it seems that the slight overrange (of 5%) already starts to compromise performance, reducing the flexibility of this approach. Similar conclusions can be derived from the average fidelity metric in Fig. 2 (right), although in this case the Givens rotations based method seems less affected by overrange, while all approaches are negatively affected when increasing matrix size. In general, we can conclude that the proposed DEP can offer good performance when allowing for some overrange, while its performance may benefit hugely when considering more advanced quantization techniques, e.g., non-uniform quantization, or even vector quantization. Analyzing the distribution of the proposed DEP under various input distributions, and designing tailored quantization for it poses a research challenge in itself, which may be considered in future work. In fact, a simple adjustment of the uniform quantization taking into account estimates of the standard deviation of the different DEP entries could greatly improve the performance, e.g., considering the “four-sigma” rule of thumb to find good overrange factors [44].

B. MIMO CSI-feedback capacity ratio

We next analyze the performance of the considered approaches in the CSI-feedback use case presented in Section V-A. We define as performance metric the achievable capacity ratio, which corresponds to the average ratio $\mathbb{E}_{\mathbf{H}}\{R_{\text{ach}}/C\}$ of the channel capacity C (achievable through perfect waterfilling [35]) and the achievable rate R_{ach} when using the reconstructed unitary matrix for precoding at the

UE, together with the waterfilling power allocation. R_{ach} may be easily computed as $\sum_i \log_2(1 + \text{SINR}_i)$ from the resulting signal-to-interference-plus-noise ratio (SINR) per stream, since the BS assumes that the precoding is capable of orthogonalizing the streams. We have averaged the results throughout 10^3 realizations of a standard independent and identically distributed (IID) Rayleigh fading channel. We consider different number of UE antennas, but the number of BS antennas is fixed to $M = 32$ antennas, and the communication SNR is fixed to 10 dB, since these have no direct impact on the considered approaches. With respect to the feedback channel characteristics, the same consideration apply as for Figs. 1 and 2 for both the AWGN and the quantization cases. For the quantization case, we have further assumed that the waterfilling power allocation values are also quantized using the same resolution per entry.

In Fig. 3a we see that the proposed DEP attains the best performance under AWGN feedback channel, requiring only around 3-4 bit AWGN channels to attain the maximum capacity. This is roughly half the feedback channel capacity required for both the naive approach, as well as the Givens rotations DEP, to achieve maximum capacity. For the case with $N = 4$ UE antennas, the Givens based DEP can only outperform the naive approach for large enough feedback channel capacity, in agreement with what was seen in Fig. 1, while its performance is further degraded as the number of UE antennas increases to $N = 8$. In Fig. 3b we see again that, under uniform quantization, the proposed DEP can suffer from some performance loss with respect to the Givens rotations DEP. However, as we allow for some quantization overrange, the proposed DEP can even outperform the Givens rotations DEP for large enough N . This hints again at the fact that the proposed approach can greatly benefit from more tailored quantization schemes, which may be an interesting direction for future work. We have not included any plots for the Givens rotations DEP with quantization overrange to avoid cluttering the figure, since the slightest overrange already incurs some degradation in its performance.







Fig. 2: Reconstruction error versus quantization resolution per entry.



(a) AWGN feedback channel.



(b) Quantized feedback.

Fig. 3: Achievable capacity ratio versus quantization resolution for CSI feedback scenario with $M = 32$ BS antennas.

C. FRIS sum capacity ratio

Next, we analyze the performance of the previous approaches in the fully-connected BD-RIS use case described in Section V-B. We consider a scenario with a FRIS, or non-reciprocal BD-RIS, where the direct link is blocked leading to $\mathbf{H}_0 = \mathbf{0}$, since in this scenario we have a closed-form expression for the FRIS configuration achieving the sum capacity, as derived in [45]. However, the presented methods are also applicable to the case of having a direct channel since we consider the transmission of the full unitary configuration matrix. Moreover, the proposed DEP can be trivially adapted to the reciprocal BD-RIS scenario with symmetric reflection matrix, reducing the required overhead as described in Remark 5. The naive approach can also trivially incorporate this symmetry restriction, but for the Givens rotations DEP this corresponds to finding a subset of Givens rotations whose product can give an arbitrary symmetric unitary matrix. This is a highly non-trivial problem, so the practicality of the Givens rotations DEP when dealing with reciprocal BD-RIS scenarios (or other physical restrictions) is questionable.

In Fig. 4 we compare the achievable sum capacity ratio, which is the counterpart of the achievable capacity ratio for the multi-user case, when feedbacking the FRIS unitary capacity achieving configuration from [45]. We have averaged the results throughout 10^4 realizations of the cascaded channels $\mathbf{H}_1$ and $\mathbf{H}_2$, both modeled as standard IID Rayleigh fading channels. The SNR is fixed to 10 dB, the number of transmit antennas to $M_T = 8$, and the number of receive antennas to $M_R = 16$, while we considered $N \in \{8, 16\}$ FRIS elements, which determine the size of the unitary matrices to be sent. As we can see from Fig. 4a, the proposed DEP outperforms the other approaches by a great margin under the AWGN feedback channel, similar to what happened in the previous CSI feedback scenario. In Fig. 4b, we also see the same trend previously observed for the uniformly quantized feedback, where in this case the proposed DEP with overrange shows the highest potential since RIS-related scenarios are characterized by a large number of passive reconfigurable elements, which favor the proposed approach. Note that the overrange has been increased manually to the point before the performance starts

to degrade, but in future work we may gain more insights by studying the distribution of the entries of the proposed DEP.

D. Decentralized hybrid beamforming sum capacity ratio

We end by analyzing the performance of the considered approaches in the generalized decentralized hybrid beamforming scenario described in Section V-C, based on [21]. We assume that the scheme from [21] for computing the decentralized analog filters is implemented at the BBU, and the resulting unitary matrices are shared with the decentralized modules using the previous approaches. In this case, the achievable sum capacity ratio is computed with respect to the channel sum capacity, as in [21], and averaged over 10^3 realizations of an IID Rayleigh fading channel. This means that, even if the unitary matrices are perfectly reconstructed, the scheme may not converge to the full capacity unless the parameter combination allows for information-lossless processing [40].

In Fig. 5a we see again that the proposed DEP is the most efficient way to transmit unitary matrices under the AWGN channel, allowing to halve the required data rate with respect to the other approaches. In the case of quantized data from Fig. 5b, we see that, even for this small matrix sizes characterizing the size of the decentralized filters, the proposed DEP accepts a reasonable amount of overrange, which allows it to outperform the Givens rotations DEP.

VII. CONCLUSIONS

In this work, we have addressed the transmission and storage of unitary matrices by mapping these matrices to a sequence of bounded real numbers with minimum dimension. To this end, we have derived a DEP, and we have characterized an interval bounding the real numbers associated to this DEP. We have presented several applications from the field of wireless communications that could benefit from employing this technique. We have also shown how to adjust the DEP when considering extra matrix constraints that may appear in some of its applications, showcasing the adaptability and practicality of our approach. Finally, we have performed a numerical analysis under AWGN channels, as well as under uniform quantization, demonstrating the potential of the proposed approach in general applications, including the mentioned ones from wireless communications. In the case of uniform quantization, the proposed approach requires adjusting the quantization range allowing some overrange for the input values with respect to the derived bounds. In this sense, an interesting direction for future work is to study alternative quantization approaches for the proposed DEP so as to better exploit its potential gains.

REFERENCES

- [1] A. F. Molisch, *Wireless communications*. John Wiley & Sons, 2012, vol. 34.
- [2] M. Akay, *Biomedical signal processing*. Academic press, 2012.
- [3] D. Pozar, *Microwave Engineering, 4th Edition*. Wiley, 2011.
- [4] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.
- [5] A. Bisio, G. Chiribella, G. M. D'Ariano, S. Facchini, and P. Perinotti, "Optimal quantum learning of a unitary transformation," *Phys. Rev. A*, vol. 81, p. 032324, Mar 2010.
- [6] M. Sedláč, A. Bisio, and M. Ziman, "Optimal probabilistic storage and retrieval of unitary channels," *Phys. Rev. Lett.*, vol. 122, p. 170502, May 2019.
- [7] R. A. Bertlmann and P. Krammer, "Bloch vectors for qudits," *J. of Phys. A: Math. and Theor.*, vol. 41, no. 23, p. 235303, may 2008.
- [8] S. L. Hyland and G. Rättsch, "Learning unitary operators with help from $u(n)$," in *Proc. 31st AAAI Conf. on Artif. Intell.*, ser. AAAI'17. AAAI Press, 2017, p. 2050–2058.
- [9] M. A. Sadrabadi, A. K. Khandani, and F. Lahouti, "Channel feedback quantization for high data rate MIMO systems," *IEEE Trans. Wireless Commun.*, vol. 5, no. 12, pp. 3335–3338, 2006.
- [10] J. C. Roh and B. D. Rao, "Efficient feedback methods for MIMO channels based on parameterization," *IEEE Trans. Wireless Commun.*, vol. 6, no. 1, pp. 282–292, 2007.
- [11] S. Schwarz and M. Rupp, "Predictive quantization on the Stiefel manifold," *IEEE Signal Process. Lett.*, vol. 22, no. 2, pp. 234–238, 2015.
- [12] K. Kim, T. Kim, D. J. Love, and I. H. Kim, "Differential feedback in codebook-based multiuser MIMO systems in slowly varying channels," *IEEE Trans. Commun.*, vol. 60, no. 2, pp. 578–588, 2012.
- [13] T. M. Cover and J. A. Thomas, *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. USA: Wiley-Interscience, 2006.
- [14] S. Shen, B. Clerckx, and R. Murch, "Modeling and architecture design of reconfigurable intelligent surfaces using scattering parameter network analysis," *IEEE Trans. Wireless Commun.*, vol. 21, no. 2, pp. 1229–1243, 2022.
- [15] H. Li, S. Shen, and B. Clerckx, "Beyond diagonal reconfigurable intelligent surfaces: From transmitting and reflecting modes to single-, group-, and fully-connected architectures," *IEEE Trans. Wireless Commun.*, vol. 22, no. 4, pp. 2311–2324, 2023.
- [16] H. Li and B. Clerckx, "Non-reciprocal beyond diagonal RIS: Multiport network models and performance benefits in full-duplex systems," *IEEE Trans. Commun.*, pp. 1–1, 2025.
- [17] C. Huang, A. Zappone, G. C. Alexandropoulos, M. Debbah, and C. Yuen, "Reconfigurable intelligent surfaces for energy efficiency in wireless communication," *IEEE Trans. Wireless Commun.*, vol. 18, no. 8, pp. 4157–4170, 2019.
- [18] E. Basar, M. Di Renzo, J. De Rosny, M. Debbah, M.-S. Alouini, and R. Zhang, "Wireless communications through reconfigurable intelligent surfaces," *IEEE Access*, vol. 7, pp. 116 753–116 773, 2019.
- [19] M. Nerini and B. Clerckx, "Analog computing for signal processing and communications – part I: Computing with microwave networks," 2025. [Online]. Available: <https://arxiv.org/abs/2504.06790>
- [20] —, "Capacity of MIMO systems aided by microwave linear analog computers (MiLACs)," 2025. [Online]. Available: <https://arxiv.org/abs/2506.05983>
- [21] J. Vidal Alegria and O. Edfors, "Decentralized multi-antenna architectures with unitary constraints," in *Proc. 58th Asilomar Conf. on Signals, Syst., and Comput.*, 2024, pp. 1615–1619.
- [22] B. Hall, *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*, ser. Graduate Texts in Mathematics. Springer, 2003.
- [23] S. Gudmundsson, *An Introduction to Riemannian Geometry - Lecture Notes in Mathematics*. Lund University, 2021.
- [24] R. Gray and D. Neuhoff, "Quantization," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2325–2383, 1998.
- [25] J. Proakis and M. Salehi, *Digital Communications*, ser. McGraw-Hill International Edition. McGraw-Hill, 2008.
- [26] T. E. Abruđan, J. Eriksson, and V. Koivunen, "Steepest descent algorithms for optimization under unitary matrix constraint," *IEEE Trans. Signal Process.*, vol. 56, no. 3, pp. 1134–1147, 2008.
- [27] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1985.
- [28] T. Gamelin, *Complex analysis*, ser. Undergraduate Texts in Mathematics. Springer Science & Business Media, 2003.
- [29] C. A. Manogue and T. Dray, "The geometry of linear algebra," 2017, [Online; accessed 2-February-2025]. [Online]. Available: <https://books.physics.oregonstate.edu/LinAlg/book-1.html>
- [30] A. Eremenko, "Spectral theorems for hermitian and unitary matrices," *Purdue University, USA*, 2017.




Fig. 4: Achievable sum capacity ratio for FRIS control feedback scenario with $M = 16$ BS antennas, $K = 8$ UEs, and $N = 16$ BD-RIS elements.




Fig. 5: Achievable sum capacity ratio for decentralized analog beamformer sharing in [21] with $M = 12$ BS antennas, $K = 6$ UEs, and $T = 8$ inputs to the BBU.

- [31] B. Hassibi and B. Hochwald, "Cayley differential unitary space-time codes," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1485–1503, 2002.
- [32] A. M. E. Br uning, H. M kel  and F. Petruccione, "Parametrizations of density matrices," *J. of Modern Opt.*, vol. 59, no. 1, pp. 1–20, 2012.
- [33] S. Ueno, T. Fujihashi, T. Koike-Akino, and T. Watanabe, "Point cloud soft multicast for untethered XR users," *IEEE Trans. Multimedia*, vol. 25, pp. 7185–7195, 2023.
- [34] A. Paulraj, R. Nabar, and D. Gore, *Introduction to Space-Time Wireless Communications*, 1st ed. USA: Cambridge Univ. Press, 2008.
- [35] E. Telatar, "Capacity of multi-antenna gaussian channels," *Eur. Trans. on Telecommun.*, vol. 10, no. 6, pp. 585–595, 1999.
- [36] E. Bj rnson, J. Hoydis, and L. Sanguinetti, "Massive MIMO networks: Spectral, energy, and hardware efficiency," *Found. and Trends® in Signal Process.*, vol. 11, no. 3-4, pp. 154–655, 2017.
- [37] J. Vidal Alegr a, J. Thunberg, and O. Edfors, "Channel orthogonalization with reconfigurable surfaces: General models, theoretical limits, and effective configuration," *IEEE Trans. Wireless Commun.*, 2025.
- [38] A. L. F. de Almeida, B. Sokal, H. Li, and B. Clerckx, "Channel estimation for beyond diagonal RIS via tensor decomposition," *IEEE Trans. Signal Process.*, pp. 1–15, 2025.
- [39] J. L. Butler and R. Lowe, "Beam-forming matrix simplifies design of electronically scanned antennas," *Electron. Des.*, vol. 9, pp. 170–173, 1961.
- [40] J. V. Alegr a, F. Rusek, and O. Edfors, "Trade-offs in decentralized multi-antenna architectures: The WAX decomposition," *IEEE Trans. Signal Process.*, vol. 69, pp. 3627–3641, 2021.
- [41] J. Vidal Alegr a and F. Rusek, "Trade-offs in decentralized multi-antenna architectures: Sparse combining modules for WAX decomposition," *IEEE Trans. Signal Process.*, vol. 71, pp. 2879–2894, 2023.
- [42] R. Cabrera, O. M. Shir, R. Wu, and H. Rabitz, "Fidelity between unitary operators and the generation of robust gates against off-resonance perturbations," *Journal of Physics A: Mathematical and Theoretical*, vol. 44, no. 9, p. 095302, feb 2011.
- [43] R. J. Spiteri, M. Schmidt, J. Ghosh, E. Zahedinejad, and B. C. Sanders, "Quantum control for high-fidelity multi-qubit gates," *New J. of Physics*, vol. 20, no. 11, p. 113009, nov 2018.
- [44] W. R. Bennett, "Spectra of quantized signals," *The Bell System Technical J.*, vol. 27, no. 3, pp. 446–472, 1948.
- [45] G. Bartoli, A. Abrardo, N. Decarli, D. Dardari, and M. Di Renzo, "Spatial multiplexing in near field MIMO channels with reconfigurable intelligent surfaces," *IET Signal Process.*, vol. 17, no. 3, p. e12195, 2023.