

Sphere decoding complexity exponent for decoding full-rate codes over the quasi-static MIMO channel

Joakim Jaldén and Petros Elia

Abstract—In the setting of quasi-static multiple-input multiple-output (MIMO) channels, we consider the high signal-to-noise ratio (SNR) asymptotic complexity required by the sphere decoding (SD) algorithm for decoding a large class of full-rate linear space-time codes. With SD complexity having random fluctuations induced by the random channel, noise and codeword realizations, the introduced *SD complexity exponent* manages to concisely describe the computational reserves required by the SD algorithm to achieve arbitrarily close to optimal decoding performance. Bounds and exact expressions for the SD complexity exponent are obtained for the decoding of large families of codes with arbitrary performance characteristics. For the particular example of decoding the recently introduced threaded cyclic division algebra (CDA) based codes – the only currently known explicit designs that are uniformly optimal with respect to the diversity multiplexing tradeoff (DMT) – the SD complexity exponent is shown to take a particularly concise form as a non-monotonic function of the multiplexing gain. To date, the SD complexity exponent also describes the minimum known complexity of any decoder that can provably achieve a gap to maximum likelihood (ML) performance that vanishes in the high SNR limit.

Index Terms—Diversity-Multiplexing Tradeoff, Sphere Decoding, Complexity, Space-Time Codes, Large Deviations.

I. INTRODUCTION

The past decade has seen the abundant use of the sphere decoding (SD) algorithm [1]–[5] as a tool for facilitating near maximum likelihood (ML) decoding for multiple-input multiple-output (MIMO) channels. The SD algorithm allows for efficient optimal or near optimal decoding of a large number of high rate space-time codes that map constituent constellation symbols linearly in space and time [1]. However, as the SD algorithm has a channel-dependent run-time, any time-limited implementation will have to declare computational outages that inevitably increase the overall probability of decoding error. This naturally raises the intriguing question of how large computational reserves are actually required for near ML performance.

While this question is hard to answer in general, or even ask in a rigorously meaningful way, we show herein that by

The research leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 228044. J. Jaldén acknowledges funding from the Swedish Foundation for Strategic Research (SSF) under the project grant ICA08-0046. P. Elia acknowledges funding by the Mitsubishi RD project Home-eNodeBS, from FP7 CELTIC SPEC-TRA project and from Agence Nationale de la Recherche project ANR-IMAGENET. A shortened version of the work was presented at the IEEE International Symposium on Information Theory (ISIT-2011).

J. Jaldén is with the ACCESS Linnaeus Center, Signal Processing Lab, School of Electrical Engineering, KTH Royal Institute of Technology, Stockholm, Sweden (email: jalden@kth.se)

P. Elia is with the Mobile Communications Department, EURECOM, Sophia Antipolis, France (email: elia@eurecom.fr)

following [6] and considering the decoding of sequences of codes in the high signal-to-noise ratio (SNR) limit, not only can the question be made rigorous: It also admits surprisingly simple explicit and general answers. Drawing from the diversity multiplexing tradeoff (DMT) – which has also previously been used to assess the reliability of reduced complexity decoders [7]–[9] – we introduce the *SD complexity exponent* as a measure of complexity of the SD algorithm. This complexity measure naturally takes into account the dependency of the SD complexity on the codeword density and the codebook size, as well as on the SNR and the fading characteristics of the wireless channel. Similar to previous work on the DMT relating the code rate and probability of decoding error, it is seen that also the complexity, although hard to characterize at any finite SNR, has mathematically tractable characterizations in the high SNR asymptote. These characterizations in turn yield valuable insights into the behavior of the algorithm.

The SD algorithm is a branch-and-bound search [4] over a regular tree and like most other works on SD complexity [2], [4], [10]–[12] we view the number of visited nodes N as the complexity of the algorithm¹. To identify an appropriate scale of interest for complexity in the DMT setting, it is useful to note that in order to achieve a multiplexing gain of r the code must in the high SNR limit have rate of² $R = r \log \rho + o(\log \rho)$ bits per channel use where ρ denotes the SNR. Consequently, the cardinality of the codebook is $|\mathcal{X}| \doteq \rho^{rT}$ where T is the codeword length and where \doteq denotes equality in the SNR exponent (cf. [6] and Section I-B). In the worst case, as will be shown later, the sphere decoder is in essence forced to perform a complete search over the entire codebook, and its complexity is in this case ρ^{rT} . However, although feasible, this event is also highly improbable. Thus, in order to quantify the probability that the sphere decoder visits a large number of nodes N , we introduce a (complexity) rate-function $\Psi(x)$ over $0 \leq x \leq rT$ such that $P(N \geq \rho^x) \doteq \rho^{-\Psi(x)}$. The function $\Psi(x)$ captures the decay-rate of the probability that the complexity exceeds a given SNR dependent threshold ρ^x . This decay-rate should be contrasted with the minimum probability of decoding error, which vanishes in the high SNR limit as $\rho^{-d(r)}$ where $d(r)$ is the diversity gain of the code under maximum likelihood (ML) decoding. We can thus judiciously argue that for any x such that $\Psi(x) > d(r)$, the probability that the complexity exceeds ρ^x is at high SNR insignificant in comparison to the overall probability of error

¹In the context of the DMT this is, as we argue later on, equivalent to measuring complexity in floating point operations (flops).

²Herein, \log denotes the base-2 logarithm and $o(\cdot)$ is the standard Landau notation where $f(\rho) = o(\phi(\rho))$ implies $\lim_{\rho \rightarrow \infty} f(\rho)/\phi(\rho) = 0$. Similarly, $f(\rho) = O(\phi(\rho))$ implies that $\limsup_{\rho \rightarrow \infty} |f(\rho)|/\phi(\rho) < \infty$ for $\phi(\rho) > 0$.

of the decoder. In other words: for x such that $\Psi(x) > d(r)$, imposing a run-time limit of ρ^x on the complexity of the algorithm – and declaring a decoding outage whenever this limit is not met – would cause a vanishing degradation in terms of the overall error probability at high SNR. This motivates us to deviate from the traditional worst-case complexity measure that fails to meaningfully describe the effective complexity, and to define the SD complexity exponent $c(r)$ as the infimum of all x for which $\Psi(x) > d(r)$. In essence, $\rho^{c(r)}$ represents the minimum computational reserves required for achieving DMT optimal performance using the SD algorithm. The main topic of this work is to give closed form expressions, and bounds, for the SD complexity exponent $c(r)$ when decoding different classes of full-rate linear codes, to be described later, including the codes proposed in [13]–[18].

Most prior works on sphere decoding complexity have considered (coherent) uncoded spatial multiplexing (SM), and asymptotic results in terms of the signal space dimension, see, e.g., [10], [11], [19], [20]. A notable exception is [21], which incorporates space-time codes into the SD complexity analysis of a multiple-symbol differentially encoded MIMO system, and give results that are asymptotic in the transmission block length. Our present work is most closely related to [12], which considers the SD complexity tail distribution for coherent SM with a fixed signal space dimension. However, unlike [12] we also incorporate the space-time codes into the analysis, as well as the SNR scalings of the rates of these codes mandated by the DMT. We also take our analysis one step further by coupling the complexity tail distribution to the DMT performance of the code in order to obtain the SD complexity exponent.

A. Outline and contributions

The general definition of the SD complexity exponent is given in Definition 1 and Theorem 1 then describes how sphere decoding and the time-out policies to be employed can guarantee a gap to ML that vanishes with increasing SNR. However, before proceeding with the statement of these results, we first consider the code-channel system, describe the basic workings of the SD algorithm, and handle different pertinent aspects that are necessary for the exposition that follows.

Following the definition of the SD complexity exponent, Theorem 2 gives, in the form of an optimization problem, a general upper bound $\bar{c}(r)$ on the SD complexity exponent $c(r)$ when decoding any full-rate code with multiplexing gain r and diversity $d(r)$. An explicit closed form expression for $\bar{c}(r)$ is then given in Theorem 3 for all DMT optimal full-rate codes. The bound $\bar{c}(r)$ is already useful in itself in that it establishes that the SD complexity exponent is much lower than the worst-case SNR exponent rT associated with a full search of the codebook. However, in the interest of also establishing the tightness of this bound, Lemma 2 provides easy-to-check sufficient conditions on the generator matrix of the code lattice, that guarantee the tightness of $\bar{c}(r)$ in the most general setting. Building on this, Theorem 4 establishes that, given any full-rate design of arbitrary DMT performance, there is always at least one non-random SD column ordering

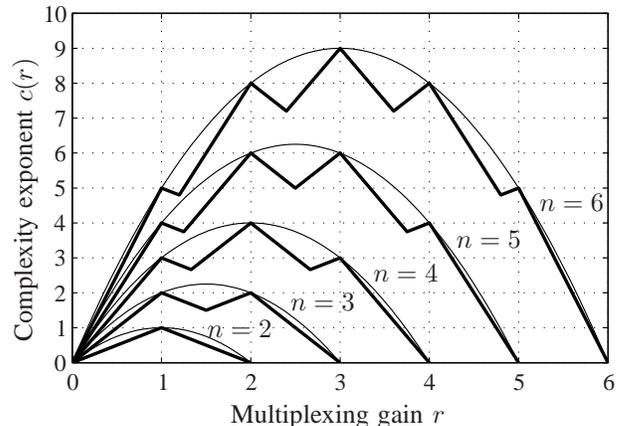


Fig. 1. The SD complexity exponent $c(r)$ for decoding threaded minimum delay DMT optimal codes with $n_T = T = n$ for $n = 2, \dots, 6$. The SD complexity exponent is illustrated by the bold lines. The same exponent also serves as an upper bound to the SD complexity exponent when decoding any minimum delay DMT optimal full-rate linear dispersive code. The thin lines show the quadratic function given by $r(n-r)$ which provides the exact complexity exponent at integer multiplexing gains.

[3], [4] for which $c(r) = \bar{c}(r)$, i.e., for which the exact $c(r)$ can be explicitly calculated from the result of Theorem 2. Theorem 5 goes one step further and establishes the exact SD complexity exponent, given any *threaded* code design and the natural column ordering, to be $c(r) = \bar{c}(r)$ while Theorem 6 provides an explicit expression for $c(r)$ for any DMT optimal threaded code design. Surprisingly, this simple expression (see Fig. 1) can also serve as an upper bound on $c(r)$ for any full-rate code, irrespective of the fading statistics. Finally, and along a different path, Theorem 7 establishes $c(r)$ for any 2×2 approximately universal code [22], irrespective of its specific structure, thus identifying the exact $c(r)$ even for possibly undiscovered code structural designs, as long as these designs are approximately universal, i.e., as long as they achieve DMT optimality for all fading statistics. Some general discussions of these results are then provided in Section V.

The SD complexity exponent for decoding the class of full-rate threaded DMT optimal codes with minimum delay, i.e., for which $n_T = T = n$ where n_T denotes the number of transmit antennas, is shown³ in Fig. 1 for $n = 2, \dots, 6$. The results shown in the figure apply to codes such as those presented in [13]–[18]. Before proving the aforementioned results, it is worth commenting on the somewhat counterintuitive result suggested by the SD complexity exponent in Fig. 1. Namely, while $c(r)$ initially increases as a function of the multiplexing gain r , it then decreases as r approaches its maximum value n_T . The initial increase can easily be explained by the fact that the cardinality (and density) of the codebook \mathcal{X} increases as a function of r . However, the decrease at high multiplexing gains can be understood in light of the coupling of the complexity with the overall probability of error: In short, at high multiplexing gains the error probability is also higher and this implies that the decoder

³A closed form expression for the complexity exponent $c(r)$ shown in Fig. 1 is given by (50) in Theorem 6.

may time out for a larger set of problem instances without seriously affecting the overall performance, leading to an overall reduction in decoding complexity. Pushing the code-decoder pair towards the maximal data rate does therefore not imply that the decoding complexity is maximized. This effect is discussed further in Section V-A, in terms of information theoretic outages.

B. Notation

We let \mathbb{Z} , \mathbb{R} , and \mathbb{C} , denote the set of integer, real and complex numbers respectively and \mathbb{F}^n and $\mathbb{F}^{m \times n}$ the set of n -vectors and $m \times n$ -matrices over $\mathbb{F} \in \{\mathbb{Z}, \mathbb{R}, \mathbb{C}\}$. Vectors are denoted by lower-case bold letters \mathbf{a} , and matrices are denoted by upper-case bold letters \mathbf{A} . We use $(\cdot)^T$ and $(\cdot)^H$ to denote the transpose and Hermitian (conjugate) transpose of vectors and matrices, and $\text{vec}(\cdot) : \mathbb{C}^{m \times n} \mapsto \mathbb{C}^{mn}$ to denote the matrix to vector operation whereby the columns of the argument are stacked on top of each other. We use $\mathbf{I}_n \in \mathbb{C}^n$ to denote the $n \times n$ identity matrix, and use $\mathbf{0}$ to denote the zero vector or matrix where the dimensions are given by the context in which $\mathbf{0}$ is used. Deviating slightly from standard usage, we refer to a tall matrix $\mathbf{U} \in \mathbb{C}^{m \times n}$ where $m \geq n$ as unitary if $\mathbf{U}^H \mathbf{U} = \mathbf{I}$. For $a \in \mathbb{C}$ we use $\Re(a)$ and $\Im(a)$ to denote the real and imaginary parts of a . For $a \in \mathbb{R}$ we use $\lfloor a \rfloor$ to denote the floor operation defined as the largest integer smaller than or equal to a , $\lceil a \rceil$ to denote the ceil operation defined as the smallest integer larger than or equal to a , and we let $(a)^+ \triangleq \max(a, 0)$.

We let $\sigma_1(\mathbf{A}) \leq \dots \leq \sigma_n(\mathbf{A})$ denote the ordered (structurally) non-zero singular values [23] of a matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ where $m \geq n$. We will on occasion use $\sigma_{\max}(\mathbf{A})$ to denote the largest singular value when the dimension of \mathbf{A} is not explicit. We make no notational difference between random variables and their realizations. We will use the notation $\Omega \triangleq \{\dots\}$ to label the stochastic event within the brackets.

Finally, in order to simplify notation we will make use of the \doteq (and \leq , \geq , $<$, $>$) notation for equalities (and inequalities) in the SNR exponent, cf. [6]. Specifically, we write $f(\rho) \leq \rho^a$ and $g(\rho) \geq \rho^b$ to denote

$$\limsup_{\rho \rightarrow \infty} \frac{\log f(\rho)}{\log \rho} \leq a \quad \text{and} \quad \liminf_{\rho \rightarrow \infty} \frac{\log g(\rho)}{\log \rho} \geq b$$

and $f(\rho) \doteq \rho^x$ when $f(\rho) \leq \rho^x$ and $f(\rho) \geq \rho^x$ simultaneously hold. The definition of $<$ and $>$ follows after replacing \leq by $<$ and \geq by $>$ in the above expressions.

II. CHANNEL MODEL AND SPACE-TIME CODES

We consider the standard block Rayleigh fading $n_T \times n_R$ quasi-static point-to-point MIMO channel model with coherence-time T given by

$$\mathbf{Y} = \mathbf{H}\mathbf{X} + \mathbf{W} \quad (1)$$

where $\mathbf{X} \in \mathbb{C}^{n_T \times T}$, where $\mathbf{Y} \in \mathbb{C}^{n_R \times T}$, and where $\mathbf{W} \in \mathbb{C}^{n_R \times T}$ denote the transmitted space-time block codeword, the block of received signals, and additive spatially and temporally white Gaussian noise. The channel gains $\mathbf{H} \in \mathbb{C}^{n_R \times n_T}$ are

assumed to be i.i.d. circularly symmetric complex Gaussian⁴ (i.e., Rayleigh fading) and constant over the duration of the transmission (i.e., quasi-static Rayleigh fading). We shall assume throughout that $n_R \geq n_T$. The transmitted codewords \mathbf{X} are assumed to be drawn uniformly from a codebook \mathcal{X} and we assume that

$$\mathbb{E}\{\|\mathbf{X}\|_{\mathbb{F}}^2\} = \frac{1}{|\mathcal{X}|} \sum_{\mathbf{X} \in \mathcal{X}} \|\mathbf{X}\|_{\mathbb{F}}^2 = \rho T, \quad (2)$$

so that the parameter ρ takes on the interpretation of an average SNR. Note also here that one use of (1) is viewed as T uses of the wireless channel in the definition of the data-rate.

We shall herein restrict our attention to full-rate (complex) linear dispersive codes [24], [25] of the form

$$\mathbf{X} = \theta \sum_{i=1}^{\kappa} s_i \mathbf{D}_i \quad (3)$$

where $s_i \in \mathbb{S} \subset \mathbb{C}$ are constellation symbols drawn from a finite alphabet \mathbb{S} , where $\{\mathbf{D}_i\}_{i=1}^{\kappa}$ is a set of linearly independent *dispersion matrices*, and where θ is a parameter regulating the transmit power. The notion of full-rate implies that each codeword $\mathbf{X} \in \mathbb{C}^{n_T \times T}$ carries $\kappa = n_T T$ constellation symbols. We will further make the additional assumption that \mathbb{S} belongs in the class of QAM-like alphabets of the form

$$\mathbb{S} = \mathbb{S}_{\eta} \triangleq \{s \mid \Re(s), \Im(s) \in \mathbb{Z} \cap [-\eta, \eta]\} \quad (4)$$

where $\eta > 0$ is a parameter regulating the size of the constellation⁵. We will use \mathbb{S}_{∞} to denote the extended (infinite) constellation obtained by letting $\eta = \infty$ in (4), and note that \mathbb{S}_{∞} is nothing but the set of Gaussian integers. QAM constellations will in general also include a translation and scaling of the underlying lattice \mathbb{S}_{∞} . However, as including such a translation would not affect the results obtained herein, and as the scaling can easily be included in the dispersion matrices \mathbf{D}_i or θ , we omit these variations and concentrate on (4) in the interest of notational simplicity.

The channel model in (1) may also be equivalently expressed in a vectorized form according to

$$\mathbf{y} = (\mathbf{I}_T \otimes \mathbf{H})\mathbf{x} + \mathbf{w} \quad (5)$$

where $\mathbf{y} = \text{vec}(\mathbf{Y})$, where $\mathbf{x} = \text{vec}(\mathbf{X})$, where $\mathbf{w} = \text{vec}(\mathbf{W})$, and where \otimes denotes the Kronecker product [26]. We shall mainly work with (5) rather than (1) directly. In the vectorized form the codewords \mathbf{x} are given by

$$\mathbf{x} = \theta \mathbf{G}\mathbf{s}$$

for $\mathbf{s} \in \mathbb{S}_{\eta}^{\kappa}$, and where the full rank matrix

$$\mathbf{G} = [\text{vec}(\mathbf{D}_1) \quad \dots \quad \text{vec}(\mathbf{D}_{\kappa})] \quad (6)$$

⁴This assumption is relaxed in Section V-B where the extension to arbitrary fading distributions is discussed.

⁵The assumption of a square constellation is made here for simplicity of exposition and in line with practical encoding schemes. This assumption though can readily be relaxed without affecting the presented results, as long as the constellation is the same for all s_i , $i = 1, \dots, \kappa$. A detailed exposition of the mathematical machinery that allows for this relaxation can be found in [9, Section III]

is referred to as the *generator matrix* of the code. The linear dispersive codes are lattice codes [25] as the codewords constitute a subset of the (complex) lattice $\theta\mathbf{G}\mathbb{S}_\infty$.

The parameters θ and η are, as noted, chosen in order to satisfy given transmit power and rate constraints. In particular, in order to ensure a multiplexing gain of

$$r \triangleq \lim_{\rho \rightarrow \infty} \frac{1}{T} \frac{\log |\mathcal{X}|}{\log \rho}, \quad (7)$$

or equivalently a rate of $R = r \log \rho + o(\log \rho)$, it must hold that $\eta \doteq \rho^{\frac{T}{2\kappa}}$ which by (2) and (4) implies that $\theta^2 \doteq \rho^{1 - \frac{T}{\kappa}}$. The code structure described above includes the codes proposed in [13]–[16], [18], as well as the QAM-based codes of [17], as special cases. Finally, we will throughout, with slight abuse of terminology but still in line with [13]–[18], use the term *code* when referring to the whole *family of codes* that is generated by a single generator matrix \mathbf{G} for different multiplexing gains and SNRs, and trust that no confusion should follow by this usage.

III. DECODING

The coherent ML decoder for (1) is well known to be

$$\hat{\mathbf{X}}_{\text{ML}} = \arg \min_{\hat{\mathbf{X}} \in \mathcal{X}} \|\mathbf{Y} - \mathbf{H}\hat{\mathbf{X}}\|_{\text{F}}^2. \quad (8)$$

The resulting diversity gain of the code, under ML decoding, is correspondingly given by (cf. [6])

$$d(r) \triangleq - \lim_{\rho \rightarrow \infty} \frac{\log \text{P}(\hat{\mathbf{X}}_{\text{ML}} \neq \mathbf{X})}{\log \rho}, \quad (9)$$

where the notation $d(r)$ accentuates the dependence of the diversity on the multiplexing gain r .

One of the main features of the linear dispersive codes, as was noted in the introduction, is that their lattice structure allows for efficient – optimal and near optimal – solutions to (8) using the sphere decoding algorithm. Using the linearity of the map from \mathbf{s} to $\mathbf{x} = \text{vec}(\mathbf{X})$ we obtain (cf. (5))

$$\mathbf{y} = \mathbf{M}\mathbf{s} + \mathbf{w} \quad (10)$$

where the code-channel generator matrix \mathbf{M} is given by

$$\mathbf{M} \triangleq \theta(\mathbf{I}_T \otimes \mathbf{H})\mathbf{G} \in \mathbb{C}^{n_{\text{R}}T \times \kappa}. \quad (11)$$

We can thus, instead of solving (8) directly, equivalently obtain an estimate of \mathbf{s} through

$$\hat{\mathbf{s}}_{\text{ML}} = \arg \min_{\hat{\mathbf{s}} \in \mathbb{S}_\eta^\kappa} \|\mathbf{y} - \mathbf{M}\hat{\mathbf{s}}\|^2, \quad (12)$$

where (12) is an optimization problem suitable for the sphere decoder, and then easily recover $\hat{\mathbf{X}}_{\text{ML}}$ from $\hat{\mathbf{s}}_{\text{ML}}$.

A. The Sphere Decoder

The sphere decoding algorithm solves (12) by a branch-and-bound search on a regular tree. Detailed descriptions of the algorithm are found in [1] and the semi-tutorial papers [2]–[4], and most implementation issues will not be repeated herein. However, in order to make our results precise and to introduce notation we need to review some of the key ideas as they apply to (12).

By the rotational invariance of the Euclidean norm it follows that (12) is equivalent to

$$\hat{\mathbf{s}}_{\text{ML}} = \arg \min_{\hat{\mathbf{s}} \in \mathbb{S}_\eta^\kappa} \|\mathbf{r} - \mathbf{R}\hat{\mathbf{s}}\|^2 \quad (13)$$

where $\mathbf{Q}\mathbf{R} = \mathbf{M}$ is the thin QR-decomposition of \mathbf{M} (i.e., $\mathbf{Q} \in \mathbb{C}^{n_{\text{R}}T \times \kappa}$ is unitary and $\mathbf{R} \in \mathbb{C}^{\kappa \times \kappa}$ is upper triangular) and where $\mathbf{r} = \mathbf{Q}^H \mathbf{y}$. The sphere decoder solves (13) by enumerating symbol vectors $\hat{\mathbf{s}} \in \mathbb{S}_\eta^\kappa$ within a given sphere of radius $\xi > 0$, i.e., $\hat{\mathbf{s}}$ that satisfy

$$\|\mathbf{r} - \mathbf{R}\hat{\mathbf{s}}\|^2 \leq \xi^2. \quad (14)$$

The set of vectors that satisfy (14) is found by recursively enumerating partial symbol vectors $\hat{\mathbf{s}}_k \in \mathbb{S}_\eta^k$ for $k = 1, \dots, \kappa$. Specifically, if $\hat{\mathbf{s}}_k$ is the vector containing the last k components of $\hat{\mathbf{s}}$, a necessary condition for (14) to be satisfied is given by

$$\|\mathbf{r}_k - \mathbf{R}_k \hat{\mathbf{s}}_k\|^2 \leq \xi^2, \quad (15)$$

where $\mathbf{r}_k \in \mathbb{C}^k$ denotes the last k components of \mathbf{r} , and where $\mathbf{R}_k \in \mathbb{C}^{k \times k}$ denotes the $k \times k$ lower right corner of \mathbf{R} .

The enumeration of partial symbol vectors \mathbf{s}_k is equivalent to the traversal of a regular tree with κ layers – one per symbol s_k where s_k is the k th component of \mathbf{s} – and $|\mathbb{S}_\eta|$ children per node [4]. There is a one-to-one correspondence between the nodes at layer k (the layers are enumerated with the root node corresponding to $k = 0$) and the partial vectors \mathbf{s}_k . We say that a node is visited by the sphere decoder if and only if the corresponding partial vector \mathbf{s}_k satisfies (15), i.e., there is a bijection between the visited nodes at layer k and the set

$$\mathcal{N}_k \triangleq \{ \hat{\mathbf{s}}_k \in \mathbb{S}_\eta^k \mid \|\mathbf{r}_k - \mathbf{R}_k \hat{\mathbf{s}}_k\|^2 \leq \xi^2 \}. \quad (16)$$

The total number of visited nodes is given by

$$N = \sum_{k=1}^{\kappa} N_k, \quad (17)$$

where $N_k \triangleq |\mathcal{N}_k|$ is the number of visited nodes at layer k . Following [2], [4], [10]–[12] we let N be the complexity of the algorithm. Note however that as the total number of flops required for evaluating the bound in (15) may be upper and lower bounded by constants that are independent of ρ [19] our results relating to the SD complexity exponent would not change if we instead considered N to be the number of flops spent by the decoder.

B. The search radius

The description of the sphere decoder is not complete without specifying how the search radius is selected. In the interest of obtaining the SD complexity exponent, we may argue that any reasonable choice of a fixed (non-random) search radius should satisfy

$$\xi \doteq \rho^0. \quad (18)$$

To see this, it is sufficient to note that the metric in (13) satisfies

$$\|\mathbf{r} - \mathbf{R}\mathbf{s}\|^2 = \|\mathbf{Q}^H \mathbf{w}\|^2$$

for the transmitted vector \mathbf{s} . Thus, if $\|\mathbf{Q}^H \mathbf{w}\|^2 > \xi^2$ the transmitted symbol vector is excluded from the search, resulting in a decoding error. By considering a radius that grows slowly with SNR, say $\xi^2 = z \log \rho \doteq \rho^0$, it can be shown that

$$P(\|\mathbf{Q}^H \mathbf{w}\|^2 \geq \xi^2) \doteq \rho^{-z}, \quad (19)$$

for $z > 0$, i.e., by selecting $z > d(r)$ the probability of excluding the transmitted vector will (for increasing ρ) vanish faster than the probability of error and cause vanishing degradation of the overall probability of error. At the same time, if the radius does not tend to infinity with increasing ρ , it will follow that $P(\|\mathbf{Q}^H \mathbf{w}\|^2 > \xi^2)$ is bounded away from zero. This implies a non-vanishing probability of error and a resulting diversity gain of zero, which is clearly undesirable. Thus, as the complexity exponent is not affected by the particular choice of z , we shall unless otherwise stated in the following for simplicity assume that $\xi^2 = z \log \rho \doteq \rho^0$, with $z > d(r)$ in order to ensure vanishing degradation to the overall probability of error. This said, the derived SD complexity exponent would be the same if we considered adaptive radius updates as used in the Schnorr-Euchner (SE) implementation [2], [3]. This may be shown by following the argument in [11], and we give a proof of this statement in the present setting in Appendix B-D.

C. Decoding Complexity

The sphere decoder complexity, or equivalently the number of visited nodes N , is as stated a random variable with a distribution that depends on a number of parameters, e.g., the system dimensions n_R, n_T and T , the SNR ρ , the multiplexing gain r , the generator matrix \mathbf{G} , and the search radius ξ . This is well known and follows by the randomness of the bound in (15). Naturally this randomness must be considered when properly analyzing the sphere decoder complexity, unless one resorts to a worst-case analysis. However, we argue that the worst-case analysis is unnecessarily pessimistic.

In order to illustrate one of the key problems with focusing on the worst-case complexity consider the event that $\mathbf{H} = \mathbf{0}$ and $\|\mathbf{w}\|^2 < \xi^2$. In this case it is easily seen that (14) and (15) are always satisfied. As a consequence, the complexity of the sphere decoder would be equal to

$$N = \sum_{k=1}^{\kappa} |\mathbb{S}_\eta|^k \doteq \sum_{k=1}^{\kappa} \rho^{\frac{rT}{\kappa} k} \doteq \rho^{rT},$$

where we have used that $\eta \doteq \rho^{\frac{rT}{2\kappa}}$ to obtain the size of \mathbb{S}_η in (4). The worst-case complexity is therefore comparable to that of a full search over \mathcal{X} as $|\mathcal{X}| \doteq \rho^{rT}$. However, there is also no point in decoding when $\mathbf{H} = \mathbf{0}$ as all codewords would yield the same ML metric which in turn implies a high probability of error. Essentially the same argument, for opting out of decoding, can be made whenever the MIMO channel is in information theoretic outage [6]. In this case it follows by Fano's inequality that the probability of decoding error will be bounded away from zero. In fact, for a code with a diversity gain of $d(r)$ any set of channel matrices \mathcal{H} for which $P(\mathbf{H} \in \mathcal{H}) < \rho^{-d(r)}$, may be neglected by the decoder with vanishing degradation of the overall probability of error.

However, rather than identifying and excluding a set of bad channel matrices directly, a more pragmatic approach is to impose a run-time constraint on the decoder and ensure that this constraint is such that the probability of it being violated is insignificant in relation to the probability of error. This leads to the following measure of the decoding complexity, which we will use throughout.

Definition 1: Let

$$\Psi(x) \triangleq - \lim_{\rho \rightarrow \infty} \frac{\log P(N \geq \rho^x)}{\log \rho} \quad (20)$$

where N is the number of nodes visited by the sphere decoder (cf. (17)). The *SD complexity exponent* is then given by

$$c(r) \triangleq \inf\{x \mid \Psi(x) > d(r)\} \quad (21)$$

where $d(r)$ (cf. (9)) is the diversity gain of the code at multiplexing gain r .

D. A vanishing gap to the ML performance

In order to illustrate the operational significance of $c(r)$, we recall that in addition to the instances where the ML decoder makes an incorrect decision, a time-limited sphere decoder can additionally make decoding errors when the search radius is selected too small, i.e., when $\mathcal{N}_\kappa = \emptyset$ (cf. (16)), or when the run-time limit of ρ^x becomes active, i.e., when $N \geq \rho^x$. These extra errors cause a gap to ML performance which can be quantified as

$$g(x) \triangleq \frac{P(\{\hat{\mathbf{X}}_{\text{ML}} \neq \mathbf{X}\} \cup \{\mathcal{N}_\kappa = \emptyset\} \cup \{N \geq \rho^x\})}{P(\hat{\mathbf{X}}_{\text{ML}} \neq \mathbf{X})}$$

describing the ratio between the probability of error of the time-limited sphere decoder and the ML decoder. With respect to $c(r)$ we then have the following.

Theorem 1: A sphere decoder with a computational constraint activated at ρ^x flops, allows for a vanishing gap to ML performance for all $x > c(r)$, i.e.,

$$\lim_{\rho \rightarrow \infty} g(x) = 1, \text{ for any } x > c(r). \quad (22)$$

The above simply states that, for any $x > c(r)$ it is possible to design a decoder based on the SD algorithm that achieves a vanishing SNR gap to the ML decoder, at a worst-case complexity of ρ^x . To see this apply the union bound to get

$$g(x) \leq \underbrace{\frac{P(\hat{\mathbf{X}}_{\text{ML}} \neq \mathbf{X})}{P(\hat{\mathbf{X}}_{\text{ML}} \neq \mathbf{X})}}_{=1} + \underbrace{\frac{P(\xi_{\text{ML}} > \xi)}{P(\hat{\mathbf{X}}_{\text{ML}} \neq \mathbf{X})}}_{\rightarrow 0} + \underbrace{\frac{P(N \geq \rho^x)}{P(\hat{\mathbf{X}}_{\text{ML}} \neq \mathbf{X})}}_{\rightarrow 0}$$

where the second and third term tend to zero with increasing ρ as the numerator tends to zero at a faster rate than the denominator cf. (19), (21). This immediately translates to a vanishing SNR gap to the ML decoder at high SNR. In short, the probabilities of the events that the search space is empty or that the complexity of the run-time-unconstrained sphere decoder exceeds ρ^x are insignificant in comparison to the probability of ML decoding errors.

Furthermore, one cannot in general time-limit the sphere decoder to ρ^x for some $x < c(r)$ and expect an arbitrary

small gap to ML performance. Specifically, one can show (cf. (68)) that $P(N \geq \rho^x) \dot{>} \rho^{-d(r)}$ for any $x < c(r)$, and as a result, it follows⁶ for $x < c(r)$ that

$$g(x) \geq \frac{P(N \geq \rho^x)}{\underbrace{P(\hat{\mathbf{X}}_{\text{ML}} \neq \mathbf{X})}_{\rightarrow \infty}}$$

implying that any attempt to significantly reduce the complexity below $\rho^{c(r)}$ will be at the expense of the vanishing SNR gap to ML decoding.

IV. THE SPHERE DECODER COMPLEXITY EXPONENT

We proceed to establish upper and lower bounds on the SD complexity exponent, in essence through the application of a principle (dating back to Gauss) which states that the number of integer lattice points within a (large) set is well approximated by the volume of the set [27], [28]. Thus, in order to approximate the number of nodes at layer k of the search tree, i.e., the size of \mathcal{N}_k defined in (16), we are primarily concerned with the volume of $([-\eta, \eta] + \sqrt{-1}[-\eta, \eta])^k \cap \mathcal{E}_k$ where \mathcal{E}_k is the elliptical set given by (cf. (16))

$$\mathcal{E}_k = \{\tilde{\mathbf{s}}_k \in \mathbb{C}^k \mid \|\mathbf{r}_k - \mathbf{R}_k \tilde{\mathbf{s}}_k\|^2 \leq \xi^2\}. \quad (23)$$

The use of the volume principle for assessing the sphere decoder complexity was previously used in [2], [12], [29] although its prior use in the communications literature is limited to the case of lattice decoding (i.e., where the constellation boundary constraint imposed by $[-\eta, \eta]$ is ignored by the decoder). Herein, we have to take the constellation boundary into account to obtain tight bounds on the SD complexity exponent.

The upper and lower bounds presented in this section are essentially obtained in three main steps: 1) The volume principle is used to obtain an expression for the number N_k of visited nodes at layer k in terms of the singular values of \mathbf{R}_k ; 2) the singular values of \mathbf{R}_k for $k = 1, \dots, \kappa$ are related to the singular values of the channel matrix \mathbf{H} ; and 3) the theory of large deviations is used similarly to [6] to identify random events likely to cause an atypically large decoding complexity. Establishing the upper bound on $c(r)$ turns out to be easier mathematically. The reason for this is primarily in the second step where the interlacing property of singular values of sub-matrices [23] can be used to lower bound the singular values of \mathbf{R}_k by the singular values of \mathbf{H} , to yield results that are universally applicable for any full rank generator matrix \mathbf{G} , cf. Theorem 2. Although the interlacing property gives both upper and lower bounds on the singular values of \mathbf{R}_k , the upper bounds are unfortunately not sufficient for establishing tight lower bounds on $c(r)$. We are therefore forced to develop tighter bounds that depend on some technical assumptions on \mathbf{G} , cf. Lemma 2. While these conditions are, at least in principle, easily verified for any given code design, they are generally hard to verify for arbitrary classes of codes.

⁶Note here that what we formally show is that under the basic technical conditions of Lemma 2 one cannot time-limit the decoder to ρ^x for any $x < c(r)$. The same statement naturally holds whenever $\Psi(x)$ is strictly decreasing in x .

Nevertheless, for some important classes discussed in Section IV-D we are able to conclude that the upper bound on $c(r)$ is tight, thereby establishing $c(r)$ exactly.

The derivation of the upper bound on $c(r)$ is given in the following sub-sections, while the derivation of the lower bound, which is similar in spirit to the upper bound but complicated by some technical details, is primarily given in Appendix B, and discussed in Section IV-D.

A. The volume principle

As noted, we begin by establishing bounds on $N_k = |\mathcal{N}_k|$ in terms of the singular values of the matrix \mathbf{R}_k in (15), the sphere radius ξ and the constellation size η . To this end, consider the following lemma, which corresponds to rigorous applications of the volume principle discussed above. The proof of the lemma is given in Appendix A.

Lemma 1: Let $\mathcal{E} \subset \mathbb{R}^n$ be the ellipsoidal set given by

$$\mathcal{E} \triangleq \{\mathbf{d} \in \mathbb{R}^n \mid \|\mathbf{c} - \mathbf{D}\mathbf{d}\|^2 \leq \xi^2\} \quad (24)$$

where $\mathbf{D} \in \mathbb{R}^{n \times n}$ and $\mathbf{c} \in \mathbb{R}^n$. Let $\mathcal{B} \subset \mathbb{R}^n$ be the hypercube given by

$$\mathcal{B} \triangleq \{\mathbf{d} \in \mathbb{R}^n \mid |d_i| \leq \eta, i = 1, \dots, n\}. \quad (25)$$

Then, the number of integer points contained in the intersection of \mathcal{E} and \mathcal{B} is upper bounded as

$$|\mathcal{E} \cap \mathcal{B} \cap \mathbb{Z}^n| \leq \prod_{i=1}^n \left[\sqrt{n} + \min \left(\frac{2\xi}{\sigma_i(\mathbf{D})}, 2\sqrt{n}\eta \right) \right], \quad (26)$$

and the number of integer points contained in \mathcal{E} is lower bounded by

$$|\mathcal{E} \cap \mathbb{Z}^n| \geq \prod_{i=1}^n \left(\frac{2\xi}{\sqrt{n}\sigma_i(\mathbf{D})} - \sqrt{n} \right)^+, \quad (27)$$

where $\sigma_i(\mathbf{D})$, $i = 1, \dots, n$ denote the singular values of \mathbf{D} .

Although Lemma 1 is phrased in terms of real valued quantities, it is easily applied to complex valued sets by considering each complex dimension as two real valued dimensions. In particular, the expression in (14) is equivalent to

$$\|\underline{\mathbf{r}}_k - \underline{\mathbf{R}}_k \hat{\underline{\mathbf{s}}}_k\|^2 \leq \xi^2$$

where

$$\underline{\mathbf{r}}_k = \begin{bmatrix} \Re(\mathbf{r}_k) \\ \Im(\mathbf{r}_k) \end{bmatrix}, \underline{\mathbf{R}}_k = \begin{bmatrix} \Re(\mathbf{R}_k) & -\Im(\mathbf{R}_k) \\ \Im(\mathbf{R}_k) & \Re(\mathbf{R}_k) \end{bmatrix},$$

and

$$\hat{\underline{\mathbf{s}}}_k = \begin{bmatrix} \Re(\hat{\mathbf{s}}_k) \\ \Im(\hat{\mathbf{s}}_k) \end{bmatrix}.$$

By noting that if $\mathbf{R}_k = \mathbf{U}\Sigma\mathbf{V}^H$ is the singular value decomposition (SVD) [23] of \mathbf{R}_k , then

$$\underline{\mathbf{R}}_k = \begin{bmatrix} \Re(\mathbf{U}) & -\Im(\mathbf{U}) \\ \Im(\mathbf{U}) & \Re(\mathbf{U}) \end{bmatrix} \begin{bmatrix} \Sigma & \mathbf{0} \\ \mathbf{0} & \Sigma \end{bmatrix} \begin{bmatrix} \Re(\mathbf{V}^H) & -\Im(\mathbf{V}^H) \\ \Im(\mathbf{V}^H) & \Re(\mathbf{V}^H) \end{bmatrix}$$

is an SVD of $\underline{\mathbf{R}}_k$, it follows that the singular values of $\underline{\mathbf{R}}_k$ are the same as those of \mathbf{R}_k albeit with a multiplicity of 2. Thus, applying (26) in Lemma 1 to \mathcal{N}_k (cf. (16)) yields an

upper bound on the number of nodes visited at layer k , which is given as

$$N_k = |\mathcal{N}_k| \leq \prod_{i=1}^k \left[\sqrt{2k} + \min \left(\frac{2\xi}{\sigma_i(\mathbf{R}_k)}, 2\sqrt{2k\eta} \right) \right]^2 \quad (28)$$

where $\sigma_i(\mathbf{R}_k)$, $i = 1, \dots, k$ denote the singular values of \mathbf{R}_k . Here, in essence, the additive $\sqrt{2k}$ term accounts for edge effects in the volume approximation, the first term in the minimum accounts for the size of the search sphere, and the second term in the minimum accounts for the constellation boundary.

The lower bound in (27) will be used later in order to assess the tightness of the upper bound on $c(r)$ developed next. The reason for providing a lower bound on $|\mathcal{E} \cap \mathbb{Z}^n|$ and not $|\mathcal{E} \cap \mathcal{B} \cap \mathbb{Z}^n|$ is that we cannot a-priori rule out that \mathcal{c} in (24) is such that $\mathcal{E} \cap \mathcal{B} = \emptyset$, a case which if not ruled out would lead to the trivial lower bound $|\mathcal{E} \cap \mathcal{B} \cap \mathbb{Z}^n| \geq 0$.

B. Singular value bounds

The interlacing theorem of singular values of sub-matrices (cf. [23, Th. 7.3.9] and [26, Corollary 3.1.3]) states that the singular values of \mathbf{R}_k are bounded by the singular values of \mathbf{R} according

$$\sigma_{i+\kappa-k}(\mathbf{R}) \geq \sigma_i(\mathbf{R}_k) \geq \sigma_i(\mathbf{R}), \quad i = 1, \dots, k, \quad (29)$$

where $\sigma_i(\mathbf{R}_k)$ and $\sigma_i(\mathbf{R})$ denotes the i th singular values of \mathbf{R}_k and \mathbf{R} respectively. As $\mathbf{R} = \mathbf{Q}^H \mathbf{M}$ where \mathbf{Q} has a set of orthogonal columns that span the range of \mathbf{M} it follows that $\sigma_i(\mathbf{R}) = \sigma_i(\mathbf{M})$. Further, by the definition of \mathbf{M} in (11) we have that $\sigma_i(\mathbf{M}) \geq \theta\gamma\sigma_i(\mathbf{I}_T \otimes \mathbf{H})$ where $\gamma \triangleq \sigma_1(\mathbf{G}) > 0$ due to the assumption that \mathbf{G} is full rank. The singular values of $\mathbf{I}_T \otimes \mathbf{H}$ are the same as those of the channel matrix \mathbf{H} in (1), albeit with a multiplicity of T , i.e.,

$$\sigma_i(\mathbf{I}_T \otimes \mathbf{H}) = \sigma_{\iota_T(i)}(\mathbf{H}), \quad i = 1, \dots, n_T,$$

where

$$\iota_T(i) \triangleq \left\lceil \frac{i}{T} \right\rceil. \quad (30)$$

This can be seen by noting that if $\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H$ is the SVD of \mathbf{H} , then

$$(\mathbf{I}_T \otimes \mathbf{U})(\mathbf{I}_T \otimes \mathbf{\Sigma})(\mathbf{I}_T \otimes \mathbf{V}^H)$$

is an SVD of $\mathbf{I}_T \otimes \mathbf{H}$ (albeit with a non-standard ordering of the singular values). Alternatively, one can apply [26, Theorem 4.2.12] to the eigenvalues of $\mathbf{I}_T \otimes \mathbf{H}^H \mathbf{H}$.

Combining the above yields a lower bound on the singular values of \mathbf{R}_k in terms of the singular values of the channel matrix \mathbf{H} according to

$$\sigma_i(\mathbf{R}_k) \geq \theta\gamma\sigma_{\iota_T(i)}(\mathbf{H}), \quad i = 1, \dots, k,$$

and an upper bound on the number of nodes visited by the sphere decoder at layer k according to

$$N_k \leq \prod_{i=1}^k \left[\sqrt{2k} + \min \left(\frac{2\xi}{\theta\gamma\sigma_{\iota_T(i)}(\mathbf{H})}, 2\sqrt{2k\eta} \right) \right]^2. \quad (31)$$

In order to bound the probability that the right hand side of (31) is atypically large in the high SNR regime, it is useful to consider the SNR dependent parameterization of the singular values (or eigenvalues) of $\mathbf{H}^H \mathbf{H}$ introduced in [6], i.e., SNR dependent random variables α_i , for $i = 1, \dots, n_T$, defined by

$$\alpha_i \triangleq -\frac{\log \sigma_i(\mathbf{H}^H \mathbf{H})}{\log \rho} \Leftrightarrow \sigma_i(\mathbf{H}^H \mathbf{H}) = \rho^{-\alpha_i}. \quad (32)$$

Note that by this definition $\sigma_i(\mathbf{H}) = \rho^{-\frac{1}{2}\alpha_i}$. The variables, α_i for $i = 1, \dots, n_T$ are referred to as the *singularity levels* of \mathbf{H} as they give an indication of how close to singular the channel \mathbf{H} is in relation to the inverse SNR ρ^{-1} . As $\xi \doteq \rho^0$, $\theta \doteq \rho^{\frac{1}{2} - \frac{rT}{2\kappa}}$ and $\eta = \rho^{\frac{rT}{2\kappa}}$ it holds that

$$\left[\sqrt{k} + \min \left(\frac{2\xi}{\theta\gamma\sigma_{\iota_T(i)}(\mathbf{H})}, 2\sqrt{k\eta} \right) \right]^2 \leq \rho^{\nu_i}$$

where

$$\nu_i \triangleq \min \left(\frac{rT}{\kappa} - 1 + \alpha_{\iota_T(i)}, \frac{rT}{\kappa} \right)^+. \quad (33)$$

By (31) it follows that

$$N_k \leq \prod_{i=1}^k \rho^{\nu_i} = \rho^{\sum_{i=1}^k \nu_i}. \quad (34)$$

However, as the SNR exponent on the right hand side of (34) is non decreasing in k it must for the total number of visited nodes N hold that $N = \sum_{k=1}^{\kappa} N_k \leq \rho^{\sum_{i=1}^{\kappa} \nu_i}$ or for any given $\delta > 0$ hold that

$$N \leq \rho^{\sum_{i=1}^{\kappa} \nu_i + \delta} \quad (35)$$

provided ρ is sufficiently large.

Consider now the set (cf. (33) and (35))

$$\mathcal{T}(x) \triangleq \left\{ \boldsymbol{\alpha} \mid \sum_{i=1}^{\kappa} \min \left(\frac{rT}{\kappa} - 1 + \alpha_{\iota_T(i)}, \frac{rT}{\kappa} \right)^+ \geq x \right\}, \quad (36)$$

where $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_{n_T})$. As (35) holds (asymptotically) for any $\delta > 0$, and since $\boldsymbol{\alpha} \notin \mathcal{T}(y)$ implies that $N < \rho^x$ for any $y < x$ by (35) and (33), it follows that

$$\lim_{\rho \rightarrow \infty} \frac{\log \mathbb{P}(N \geq \rho^x)}{\log \rho} \leq \lim_{\rho \rightarrow \infty} \frac{\log \mathbb{P}(\boldsymbol{\alpha} \in \mathcal{T}(y))}{\log \rho}.$$

Equivalently (cf. (20))

$$\Psi(x) \geq -\lim_{\rho \rightarrow \infty} \frac{\log \mathbb{P}(\boldsymbol{\alpha} \in \mathcal{T}(y))}{\log \rho} \quad (37)$$

for $y < x$. The value of the bound in (37) is that the right hand side is readily computed using large deviation theory [30].

C. Large deviations

A sequence of random vectors $\boldsymbol{\beta}_\rho \in \mathbb{R}^n$ parameterized by ρ is said to satisfy the *large deviation principle* [30] with *rate function* I ,

$$I : \mathbb{R}^n \mapsto \{\mathbb{R}_+, \infty\},$$

if for every open set $\mathcal{G} \subseteq \mathbb{R}^n$ it holds that

$$\liminf_{\rho \rightarrow \infty} \frac{\log \mathbb{P}(\boldsymbol{\beta}_\rho \in \mathcal{G})}{\log \rho} \geq -\inf_{\boldsymbol{\beta} \in \mathcal{G}} I(\boldsymbol{\beta}) \quad (38)$$

and if for every closed set $\mathcal{F} \subseteq \mathbb{R}^n$ it holds that

$$\liminf_{\rho \rightarrow \infty} \frac{\log P(\beta_\rho \in \mathcal{F})}{\log \rho} \leq - \inf_{\beta \in \mathcal{F}} I(\beta). \quad (39)$$

Although not stated formally, one of the central results of [6] is that the sequence of random variables given by $\alpha_\rho = \alpha = (\alpha_1, \dots, \alpha_{n_T})$ (cf. (32)) satisfies the large deviation principle with rate function (see the proof of Theorem 4 in [6])

$$I(\alpha) = \sum_{i=1}^{n_T} (n_R - n_T + 2i - 1)\alpha_i \quad (40)$$

if $\alpha_1 \geq \dots \geq \alpha_{n_T} \geq 0$ and $I(\alpha) = \infty$ otherwise. This observation was key in establishing the DMT in [6].

By combining (37) with (39), and noting that $\mathcal{T}(y)$ is a closed set, it follows that

$$\Psi(x) \geq f(y) \triangleq \inf_{\alpha \in \mathcal{T}(y)} I(\alpha) \quad (41)$$

for any $y < x$. As $\mathcal{T}(x) \subseteq \mathcal{T}(y)$ for all $y < x$ it follows that $f(y)$ is non-decreasing and it can additionally be verified that $f(y)$ is left-continuous, i.e.,

$$\sup_{y < x} f(y) = f(x),$$

which implies that

$$\Psi(x) \geq f(x) = \inf_{\alpha \in \mathcal{T}(x)} I(\alpha). \quad (42)$$

From (42) it follows that the complexity exponent $c(r)$ is upper bounded by $\bar{c}(r)$ where

$$\bar{c}(r) \triangleq \inf\{x \mid f(x) > d(r)\} = \sup\{x \mid f(x) \leq d(r)\} \quad (43)$$

and where the last equality follows as $f(x)$ is non-decreasing. Further, by the left-continuity of $f(x)$ it follows that the supremum on the right is attained, i.e., the supremum can be replaced by a maximum.

Note however that the condition that $f(x) \leq d(r)$ is satisfied if and only if there exist an $\alpha \in \mathcal{T}(x)$ such that $I(\alpha) \leq d(r)$. Thus, $\bar{c}(r)$ in (43) could alternatively be obtained as the solution to a constrained maximization problem according to

$$\max_{\alpha, x} x \quad (44a)$$

$$\text{s.t.} \quad \sum_{i=1}^{\kappa} \min\left(\frac{rT}{\kappa} - 1 + \alpha_{\nu_T(i)}, \frac{rT}{\kappa}\right)^+ \geq x \quad (44b)$$

$$\sum_{i=1}^{n_T} (n_R - n_T + 2i - 1)\alpha_i \leq d \quad (44c)$$

$$\alpha_1 \geq \dots \geq \alpha_{n_T} \geq 0, \quad (44d)$$

where (44b) follows from the constraint $\alpha \in \mathcal{T}(x)$, and where (44c) and (44d) follows from $I(\alpha) \leq d(r)$. It is straightforward to show that the optimal x in (44) must be such that (44b) is satisfied with equality. By further noting that the sum in (44b) contains only n_T distinct terms, each with multiplicity T , it can be seen that

$$\begin{aligned} & \sum_{i=1}^{\kappa} \min\left(\frac{rT}{\kappa} - 1 + \alpha_{\nu_T(i)}, \frac{rT}{\kappa}\right)^+ \\ &= \sum_{i=1}^{n_T} T \min\left(\frac{r}{n_T} - 1 + \alpha_i, \frac{r}{n_T}\right)^+, \end{aligned}$$

where we have also used the full-rate assumption that $\kappa = n_T T$. We summarize the above in the following theorem.

Theorem 2: The SD complexity exponent $c(r)$ of decoding any full-rate linear dispersive code with multiplexing gain r and diversity $d(r)$ is upper bounded as $c(r) \leq \bar{c}(r)$ where

$$\bar{c}(r) \triangleq \max_{\alpha} \sum_{i=1}^{n_T} T \min\left(\frac{r}{n_T} - 1 + \alpha_i, \frac{r}{n_T}\right)^+ \quad (45a)$$

$$\text{s.t.} \quad \sum_{i=1}^{n_T} (n_R - n_T + 2i - 1)\alpha_i \leq d(r) \quad (45b)$$

$$\alpha_1 \geq \dots \geq \alpha_{n_T} \geq 0. \quad (45c)$$

The upper bound given by Theorem 2 can naturally be computed given explicit values for the multiplexing gain r and diversity gain⁷ $d(r)$. However, it is also possible in some cases to give general solutions as a function of r when the DMT curve $d(r)$ of the code is known explicitly. In particular, DMT optimal codes such as those presented in [13]–[18] have a diversity gain of $d(k) = (n_T - k)(n_R - k)$ at any integer multiplexing gain $r = k$ [6]. In this case it is straightforward to verify that an⁸ optimal α in (45) is given by

$$\alpha_i^* = 1, \quad \text{for } i = 1, \dots, n_T - k$$

and

$$\alpha_i^* = 0, \quad \text{for } i = n_T - k + 1, \dots, n_T.$$

To see this, note that the objective function in (45a) is symmetric with respect to permutations of the set of α_i for $i = 1, \dots, n_T$. As the sum in the diversity constraint (45b) places more weight on α_j than on α_i , for $j > i$, it is optimal to increase α_1 until the term in (45a) containing α_1 saturates (i.e., when $\alpha_1 = 1$), then to increase α_2 etcetera, until the constraint is satisfied with equality. This yields the aforementioned solution. Note also that $\alpha_1^*, \dots, \alpha_{n_T}^*$ are the same singularity levels that give the typical outages in [6], (cf. Section V-A). Inserting the optimal solution into (45a) yields

$$\bar{c}(k) = \frac{Tk(n_T - k)}{n_T},$$

which is a remarkably simple upper bound on the SD complexity exponent of decoding any DMT optimal code at an integer multiplexing gain $r = k$.

For a DMT optimal code at a possibly non-integer value of r , let k be the integer such that $r \in [k, k + 1)$, i.e., $k = \lfloor r \rfloor$. The optimal solution is in this case given by

$$\alpha_i^* = 1, \quad \text{for } i = 1, \dots, n_T - k - 1,$$

$$\alpha_i^* = 0, \quad \text{for } i = n_T - k + 1, \dots, n_T,$$

and

$$\alpha_{n_T - k}^* = k + 1 - r.$$

Substituting the above solution back into (45a) yields

$$\bar{c}(r) = \frac{T}{n_T} \left(r(n_T - k - 1) + (n_T k - r(n_T - 1))^+ \right).$$

⁷Note that Theorem 2 does not assume a diversity optimal code.

⁸In general, (45) does not have a unique optimal point as $\min(a, b)^+$ is constant in a for $a \leq 0$ and $a \geq b$.

We summarize the above in the following theorem.

Theorem 3: The SD complexity exponent $c(r)$ of decoding any DMT optimal full-rate linear dispersive code with integer multiplexing gain $r = k$ is upper bounded as

$$c(k) \leq \bar{c}(k) = \frac{Tk(n_T - k)}{n_T}. \quad (46)$$

For general r where $0 \leq r \leq n_T$ the SD complexity exponent $c(r)$ is upper bounded as $c(r) \leq \bar{c}(r)$ where

$$\bar{c}(r) = \frac{T}{n_T} \left(r(n_T - \lfloor r \rfloor - 1) + (n_T \lfloor r \rfloor - r(n_T - 1))^+ \right). \quad (47)$$

The function $\bar{c}(r)$ in (47) is a piecewise linear function in r , although slightly more involved than the set of straight lines describing the optimal DMT $d(r)$. For $n_T = T = n$ the function in (47) coincides with the curve for $c(r)$ shown in Fig. 1.

D. Establishing the exact SD complexity exponent

We now turn to specific cases where we can exactly establish the SD complexity exponent $c(r)$ by establishing that the upper bound $c(r) \leq \bar{c}(r)$ is in fact tight. To this end, we begin with the following lemma, which provides a sufficient condition for $c(r) = \bar{c}(r)$, i.e., for the tightness of the upper bound.

Lemma 2: Let $\mathbf{G}_{|p} \in \mathbb{C}^{\kappa \times Tp}$ be the matrix consisting of the first Tp columns of the generator matrix $\mathbf{G} \in \mathbb{C}^{\kappa \times \kappa}$. If there exists, for $p = 1, \dots, n_T$, unitary matrices $\mathbf{U}_p \in \mathbb{C}^{n_T \times p}$ such that

$$\text{rank}((\mathbf{I}_T \otimes \mathbf{U}_p^H) \mathbf{G}_{|p}) = pT, \quad (48)$$

then $c(r) = \bar{c}(r)$ for all $r \in [0, n_T]$, where $\bar{c}(r)$ is given by (45) in Theorem 2.

The proof of Lemma 2 is similar in spirit to the proof of Theorem 2, although riddled with technical details, and therefore relegated to Appendix B. In essence, the condition posed in (48) implies that there are certain orientations of the right singular vectors of the channel \mathbf{H} (in relation to the code generated by \mathbf{G}) for which the lower bound in (29) is sufficiently tight. Details are provided in Appendix B, and some additional discussions of (48) and the general applicability of the lemma can be found in Section V. However, we first apply Lemma 2 to find $c(r)$ in some very important special cases.

To this end, it is useful to first note that permuting the columns of \mathbf{G} , i.e., replacing \mathbf{G} with $\mathbf{G}\mathbf{\Pi}$ where $\mathbf{\Pi} \in \mathbb{R}^{\kappa \times \kappa}$ is a permutation matrix, does not change the code \mathcal{X} . Instead, the effect such a permutation would have is that it would change the order in which the symbols in \mathbf{s} are enumerated by the sphere decoder described in Section III-A (cf. [3, Section IV]). In the present context, the first pT columns of $\mathbf{G}\mathbf{\Pi}$, i.e., $[\mathbf{G}\mathbf{\Pi}]_{|p}$, may differ from those of \mathbf{G} . Thus, we see that (48) depends not only on the code itself, but also on the order in which the constituent symbols s_i are enumerated by the sphere decoder (cf. [3], [4] where the topic of column ordering is discussed in detail).

In the context of Lemma 2, it can be seen that $(\mathbf{I}_T \otimes \mathbf{U}_p^H) \mathbf{G} \in \mathbb{C}^{n_T \times \kappa}$ has rank pT for any unitary \mathbf{U}_p due to the

full rank assumption on \mathbf{G} . One can therefore select pT linearly independent columns, or equivalently, find a permutation matrix $\mathbf{\Pi}$ such that $(\mathbf{I}_T \otimes \mathbf{U}_p^H)[\mathbf{G}\mathbf{\Pi}]_{|p}$ has full rank. Using a similar argument, we can recursively construct a (single) $\mathbf{\Pi}$ for which there are \mathbf{U}_p for $p = 1, \dots, n_T$ satisfying

$$\text{rank}((\mathbf{I}_T \otimes \mathbf{U}_p^H)[\mathbf{G}\mathbf{\Pi}]_{|p}) = pT$$

by constructing \mathbf{U}_{p-1} from \mathbf{U}_p by removing a column, selecting the appropriate columns from \mathbf{G} , and starting the recursion with an arbitrary \mathbf{U}_{n_T} . Interpreting the above in light of Theorem 2 and Lemma 2 we can thus establish that for any full-rate linear dispersive code design, $\bar{c}(r)$ as defined in Theorem 2 and given in Theorem 3 for DMT optimal codes, is the tightest upper bound on the SD complexity exponent that can possibly hold under arbitrary column orderings. This is formalized in the following.

Theorem 4: Given any full-rate linear dispersive code achieving diversity $d(r)$, there is always at least one column ordering for which $c(r) = \bar{c}(r)$.

However, while Theorem 4 is useful in the sense that it tells us that one could not improve upon the tightness of $\bar{c}(r)$ without introducing further assumptions regarding the particular code design considered, it is obviously not of practical interest to use the worst possible column ordering. Therefore, we turn our attention to the important class of threaded codes [31] for which we will show that the *natural* column ordering $\mathbf{\Pi} = \mathbf{I}_\kappa$ implies $c(r) = \bar{c}(r)$.

1) *Threaded codes:* The class of threaded code designs is of particular interest, as it includes full-rate codes that perform very well in a variety of settings. The threaded algebraic space-time (TAST) codes [31], codes constructed from cyclic division algebras (CDAs) [32], [33], and specifically modified CDA codes [15]–[18] that were shown (cf. [16]) to achieve the entire DMT, are prime examples. The CDA based threaded designs are also the only currently known explicit constructions capable of achieving the DMT for all values of n_T and simultaneously over all $r \in [0, n_T]$. All these codes have a common threaded structure. Specifically an $n \times n$ threaded code is built from n component codes mapped cyclically in threads (or layers) to the codewords \mathbf{X} . For example, in the special case of $n = n_T = T = 4$, the thread structure is given by

$$\begin{bmatrix} 1 & 4 & 3 & 2 \\ 2 & 1 & 4 & 3 \\ 3 & 2 & 1 & 4 \\ 4 & 3 & 2 & 1 \end{bmatrix}$$

where the numbers 1, 2, 3, 4 indicate the thread to which a particular entry of \mathbf{X} belongs. In general, symbol j in thread l is mapped to $[\mathbf{X}]_{j,k}$ where $k = \text{mod}(j-l, n) + 1$ and where $\text{mod}(\cdot, n)$ denotes the modulo n operation. For example, in the case of perfect codes [17], [18] which also employ a threaded structure, the code follows from

$$\text{lay}(\mathbf{X}) = \theta \underbrace{\begin{bmatrix} \mathbf{B}_0 \mathbf{C} & & & \\ & \ddots & & \\ & & \mathbf{B}_{n-1} \mathbf{C} & \\ & & & \mathbf{I} \end{bmatrix}}_{\mathbf{Y}} \underbrace{\begin{bmatrix} \mathbf{s}^{(1)} \\ \vdots \\ \mathbf{s}^{(n)} \end{bmatrix}}_{\mathbf{s}} \quad (49)$$

where

$$B_i \triangleq \text{Diag}(\underbrace{1, 1, \dots, 1}_{n-i \text{ entries}}, \underbrace{\gamma, \gamma, \dots, \gamma}_i), \quad i = 0, \dots, n-1$$

are full rank diagonal matrices incorporating a properly chosen thread-separating scalar $\gamma \in \mathbb{C}$, where $C \in \mathbb{C}^{n \times n}$ is a (unitary) full rank generator matrix for the component code of each thread, $s^{(l)} \in \mathbb{S}_\eta^n$ are the constellation symbols of thread l , and where $\text{lay}(\mathbf{X})$ denotes the matrix to vector operation obtained by stacking the elements of \mathbf{X} according to their thread (cf. the column based stacking of the $\text{vec}(\cdot)$ operation).

Regarding the cost of decoding by such codes, we note that the corresponding generator matrix $G \in \mathbb{C}^{n^2 \times n^2}$ is related to \mathbf{Y} through a permutation of the rows (cf. (49)) in such a way that the (i, j) th block $G_{ij} \in \mathbb{C}^{n \times n}$ of G contains exactly one non zero row which itself is one of the rows of $B_{j-1}C$. Consequently, $G_{|ip} \triangleq [G_{i1} \ \dots \ G_{ip}] \in \mathbb{C}^{n \times np}$ has rank p and contains exactly p non-zero rows. This holds for any n and $p \leq n$. Now, let $U_p \in \mathbb{C}^{n \times p}$ be a unitary matrix with the property that any p rows of U_p are linearly independent. Such matrices can clearly be constructed, and an example is the matrix that contains the first p discrete Fourier transform (DFT) vectors of length n . Let $\tilde{G}_{|ip} \in \mathbb{C}^{p \times np}$ be the matrix containing only the non-zero rows of $G_{|ip} \in \mathbb{C}^{p \times np}$, and let $\tilde{U}_{ip} \in \mathbb{C}^{p \times p}$ be the full rank matrix consisting of the rows of U_p matching the non-zero rows of $G_{|ip}$. It follows that

$$(\mathbf{I}_n \otimes U_p^H)G_{|p} = \underbrace{\begin{bmatrix} \tilde{U}_{1p}^H & & \\ & \ddots & \\ & & \tilde{U}_{np}^H \end{bmatrix}}_{\tilde{U}_p} \underbrace{\begin{bmatrix} \tilde{G}_{|1p} \\ \vdots \\ \tilde{G}_{|np} \end{bmatrix}}_{\tilde{G}_{|p}} \in \mathbb{C}^{np \times np}$$

is full rank as both $\tilde{U}_p \in \mathbb{C}^{np \times np}$ and $\tilde{G}_{|p} \in \mathbb{C}^{np \times np}$ are full rank and square matrices. Note also that the same argument can be made regardless of the ordering of the threads, and for any other code with a threaded structure, provided the symbols in s are grouped into layers as in (49). This is stated in the following.

Theorem 5: The SD complexity exponent, given any threaded code with $n = n_T = T$ that is decoded with the natural column ordering or under any other thread-wise grouping, is $c(r) = \bar{c}(r)$ where $\bar{c}(r)$ is given in Theorem 2.

Consequently directly from Theorems 3 and 5, we have the following result for DMT optimal threaded codes.

Theorem 6: Sphere decoding with thread-wise grouping of any DMT optimal threaded code with $n = n_T = T$, achieves DMT optimality with an SD complexity exponent of

$$c(r) = r(n - \lfloor r \rfloor - 1) + (n \lfloor r \rfloor - r(n - 1))^+ \quad (50)$$

which, for integer values of $r = k$, simplifies to

$$c(k) = k(n - k). \quad (51)$$

We briefly note that as expected the complexity increases with increasing $n = n_T = T$ for any fixed r which is quite natural as the size of the codebook \mathcal{X} and the signal space dimension increase. One can however also note that $c(r)$ is

independent of the number of receive antennas n_R (provided $n_R \geq n_T$). This is specific to the DMT optimal behavior and threaded structure of the codes, and may be explained by the fact that even though the channel quality generally improves by adding receive antennas - thus generally reducing complexity - the same improvement also occurs in the error probability performance of the code, and these two effects cancel each other in the SD complexity exponent.

2) *2 × 2 approximately universal codes:* We here go one step further and identify a class of codes for which we can state, without limitations on the actual code structure, that $c(r) = \bar{c}(r)$ for any column ordering. In particular, we establish this for the class of all 2 × 2 approximately universal codes, i.e., all minimum delay codes that can achieve DMT optimality over the 2 × n_R channel irrespective of fading statistics [22]. This is accomplished, albeit only for the specific case of $n_T \times T = 2 \times 2$, by proving that (48) follows from the so called non-vanishing determinant (NVD) condition [33] which is well known to be a necessary and sufficient condition for approximate universality. We consequently have the following theorem, proven in Appendix D.

Theorem 7: Any 2 × 2 full-rate approximately universal linear dispersive code, irrespective of its structure, introduces an SD complexity exponent of

$$c(r) = \min(r, 2 - r).$$

As the NVD property does not depend on the ordering of the columns of G , it also follows that the conclusion of Theorem 7 holds irrespective of the column ordering.

V. IMPLICATIONS AND DISCUSSIONS

A. Decoding complexity and information theoretic outages

We recall that the claim of Theorem 2 may be expressed in terms of the function (cf. (45a))

$$\bar{c}(r : \alpha) \triangleq \sum_{i=1}^{n_T} T \min\left(\frac{r}{n_T} - 1 + \alpha_i, \frac{r}{n_T}\right)^+, \quad (52)$$

which provides a conditional, asymptotic, upper bound on the sphere decoding complexity according to $N \leq \rho^{\bar{c}(r:\alpha)}$ in terms of the multiplexing gain r and the singularity level α . The final upper bound $\bar{c}(r)$ in (45) is then given as the worst-case $\bar{c}(r : \alpha)$ over all singularity levels that occur with a probability that is larger than or equal to the probability of error, given asymptotically by the diversity gain $d(r)$ of the code.

The characterization of the DMT in [6] relies on the asymptotic probability of outages at high SNR, i.e., the probability that the i.i.d. Rayleigh fading AWGN channel given by

$$\mathbf{y}_t = \mathbf{H}\mathbf{x}_t + \mathbf{w}_t$$

with a power constraint $\mathbb{E}\{\|\mathbf{x}_t\|^2\} \leq \rho$ cannot support an asymptotic data-rate of $R = r \log \rho + o(\log \rho)$. As was shown in [6], this occurs when the singularity levels belong to the outage set $\mathcal{A}(r) = \{\alpha \mid \sum_i (1 - \alpha_i)^+ < r\}$, and the diversity of the outage event is given by the most likely set of singularity levels that satisfy this condition, i.e., $d(r) = \inf_{\alpha \in \mathcal{A}(r)} I(\alpha)$ where $I(\alpha)$ is given in (40).

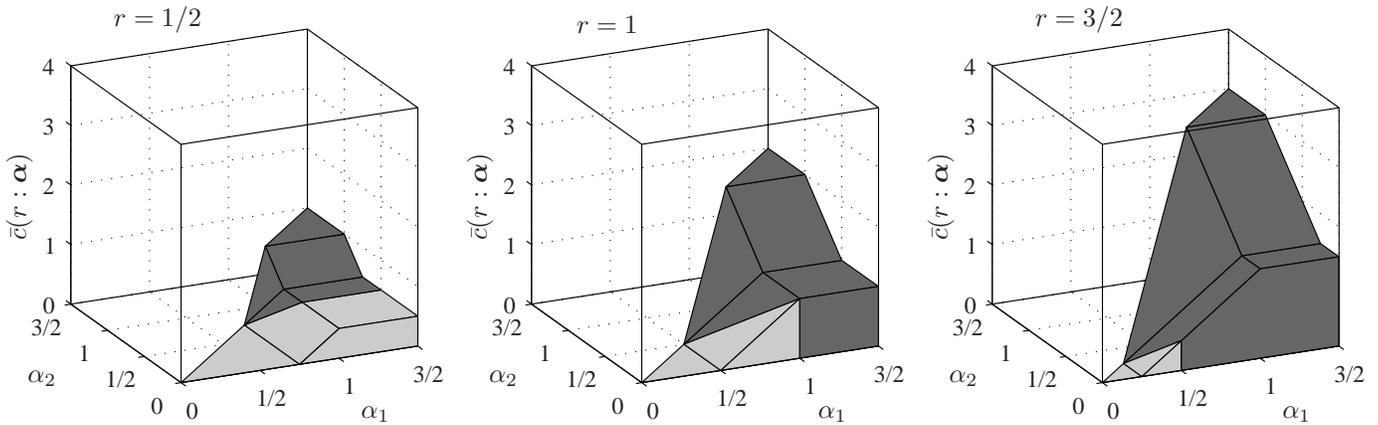


Fig. 2. Conditional SD complexity exponent bound $\bar{c}(\alpha)$ as a function of singularity levels $\alpha = (\alpha_1, \alpha_2)$. Singularity levels that corresponds to outage events at the target multiplexing gain r are shown in dark grey, while singularity levels capable of supporting the target multiplexing gain are shown in light grey.

If we restrict attention to the set of singularity levels whose probability of occurring does not vanish exponentially fast, i.e., for which $I(\alpha) < \infty$ or equivalently $\alpha_{n_T} \geq \dots \geq \alpha_1 \geq 0$, we can for DMT optimal codes⁹ make an interesting connection between the decoding complexity and information theoretic outages. In particular, as $d(r) = \inf_{\alpha \in \mathcal{A}(r)} I(\alpha)$ it follows that $\sum_i (n_R - n_T + 2i - 1)\alpha_i \leq d(r)$ if and only if $\sum_i (1 - \alpha_i)^+ \geq r$. We can thus, for DMT optimal codes, equivalently express (45) according to

$$\bar{c}(r) = \max_{\alpha} \bar{c}(r : \alpha) \quad (53a)$$

$$\text{s.t. } \sum_{i=1}^{n_T} (1 - \alpha_i)^+ \geq r \quad (53b)$$

$$\alpha_1 \geq \dots \geq \alpha_{n_T} \geq 0, \quad (53c)$$

which may be interpreted as the worst-case complexity (bound) over all channels that are not in outage. This significantly strengthens the connection between channel and decoding outages touched upon in Section III-C.

The concept is illustrated for $n_T = T = 2$ in Fig. 2 where $\bar{c}(r : \alpha)$ is plotted as a function of $\alpha = (\alpha_1, \alpha_2)$ over $\alpha_1 \geq \alpha_2 \geq 0$. In this case, $\bar{c}(r) = \min(r, 2 - r)$. Note also here that we know by Theorem 7 that $c(r) = \bar{c}(r)$. Singularity levels that are in the outage region $\sum_i (1 - \alpha_i)^+ < r$ are shown in a darker shade. It can be seen that increasing the multiplexing gain r increases the codeword density and codebook size and consequently broadens the set of singularity levels that can potentially lead to higher complexity. However, increasing the multiplexing gain also reduces the set of channels that support the data-rate, thus limiting the set of singularity levels for which the decoder needs to be applied, leading to an overall reduction in the SD complexity exponent as r approaches its maximum value.

Further, the connection to information theoretic outages allows for an intuitive explanation of the result of Theorem 6,

⁹To be precise, we are assuming here that we are working with approximately universal codes for which it is known that errors are only likely when the channel is in outage [22]. All explicitly constructed full-rate DMT-optimal codes known to date, are also approximately universal.

and in particular (51). To this end, it is illustrative to consider a heuristic argument involving low rank channel matrices \mathbf{H} . As noted in [6], the typical outages at integer multiplexing gains $r = k$ are caused by channels that are close to the set of rank k matrices, i.e., that have $n - k$ small singular values. If we for the purpose of illustration assume that \mathbf{H} has rank k , it follows that $\mathbf{I}_T \otimes \mathbf{H}$ for $T = n$, and \mathbf{M} , has rank nk , and consequently a null-space of dimension $n(n - k)$. This implies that the $n(n - k) \times n(n - k)$ lower right block of \mathbf{R} is¹⁰ identically equal to zero, and the sphere decoder pruning criteria become totally ineffective up to and including layer $n(n - k)$. As the size of \mathbb{S}_η is $|\mathbb{S}_\eta| \doteq \rho^{\frac{k}{n}}$ for $r = k$, we see that the number of nodes searched at layer $n(n - k)$ of the SD search tree is $|\mathbb{S}_\eta^{n(n-k)}| \doteq \rho^{k(n-k)}$ (cf. (51)). In order to ensure close to optimal performance, the sphere decoder must be able to decode for channels where $n - k$ singular values are close to zero. However, channels with even more singular values close to zero occur with a probability that is small in relation to the outage probability or the probability of ML decoder error, and can thus be safely ignored by the decoder.

B. A complexity bound that holds for all fading statistics

It is perfectly conceivable that the sphere decoding complexity may rise under specific codes and under specific fading statistics that tend to regularly introduce channel instances that are difficult to decode for. A natural question is then whether one can bound the complexity, irrespective of the code and of the statistical characterization of the channel. Considering Theorem 2, and the proof of this theorem, we can see that the i.i.d. Rayleigh fading assumption only enters through the rate function $I(\alpha)$. Consequently, we may directly restate Theorem 2 for other fading distributions after updating (45b) and (45c) with the appropriate rate-function $I(\alpha)$. Some relevant

¹⁰This also requires that the first nk columns of \mathbf{R} are linearly independent. In fact, the rigorous treatment of this technical detail is largely responsible for much of the difficulty in establishing the lower bounds on $c(r)$. In particular, condition (48) in Lemma 2 guarantees that this happens with sufficiently high probability, while Lemma 3 provides a perturbation analysis that allows us to extend this intuitive reasoning to channels that are close to the set of rank deficient matrices.

examples of rate-functions for other fading distributions are given in [34].

Further, regardless of which $I(\alpha)$ applies, the upper bound $\bar{c}(r)$ in Theorem 2 is non-decreasing in $d(r)$ and maximized when $d(r)$ corresponds to the outage exponent. In this case $\bar{c}(r)$ is, again, given by (53), which does not explicitly depend the fading distribution other than through the assumption that $P(\alpha_{n_T} < 0)$ vanishes exponentially fast; an assumption that holds for all reasonable distributions. This implies that SD complexity exponent is universally upper bounded as (cf. Theorem 3)

$$c(r) \leq \frac{T}{n_T} \left(r(n_T - \lfloor r \rfloor - 1) + (n_T \lfloor r \rfloor - r(n_T - 1))^+ \right)$$

for any full-rate code and statistical characterization of the channel. This is also clearly the tightest upper bound that can hold for all (full-rate) codes and fading statistics.

C. Fast decodable codes

In [35]–[37] a family of DMT optimal $n_T \times T = 2 \times 2$ space-time codes called fast decodable codes [38] were constructed. The SD complexity exponent (and also its upper bound) provides an interesting approach for comparing the complexity of decoding regular codes and fast decodable codes. Before doing so it should be noted that these fast decodable codes are not, strictly speaking, of the form in (3) as the real and imaginary part of each constituent symbol is dispersed separately. Nevertheless, the fast decodable codes may be decoded by an equivalent *real valued* sphere decoder that performs a search over a tree with 2κ layers and $|\mathbb{S}_\eta|^{\frac{1}{2}}$ branches per node, and we can compare the reported worst-case complexity of this real valued sphere decoder to the complexity of the complex valued sphere decoder considered herein.

The fast decodable codes have the appealing property that the upper right 4×4 block of the real valued $\mathbf{R} \in \mathbb{R}^{8 \times 8}$ (cf. (13)) is always a diagonal matrix, regardless of the particular realization of \mathbf{H} . While the regular real valued sphere decoder for a 2×2 full-rate code would perform a (bounded) search over the entire tree, it is sufficient for the fast decodable codes to (without loss of optimality) perform a search over only the 4 first layers, and extend each node at layer 4 to a valid codeword through a faster, linear, ML decoding. This simplified version of the real valued sphere decoder can be viewed as a search over a regular tree where each node has $|\mathbb{S}_\eta|^{\frac{1}{2}}$ children up to layer 4, but only one child per node for the 4 remaining layers. Consequently, the worst-case number of nodes visited by the simplified sphere decoder is $5|\mathbb{S}_\eta|^{\frac{1}{2}} + \sum_{k=1}^3 |\mathbb{S}_\eta|^{\frac{k}{2}} \doteq |\mathbb{S}_\eta|^2 \doteq \rho^r$ as opposed to $\sum_{k=1}^8 |\mathbb{S}_\eta|^{\frac{k}{2}} \doteq \rho^{2r}$ for the regular real valued sphere decoder, cf. [38]. Thus, fast decodability implies a reduction by a factor of 2 in the worst-case SNR exponent, which is significant at high SNR.

However, this worst-case SNR exponent of the simplified SD algorithm should be viewed in light of the SD complexity exponent induced by any 2×2 approximately universal code as given by Theorem 7, i.e., $c(r) = \min(r, 2 - r)$. The worst-case SNR exponent of the regular sphere decoder and of the simplified sphere decoder, and the SD complexity exponent

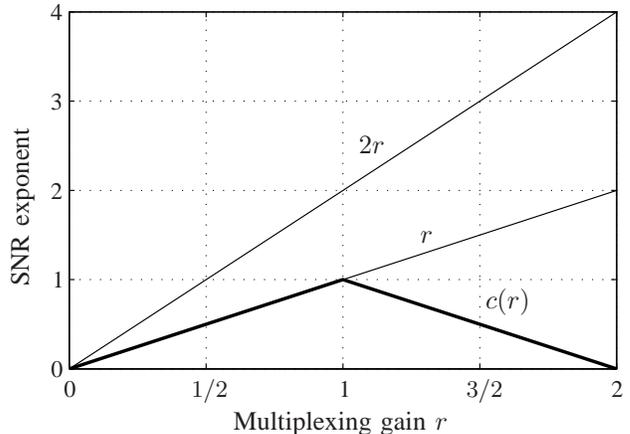


Fig. 3. Comparing the SD complexity exponent $c(r) = \min(r, 2 - r)$, with the (worst-case) SNR exponent ($2r$) of the regular sphere decoder, and with the maximal SNR exponent (r) of the simplified sphere decoder.

$c(r)$ are shown in Fig. 3. For multiplexing gains lower than or equal to 1, the SD complexity exponent and the worst-case SNR exponent of the simplified sphere decoder actually coincide, while the SD complexity exponent is strictly lower for multiplexing gains higher than one. The interpretation is that a run-time constrained sphere decoder will yield asymptotic ML performance for *any* 2×2 approximately universal code at a complexity that is comparable to the reported worst-case complexity of the fast decodable codes at low multiplexing gains, and significantly better at high multiplexing gains. Thus, in a sense, all approximately universal codes are fast decodable at high SNR. However, we hasten to add that the fast decodable structure can naturally still be desirable in many cases of practical interest.

D. The applicability of Lemma 2

Finally, we discuss the application of Lemma 2 to codes not considered herein. To this end, note that for any given generator matrix \mathbf{G} of some code not covered by Section IV-D, it should be clear that if (48) holds then Lemma 2 could be used to establish that $c(r) = \bar{c}(r)$. This said, we also wish to caution the reader that (48) only represents a sufficient condition for $c(r) = \bar{c}(r)$. It does not necessarily follow that $c(r) < \bar{c}(r)$ if (48) is not true. In other words, the question of if there are code designs that improve upon the bound $\bar{c}(r)$ is not answered in the positive by finding code designs for which (48) does not hold.

As for testing (48) it should also be noted that one does not have to restrict the search for \mathbf{U}_p to the set of unitary matrices. Any full rank matrix $\mathbf{A} \in \mathbb{C}^{n_T \times p}$ can be factored, e.g., by the QR decomposition, as $\mathbf{U}_p \mathbf{T} = \mathbf{A}$ where $\mathbf{T} \in \mathbb{C}^{p \times p}$ has rank p and where \mathbf{U}_p is unitary. Hence, as $(\mathbf{I}_T \otimes \mathbf{T}^H)$ is full rank, it follows that

$$(\mathbf{I}_T \otimes \mathbf{A}^H) \mathbf{G}_{|p} = (\mathbf{I}_T \otimes \mathbf{T}^H) (\mathbf{I}_T \otimes \mathbf{U}_p^H) \mathbf{G}_{|p}$$

is rank deficient if and only if (48) fails to hold. Hence, the statement of Lemma 2 could be phrased in terms of the existence of any \mathbf{U}_p , not necessarily unitary.

Further, let

$$p(\mathbf{A}) \triangleq |(\mathbf{I}_T \otimes \mathbf{A}^H)\mathbf{G}|_p,$$

where $|\cdot|$ denotes the determinant, and note that $p(\mathbf{A})$ is a polynomial in the elements of \mathbf{A} . It can thus be seen that if $p(\mathbf{A}) \neq 0$ for some $\mathbf{A} \in \mathbb{C}^{n_T \times p}$, i.e., $p(\mathbf{A})$ is not the zero polynomial, it follows that the set of \mathbf{A} for which $p(\mathbf{A}) = 0$ has zero Lebesgue measure. It is then a straightforward extension to show that (48) holds either never or almost always with respect to the set of unitary matrices \mathbf{U}_p over the Stiefel manifold (i.e., the set of all unitary $n \times p$ matrices) endowed with the Haar (uniform) measure. This suggests a rather interesting conceptual method for verifying (48). Given a specific generator matrix one could at least in theory test the condition of Lemma 2 by selecting \mathbf{U}_p (or \mathbf{A}) uniformly at random, and the condition of Lemma 2 would be proven with probability one if true. However, finite precision computations will limit the practical applicability of such an approach, although symbolic computations could potentially be a way to test a specific code design.

VI. CONCLUSION

The work addressed the open question of identifying the computational cost of near-ML sphere decoding. In the high-SNR high-rate regime, the introduced SD complexity exponent asymptotically described this cost, concisely revealing the cost's natural dependencies to the codeword density, the codebook size, as well as to the SNR, dimensions and fading characteristics of the wireless channel. This exponent currently sets the bar with respect to the computational reserves required for decoding with arbitrarily close to ML performance, and the clear challenge is now to identify transceivers with a lesser complexity exponent that can still guarantee a vanishing ML gap.

The simplicity of the provided complexity guarantees can offer insight into designing robust encoders, decoders, and time-out policies, as well as guidelines for network planning in settings where rate, reliability, and computational complexity are principal intertwined concerns. Such guarantees can apply towards substantial savings in energy, processing power, and hardware.

APPENDIX A PROOF OF LEMMA 1

In the following, we provide a proof of Lemma 1, starting with the upper bound in (26) and then establishing the lower bound in (27). To this end, note that the length of the i th semi-axis of \mathcal{E} , denoted e_i , is given by

$$e_i \triangleq \frac{2\xi}{\sigma_i(\mathbf{D})}.$$

Let \mathcal{C}_1 be the smallest orthotope (box), aligned with and containing \mathcal{E} , i.e., \mathcal{C}_1 is an orthotope with side lengths e_i (see Fig. 4). Let \mathcal{C}_2 be a hypercube with side-length $2\sqrt{n}\eta$, centered at the origin and aligned with \mathcal{C}_1 (see Fig. 4). As the diagonal of \mathcal{B} is $2\sqrt{n}\eta$ it follows that $\mathcal{B} \subset \mathcal{C}_2$, regardless of the orientation of \mathcal{C}_2 . Let \mathcal{C}_3 be given by $\mathcal{C}_3 = \mathcal{C}_1 \cap \mathcal{C}_2$ and note that

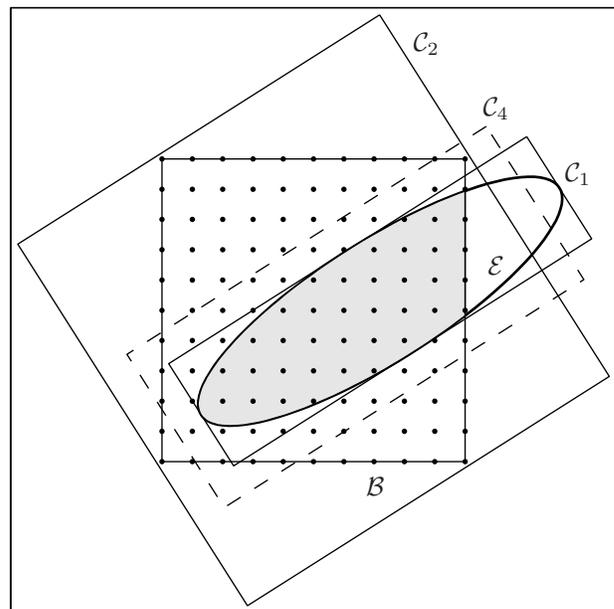


Fig. 4. Illustration of the proof of (26) in Lemma 1 in the case of $n = 2$. The lemma provides an upper bound on the number of integer points within the shaded area, corresponding to the intersection of the ellipsoid and the constellation boundary.

$\mathcal{E} \cap \mathcal{B} \subset \mathcal{C}_3$ as $\mathcal{E} \subset \mathcal{C}_1$ and $\mathcal{B} \subset \mathcal{C}_2$. As \mathcal{C}_1 and \mathcal{C}_2 are aligned, it follows that \mathcal{C}_3 is also an orthotope. Let l_1, \dots, l_n denote the side-lengths of \mathcal{C}_3 and note that $l_i \leq \min(e_i, 2\sqrt{n}\eta)$.

The mean value theorem [28] states that for any convex body (or set) $\mathcal{C} \subset \mathbb{R}^n$ it holds that

$$\text{Vol}(\mathcal{C}) = \int_{\mathcal{U}} |\mathbb{Z}^n \cap \mathcal{C} + \mathbf{u}| d\mathbf{u} \quad (54)$$

where

$$\mathcal{U} \triangleq \left[-\frac{1}{2}, \frac{1}{2}\right]^n$$

denotes the unit cube in \mathbb{R}^n , i.e., the volume of the set equals the average number of integer points in the set when perturbed by a uniform random perturbation over the unit cube [28]. This statement is a rigorous version of the intuitive notion that the number of integer points in \mathcal{C} may be approximated by its volume, and it can be used to obtain a non-random upper bound on the number of integer points in \mathcal{C}_3 . To this end, let \mathcal{C}_4 be the orthotope, aligned with and centered around \mathcal{C}_3 , with side lengths $l_i + \sqrt{n}$ (see Fig. 4). By construction, it follows that

$$\mathcal{C}_3 \subset \mathcal{C}_4 + \mathbf{u}$$

for any $\mathbf{u} \in \mathcal{U}$. It therefore follows by (54) that

$$|\mathcal{C}_3 \cap \mathbb{Z}^n| \leq \text{Vol}(\mathcal{C}_4) = \prod_{i=1}^n [\sqrt{n} + l_i],$$

where

$$l_i \leq \min\left(\frac{2\xi}{\sigma_i(\mathbf{D})}, 2\sqrt{n}\eta\right).$$

As $\mathcal{E} \cap \mathcal{B} \subset \mathcal{C}_3$ the upper bound in (26) follows.

In order to establish the lower bound in (27) we may redefine \mathcal{C}_1 to be the orthotope, aligned with the semiaxes of \mathcal{E} and with the same center, having side-lengths

$$b_i = \frac{2\xi}{\sqrt{n}\sigma_i(\mathbf{D})},$$

for $i = 1, \dots, n$. Now, by construction $\mathcal{C}_1 \subset \mathcal{E}$ which implies that $|\mathcal{C}_1 \cap \mathbb{Z}^n| \leq |\mathcal{E} \cap \mathbb{Z}^n|$. Let \mathcal{C}_2 be another orthotope, aligned with \mathcal{C}_1 and with side-lengths $\max(b_i - \sqrt{n}, 0)$. It follows that $\mathcal{C}_2 + \mathbf{u} \subset \mathcal{C}_1$ for any $\mathbf{u} \in \mathcal{U}$. By reasoning similar to what is used in the proof of the upper bound (cf. (54)) it follows that

$$|\mathcal{E} \cap \mathbb{Z}^n| \geq \text{vol}(\mathcal{C}_2) = \prod_{i=1}^n \max(b_i - \sqrt{n}, 0),$$

which establish the lower bound in (27).

APPENDIX B PROOF OF LEMMA 2

Let $\boldsymbol{\alpha}^* = (\alpha_1^*, \dots, \alpha_{n_T}^*)$ be an optimal point of (45) and let q be the largest integer for which (cf. (45a))

$$\frac{r}{n_T} - 1 + \alpha_q^* > 0. \quad (55)$$

Note that we can without loss of generality assume that $q \geq 1$ as otherwise $\bar{c}(r) = 0$ (cf. (45a)) and $c(r) = \bar{c}(r)$ would be a trivial statement. It follows that $\alpha_i^* > 0$ for $i = 1, \dots, q$ and we may also without loss of generality assume that $\alpha_i^* \leq 1$ for $i = 1, \dots, n_T$ as the objective in (45) does not increase in α_i beyond $\alpha_i = 1$. The goal will be to show that layer $k = qT$ of the sphere decoder contains close to $\rho^{\bar{c}(r)}$ nodes with a probability that is large with respect to the probability of decoding error $P(\mathbf{X}_{\text{ML}} \neq \mathbf{X}) \doteq \rho^{-d(r)}$.

To this end, let $\mathbf{H} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^H$ be the singular value decomposition of \mathbf{H} , where

$$\boldsymbol{\Sigma} \triangleq \text{Diag}(\sigma_1(\mathbf{H}), \dots, \sigma_{n_T}(\mathbf{H}))$$

and where $\mathbf{U}^H\mathbf{U} = \mathbf{I}$. Let \mathbf{U}_p denote the last $p \triangleq n_T - q$ columns of \mathbf{U} (corresponding to the p largest singular values) and let $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_{n_T})$ be the random vector of singularity levels given by (32). Now, consider the set of conditions (or events) given by

$$\Omega_1 \triangleq \{ \alpha_i^* - 2\delta < \alpha_i < \alpha_i^* - \delta, \quad i = 1, \dots, q, \\ 0 < \alpha_i < \delta, \quad i = q + 1, \dots, n_T \}, \quad (56a)$$

for some given (small) $\delta > 0$,

$$\Omega_2 \triangleq \{ \sigma_1((\mathbf{I}_T \otimes \mathbf{U}_p^H)\mathbf{G}_{|p}) \geq u \}, \quad (56b)$$

for some given $u > 0$,

$$\Omega_3 \triangleq \{ \|\mathbf{Q}^H \mathbf{w}\| \leq 1 \}, \quad (56c)$$

and

$$\Omega_4 \triangleq \{ \|\mathbf{s}\| \leq \frac{1}{2}\eta \}. \quad (56d)$$

Note also that by choosing δ sufficiently small, we may without loss of generality assume that Ω_1 implies that $\alpha_i > 0$ for all $i = 1, \dots, n_T$.

The following proof is structured as follows: First, in Sections B-A and B-B, it is established that (56) represent

sufficient conditions for the number of nodes N_k visited in layer $k = qT$ to be close to $\rho^{\bar{c}(r)}$. Then, in Section B-C, it is established that the set of conditions in (56) are simultaneously satisfied with a probability that is large with respect to the probability of error.

A. The Constellation boundary

We begin by proving that, given (56), the constellation boundary may be ignored, i.e., that \mathbb{S}_η may be replaced by \mathbb{S}_∞ in (16) without changing the result, thus making the lower bound (27) in Lemma 1 applicable. To this end, let $\hat{\mathbf{s}}_k \in \mathbb{S}_\infty^k$ be an arbitrary point in the k -dimensional infinite constellation (i.e., the Gaussian integer lattice) and assume that $\hat{\mathbf{s}}_k$ satisfies the sphere constraint at layer k , i.e.

$$\|\mathbf{r}_k - \mathbf{R}_k \hat{\mathbf{s}}_k\| \leq \xi.$$

Note that $\mathbf{r}_k = \mathbf{R}_k \mathbf{s}_k + \mathbf{v}_k$, where \mathbf{s}_k denotes the last k components of the transmitted symbol vector $\mathbf{s} \in \mathbb{S}_\eta^\kappa$ and where \mathbf{v}_k denotes the last k components of $\mathbf{v} \triangleq \mathbf{Q}^H \mathbf{w}$. It follows that

$$\|\mathbf{r}_k - \mathbf{R}_k \hat{\mathbf{s}}_k\| = \|\mathbf{R}_k(\mathbf{s}_k - \hat{\mathbf{s}}_k) + \mathbf{v}_k\| \\ \geq \sigma_1(\mathbf{R}_k) \|\hat{\mathbf{s}}_k - \mathbf{s}_k\| - \|\mathbf{v}_k\|$$

which implies that

$$\|\hat{\mathbf{s}}_k - \mathbf{s}_k\| \leq \frac{1}{\sigma_1(\mathbf{R}_k)} (\xi + \|\mathbf{v}_k\|)$$

and

$$\|\hat{\mathbf{s}}_k\| \leq \frac{1}{\sigma_1(\mathbf{R}_k)} (\xi + \|\mathbf{v}_k\|) + \|\mathbf{s}_k\|. \quad (57)$$

By the interlacing property of singular values (cf. (29)) it further follows that

$$\sigma_1(\mathbf{R}_k) \geq \theta \gamma \sigma_1(\mathbf{H}) \doteq \rho^{\frac{1}{2} - \frac{rT}{2\kappa} - \frac{1}{2}\alpha_1} \geq \rho^{\frac{1}{2}\delta - \frac{rT}{2\kappa}}$$

where we recall that $\theta \doteq \rho^{\frac{1}{2} - \frac{rT}{2\kappa}}$ is the power scaling and $\gamma = \sigma_1(\mathbf{G}) > 0$, and where the last inequality is implied by (56a) and $\alpha_1^* \leq 1$. As $\xi \doteq \rho^0$ and $\|\mathbf{v}_k\| \leq \|\mathbf{Q}^H \mathbf{w}\| \leq 1$ by (56c) it follows that

$$\frac{1}{\sigma_1(\mathbf{R}_k)} (\xi + \|\mathbf{v}_k\|) \leq \rho^{\frac{rT}{2\kappa} - \frac{1}{2}\delta}.$$

By (57), (56d), $\|\mathbf{s}_k\| \leq \|\mathbf{s}\| \leq \eta$ and since $\rho^{\frac{rT}{2\kappa} - \frac{1}{2}\delta} < \frac{1}{2}\eta$, it follows that

$$\|\hat{\mathbf{s}}_k\| \leq \eta$$

given that ρ is sufficiently large. This implies that $\hat{\mathbf{s}}_k \in \mathbb{S}_\eta^k$. Thus, any integer point that satisfies the sphere constraint must also belong to the constellation, and we can proceed using (27) to lower bound the complexity.

B. Singular value bounds

We proceed to provide bounds on the singular values of \mathbf{R}_k in order to lower bound the number of nodes visited in layer $k = qT$. However, as stated previously the interlacing theorem is, unlike in the derivation of the upper bound, not sufficient for our purposes. Instead, we consider the following lemma, proven in Appendix C.

Lemma 3: Let $\mathbf{A} \in \mathbb{C}^{m \times n}$, $m \geq n$ be an arbitrary matrix and $\mathbf{QR} = \mathbf{A}$ be the QR decomposition of \mathbf{A} . Partition \mathbf{A} , \mathbf{Q} and \mathbf{R} according to

$$[\mathbf{A}_1 \quad \mathbf{A}_2] = [\mathbf{Q}_1 \quad \mathbf{Q}_2] \begin{bmatrix} \mathbf{R}_{11} & \mathbf{R}_{12} \\ \mathbf{0} & \mathbf{R}_{22} \end{bmatrix},$$

where $\mathbf{A}_1 \in \mathbb{C}^{m \times n-k}$ and $\mathbf{R}_{22} \in \mathbb{C}^{k \times k}$. Then, assuming that $\sigma_i(\mathbf{A}) < \sigma_1(\mathbf{A}_1)$ for $i = 1, \dots, k$, it holds that

$$\sigma_i(\mathbf{R}_{22}) \leq \left[\frac{\sigma_n(\mathbf{A})}{\sigma_1(\mathbf{A}_1)} + 1 \right] \sigma_i(\mathbf{A}). \quad (58)$$

Applied to the effective channel matrix \mathbf{M} it follows that

$$\sigma_i(\mathbf{R}_k) \leq \left[\frac{\sigma_\kappa(\mathbf{M})}{\sigma_1(\mathbf{M}_1)} + 1 \right] \sigma_i(\mathbf{M}) \quad (59)$$

where \mathbf{M}_1 contains the first pT columns ($p = n_T - q$) of \mathbf{M} , assuming that $\sigma_i(\mathbf{M}) < \sigma_1(\mathbf{M}_1)$ as will be shown for $i = 1, \dots, qT$ later. In order to lower bound $\sigma_1(\mathbf{M}_1)$ note that

$$\mathbf{M}_1 = \theta(\mathbf{I}_T \otimes \mathbf{H})\mathbf{G}_{|p},$$

where $\mathbf{G}_{|p}$ denotes the first pT columns of \mathbf{G} and

$$\mathbf{M}_1^H \mathbf{M}_1 = \theta^2 \mathbf{G}_{|p}^H (\mathbf{I}_T \otimes \mathbf{H}^H \mathbf{H}) \mathbf{G}_{|p}.$$

As

$$\mathbf{H}^H \mathbf{H} \succeq \sigma_{q+1}(\mathbf{H}^H \mathbf{H}) \mathbf{U}_p \mathbf{U}_p^H$$

where \mathbf{U}_p denotes the matrix containing the p singular vectors corresponding to the p largest singular values, and where $\mathbf{A} \succeq \mathbf{B}$ denotes that $\mathbf{A} - \mathbf{B}$ is positive semi-definite, it follows that

$$\mathbf{M}_1^H \mathbf{M}_1 \succeq \theta^2 \sigma_{q+1}(\mathbf{H}^H \mathbf{H}) \mathbf{G}_{|p}^H (\mathbf{I}_T \otimes \mathbf{U}_p \mathbf{U}_p^H) \mathbf{G}_{|p}.$$

Considering the smallest singular value of $\mathbf{M}_1^H \mathbf{M}_1$ yields

$$\sigma_1(\mathbf{M}_1^H \mathbf{M}_1) \geq \theta^2 \sigma_{q+1}(\mathbf{H}^H \mathbf{H}) \sigma_1(\mathbf{G}_{|p}^H (\mathbf{I}_T \otimes \mathbf{U}_p \mathbf{U}_p^H) \mathbf{G}_{|p})$$

and consequently

$$\sigma_1(\mathbf{M}_1) \geq \theta \sigma_{q+1}(\mathbf{H}) = \theta \rho^{-\frac{1}{2}a_{q+1}} \geq \rho^{\frac{1}{2} - \frac{rT}{2\kappa} - \frac{1}{2}\delta}, \quad (60)$$

where the first inequality follows by (56b) and the last inequality follows by (56a) together with $\theta \doteq \rho^{\frac{1}{2} - \frac{rT}{2\kappa}}$ and as $u > 0$ is fixed (independent of ρ). Further,

$$\sigma_i(\mathbf{M}) = \theta \sigma_i((\mathbf{I}_T \otimes \mathbf{H})\mathbf{G}) \leq \theta \Gamma \sigma_{\iota_T(i)}(\mathbf{H})$$

where $\Gamma \triangleq \sigma_{\max}(\mathbf{G}) = \sigma_\kappa(\mathbf{G})$, and it follows from (56a) that

$$\sigma_i(\mathbf{M}) \leq \rho^{\frac{1}{2} - \frac{rT}{2\kappa} - \frac{1}{2}\alpha_{\iota_T(i)}^* + \delta} \quad (61)$$

for $i = 1, \dots, qT$. As $\alpha_i^* > 0$ for $i = 1, \dots, q$, it follows by comparing (60) and (61) that $\sigma_i(\mathbf{M}) \leq \sigma_1(\mathbf{M}_1)$ for $i = 1, \dots, qT$, given that δ is sufficiently small and that ρ is sufficiently large, making Lemma 3 applicable for $k = qT$.

For the maximal singular value of \mathbf{M} we have (cf. (60))

$$\sigma_\kappa(\mathbf{M}) \leq \rho^{\frac{1}{2} - \frac{rT}{2\kappa} - \frac{1}{2}\alpha_{n_T}} \leq \rho^{\frac{1}{2} - \frac{rT}{2\kappa}}$$

where the last inequality follows as $\alpha_{n_T} > 0$. Combined with (60) it follows that

$$\left[\frac{\sigma_\kappa(\mathbf{M})}{\sigma_1(\mathbf{M}_1)} + 1 \right] \leq \rho^{\frac{1}{2}\delta},$$

and from (59) and (61) that

$$\sigma_i(\mathbf{R}_k) \leq \rho^{\frac{1}{2} - \frac{rT}{2\kappa} - \frac{1}{2}\alpha_{\iota_T(i)}^* + \frac{3}{2}\delta}$$

for $i = 1, \dots, qT$. Consequently (cf. (27)),

$$\frac{2\xi}{\sqrt{2k}\sigma_i(\mathbf{R}_k)} \geq \rho^{\frac{rT}{2\kappa} + \frac{1}{2}\alpha_{\iota_T(i)}^* - \frac{1}{2} - \frac{3}{2}\delta}, \quad (62)$$

given that δ is sufficiently small. Further, as

$$\frac{rT}{2\kappa} + \frac{1}{2}\alpha_{\iota_T(i)}^* - \frac{1}{2} = \frac{1}{2} \left(\frac{r}{n_T} - 1 + \alpha_{\iota_T(i)}^* \right) > 0$$

for $i = 1, \dots, k$ where $k = 2qT$ by the condition for q in (55), it follows that the lower bound in (62) tends to infinity with increasing ρ provided that δ is small, and we may conclude that

$$\left[\frac{2\xi}{\sqrt{2k}\sigma_i(\mathbf{R}_k)} - \sqrt{2k} \right]^2 \geq \rho^{\frac{rT}{\kappa} + \alpha_{\iota_T(i)}^* - 1 - 3\delta} > 0 \quad (63)$$

where the last inequality holds, again, provided δ is small.

Combining (27) in Lemma 1 and (63), and making the real valued expansion as we did for Theorem 2, yields a lower bound on the number of nodes visited by the sphere decoder in layer $k = qT$ given by

$$N_k \geq \prod_{i=1}^k \rho^{\frac{rT}{\kappa} + \alpha_{\iota_T(i)}^* - 1 - 3\delta} = \rho^{v - 3k\delta} \quad (64)$$

where

$$\begin{aligned} v &\triangleq \sum_{i=1}^k \frac{rT}{\kappa} + \alpha_{\iota_T(i)}^* - 1 \\ &= \sum_{i=1}^q T \left(\frac{r}{n_T} + \alpha_i^* - 1 \right). \end{aligned} \quad (65)$$

By noting that

$$0 \leq \frac{r}{n_T} + \alpha_i^* - 1 \leq \frac{r}{n_T}$$

for $i = 1, \dots, q$ by the assumption that $\alpha_i^* \leq 0$ and the definition of q , it follows that

$$T \left(\frac{r}{n_T} + \alpha_i^* - 1 \right) = T \min \left(\frac{r}{n_T} - 1 + \alpha_i^*, \frac{r}{n_T} \right)^+ \quad (66)$$

for $i = 1, \dots, q$. Further, as

$$\frac{r}{n_T} + \alpha_i^* - 1 \leq 0$$

for $i > q$, also by the definition of q , the right hand side of (66) is equal to 0 for $i > q$. Thus, it follows that

$$v = \sum_{i=1}^{n_T} T \min \left(\frac{r}{n_T} - 1 + \alpha_i^*, \frac{r}{n_T} \right)^+ = \bar{c}(r)$$

where the last equality follows due to the optimality of α^* in (45), and we then obtain from (64) that

$$N \geq N_k \geq \rho^{\bar{c}(r) - 3k\delta},$$

given that ρ is sufficiently large and that $\delta > 0$ is small. However, as $\delta > 0$ can be chosen arbitrarily small it is concluded that (56) represents sufficient conditions under which the number of nodes visited is arbitrarily close to the upper bound of $\rho^{\bar{c}(r)}$ given by Theorem 2.

C. Probabilities

We now turn to the probability that the conditions imposed by (56) are simultaneously satisfied. The events in (56) are independent¹¹. As (56) imply $N \geq \rho^{\bar{c}(r)-\frac{3}{2}k\delta}$ it follows that

$$P\left(N \geq \rho^{\bar{c}(r)-3k\delta}\right) \geq \prod_{i=1}^4 P(\Omega_i),$$

given that ρ is sufficiently large.

The assumption made in Lemma 2, i.e., condition (48), guarantees that

$$\sigma_1((\mathbf{I} \otimes \mathbf{U}_p^H)\mathbf{G}|_p) > 0$$

for some \mathbf{U}_p . However, by the continuity of singular values [23] it follows for sufficiently small $u > 0$ (cf. (56b)) that $P(\Omega_2) > 0$, which implies $P(\Omega_2) \doteq \rho^0$ as Ω_2 is independent of ρ . The same is true for Ω_3 , i.e., $P(\Omega_3) \doteq \rho^0$. It may also be shown that $P(\Omega_4)$ converges to a strictly positive limit¹² and that therefore $P(\Omega_4) \doteq \rho^0$. It follows that

$$P(N \geq \rho^{\bar{c}(r)-3k\delta}) \doteq P(\Omega_1).$$

The probability of Ω_1 may again be assessed by using large deviation techniques as in [6]. In particular, it is noted that the condition imposed by Ω_1 (cf. (56a)) specifies an open set of admissible α . Applying (38) and (40) yields

$$\begin{aligned} -\lim_{\rho \rightarrow \infty} \frac{\log P(\Omega_1)}{\log \rho} &\leq \sum_{i=1}^q (n_R - n_T + 2i - 1)(\alpha_i^* - 2\delta) \\ &\leq d(r) - 2(n_R - n_T + q)q\delta < d(r), \end{aligned} \quad (67)$$

where the second inequality follows from (45b) and the feasibility of α^* . Thus,

$$-\lim_{\rho \rightarrow \infty} \frac{\log P(N \geq \rho^{\bar{c}(r)-3k\delta})}{\log \rho} < d(r). \quad (68)$$

By the definition of the SD complexity exponent $c(r)$ (cf. (21)) it follows by (68) that $c(r) \geq \bar{c}(r) - 3k\delta$. As the bound holds for arbitrarily small $\delta > 0$, it follows that $c(r) = \bar{c}(r)$, establishing the tightness of (45) and Lemma 2.

D. The extension to adaptive radius updates

The derivations above make the assumption that the search radius ξ is a non-random function of ρ that satisfies $\xi \doteq \rho^0$. It is thus natural to ask if the SD complexity exponent could potentially be improved by choosing ξ adaptively based on the problem data \mathbf{H} and \mathbf{Y} , as is done when using, e.g., the Schnorr-Euchner SD algorithm implementation [2], [3]. However, we will show here that it can not, and therefore that the assumption of a non-adaptive radius is made without loss of generality. The argument is similar to the one in [11].

¹¹The independence of Ω_1 and Ω_2 follows by the i.i.d. Rayleigh assumption on \mathbf{H} , which make the singular values and singular vectors of $\mathbf{H}^H\mathbf{H}$ independent [39].

¹²This is provided that $r > 0$ in which case the the subset of the constellation defined by Ω_4 contains an asymptotically deterministic and strictly positive fraction of the full constellation, cf. the proof of Lemma 1 in [9]. When $r = 0$ the statement that $\bar{c}(r)$ is tight is trivial as $\bar{c}(0) = 0$.

To this end, we note that even if ξ is adaptively chosen, it cannot be chosen smaller than the distance to the closest codeword, i.e., the (square root of the) minimum metric in (13), because otherwise no codeword will be chosen by the search. As was already argued in Section III-B, the distance to the *transmitted* codeword is $\|\mathbf{Q}^H\mathbf{w}\|$ and $P(\|\mathbf{Q}^H\mathbf{w}\| \geq \epsilon)$ can be made arbitrarily close to one by appropriately choosing $\epsilon > 0$. Consequently, whenever the transmitted codeword s is the ML decision, i.e., yields the minimum metric in (13), we could use $\xi \geq \epsilon$ where $\epsilon \doteq \rho^0$ as an (arbitrarily likely) probabilistic lower bound on an adaptive search radius ξ in the proof of the lower bounds on the complexity (e.g., in (63)) and obtain the same SNR exponents in these bounds.

Thus, to complete the argument we must only rule out the possibility that the probability that s yields the minimum metric in (13) under the conditions of Ω imposed in (56) is small, as the lower bound is derived explicitly under Ω . To this end, assume that it is false, i.e., that $P(\mathbf{s}_{\text{ML}} \neq s|\Omega) \geq \epsilon$ for any (fixed and SNR independent) $\epsilon > 0$. In this case we could lower bound the error probability of the ML decoder according to (cf. (67))

$$P(\mathbf{s}_{\text{ML}} \neq s) \geq P(\mathbf{s}_{\text{ML}} \neq s|\Omega)P(\Omega) \geq \epsilon P(\Omega) \doteq \rho^{-d(r)},$$

which would violate the definition of $d(r)$ as the diversity order of the ML decoder. Consequently, at sufficiently high SNR it must hold that $P(\mathbf{s}_{\text{ML}} = s|\Omega) \geq 1 - \epsilon$ for any $\epsilon > 0$. We can thus choose $\epsilon > 0$ such that, under Ω , it follows that $\mathbf{s}_{\text{ML}} = s$ and $\|\mathbf{Q}^H\mathbf{w}\| \geq \epsilon$ with arbitrary high probability, implying that $\xi_{\text{ML}} > \epsilon$ where ξ_{ML}^2 is the minimum metric in (13). In other words, there is some $\epsilon > 0$ for which

$$P(\Omega \cup \{\xi_{\text{ML}} \geq \epsilon\}) \doteq \rho^{-d(r)}.$$

Completing the proof of Lemma 2 with $\epsilon \doteq \rho^0$ in place of ξ (as $\xi \geq \xi_{\text{ML}} \geq \epsilon$ throughout the search) proves that $c(r) = \bar{c}(r)$ under (48) also if we allow for SD implementations that adaptively choose and update the search radius ξ . Finally, it should be noted that what is shown here is not that adaptively choosing the search radius cannot reduce complexity – it does – but only that this reduction is not significant enough to reduce the complexity exponent.

APPENDIX C PROOF OF LEMMA 3

Consider the matrix, $\underline{\mathbf{A}}$, given by

$$\underline{\mathbf{A}} \triangleq \mathbf{U}\underline{\Sigma}\mathbf{V}^H,$$

where $\underline{\Sigma} = \text{Diag}(0, \dots, 0, \sigma_{i+1}(\mathbf{A}), \dots, \sigma_n(\mathbf{A}))$, and where \mathbf{U} and \mathbf{V} denote the right and left singular vector of \mathbf{A} respectively. Partition $\underline{\mathbf{A}} \in \mathbb{C}^{m \times n}$ according to

$$\underline{\mathbf{A}} = [\underline{\mathbf{A}}_1 \quad \underline{\mathbf{A}}_2]$$

where $\underline{\mathbf{A}}_1 \in \mathbb{C}^{m \times n-k}$ and $\underline{\mathbf{A}}_2 \in \mathbb{C}^{m \times k}$.

By the nature of the QR decomposition, it holds that

$$\mathbf{P} \triangleq \mathbf{\Pi}_{\underline{\mathbf{A}}_1}^\perp \underline{\mathbf{A}}_2 = \mathbf{Q}_2 \mathbf{R}_{22}$$

where $\Pi_{\mathbf{A}_1}^\perp$ denotes the projection onto the orthogonal complement of the range of \mathbf{A}_1 (i.e. the nullspace of \mathbf{A}_1^H). Additionally, let

$$\underline{\mathbf{P}} \triangleq \Pi_{\mathbf{A}_1}^\perp \mathbf{A}_2.$$

As \mathbf{Q}_2 is a unitary matrix it follows that

$$\sigma_i(\mathbf{R}_{22}) = \sigma_i(\mathbf{P}).$$

In what follows, we will consider the singular values of \mathbf{P} in order to establish the lemma. To this end, we will make use of two results due to Weyl and Stewart. For a modern proof of Theorem 8, see e.g. [23, Corollary 7.3.8]. The statement in Theorem 9 follows by combining [40, Theorem 2.3] and [40, Theorem 2.4]. In the following, $\|\mathbf{B}\| = \sigma_{\max}(\mathbf{B})$ denotes the spectral matrix norm.

Theorem 8 (Weyl): For arbitrary $\mathbf{B}, \mathbf{C} \in \mathbb{C}^{p \times q}$ it holds that

$$|\sigma_i(\mathbf{B}) - \sigma_i(\mathbf{C})| \leq \|\mathbf{B} - \mathbf{C}\|. \quad (69)$$

Theorem 9 (Stewart): For $\mathbf{B}, \mathbf{C} \in \mathbb{C}^{p \times q}$ such that $\text{rank}(\mathbf{B}) = \text{rank}(\mathbf{C})$,

$$\|\Pi_{\mathbf{B}}^\perp - \Pi_{\mathbf{C}}^\perp\| \leq \min(\|\mathbf{B}^\dagger\|, \|\mathbf{C}^\dagger\|) \|\mathbf{B} - \mathbf{C}\|, \quad (70)$$

where $(\cdot)^\dagger$ denotes the Moore-Penrose pseudo inverse [23].

By noting that

$$\|\mathbf{A}_l - \underline{\mathbf{A}}_l\| \leq \|\mathbf{A} - \underline{\mathbf{A}}\| = \sigma_i(\mathbf{A}), \quad (71)$$

for $l = 1, 2$ and using the assumption that $\sigma_1(\mathbf{A}_1) > \sigma_i(\mathbf{A})$ it follows by Theorem 8 that

$$\sigma_1(\underline{\mathbf{A}}_1) \geq \sigma_1(\mathbf{A}_1) - \|\mathbf{A}_1 - \underline{\mathbf{A}}_1\| \geq \sigma_1(\mathbf{A}_1) - \sigma_i(\mathbf{A}) > 0$$

implying that $\underline{\mathbf{A}}_1$ is full rank. As $\sigma_1(\mathbf{A}_1) > 0$ is directly implied by $\sigma_1(\mathbf{A}_1) > \sigma_i(\mathbf{A})$ it follows that $\text{rank}(\underline{\mathbf{A}}_1) = \text{rank}(\mathbf{A}_1)$ which makes Theorem 9 applicable to $\Pi_{\mathbf{A}_1}^\perp - \Pi_{\underline{\mathbf{A}}_1}^\perp$. As

$$\begin{aligned} \mathbf{P} - \underline{\mathbf{P}} &= \Pi_{\mathbf{A}_1}^\perp \mathbf{A}_2 - \Pi_{\underline{\mathbf{A}}_1}^\perp \underline{\mathbf{A}}_2 \\ &= (\Pi_{\mathbf{A}_1}^\perp - \Pi_{\underline{\mathbf{A}}_1}^\perp) \underline{\mathbf{A}}_2 + \Pi_{\underline{\mathbf{A}}_1}^\perp (\mathbf{A}_2 - \underline{\mathbf{A}}_2) \end{aligned}$$

it follows that

$$\|\mathbf{P} - \underline{\mathbf{P}}\| \leq \|\mathbf{A}_1^\dagger\| \|\mathbf{A}_1 - \underline{\mathbf{A}}_1\| \|\underline{\mathbf{A}}_2\| + \|\mathbf{A}_2 - \underline{\mathbf{A}}_2\|,$$

where we used Theorem 9 and the fact that $\|\mathbf{BC}\| \leq \|\mathbf{B}\| \|\mathbf{C}\|$ and $\|\Pi_{\mathbf{B}}^\perp\| \leq 1$ [23]. By noting that $\|\mathbf{A}_1^\dagger\| = 1/\sigma_1(\mathbf{A}_1)$, that $\|\underline{\mathbf{A}}_2\| \leq \|\mathbf{A}_2\| = \sigma_n(\mathbf{A})$, that $\|\mathbf{A}_1 - \underline{\mathbf{A}}_1\| \leq \sigma_i(\mathbf{A})$ and that $\|\mathbf{A}_2 - \underline{\mathbf{A}}_2\| \leq \sigma_i(\mathbf{A})$ (cf. (71)), it follows that

$$\mu \triangleq \left[\frac{\sigma_n(\mathbf{A})}{\sigma_1(\mathbf{A}_1)} + 1 \right] \sigma_i(\mathbf{A}) \geq \|\mathbf{P} - \underline{\mathbf{P}}\|. \quad (72)$$

By again applying Theorem 8 to (72) it follows that $\sigma_i(\mathbf{P}) \leq \sigma_i(\underline{\mathbf{P}}) + \mu$. Note however that

$$\text{rank}(\underline{\mathbf{A}}) = \text{rank}(\underline{\mathbf{A}}_1) + \text{rank}(\underline{\mathbf{P}})$$

where $\underline{\mathbf{P}} \triangleq \Pi_{\underline{\mathbf{A}}_1}^\perp \underline{\mathbf{A}}_2 \in \mathbb{C}^{m \times k}$. As $\text{rank}(\underline{\mathbf{A}}_1) = n - k$ and $\text{rank}(\underline{\mathbf{A}}) \leq n - i$ it follows that

$$\text{rank}(\underline{\mathbf{P}}) \leq k - i$$

and $\sigma_i(\underline{\mathbf{P}}) = 0$. Thus, $\sigma_i(\mathbf{R}_{22}) = \sigma_i(\mathbf{P}) \leq \mu$ establishing the lemma.

APPENDIX D PROOF OF THEOREM 7

Let $\underline{\mathcal{X}}$ be the un-normalized extended codebook corresponding to the un-normalized lattice points $\mathbf{G}\mathbb{S}_\infty$, i.e., where $\mathcal{X} \subseteq \theta \underline{\mathcal{X}}$. A space-time code is said to satisfy the non-vanishing determinant (NVD) condition if [41]

$$\inf_{\underline{\mathbf{x}} \in \underline{\mathcal{X}} \setminus \mathbf{0}} |\underline{\mathbf{x}}| > 0, \quad (73)$$

i.e., if there are no non-zero un-normalized (difference) codewords with arbitrarily small determinants. The proof of the theorem draws from the well known fact that the NVD condition is a necessary condition for achieving approximate universality, and divides the problem into a few (exhaustive) cases where either the condition in Lemma 2 is shown to hold, or the NVD property is shown to be violated, thus eliminating the possibility of NVD codes that would violate the rank condition in Lemma 2. To this end, consider a partitioning of the 4×4 generator matrix according to

$$\mathbf{G} = \begin{bmatrix} \mathbf{G}_{11} & \mathbf{G}_{12} \\ \mathbf{G}_{21} & \mathbf{G}_{22} \end{bmatrix}$$

where $\mathbf{G}_{ij} \in \mathbb{C}^{2 \times 2}$. First of all, let us note that the case where $p = 2$ is trivially satisfied as \mathbf{G} is full rank, i.e., the matrix

$$(\mathbf{I}_2 \otimes \mathbf{U}_2^H) \mathbf{G}$$

is full rank for any unitary $\mathbf{U}_2 \in \mathbb{C}^{2 \times 2}$. We can thus restrict attention to the case of $p = 1$ and consider the rank of

$$(\mathbf{I}_2 \otimes \mathbf{u}^H) \mathbf{G}_{|1} = \begin{bmatrix} \mathbf{u}^H \mathbf{G}_{11} \\ \mathbf{u}^H \mathbf{G}_{21} \end{bmatrix} \in \mathbb{C}^{2 \times 2} \quad (74)$$

where $\mathbf{u} = \mathbf{U}_1 \in \mathbb{C}^{2 \times 1}$. In the cases where the NVD property is shown to not hold, it is sufficient to consider non-zero (unnormalized) codewords of the form

$$\underline{\mathbf{x}} = \begin{bmatrix} \mathbf{G}_{11} & \mathbf{G}_{12} \\ \mathbf{G}_{21} & \mathbf{G}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{s}_1 \\ \mathbf{s}_2 \end{bmatrix}$$

where $\mathbf{s}_1 \in \mathbb{S}_\infty^2$, and $\mathbf{s}_1 \neq \mathbf{0}$, and where $\mathbf{s}_2 = \mathbf{0}$. The (un-normalized) codewords in matrix form are in this case given by $\underline{\mathbf{X}} = [\mathbf{G}_{11} \mathbf{s}_1 \quad \mathbf{G}_{21} \mathbf{s}_1]$ and we have $\mathbf{u}^H \underline{\mathbf{X}} = [\mathbf{u}^H \mathbf{G}_{11} \mathbf{s}_1 \quad \mathbf{u}^H \mathbf{G}_{21} \mathbf{s}_1]$. All codewords discussed in what follows are assumed to have this structure.

We will now consider different cases depending on the rank of \mathbf{G}_{11} and \mathbf{G}_{21} . However, as it is straightforward to see that $\underline{\mathbf{X}}$ has zero determinant (for any \mathbf{s}_1) if either $\mathbf{G}_{11} = \mathbf{0}$ or $\mathbf{G}_{21} = \mathbf{0}$, the cases that need consideration are those when the rank of both \mathbf{G}_{11} and \mathbf{G}_{21} is equal to one (case a), when the rank of both \mathbf{G}_{11} and \mathbf{G}_{21} is equal to two (case b), and when the rank of either \mathbf{G}_{11} or \mathbf{G}_{21} is equal to one and the rank of the other is equal to two (case c).

A. Case a

Consider the case where the rank of both \mathbf{G}_{11} and \mathbf{G}_{21} is one, i.e., where $\mathbf{G}_{11} = \mathbf{b}_1 \mathbf{a}_1^H$ and $\mathbf{G}_{21} = \mathbf{b}_2 \mathbf{a}_2^H$. If \mathbf{a}_1 and \mathbf{a}_2 are not linearly dependent, the condition in (74) is satisfied for any \mathbf{u} such that $\mathbf{u}^H \mathbf{b}_1 \neq 0$ and $\mathbf{u}^H \mathbf{b}_2 \neq 0$, and we can thus restrict attention to the case where \mathbf{a}_1 and \mathbf{a}_2 are linearly dependent. Here, we may without loss of generality assume

that $\mathbf{a}_1 = \mathbf{a}_2 = \mathbf{a}$ by absorbing any complex scalars into \mathbf{b}_1 and \mathbf{b}_2 .

Note however that given any $\epsilon > 0$ we can always find a point $\mathbf{s}_1 \in \mathbb{S}_\infty^2$, where $\mathbf{s}_1 \neq \mathbf{0}$ such that¹³ $\|\mathbf{a}^H \mathbf{s}_1\|^2 < \epsilon$. For any such \mathbf{s}_1 it follows that (cf. [23, Theorem 7.3.10])

$$\sigma_{\max}^2(\underline{\mathbf{X}}) = \max_{\mathbf{u} \in \mathbb{C}^2: \|\mathbf{u}\|=1} \|\mathbf{u}^H \underline{\mathbf{X}}\|^2 \leq (\|\mathbf{b}_1\|^2 + \|\mathbf{b}_2\|^2)\epsilon,$$

i.e., the maximal singular value of $\underline{\mathbf{X}}$ can be made arbitrarily small. However, this violates the assumed NVD property of the code as a small maximal singular value implies a small determinant, and concludes case a.

B. Case b

When \mathbf{G}_{11} and \mathbf{G}_{21} are full rank we can always find a vector \mathbf{u} such that $\mathbf{u}^H \mathbf{G}_{11}$ and $\mathbf{u}^H \mathbf{G}_{21}$ are linearly independent (thus satisfying the condition of Lemma 2) unless \mathbf{G}_{11} and \mathbf{G}_{21} are linearly dependent, i.e., when $\mathbf{G}_{11} = a\mathbf{G}_{21}$ for some $a \in \mathbb{C}$. However, in this case we have that $\mathbf{G}_{11}\mathbf{s}_1 = a\mathbf{G}_{21}\mathbf{s}_1$ for any \mathbf{s}_1 which implies that the columns of $\underline{\mathbf{X}}$ are linearly dependent, and the rank of $\underline{\mathbf{X}}$ is zero. This concludes case b.

C. Case c

In this case we may assume that $\mathbf{G}_{11} = \mathbf{b}_1 \mathbf{a}_1^H$ has rank one and \mathbf{G}_{21} has rank two (the opposite case is handled equivalently). Here, as both the set of \mathbf{u} for which $\mathbf{u}^H \mathbf{b}_1 = 0$ and where $\mathbf{u}^H \mathbf{G}_{21}$ is linearly dependent of \mathbf{a}_1^H have zero measure, we may pick \mathbf{u} such that $\mathbf{u}^H \mathbf{b}_1 \neq 0$ and such that $\mathbf{u}^H \mathbf{G}_{21}$ is linearly independent of \mathbf{a}_1^H , thus satisfying the conditions of Lemma 2. This concludes the proof of Theorem 7.

REFERENCES

- [1] E. Viterbo and J. Boutros, "A universal lattice code decoder for fading channels," *IEEE Trans. Inform. Theory*, vol. 45, no. 5, pp. 1639–1642, July 1999.
- [2] E. Agrell, T. Eriksson, A. Vardy, and K. Zeger, "Closest point search in lattices," *IEEE Trans. Inform. Theory*, vol. 48, no. 8, pp. 2201–2214, Aug. 2002.
- [3] M. O. Damen, H. E. Gamal, and G. Caire, "On maximum-likelihood detection and the search for the closest lattice point," *IEEE Trans. Inform. Theory*, vol. 49, no. 10, pp. 2389–2401, Oct. 2003.
- [4] A. D. Murugan, H. E. Gamal, M. O. Damen, and G. Caire, "A unified framework for tree search decoding: Rediscovering the sequential decoder," *IEEE Trans. Inform. Theory*, vol. 52, no. 3, pp. 933–953, Mar. 2006.
- [5] A. Ghasemmehdi and E. Agrell, "Faster recursions in sphere decoding," *IEEE Trans. Inform. Theory*, vol. 57, no. 6, pp. 3530–3536, June 2011.
- [6] L. Zheng and D. N. C. Tse, "Diversity and multiplexing: A fundamental tradeoff in multiple-antenna channels," *IEEE Trans. Inform. Theory*, vol. 49, no. 5, pp. 1073–1096, May 2003.
- [7] K. Kumar, G. Caire, and A. Moustakas, "Asymptotic performance of linear receivers in MIMO fading channels," *IEEE Trans. Inform. Theory*, vol. 55, no. 10, pp. 4398–4418, Oct. 2009.
- [8] M. Taherzadeh and A. K. Khandani, "On the limitations of the naive lattice decoding," *IEEE Trans. Inform. Theory*, vol. 56, no. 10, pp. 4820–4826, Oct. 2010.
- [9] J. Jaldén and P. Elia, "DMT optimality of LR-aided linear decoders for a general class of channels, lattice designs, and system models," *IEEE Trans. Inform. Theory*, vol. 56, no. 10, pp. 4765–4780, Oct. 2010.
- [10] J. Jaldén and B. Ottersten, "On the complexity of sphere decoding in digital communications," *IEEE Trans. Signal Processing*, vol. 53, no. 4, pp. 1474–1484, Apr. 2005.
- [11] J. Jaldén and B. Ottersten, "On the limits of sphere decoding," in *Proc. IEEE International Symposium on Information Theory, ISIT*, Adelaide, Australia, Sept. 2005, pp. 1691–1695.
- [12] D. Seethaler, J. Jaldén, C. Studer, and H. Bölcskei, "On the complexity distribution of sphere decoding," *IEEE Trans. Inform. Theory*, vol. 57, no. 9, pp. 5754–5768, Sept. 2011.
- [13] H. Yao and G. W. Wornell, "Structured space-time block codes with optimal diversity-multiplexing tradeoff and minimum delay," in *Proc. IEEE Global Telecommunications Conference, GLOBECOM*, San Francisco, CA, Dec. 2003.
- [14] J.-C. Belfiore, G. Rekaya, and E. Viterbo, "The golden code: A 2 x 2 full-rate space-time code with non-vanishing determinants," *IEEE Trans. Inform. Theory*, vol. 51, no. 4, pp. 1432–1436, Apr. 2005.
- [15] T. Kiran and B. Sundar Rajan, "STBC-schemes with non-vanishing determinant for certain number of transmit antennas," *IEEE Trans. Inform. Theory*, vol. 51, no. 8, pp. 2984–2992, Aug. 2005.
- [16] P. Elia, K. R. Kumar, S. A. Pawar, P. Vijay Kumar, and H.-F. Lu, "Explicit space-time codes achieving the diversity-multiplexing gain tradeoff," *IEEE Trans. Inform. Theory*, vol. 52, no. 9, pp. 3869–3884, Sept. 2006.
- [17] F. Oggier, G. Rekaya, J.-C. Belfiore, and E. Viterbo, "Perfect space-time block codes," *IEEE Trans. Inform. Theory*, vol. 52, no. 9, pp. 3885–3902, Sept. 2006.
- [18] P. Elia, B. A. Sethuraman, and P. Vijay Kumar, "Perfect space-time codes for any number of transmit antennas," *IEEE Trans. Inform. Theory*, vol. 53, no. 11, pp. 3853–3868, Nov. 2007.
- [19] B. Hassibi and H. Vikalo, "On the sphere-decoding algorithm I. Expected complexity," *IEEE Trans. Signal Processing*, vol. 53, no. 8, pp. 2806–2818, Aug. 2005.
- [20] —, "On the sphere-decoding algorithm II. Generalizations, second-order statistics, and applications to communications," *IEEE Trans. Signal Processing*, vol. 53, no. 8, pp. 2819–2834, Aug. 2005.
- [21] V. Pauli and L. Lampe, "On the complexity of sphere decoding for differential detection," *IEEE Trans. Inform. Theory*, vol. 53, no. 4, pp. 1595–1603, Apr. 2007.
- [22] S. Tavildar and P. Viswanath, "Approximately universal codes over slow-fading channels," *IEEE Trans. Inform. Theory*, vol. 52, no. 7, pp. 3233–3258, July 2006.
- [23] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1985.
- [24] B. Hassibi and B. M. Hochwald, "High-rate codes that are linear in space and time," *IEEE Trans. Inform. Theory*, vol. 48, no. 7, pp. 1804–1824, July 2002.
- [25] H. El Gamal, G. Caire, and M. O. Damen, "Lattice coding and decoding achieve the optimal diversity-multiplexing tradeoff of MIMO channels," *IEEE Trans. Inform. Theory*, vol. 50, no. 6, pp. 968–985, June 2004.
- [26] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*. Cambridge University Press, 1991.
- [27] E. Krätzel, *Lattice Points*. Berlin: Kluwer Academic Publishers, 1988.
- [28] P. Gritzmann and J. M. Wills, "Lattice points," in *Handbook of Convex Geometry*, P. M. Gruber and J. M. Wills, Eds. North-Holland, 1993, vol. B, ch. 3.2.
- [29] A. H. Banihashemi and A. K. Khandani, "On the complexity of decoding lattices using the Korkin-Zolotarev reduced basis," *IEEE Trans. Inform. Theory*, vol. 44, no. 1, pp. 162–171, Jan. 1998.
- [30] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, 2nd ed. Springer-Verlag New York Inc., 1998.
- [31] H. El Gamal and M. Damen, "Universal space-time coding," *IEEE Trans. Inform. Theory*, vol. 49, no. 5, pp. 1097–1119, May 2003.
- [32] B. A. Sethuraman, B. Sundar Rajan, and V. Shashidhar, "Full-diversity, high-rate, space-time block codes from division algebras," *IEEE Trans. Inform. Theory*, vol. 49, no. 10, pp. 2596–2616, Oct. 2003.
- [33] J.-C. Belfiore and G. Rekaya, "Quaternionic lattices for space-time coding," in *Proc. IEEE Information Theory Workshop, ITW*, Paris, France, Mar. 2003.
- [34] L. Zhao, W. Mo, Y. Ma, and Z. Wang, "Diversity and multiplexing tradeoff in general fading channels," *IEEE Trans. Inform. Theory*, vol. 53, no. 4, pp. 1549–1557, Apr. 2007.
- [35] O. Tirkkonen and R. Kashaev, "Combined information and performance optimization of linear MIMO modulations," in *Proc. IEEE International*

¹³Since $\mathbf{a}^H \mathbb{S}_\infty^2$ is the projection of the two dimensional Gaussian integer lattice onto a one dimensional space, one can also view the statement as an application of Dirichlet's box principle, cf. [8]. In general, the problem of finding non-zero integer vectors that are approximately orthogonal to a given vector is known as approximate integer relations (IRs), and is related to simultaneous Diophantine approximations [42].

- Symposium on Information Theory, ISIT*, Lausanne, Switzerland, June 2002.
- [36] J. Paredes, A. B. Gershman, and M. Gharavi-Alkhansari, "A 2×2 space-time code with non-vanishing determinant and fast maximum likelihood decoding," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP*, Honolulu, Hawaii, Apr. 2007, pp. 877–880.
 - [37] M. Samuel and M. P. Fitz, "Reducing the detection complexity by using 2×2 multi-strata space-time codes," in *Proc. IEEE International Symposium on Information Theory, ISIT*, Nice, France, June 2007.
 - [38] E. Biglieri, Y. Hong, and E. Viterbo, "On fast-decodable space-time block codes," *IEEE Trans. Inform. Theory*, vol. 55, no. 2, pp. 524–530, Feb. 2009.
 - [39] A. Tulino and S. Verdú, "Random matrix theory and wireless communications," *Foundations and Trends in Communications and Information Theory*, vol. 1, no. 1, June 2004.
 - [40] G. W. Stewart, "On the perturbation of pseudo-inverses, projections and linear least squared problems," *SIAM Review*, vol. 19, no. 4, pp. 634–662, Oct. 1977.
 - [41] F. Oggier, J.-C. Belfiore, and E. Viterbo, "Cyclic division algebras: A tool for space-time coding," *Foundations and Trends in Communications and Information Theory*, vol. 4, no. 1, pp. 1–95, 2007.
 - [42] I. V. L. Clarkson, "Approximation of linear forms by lattice points with applications to signal processing," Ph.D. dissertation, The Australian National University, 1997.