

Barriers for Recent Methods in Geodesic Optimization

W. Cole Franks  

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA

Philipp Reichenbach  

Institut für Mathematik, Technische Universität Berlin, Germany

Abstract

We study a class of optimization problems including matrix scaling, matrix balancing, multidimensional array scaling, operator scaling, and tensor scaling that arise frequently in theory and in practice. Some of these problems, such as matrix and array scaling, are convex in the Euclidean sense, but others such as operator scaling and tensor scaling are *geodesically convex* on a different Riemannian manifold. Trust region methods, which include box-constrained Newton’s method, are known to produce high precision solutions very quickly for matrix scaling and matrix balancing (Cohen et. al., FOCS 2017, Allen-Zhu et. al. FOCS 2017), and result in polynomial time algorithms for some geodesically convex problems like operator scaling (Garg et. al. STOC 2018, Bürgisser et. al. FOCS 2019). One is led to ask whether these guarantees also hold for multidimensional array scaling and tensor scaling.

We show that this is not the case by exhibiting instances with exponential *diameter bound*: we construct polynomial-size instances of 3-dimensional array scaling and 3-tensor scaling whose approximate solutions all have doubly exponential condition number. Moreover, we study convex-geometric notions of complexity known as margin and gap, which are used to bound the running times of all existing optimization algorithms for such problems. We show that margin and gap are exponentially small for several problems including array scaling, tensor scaling and polynomial scaling. Our results suggest that it is impossible to prove polynomial running time bounds for tensor scaling based on diameter bounds alone. Therefore, our work motivates the search for analogues of more sophisticated algorithms, such as interior point methods, for geodesically convex optimization that do not rely on polynomial diameter bounds.

2012 ACM Subject Classification Theory of computation → Algebraic complexity theory; Theory of computation → Continuous optimization; Mathematics of computing → Combinatorics

Keywords and phrases Geodesically Convex Optimization, Weight Margin, Moment Polytope, Diameter Bounds, Tensor Scaling, Matrix Scaling

Digital Object Identifier 10.4230/LIPIcs.CCC.2021.13

Related Version *Long Version*: <https://arxiv.org/abs/2102.06652>

Funding *Philipp Reichenbach*: acknowledges funding by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 787840).

Acknowledgements The authors thank Jason Altschuler, Peter Bürgisser, Visu Makam, Adam Sawicki and Michael Walter for helpful discussions. Furthermore, the authors thank the anonymous referees for helpful comments and suggestions.

1 Introduction

We study a class of optimization problems ubiquitous in theoretical computer science, machine learning, quantum information theory and statistics. The programs we consider are continuous optimization problems over matrix groups. More precisely, they can be posed as Euclidean norm minimization over the closure of a group orbit. The programs span two historically distinct contexts: In one context, the optimization problems are convex, and in the other they are not convex but rather *geodesically convex* on a suitable manifold.



© W. Cole Franks and Philipp Reichenbach;
licensed under Creative Commons License CC-BY 4.0
36th Computational Complexity Conference (CCC 2021).

Editor: Valentine Kabanets; Article No. 13; pp. 13:1–13:54

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



The *commutative setting*, in which the underlying group is Abelian, captures matrix scaling, matrix balancing and array scaling, which arise in scientific computing and optimal transport [17, 46]. Such problems fall into the framework of unconstrained geometric programming. Though these problems are convex, there are at least two reasons to study them further. Firstly, they are of such practical importance that speed matters. Naïvely applying powerful algorithms like ellipsoid and interior point methods can be impractically slow. Hence, it is important to understand when faster methods can succeed. Matrix scaling and balancing, in particular, have enjoyed some success stories - there are fast algorithms to obtain high precision solutions [16, 3], and there are more general upper bounds [14]. Secondly, the algorithms developed for the commutative setting are candidates for generalization to our second setting, which takes place in the less well-understood arena of geodesically convex optimization.

The second context, which we call the *noncommutative setting*, arises when the underlying group is non-Abelian. The noncommutative setting captures problems like operator and tensor scaling [25, 13], the quantum marginal problem [11] and statistical estimators such as Tyler's M estimator [23] and maximum likelihood estimates for matrix and tensor normal models [6]. Deciding whether the value of the optimization problem is zero or not is equivalent to deciding a central polynomial identity testing (P.I.T.) problem in invariant theory known as the *null cone* problem. It is hoped that efficient optimization algorithms will result in efficient algorithms for the null-cone problem. One approach to complexity lower bounds, geometric complexity theory, suggests that these P.I.T. problems should be in P [43, 26], and the optimization approach has resulted in polynomial time algorithms in some cases [25, 2]. The optimization problems that arise in the noncommutative setting are not convex in the Euclidean sense, but rather *geodesically convex*, a notion of convexity on a Riemannian manifold. Currently, the only implementable algorithms for geodesically convex optimization are analogues of gradient descent and trust region methods [1, 53, 2]. There are, as of yet, no efficiently implementable geodesically convex counterparts to the interior point or cutting plane methods.

In both the commutative and noncommutative settings, algorithms are typically analysed using two quantities. One is *diameter*, or how far approximate minimizers can be from the origin. The other is a geometric measure of well-conditionedness known as *margin* (or *gap* in the noncommutative case), which has several variants in the literature and appears in two primary ways. Firstly, the smaller the margin, the higher the degree of precision required to decide if the value of the optimization problem is zero or not [12, 30]. Secondly, the larger the margin, the smaller the diameter [48, 50, 12, 14]. In this paper we show the following:

- i) In the commutative setting, and in particular for array scaling, approximate minimizers for the functions we study can have doubly exponential condition number. That is, the problems have exponential diameter. As a consequence, popular classes of algorithms such as gradient descent and trust region methods cannot produce high-precision solutions in polynomial time in general. This result applies in the noncommutative setting as well, which provides evidence that even cutting plane methods are unlikely to produce high-precision solutions in polynomial time. This shows it is necessary to develop powerful methods like the interior point method in the geodesically convex setting.
- ii) In the commutative and noncommutative settings, we study the margin and gap, respectively, which appear in running time bounds for all existing algorithms. We prove that these measures can be exponentially small in the input size for several problems including array scaling and tensor scaling. In the commutative case, this gives

evidence that existing algorithms for array scaling do not run in near-linear time. In the noncommutative case, our results show that margin-based analyses like [12] cannot prove polynomial time guarantees for deciding the null cone problem for tensor scaling using trust region methods.

We use the remainder of the introduction to describe both settings in more detail, state our main results precisely, and discuss previous work. For both the commutative and noncommutative settings, we proceed in the following order. We start with an introduction and motivation of the setting, continue with diameter bounds and afterwards treat bounds on the margin and gap, respectively. We end each setting with a short discussion of the main proof techniques.

1.1 The commutative setting: matrix scaling and its relatives

Matrix scaling and array scaling

Consider the matrix scaling problem: given a nonnegative matrix A , find nonnegative diagonal matrices X, Y such that XAY is doubly stochastic (i.e. has row and column sums equal to one). The matrices, if they exist, can be found by the exceedingly simple and fast alternating minimization method known as *Sinkhorn's algorithm*. It is frequently used in practice, e.g. for quickly approximating the solution to optimal transport problems [17].

Like all other algorithms for matrix scaling, Sinkhorn's algorithm is typically analyzed through optimization. One finds that X and Y are $e^{\text{diag}(x)}, e^{\text{diag}(y)}$, where $x, y \in \mathbb{R}^n$ are solutions to the following optimization problem:

$$\inf_{x, y \in \mathbb{R}^n} \sum A_{ij} e^{x_i + y_j - \bar{x} - \bar{y}} \quad (1)$$

for $\bar{z} := \frac{1}{n} \sum z_i$ (c.f. [35]). Moreover, the infimum is greater than zero if and only if A is *approximately scalable*, i.e. the row and column sums of XAY can be made arbitrarily close to one for X, Y nonnegative, diagonal.

More generally, given a finite set $\Omega \subseteq \mathbb{R}^m$ and a nonnegative function $p : \Omega \rightarrow \mathbb{R}_{\geq 0}$, define the *capacity* [30] as the value of the unconstrained geometric program

$$\text{cap}(p) := \inf_{x \in \mathbb{R}^m} f_p(x) := \inf_{x \in \mathbb{R}^m} \sum_{\omega \in \Omega} p_{\omega} e^{\omega \cdot x}. \quad (2)$$

The capacity is positive if and only if zero is in the *Newton polytope* $\text{conv}(\text{supp } p)$. Matrix scaling arises when $m = 2n$ and $\Omega = \{(\varepsilon_i, \varepsilon_j) : i, j \in [n]\}$ for $\varepsilon_k := e_k - \frac{1}{n} \mathbb{1}_n$, where $e_k \in \mathbb{R}^n$ is the k^{th} canonical unit vector and $\mathbb{1}_n \in \mathbb{R}^n$ denotes the all-ones vector. In this case Equation (2) reduces to precisely Equation (1), and $\|\nabla \log f_p(x)\|$ measures the deviation of p from doubly stochastic.

Matrix balancing, in which we instead wish to find a scaling for which the i^{th} row and column sum match, arises when $m = n$ and $\Omega = \{e_i - e_j : i \neq j \in [n]\}$. When $m = 3n$ and $\Omega = \{(\varepsilon_i, \varepsilon_j, \varepsilon_k) : i, j, k \in [n]\}$ we obtain the 3-dimensional *array scaling problem*. In analogy to matrix scaling, in array scaling one has an array p of numbers in $(\mathbb{R}_{\geq 0}^n)^{\otimes 3}$ and seeks positive vectors $X, Y, Z \in \mathbb{R}_{\geq 0}^n$ so that the array q with entries $q_{ijk} = p_{ijk} X_i Y_j Z_k$ is *tristochastic*. That is, the sum over every *slice* is equal to one, i.e. $\sum_{j,k} q_{i_0,j,k} = \sum_{i,k} q_{i,j_0,k} = \sum_{i,j} q_{i,j,k_0} = 1$ for all $i_0, j_0, k_0 \in [n]$. If it is possible to satisfy these equations to arbitrary precision we say p is approximately scalable. As for matrix scaling, p is approximately scalable if and only if $\text{cap}(p) > 0$. In the same manner, we obtain d -dimensional array scaling for $m = dn$ and

$$\Omega = \Omega_{n,d} := \{\varepsilon_i : i \in [n]\}^d \subseteq (\mathbb{R}^n)^d. \quad (3)$$

We can think of subsets of $\Omega_{n,d}$ as d -uniform, d -partite hypergraphs. Up to an additive shift by $-\frac{1}{n}\mathbb{1}_{nd}$, the elements of $\Omega_{n,d}$ are indicator vectors of the edges in such hypergraphs. For $d = 2$, the matrix p is scalable if and only if the bipartite graph corresponding to $\text{supp } p$ contains a perfect matching, but this is not the case for $d \geq 3$ (indeed, d -partite hypergraph matching is NP-hard).

Algorithms for array scaling

Array scaling serves the same role for speeding up multimarginal transport as matrix scaling for optimal transport, and yet again there is a simple and fast alternating minimization algorithm that produces ε -tristochastic scalings in time $O(1/\varepsilon^2)$ [5, 39]. Moreover, algorithms to approximate the capacity arise in varied settings including radial isotropic position [33], entropy maximization [50], and approximate counting [7].

It is natural to ask if there are *high-precision* algorithms for array scaling with $\log(1/\varepsilon)$ dependence on the error and linear or mild dependence on the number of nonzero entries. For matrix scaling and matrix balancing, several works have shown that trust regions and interior point methods *can* obtain such guarantees [16, 3]. Our work is concerned with whether the performance of such algorithms carries over to array scaling and the computation of the capacity in general.

1.1.1 Diameter lower bounds

Guarantees for many iterative algorithms in convex optimization require *diameter bounds*, or bounds on the distance R from the starting point to an ε -approximate solution. Trust region methods, also called *box-constrained Newton's method*, are iterative algorithms that, at each step, move to the best solution within a typically small distance D of the previous solution. By their nature, trust region methods take at least R/D steps to produce an ε -approximate solution. Gradient descent for Lipschitz functions also depends quadratically on a diameter bound, and cutting plane methods typically use diameter bounds to control the volume of a starting region.

Known diameter upper and lower bounds

For matrix scaling and matrix balancing, it has been shown in [16] that one may take $R = O(n \log(w_A/\varepsilon))$, where w_A is the ratio between the sum of the entries of the matrix and the least nonzero entry. For 3-dimensional array scaling, the best upper bound of which we are aware is $R = O(n^{3/2}2^{6n} \log(1/\varepsilon))$, which follows from the general upper bound of [50] on diameter bounds for unconstrained geometric programming. There is also a diameter bound for array scaling in the multimarginal transport context that is polynomial in the input size assuming the tensor has no nonzero entries [39].

Regarding diameter *lower* bounds, in the context of computing maximum entropy distributions it was shown that there is some bounded set $\Omega \subset \mathbb{Z}^m$ in a $\text{poly}(m)$ size ball such that there are *no* ε -approximate minimizers of norm $\text{poly}(m, \log 1/\varepsilon)$ for f_p as in Equation (2) [50].

Main theorem

Where do the polynomial diameter bounds for matrix scaling (i.e. 2-dimensional array scaling) transition to the superpolynomial diameter bounds for general Ω ? We show that this transition takes place in the next simplest problem, the 3-dimensional array scaling problem.

► **Theorem 1.1.** *There is an absolute constant $C > 0$ and an array $p_{ijk} \in (\mathbb{R}_{\geq 0}^n)^{\otimes 3}$ with $O(n)$ nonzero entries, each of bit-complexity $O(n)$, that satisfies the following property. For all $0 < \varepsilon \leq \exp(-Cn^2 \log n)$ and $(x, y, z) \in \mathbb{R}^{3n}$, if*

$$f_p(x, y, z) \leq \text{cap}(p) + \varepsilon$$

then $\|(x, y, z)\|_2 = \Omega(2^{n/3} \log(1/\varepsilon))$.

To emphasize that the difficulties do not lie in an additive vs multiplicative approximation, we remark that our array p has unit sum and $\text{cap}(p) = 1/2$. By a simple duplication trick, the same bound holds for d -dimensional array scaling with $d \geq 3$; see Corollary 3.7.

Implications of Theorem 1.1 and relation to the literature

Theorem 1.1 shows that trust region methods for array scaling with polynomial step size cannot provide high-precision solutions in $\text{poly}(n, \log(1/\varepsilon))$ time for $d \geq 3$. Moreover, gradient descent on the Lipschitz convex function $\log f_p$ has a bounded step size, and so also cannot provide high precision solutions in polynomial time.

In [50, Section 2.1] the authors ask whether there is Ω whose elements are Boolean (up to an additive shift) with a superpolynomial diameter lower bound. As subsets of $\Omega_{n,d}$ are automatically of this form, we answer their open problem in the affirmative. Our lower bound on $\log R$ is tight up to constant factors by the diameter upper bound from [50] mentioned above; moreover the logarithmic dependence on ε is best possible. Determining the correct constant in the exponent is an interesting open direction. We believe that that the requirement that ε is very small is an artifact of our specific construction and proof strategy, and thus can probably be relaxed significantly.

Lastly, we remark that [14] bounds the diameter for f_p by a polynomial in the *facet gap*, i.e. the minimum distance between an element of $\text{supp } p$ and an affine hull of a facet of the Newton polytope. The construction in Theorem 1.1 has exponentially small facet gap; see Corollary 3.6.

1.1.2 Margins: the geometry of scaling problems

Many computational aspects of the capacity rely on the convex geometry of the finite set $\Omega \subseteq \mathbb{R}^m$. Consider the following quantity, which we call the *margin* of Ω . The margin is the minimum *positive* distance from a convex hull of a subset of Ω to the origin. Formally,

► **Definition 1.2 (Margin).** *For a finite set $\Omega \subseteq \mathbb{R}^m$, define the margin $\gamma(\Omega)$ by*

$$\gamma(\Omega) := \min \{ \text{dist}(0, \text{conv}(S)) \mid S \subseteq \Omega, 0 \notin \text{conv}(S) \}.$$

We point out that for all considered capacity problems in this paper, the margin is actually the *weight margin* (c.f. [12] and our Definition 4.3) of a certain group representation. For example, the margin for array scaling is the weight margin for tensor scaling. We now discuss how the margin enters in decision problems and diameter bounds.

Margin as a precision parameter for the decision problem

To illustrate how the margin enters the decision problem of whether $\text{cap}(p) > 0$, consider matrix scaling. To certify that the capacity of a matrix is nonzero, we compute ε -doubly stochastic scalings for some ε smaller than the distance to doubly stochastic attained by any matrix that is *not* approximately scalable. This turns out to be precisely $\gamma(\Omega_{n,2})$. More

generally, it is a classical fact that for p with support contained in Ω , the gradient $\nabla \log f_p(x)$ can take any value in the Newton polytope of p . Thus, $\text{cap}(p) > 0$ if and only if there is some x with $\|\nabla \log f_p(x)\| \leq \gamma(\Omega)$.

For matrix scaling and matrix balancing, it is known that $\gamma(\Omega)$ is on the order of $n^{-3/2}$, despite the exponential number of subsets $S \subseteq \Omega$! This luck can be attributed to the extraordinary geometry of Ω in these cases, whose elements form the rows of a totally unimodular matrix (up to a shift). On the other hand, for d -dimensional array scaling for $n = 2$, the margin $\gamma(\Omega_{2,d})$ is on the order of the margin of the d -dimensional hypercube $\{\pm 1\}^d$, which satisfies $\gamma(\{\pm 1\}^d) = d^{-\frac{d}{2}(1+o(1))}$ by [4]. However, between the extreme cases $\Omega_{n,2}$ (matrix scaling) and $\Omega_{2,d}$ (the hypercube), very little is known.

Margin and related quantities for diameter bounds

In addition to their role in the decision problem, margins and related quantities can be used to prove diameter bounds for Equation (2). The work [12] proves the diameter bound $\text{poly}(\gamma(\Omega)^{-1}, \log(1/\varepsilon))$. In [50] it is shown that the diameter is polynomial in the logarithm of the minimum nonzero p_ω and a quantity called the *unary facet complexity*. The latter is defined as the maximal length of an integer normal vector of a face of the Newton polytope $\text{conv}(\text{supp } p)$. In the case of d -dimensional arrays, one can use Cramer's rule to crudely bound the unary facet complexity by $(d+1)^{dn}$. In the case when 0 is in the relative interior of the Newton polytope, [48] has shown that there is a minimizer with Euclidean norm $O(\log|\text{supp } p|/\eta)$, where η is the distance from 0 to the boundary of the Newton polytope. The diameter bounds in [48, 50] were used to design ellipsoid methods that are tractable even for $|\text{supp } p|$ very large, and in [14] they were used to bound the running time of interior point methods.

Main theorem

One is led to ask if the margin remains large for array scaling when $d \geq 3$. We show that this is not the case. In fact, the margin becomes exponentially small in nd for $d \geq 3$. What follows is stated in more detail later in Theorem 2.1.

► **Theorem 1.3.** *Let $d \geq 3$ and $n \geq 2$. Let $\Omega_{n,d} = \{\varepsilon_i : i \in [n]\}^d \subseteq (\mathbb{R}^n)^d$, where $\varepsilon_j := e_j - \frac{1}{n}\mathbb{1}_n$. There exists a constant $C > 0$, independent of n and d , such that $\gamma(\Omega_{n,d}) \leq 2^{-Cnd}$.*

That is, there are d -dimensional arrays $p \in (\mathbb{R}_{\geq 0}^n)^{\otimes d}$ such that the d -tuple of marginals of p is at distance at most 2^{-Cnd} from $\frac{1}{n}(\mathbb{1}_n, \dots, \mathbb{1}_n)$, yet the support of p does not admit an array with uniform marginals, i.e. $\text{cap}(p) = 0$. We note that the support of the array p we construct has $O(nd)$ elements.

Implications of Theorem 1.3 and relation to the literature

We remark that the construction yields a tensor whose Newton polytope has a facet exponentially close to the origin. Therefore, the bound proved in [14] on the number of iterations for interior point methods on 3-tensors is $\Omega(k^{3/2} + k^{1/2} \log(1/\varepsilon))$ for tensors with $O(k)$ nonzero entries.

Theorem 1.3 aligns with existing results showing that the $d > 2$ array case is more complex than the matrix case. Indeed, it is known that the polytope of arrays with uniform marginals, known as the *d -index axial assignment polytope*, has many more vertices when $d \geq 3$ and that the vertices can have exponential entries [40]. In contrast, for $d = 2$ this polytope (known as the Birkhoff-von Neumann polytope) has integral vertices by the Birkhoff-von Neumann theorem.

The exponential rate of decay in Theorem 1.3 is tight up to log factors: [12, Theorem 6.10 Item 3] shows that the margin for d -dimensional array scaling is at least $(n\sqrt{d})^{-dn-1}$. It is interesting to ask whether the true bound is $2^{-\Theta(nd)}$ as in our upper bound or $2^{-\Theta(nd(\log n + \log d))}$ as in the lower bound. [4] shows that the latter is correct in the case $n = 2$.

1.1.3 Proof techniques for the commutative setting

We first discuss the techniques for proving our margin bounds. Theorem 1.3 is proven by explicit construction of witness sets $\Gamma_{n,d} \subseteq \Omega_{n,d} := \{\varepsilon_i : i \in [n]\}^d$, i.e. $0 \notin \text{conv}(\Gamma_{n,d})$ but zero is exponentially close to $\text{conv}(\Gamma_{n,d})$. This is done by using that $\sum_i n^{-1}\varepsilon_i$ is the unique way to express zero as a convex combination of the ε_i , compare Lemma 2.2, and by heavily exploiting the combinatorics of $\Omega_{n,d}$. For example, in the case $d = 3$ and $n \geq 3$ the key combinatorial idea builds on a construction by Kravtsov in [38]. Kravtsov's motivation is to characterize the non-integer vertices of the 3-index axial assignment polytope. He explicitly constructs a certain non-integer vertex with maximal support [38, Theorem 1 with $k = 0$] which has an exponentially small entry.

By definition of the 3-index axial assignment polytope, the support of this vertex corresponds to a subset $S \subseteq \Omega_{n,3}$ with $0 \in \text{conv}(S)$. Removing the element of S corresponding to the small entry in Kravtsov's vertex yields our witness set $\Gamma_{n,3}$ with a convex hull very close to zero. In fact, the whole idea generalizes (in a technical way) whenever $d = 6r - 3$, $r \geq 1$ and $n \geq 3$, see section 2.3. For $n = 2$ and $d \geq 3$, the bound follows from the existing work [4], as mentioned before. While the construction in that work via $\{-1, 1\}$ matrices yields a stronger bound, we provide a different construction of $\{-1, 1\}$ matrices¹, which has the additional property of *freeness*. The latter will prove useful when we adapt Theorem 1.3 to the noncommutative case.

We now discuss the proof of the diameter lower bound, Theorem 1.1. The high level idea is as follows. We first construct a subset $\Omega_0 \subseteq \Omega_{n,3}$ with $0 \in \text{conv}(\Omega_0)$ such that there is another element $\omega \in \Omega_{n,3}$ exponentially close to $\text{conv}(\Omega_0)$, much like our construction of the witness set for small margin discussed above. We then choose an appropriate array p supported on $\Omega_0 \cup \omega$. This suggests that the only approximate minimizers of f_p have a very large component in the direction x from ω to $\text{conv}(\Omega_0)$, because as $y \in \mathbb{R}^m$ tends to a minimizer of f_p the term $e^{y \cdot \omega}$ should vanish compared to the others. This reasoning requires that y is approximately a multiple of x ; to enforce this we also ensure that zero is far into the relative interior of $\text{conv}(\Omega_0)$.

The structure of this argument bears some similarity to that in [50], which uses the construction of [4]. The main difference is that the set $\Omega_{n,3}$ in the 3-dimensional array scaling problem consists of vectors of very specific structure: up to an additive shift of $-\frac{1}{n}\mathbb{1}_{3n}$, they are Boolean vectors in \mathbb{R}^{3n} with exactly one nonzero entry among indices in the intervals $[1, n]$, $[n + 1, 2n]$, $[2n + 1, 3n]$. Thus, our construction of Ω_0 must consist of vectors of this special form and not simply bounded integral vectors as in [50]. This is the main additional technical contribution of our construction.

¹ The $(-1, 1)$ matrices from our construction are obtained by replacing all two's in the entries of A_{2r} (6) with -1 .

1.2 The noncommutative setting

In the noncommutative setting, we consider a group G acting on \mathbb{C}^m .² The optimization problem we investigate is given by the *capacity* of a vector $v \in \mathbb{C}^m$ (c.f. [12]):

$$\text{cap}(v) := \inf_{g \in G} f_v(g) := \inf_{g \in G} \|g \cdot v\|^2. \quad (4)$$

For the majority of this paper we work with the *tensor scaling action*, in which $G = \text{SL}(n, \mathbb{C})^d$, the group of d -tuples of complex matrices with determinant one, acts on $v \in (\mathbb{C}^n)^{\otimes d}$ by $(g_1, \dots, g_d) \cdot v = (g_1 \otimes \dots \otimes g_d)v$. The corresponding representation is always denoted by $\pi_{n,d}$. Sometimes we also consider the *operator scaling action*, in which $\text{SL}(n)^2$ acts on $v \in (\mathbb{C}^n)^{\otimes 2} \otimes \mathbb{C}^k$ by $(g_1, g_2) \cdot v = (g_1 \otimes g_2 \otimes I_k)v$.

Though Equation (4) looks quite different from Equation (2), one can show that restricting Equation (4) to a certain Abelian subgroup of G (a torus) and making a change of variables yields an instance of Equation (2) (c.f. [12]). For example, restricting the tensor scaling action to the diagonal matrices in G amounts precisely to the array scaling problem from the previous subsection. Likewise, restricting to diagonal matrices in the operator scaling action yields an instance of matrix scaling.

Relation to null cone problem and Geometric Complexity Theory

We study Equation (4) because it is deeply connected to invariant theory through a well-known connection between group orbits and invariant polynomials: zero is in the closure of an orbit of a vector v if and only if every non-constant homogeneous G -invariant polynomial vanishes on v , i.e. if v is in the null-cone. Null-cone membership is a well-studied polynomial identity testing (P.I.T.) problem. One approach to complexity lower bounds, geometric complexity theory, suggests that null-cone membership should be in P [43, 26].

Solving Equation (4) directly allows one to study the null-cone problem through optimization: one notes that $\text{cap}(v) = 0$ if and only if v is in the null cone. In fact, Equation (4) is a geodesically convex optimization problem over a certain Riemannian manifold. Algebraic and optimization-based algorithms have, independently and nearly concurrently, resulted in polynomial time algorithms for nearly the same set of P.I.T. problems arising in invariant theory [22, 43, 25, 34, 20, 2], including the null-cone problem for the operator scaling and simultaneous conjugation action. However, neither approach has succeeded in solving the null-cone problem for the 3-tensor action. Recent degree lower bounds for invariant polynomials for the 3-tensor action pose significant challenges for the algebraic approach [21]. It is natural to ask whether the optimization approach can overcome these challenges.

Algorithms for computing the capacity

A nonzero tensor $w = g \cdot v$ attains the capacity when w has all *quantum marginals* equal to I_n/n . The quantum marginals of a tensor w , analogous to the sums along slices of an array, are the three $n \times n$ matrices $M_1 M_1^\dagger, M_2^\dagger, M_3 M_3^\dagger$ for the $n \times n^2$ matrices M_1, M_2, M_3 known as *flattenings* of $w/\|w\|$. For operator scaling, the capacity is attained when the first two quantum marginals are I_n/n . To compute the capacity, existing algorithms attempt to find g such that the quantum marginals of $g \cdot v$ are all close to I_n/n . There are alternating minimization algorithms that can attain distance ε in time $\text{poly}(n, 1/\varepsilon)$ [25, 13], and for the

² Technically we require that G is a reductive group over \mathbb{C} which acts rationally on \mathbb{C}^m . All the group actions in this paper satisfy this assumption.

operator scaling this is possible in $\text{poly}(n, \log(1/\varepsilon))$ time [2]. However, for 3-tensor scaling, running time $\text{poly}(1/\varepsilon)$ is not sufficient to efficiently decide null-cone membership, and the only algorithms with $\log(1/\varepsilon)$ dependence on ε have an exponential dependence on n [12].

To explain the increased complexity, we discuss a noncommutative analogue of the Newton polytope known as the *moment polytope*, denoted $\Delta_G(v)$. In particular, $0 \notin \Delta_G(v)$ if and only if v is in the null-cone (i.e. $\text{cap}(v) = 0$).³ For tensor scaling, the moment polytope is the set of tuples of spectra of the quantum marginals as w ranges over $\overline{G \cdot v}$, shifted by $-\frac{1}{n}(\mathbb{1}_n, \mathbb{1}_n, \mathbb{1}_n)$. The *gap* of the action of G , i.e. the minimum positive distance from 0 to a moment polytope $\Delta_G(v)$, is a noncommutative generalization of the margin. Whereas the operator scaling and simultaneous conjugation actions have polynomially large gaps, we show that the gap for the tensor scaling action is exponentially small. Scaling algorithms amount to outer ε -approximation algorithms for $\Delta_G(v)$, which is why $\text{poly}(1/\varepsilon)$ -time algorithms do not suffice to decide null-cone membership. Like for the margin, the smaller gap corresponds to a larger diameter, which is why so far no algorithm has had running time $\text{poly}(n, \log(1/\varepsilon))$.

1.2.1 Diameter lower bound for noncommutative scaling

Here we describe how diameter bounds cause the state-of-the-art algorithms to be slow for the tensor scaling action. We begin by discussing geodesically convex optimization. In general Equation (4) is not convex, but rather *geodesically convex*. That is, G can be viewed as a manifold in such a way that the function $g \mapsto \|g \cdot v\|^2$ is convex along “geodesics” of the form $\gamma(t) = e^{tH}g$ for H Hermitian. The manifold we consider is not exactly G but rather a quotient P of it; we will make this more precise later in Section 4.5. For $G = \text{SL}(n)^d$, the manifold P is the set of tuples of positive-definite matrices with determinant one. P is equipped with the geometry on positive-definite matrices known in statistics as the Fisher-Rao metric, and studied in depth in e.g. [9]. Though we do not need many details of this geometry here, one can think of the distance between $g, h \in G$ as a bound on the logarithms of the singular values of $g^{-1}h$. In particular, the geodesic “ball” of radius R about the identity in G is the intersection of G with the set $\{U \exp(A) : A \text{ Hermitian}, \|A\|_F \leq R, U \text{ unitary}\}$. Note that the ball of radius $\sqrt{n}R$ includes all elements of G whose singular values are in $[e^{-R}, e^R]$.⁴

The existing algorithms to compute Equation (4) adapt simple first order methods, such as gradient descent, and second order methods, such as trust regions, to the geodesically convex setting [53, 2, 12]. As in the commutative case, to run in polynomial time such algorithms require that an ε -approximate solution is contained in a geodesic ball of radius $\text{poly}(n^d, \log(1/\varepsilon))$. However, for 3-tensors we have the following diameter lower bound.

► **Theorem 1.4** (Noncommutative diameter lower bound). *There is a constant $C > 0$ such that the following holds. For all $\varepsilon \leq \exp(-Cn^2 \log n)$, there is a tensor $v = v(\varepsilon) \in (\mathbb{C}^n)^{\otimes 3}$ with $O(n)$ nonzero entries of bit complexity $O(\log n + \log(1/\varepsilon))$, and a geodesic ball $B = B(\varepsilon)$ of radius $\Omega(2^{n/3} \log(1/\varepsilon))$ about the identity in $\text{SL}(n)^3$, such that*

$$\inf_{g \in B} \|g \cdot v\|^2 \geq \text{cap}(v) + \varepsilon.$$

To emphasize that the difficulties are not caused by requiring additive approximation, we remark that the vector v satisfies $1/4 \leq \text{cap}(v) \leq 1$ and $1/2 \leq \|v\| \leq 1$. A duplication trick analogous to Corollary 3.7 yields the same diameter bound for $d \geq 3$, but for the action of G simultaneously on a tuple of tensors rather than on a single one. See Corollary 4.24.

³ Moment polytope membership is an interesting problem in and of itself; for $d = 3$, for generic $v \in (\mathbb{C}^n)^{\otimes 3}$, $\Delta_G(v)$ is the *Kronecker polytope* arising in representation theory and quantum information theory. Deciding membership in this polytope is known to be in $\text{NP} \cap \text{coNP}$ but not known to be in P [10].

⁴ We define exponentials, Hermitian-ness, and Frobenius norm on tuples by treating them as block diagonal matrices.

Implications of Theorem 1.4 and relation to the literature

Theorem 1.4 shows that trust region methods with constant step size cannot ε -approximate the capacity in $\text{poly}(n, 1/\varepsilon)$ time for 3-tensors. It also shows that cutting plane methods are unlikely to do so. Cutting plane methods, such as ellipsoid, require an exponential bound on the volume of a known region containing an approximate optimizer. This is the case for Rusciano’s non-constructive query upper bound for cutting plane methods on manifolds of non-positive curvature [47], which is essentially tight [32]⁵. The volume of a ball in the manifold we consider grows exponentially in the radius (see Section 4.5), so this query bound will be exponential. Regarding tightness, the best upper bound known to the authors for the diameter bound in the noncommutative case is $O(n(\sqrt{3}n)^{1+3n} \log(1/\varepsilon))$, which can be deduced from the diameter and margin bounds [12, Proposition 5.6, Theorem 6.10]. This matches our lower bound up to logarithmic factors in the exponent. As with Theorem 1.1, Theorem 1.4 holds only values of ε that are very small (though still of polynomial bit-complexity). It would be very interesting to prove a version of Theorem 1.1 for ε larger than the gap, which is $\exp(-O(n))$. This would imply that trust region methods cannot solve the null-cone problem for the 3-tensor action in polynomial time.

1.2.2 Gaps: the geometry of noncommutative scaling problems

In analogy to the commutative case, one typically attempts to certify $\text{cap}(v) > 0$, i.e. $0 \in \Delta_G(v)$, by finding a tensor $g \cdot v$ such that all the quantum marginals are close to $\frac{1}{n}I_n$. In order to certify $\text{cap}(v) > 0$ their distance to $\frac{1}{n}(I_n, I_n, \dots, I_n)$ must be at most a certain quantity, which we call the *gap*.

► **Definition 1.5 (Gap).** *The gap⁶ for the d -tensor scaling problem is*

$$\gamma_G(\pi_{n,d}) := \min \{ \text{dist}(0, \Delta_G(v)) \mid v \in (\mathbb{C}^n)^{\otimes d}, v \neq 0, 0 \notin \Delta_G(v) \}.$$

If the gap is exponentially small, high-precision algorithms will be necessary to decide if $\text{cap}(v) > 0$. In operator scaling, the gap is known to be $\Omega(n^{-3/2})$ [29], which explains why we do not need high-precision algorithms for the decision problem in that case. In addition to its role in the decision problem, the inverse of the gap⁷ is used to control the diameter bound [12]! In that sense, the presence of a small gap can explain both the need for high precision algorithms and the slowness of existing high-precision algorithms. We show that, indeed, the tensor scaling action has an exponentially small gap for $d \geq 3$.

► **Theorem 1.6.** *There is a constant $C > 0$ such that for all $d \geq 3$ and $n \geq 2$, there are non-zero tensors $v \in (\mathbb{C}^n)^{\otimes d}$ such that $0 \notin \Delta_G(v)$ but $\text{dist}(0, \Delta_G(v)) \leq 2^{-Cdn}$. That is, the gap for d -tensor scaling satisfies*

$$\gamma_G(\pi_{n,d}) \leq 2^{-Cdn}.$$

A detailed statement on bounds for the gap can be found in Theorem 4.11, and we show in Appendix C how to fill in the missing values of n, d to obtain Theorem 1.6. Since the gap is larger than the margin (c.f. Proposition 4.6), Theorem 1.6 is at least as tight as Theorem 1.3, i.e. the exponent Cnd is tight up to an $O(\log n + \log d)$ factor.

⁵ [32] applies to the hyperbolic plane, which is a totally geodesic submanifold of the manifold P we consider.

⁶ This notion can be defined similarly for any rational representation π of a reductive group G , see Definition 4.3. This definition of the gap is already described in [12].

⁷ actually, a smaller quantity known as *weight margin*

Interestingly, for local dimension $n = 2$ [42, Main result] shows that $\text{dist}(0, \Delta_G(v))^2$ for some moment polytope $\Delta_G(v) \neq 0$ tends for $d \rightarrow \infty$ to the Gamma distribution $\Gamma(1/2, 2d)$, where $2d$ is the rate parameter. Therefore, the witnesses of the exponential behaviour in Theorem 4.11(a) are quite rare. Moreover, the authors numerically found several tensors of format $(\mathbb{C}^2)^{\otimes d}$ with $\text{dist}(0, \Delta_G(v))$ at most $\exp(-d)$; Theorem 1.6 confirms that this exponential behavior is the case for all n and d .

Margin and gap results for other group actions

In addition to the tensor scaling action, we also consider some other actions of groups G of interest in computational invariant theory. The first is the action of the special linear group on the space of homogeneous d -forms $\mathbb{C}[x_1, \dots, x_n]_d$, in which $G = \text{SL}(n)$ acts by $g \cdot p(x) = p(g^{-1}x)$ for $p \in \mathbb{C}[x_1, \dots, x_n]_d$. Homogeneous d -forms were among the objects studied earliest in computational invariant theory, and much of the theory was developed to catalogue invariants of the $\text{SL}(n)$ action on forms [52]. Still, deciding null-cone membership for $d = 3$ seems challenging. After extending the definition of the gap to other group actions in Section 4, we explain the difficulty by showing that the gap for this action is also inverse exponential in n as soon as $d \geq 3$, see Theorem 4.17. This shows that the diameter bound in [12] becomes exponentially large in n .

The other group action we consider is the action of $\text{SL}(n)^d$ on *quivers* with d vertices. A quiver is a directed multigraph, and a quiver representation is a labelling of the vertex set Q_0 of the quiver with finite-dimensional vector spaces and the edge set Q_1 with a linear map from the vector space at the tail of the edge to the vector space at the head of the edge. Given a quiver representation A with vertices labeled by \mathbb{C}^{n_x} for $x \in Q_0$ and edges $e : x \rightarrow y$ labeled with matrices A_e , the group $G = \prod_{x \in Q_0} \text{SL}(n_x)$ acts on A by $(g \cdot A)_e = g_y A_e g_x^{-1}$. Quiver representations include the operator scaling action, and an action used to bound the Brascamp-Lieb constant in analysis. In Section 4.6 we show that the (*weight*) *margin* can become exponentially small as the number of vertices grows. For this, we exhibit a quiver with $d - 1$ arrows, d vertices of dimension n and weight margin $O(n^{-d})$, see Theorem 4.25. This bound shows that the diameter bound computed in [12] can become exponentially large in d . Furthermore, when allowing n copies of each arrow in the constructed quiver, i.e. $n(d - 1)$ arrows in total, we can ensure the same bound for the gap, Theorem 4.25.

1.2.3 Proof technique in the noncommutative case: Freeness

Regarding the idea of the proof, we may transfer both the diameter lower bound and the gap upper bound to the commutative case by virtue of the tensors we construct having *free support*.

A tensor has free support if any two distinct $(d - 1)$ -dimensional slices of the tensor have disjoint support. This condition ensures that, even after being acted on by any diagonal group elements, the tensor's quantum marginals are all diagonal. This allows us to restrict to the action of the diagonal matrices and thereby reduce to the commutative (array scaling) case. Thus, we may obtain the same bounds on the tensor gap as for the array margin. However, this requires additional care to ensure freeness of our constructions. This is why we cannot naïvely use the construction of [4] for d -tensors with $n = 2$. Regarding the noncommutative diameter bound, we show that for tensors with free support the diameter bound matches that of the commutative problem obtained by restricting to the diagonal. To do this, we project the group elements to the set of diagonal elements, and use the properties of spaces of non-positive curvature to show that this projection moves the point nearer to the origin and decreases the function value.

The idea and the concept of freeness generalize to rational representations of reductive groups [24].⁸ The key statement is given in full generality in Proposition 4.8. This proposition is needed to prove bounds on the gap for the action on homogeneous polynomials and for the action on quivers. Interestingly, in [21] the concept of freeness is used in a similar way⁹ to prove exponential lower bounds on the degree of invariants for actions on cubic forms and 3-tensors. There, *free* is called *uncramped* and it is used crucially to prove closedness of certain orbits.

Freeness also played a role in the numerical results by Sawicki and Maciążek, which were obtained by applying the algorithm of [41] to several free tensors of local dimension two.

1.3 Organization of the paper

We begin with the commutative case, which is split into the study of the margin in Section 2 and diameter bounds in Section 3. Then we move to the noncommutative case in Section 4. The appendix contains some representation-theoretic background and proofs of technical lemmas, as well as a glossary of notation.

2 The geometry of commutative scaling problems

The purpose of this section is to show the following theorem on the margin of d -dimensional array scaling. Recall that the latter arises for $\Omega_{n,d} := \{\varepsilon_i : i \in [n]\}^d \subseteq (\mathbb{R}^n)^d$.

► **Theorem 2.1** (Margin for array scaling). *The margin of $\Omega_{n,d} \subseteq (\mathbb{R}^n)^d$ is bounded as follows.*

- (a) *If $n = 2$ and $d \geq 3$, then $\gamma(\Omega_{2,d}) \leq 2^{-\frac{d}{2}+1}$.*
- (b) *If $n \geq 3$ and $d = 3$, then $\gamma(\Omega_{n,3}) \leq 2^{-n+1}$.*
- (c) *If $n \geq 3$ and $d = 6r - 3$ for some integer $r \geq 2$, then*

$$\gamma(\Omega_{n,d}) \leq \frac{\sqrt{6}}{(n-1)\sqrt{r}} 2^{-r(n-1)+1} \leq 2^{-r(n-1)+1} = 2^{-\frac{(d+3)(n-1)}{6}+1}.$$

By “padding” the tensors appropriately, one sees that a bound for $\gamma(\Omega_{n,d})$ also applies to $\gamma(\Omega_{n,d+1})$ (see Proposition C.1). Combining this result with Theorem 2.1 above implies Theorem 1.3 from the introduction. The next three subsections each prove one of the parts of Theorem 2.1; the construction for part (a) with $n = 2$ is slightly different and the construction for part (c), $d > 3$ builds on the one for part (b), $d = 3$.

To prove the results, we will frequently use the following simple lemma. Recall that an *affine linear combination* of $v_1, \dots, v_k \in \mathbb{R}^m$ is $\lambda_1 v_1 + \dots + \lambda_k v_k$ for $\lambda_i \geq 0$, $\sum_{i=1}^k \lambda_i = 1$. The affine hull $\text{Aff}(S)$ of a set $S \subset \mathbb{R}^m$ is the set of all affine linear combinations of finite subsets of S , or equivalently the affine space (i.e. translate of a subspace) of lowest dimension containing S .

► **Lemma 2.2.** *In \mathbb{R}^n we have*

$$\sum_{i=1}^n \frac{1}{n} \varepsilon_i = 0_n \tag{5}$$

and this is the only affine linear combination of $\varepsilon_1, \dots, \varepsilon_n$ giving zero.

⁸ This concept is also implicitly contained in [49, Lemma 7.1] and can at least be traced back to [18] as *strong orthogonality*.

⁹ Indeed, [21, Theorem 6.5] is used to show the vanishing of the moment map at a vector. First, freeness is used as in Proposition 4.8 to ensure that one can restrict to the moment map for the maximal torus. Second, condition (2) of [21, Theorem 6.5] just states that the moment map for the torus action vanishes at the vector.

$$A_4 = \begin{pmatrix} * & * & * & * \\ & * & & \\ * & & * & * \\ * & & * & * \end{pmatrix}, \quad A_6 = \begin{pmatrix} * & * & * & * & * & * \\ & * & & & * & * \\ * & & * & * & * & * \\ * & & * & * & * & * \\ * & & * & * & * & * \\ * & & * & * & * & * \end{pmatrix}$$

■ **Figure 1** The positions of the ones in A_4 and A_6 are marked by $*$ in the following figure and the cells are colored according to whether they belong to A_2, B_1, B_2 or B_3 .

Proof. One calculates directly that $\sum_i \frac{1}{n} \varepsilon_i = 0_n$. To show uniqueness of this affine combination, we note that the vectors $e_2, \dots, e_n, \mathbb{1}_n$ are linearly independent. Thus, $\varepsilon_2, \dots, \varepsilon_n$ are linearly independent. On the other hand, $\varepsilon_1, \dots, \varepsilon_n$ are linearly dependent. Therefore $\{(\lambda_1, \dots, \lambda_n) \in \mathbb{R}^n \mid \sum_i \lambda_i \varepsilon_i = 0_n\}$ is a one-dimensional subspace of \mathbb{R}^n , which yields the uniqueness of the affine linear combination. ◀

2.1 Local dimension two: the hypercube

In this subsection we prove part (a) of Theorem 2.1 by showing that the margin of $\Omega_{2,d}$ is exponentially small in d . This follows from [4], but we present a new construction which has the additional property of *freeness*, which we discuss later in Section 4. Recall that

$$\Omega_{2,d} = \{(\varepsilon_{i_1}, \dots, \varepsilon_{i_d}) \mid i_1, \dots, i_d \in [2]\} \subseteq (\mathbb{R}^2)^d.$$

In the following we construct a subset of $\Omega_{2,d}$, which witnesses the exponentially small margin. For this, we construct a matrix with entries in $[2]$, and each row of the matrix will correspond to an element of $\Omega_{2,d}$. For example, the row $(1, 2, 2)$ would correspond to $(\varepsilon_1, \varepsilon_2, \varepsilon_2) \in \Omega_{2,3}$. To do so, we begin with the matrices

$$A_2 := \begin{pmatrix} 1 & 1 \\ 2 & 1 \end{pmatrix}, \quad B_1 := \begin{pmatrix} 1 & 1 \\ 2 & 2 \end{pmatrix}, \quad B_2 := \begin{pmatrix} 1 & 2 \\ 2 & 2 \end{pmatrix}, \quad B_3 := \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix},$$

and define recursively

$$A_{2r+2} := \begin{pmatrix} & & & B_1 \\ & A_{2r} & & \vdots \\ & & & B_1 \\ B_2 & \cdots & B_2 & B_3 \end{pmatrix} = \begin{pmatrix} A_2 & B_1 & \cdots & B_1 \\ B_2 & B_3 & \ddots & \vdots \\ \vdots & \ddots & \ddots & B_1 \\ B_2 & \cdots & B_2 & B_3 \end{pmatrix} \tag{6}$$

for $r \geq 1$. Figure 1 is supplied as a visualization aid.

We remark that the entry of A_{2r} at position (i, j) is independent of r and denote it by $a(i, j)$. We set for $r \geq 1$

$$\Gamma_{2,2r} := \{(\varepsilon_{a(i,1)}, \varepsilon_{a(i,2)}, \dots, \varepsilon_{a(i,2r)}) \mid i \in [2r]\} \subseteq \Omega(\pi_{2,2r}) \subseteq (\mathbb{R}^2)^{2r},$$

$$\Gamma_{2,2r+1} := \{(\varepsilon_{a(i,1)}, \varepsilon_{a(i,2)}, \dots, \varepsilon_{a(i,2r)}, \varepsilon_{\chi(i)}) \mid i \in [2r]\} \subseteq \Omega(\pi_{2,2r+1}) \subseteq (\mathbb{R}^2)^{2r+1},$$

where $\chi: \mathbb{N} \rightarrow \{1, 2\}$, $i \mapsto i \bmod 2$. That is, $\Gamma_{2,2r}$ is the subset of $\Omega_{2,2r}$ induced by the rows of A_{2r} and $\Gamma_{2,2r+1}$ is obtained by alternatingly appending ε_1 or ε_2 to the $2r$ -many elements of $\Gamma_{2,2r}$.

13:14 Barriers in Geodesic Optimization

► **Lemma 2.3.** *For $r \geq 1$ it holds that $0 \notin \text{Aff}(\Gamma_{2,2r})$ and $0 \notin \text{Aff}(\Gamma_{2,2r+1})$.*

Proof. By construction, $0 \in \text{Aff}(\Gamma_{2,2r+1})$ implies $0 \in \text{Aff}(\Gamma_{2,2r})$, so it suffices to prove $0 \notin \text{Aff}(\Gamma_{2,2r})$. We proceed by induction on $r \geq 1$. For $r = 1$, it is clear that $0 \notin \text{Aff}(\Gamma_{2,2}) \subseteq \mathbb{R}^2 \times \{\varepsilon_1\}$. Now assume that $0 \notin \text{Aff}(\Gamma_{2,2r})$. For the sake of contradiction, let

$$\sum_{i=1}^{2r+2} \lambda_i (\varepsilon_{a(i,1)}, \varepsilon_{a(i,2)}, \dots, \varepsilon_{a(i,2r+2)}) = 0 \in (\mathbb{R}^2)^{2r+2} \quad (7)$$

be an affine linear combination of $\Gamma_{2,2r+2}$. Then equation (7) gives in each of the $(2r+2)$ -many \mathbb{R}^2 -components the affine linear combination $2^{-1}(\varepsilon_1 + \varepsilon_2) = 0$, by Lemma 2.2. Considering the scalar factor of ε_1 in the first, the penultimate and the last \mathbb{R}^2 -component respectively, we conclude

$$\underbrace{\sum_{j=1}^{r+1} \lambda_{2j-1}}_{\text{first}} = \frac{1}{2} = \underbrace{\lambda_{2r+2} + \sum_{j=1}^r \lambda_{2j-1}}_{\text{penultimate}} = \frac{1}{2} = \underbrace{\lambda_{2r+2} + \sum_{j=1}^{r+1} \lambda_{2j-1}}_{\text{last}}$$

by construction of A_{2r+2} . Hence, $\lambda_{2r+2} = 0$ using the first and last component. Furthermore, the first and penultimate column give $\lambda_{2r+1} = \lambda_{2r+2} = 0$. Therefore, the first $2r$ -many components in Equation (7) show $0 \in \text{Aff}(\Gamma_{2,2r})$, which contradicts our induction hypothesis. ◀

► **Lemma 2.4.** *For $r \geq 1$ it holds that $\text{dist}(0, \text{conv}(\Gamma_{2,2r})) \leq 2^{-r+\frac{1}{2}}$ and $\text{dist}(0, \text{conv}(\Gamma_{2,2r+1})) \leq 2^{-r+\frac{1}{2}}$.*

Proof. We first prove the inequality for $\text{conv}(\Gamma_{2,2r})$. For $i \in [2r]$ let $\omega_i := (\varepsilon_{a(i,1)}, \dots, \varepsilon_{a(i,2r)}) \in (\mathbb{R}^2)^{2r}$ be the weight in $\Gamma_{2,2r}$ that corresponds to the i^{th} row of A_{2r} . Consider the convex combination

$$(x_1, \dots, x_{2r}) := 2^{-r}(\omega_{2r-1} + \omega_{2r}) + \sum_{l=1}^{r-1} 2^{-l-1}(\omega_{2l-1} + \omega_{2l}) \in (\mathbb{R}^2)^{2r}. \quad (8)$$

Note that $x_i \in \mathbb{R}^2$. We will argue that $(x_1, \dots, x_{2r}) = 2^{-r+1}(0_2, \dots, 0_2, \varepsilon_1)$. Since x is a convex combination of the elements in $\Gamma_{2,2r}$, the statement then follows from $\|\varepsilon_1\| = \sqrt{2}^{-1}$.

We consider A_{2r} like in its construction (6) as a $r \times r$ block matrix with block entries being 2×2 matrices. For $m \in [r]$ the two weights ω_{2m-1} and ω_{2m} correspond to the m^{th} block row of A_{2r} and have the same scalar factor in (8). Hence, whenever for $i \in [2r]$ the i^{th} column of the m^{th} block row of A_{2k} contains exactly one entry equal to one (and so the other entry equals two), then the contribution of ω_{2m-1} and ω_{2m} to x_i cancels due to $\varepsilon_1 + \varepsilon_2 = 0_2$. In particular, in (8) all contributions of block entries equal to B_1 cancel. Therefore the last column of A_{2r} gives

$$x_{2r} = 2^{-r}(\varepsilon_1 + \varepsilon_1) = 2^{-r+1}\varepsilon_1.$$

Furthermore, $x_1 = x_3 = \dots = x_{2r-1} = 0_2$ using that also the first columns of A_2 , of B_2 and of B_3 contain exactly one entry equal to one. For $r = 1$ we are done. If $r \geq 2$, then reading off the second column of A_{2r} , we find

$$x_2 = \underbrace{2^{-2}(\varepsilon_1 + \varepsilon_1)}_{\text{first block row}} + \underbrace{2^{-r}(\varepsilon_2 + \varepsilon_2)}_{\text{last block row}} + \sum_{l=2}^{r-1} \underbrace{2^{-l-1}(\varepsilon_2 + \varepsilon_2)}_{\text{middle rows}} = 2^{-1}(\varepsilon_1 + \varepsilon_2) = 0_2.$$

Analogously, as B_1 does not contribute we compute for $j = 2, 3, \dots, r - 1$ that

$$x_{2j} = \underbrace{2^{-j-1}(\varepsilon_1 + \varepsilon_1)}_{j^{\text{th}} \text{ block row}} + \underbrace{2^{-r}(\varepsilon_2 + \varepsilon_2)}_{\text{last block row}} + \sum_{l=j+1}^{r-1} \underbrace{2^{-l-1}(\varepsilon_2 + \varepsilon_2)}_{\text{in between rows}} = 2^{-j}(\varepsilon_1 + \varepsilon_2) = 0_2,$$

because the second columns of B_2 and B_3 are, respectively, $(2, 2)^T$ and $(1, 1)^T$. This proves the inequality in the case $\Gamma_{2,2r}$.

By construction, for $\Gamma_{2,2r+1}$ the same convex combination works, because the last \mathbb{R}^2 -component does not contribute as the entries of the weights alternate between ε_1 and ε_2 . ◀

Finally, Lemma 2.3 and Lemma 2.4 together yield Theorem 2.1(a), noting that for odd $d = 2r + 1$ one has $-r + 1/2 = -(d/2) + 1$.

2.2 3-tensors

The main goal of this section is to show that the margin of $\Omega_{n,3}$ is exponentially small in n , i.e. to show Theorem 2.1(b). To do so, we set

$$\mathfrak{W}_n := \bigcup_{s=2}^n \{(s, 1, s), (s, s, 1), (s - 1, s, s)\} \subseteq [n] \times [n] \times [n] \tag{9}$$

and consider the corresponding subset

$$\Gamma_{n,3} := \{(\varepsilon_i, \varepsilon_j, \varepsilon_k) \mid (i, j, k) \in \mathfrak{W}_n\} \subseteq \Omega_{n,3}. \tag{10}$$

The key combinatorial idea, which is presented in the following lemma, is due to [38, Theorem 1 with $k = 0$].¹⁰ According to [38] the special case $k = 0$ is already contained in [37, Theorem 9].

► **Lemma 2.5.** *Let $n \geq 3$. For $(i, j, k) \in [n]^3 \setminus (\mathfrak{W}_n \cup \{(1, 1, 1)\})$ set $\lambda_{i,j,k} := 0$. Moreover, define*

$$\lambda_{1,1,1} := 2^{-n+1}, \quad \lambda_{1,2,2} := 1 - 2^{-n+1}, \quad \lambda_{n,1,n} = \lambda_{n,n,1} := 2^{-1}$$

and for $s = 2, 3, \dots, n - 1$

$$\lambda_{s,1,s} = \lambda_{s,s,1} := 2^{-n+s-1}, \quad \lambda_{s,s+1,s+1} := 1 - 2^{-n+s}.$$

Then the following equations hold:

$$\left(\forall i \in [n] : \sum_{j,k=1}^n \lambda_{i,j,k} = 1 \right), \left(\forall j \in [n] : \sum_{i,k=1}^n \lambda_{i,j,k} = 1 \right), \left(\forall k \in [n] : \sum_{i,j=1}^n \lambda_{i,j,k} = 1 \right). \tag{11}$$

In particular, $\sum_{i,j,k} \lambda_{i,j,k} = n$.

Proof. This is [38, Theorem 1 with $k = 0$]. Alternatively, the statement can be checked by straightforward computation. ◀

¹⁰In [38] Kravtsov extensively studies so-called complete r -noninteger vertices (r -CNVs) of the three-index axial assignment polytope. For $k \in \{0, 1, \dots, n - 2\}$, [38, Theorem 1] states explicitly a $(3n - 2 - k)$ -CNV, among these we use the $(3n - 2)$ -CNV (i.e. $k = 0$). Moreover, [38, Theorem 2] states that such r -CNVs of the three-index axial assignment polytope actually only occur for $r \in \{2n, 2n + 1, \dots, 3n - 2\}$, and the later theorems in [38] fully characterize the r -CNVs and study their combinatorial properties.

13:16 Barriers in Geodesic Optimization

► **Example 2.6.** To visualize Lemma 2.5 it is helpful to consider the slices Λ_i given by $(\Lambda_i)_{j,k} = \lambda_{i,j,k}$. For $n = 4$ one has

$$\Lambda_1 = \frac{1}{8} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 7 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \Lambda_2 = \frac{1}{8} \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 6 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

$$\Lambda_3 = \frac{1}{8} \begin{pmatrix} 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4 \end{pmatrix}, \quad \Lambda_4 = \frac{1}{8} \begin{pmatrix} 0 & 0 & 0 & 4 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 4 & 0 & 0 & 0 \end{pmatrix}.$$

For $n = 5$ one has

$$\Lambda_1 = \frac{1}{16} \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 15 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad \Lambda_2 = \frac{1}{16} \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 14 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix},$$

$$\Lambda_3 = \frac{1}{16} \begin{pmatrix} 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 12 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad \Lambda_4 = \frac{1}{16} \begin{pmatrix} 0 & 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 4 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 8 \end{pmatrix}, \quad \Lambda_5 = \frac{1}{16} \begin{pmatrix} 0 & 0 & 0 & 0 & 8 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 8 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

► **Lemma 2.7.** For $n \geq 3$, it holds that $\text{dist}(0, \text{conv}(\Gamma_{n,3})) \leq 2^{-n+1}$.

Proof. Define $\lambda_{i,j,k} \geq 0$ for all $i, j, k \in [n]$ as in Lemma 2.5. Note that $\sum_{i=1}^n \varepsilon_i = 0$; thus Lemma 2.5 implies

$$\sum_{i,j,k} \lambda_{i,j,k}(\varepsilon_i, \varepsilon_j, \varepsilon_k) = 0_{3n}, \quad \text{equivalently} \quad -2^{-n+1}(\varepsilon_1, \varepsilon_1, \varepsilon_1) = \sum_{(i,j,k) \in \mathfrak{W}_n} \lambda_{i,j,k}(\varepsilon_i, \varepsilon_j, \varepsilon_k).$$

Normalizing the latter equation we obtain

$$x := -\frac{1}{c 2^{n-1}}(\varepsilon_1, \varepsilon_1, \varepsilon_1) \in \text{conv}(\Gamma_{n,3}), \quad \text{where } c := \sum_{(i,j,k) \in \mathfrak{W}_n} \lambda_{i,j,k} = n - 2^{-n+1} \geq \sqrt{3}.$$

Finally, $\|\varepsilon_1\|^2 \leq 1$ implies $\|x\| \leq c^{-1} 2^{-n+1} \sqrt{3} \leq 2^{-n+1}$. ◀

To finish the proof of Theorem 2.1(b) we are left to show $0 \notin \text{conv}(\Gamma_{n,3})$. We actually prove the stronger statement $0 \notin \text{Aff}(\Gamma_{n,3})$.

► **Lemma 2.8.** The zero vector is not contained in the affine hull of $\Gamma_{n,3}$.

Proof. For a proof by contradiction we assume $0 \in \text{Aff}(\Gamma_{n,3})$. Then there exist $a_s, b_s, c_s \in \mathbb{R}$ for $s = 2, 3, \dots, n$ such that $\sum_s a_s + b_s + c_s = 1$ and

$$\sum_{s=2}^n (a_s(\varepsilon_s, \varepsilon_1, \varepsilon_s) + b_s(\varepsilon_s, \varepsilon_s, \varepsilon_1) + c_s(\varepsilon_{s-1}, \varepsilon_s, \varepsilon_s)) = (0_n, 0_n, 0_n) \in (\mathbb{R}^n)^3.$$

In each of the three \mathbb{R}^n -components we obtain 0_n as an affine linear combination of $\varepsilon_1, \dots, \varepsilon_n$. Applying Lemma 2.2 to the coefficient of ε_{s-1} in the first component, respectively to the coefficient of ε_s in the second and third component yields

$$a_{s-1} + b_{s-1} + c_s = \frac{1}{n} \quad \text{for } s = 2, 3, \dots, n \quad (12)$$

$$\text{respectively} \quad b_s + c_s = a_s + c_s = \frac{1}{n} \quad \text{for } s = 2, 3, \dots, n \quad (13)$$

where we necessarily set $a_1 = b_1 := 0$. Equation (12) for $s = 2$ is $c_2 = n^{-1}$ and hence $a_2 = b_2 = 0$ by (13) for $s = 2$. But now (12) for $s = 3$ gives $c_3 = n^{-1}$ and we can proceed inductively to conclude $c_s = n^{-1}$ and $a_s = b_s = 0$ for all $s = 2, 3, \dots, n$. This gives the contradiction $1 = \sum_{s=2}^n (a_s + b_s + c_s) = \frac{n-1}{n}$, so we must have $0 \notin \text{Aff}(\Gamma_{n,3})$. Another contradiction arises when one applies Lemma 2.2 to the coefficient ε_n in the first component, which yields $a_n + b_n = n^{-1}$. ◀

2.3 d -tensors

In this subsection we show that the margin of $\Omega_{n,d}$ is inverse exponential in nd for $n, d \geq 3$, proving part (c) of Theorem 2.1.

Let us give some intuition for our construction. The main idea is to recycle the construction from the previous subsection for some multiple of n , i.e. considering \mathfrak{W}_{rn} for $r \geq 2$. Thereby, the main challenge is to ensure that the constructed subset of $\Omega_{n,d}$ does not contain zero in its convex hull. We can try to extend the elements of $\Omega_{n,3}$ to elements of $\Omega_{n,d}$. One natural idea is duplicate each component $d/3$ times, i.e. when $d = 6$ the vector $(\varepsilon_i, \varepsilon_j, \varepsilon_k) \in \Omega_{n,3}$ becomes $(\varepsilon_i, \varepsilon_i, \varepsilon_j, \varepsilon_j, \varepsilon_k, \varepsilon_k) \in \Omega_{n,6}$. However, we need a subset of $\Omega_{n,d}$ with rn many elements to imitate the construction from the previous subsection. We still extend the elements of $\Omega_{n,3}$ in this way, but will additionally “shift” and “twist” by some functions $\sigma_1, \dots, \sigma_{2r-1}: [rn] \rightarrow [n]$, so that the elements of our set will look like

$$\left(\varepsilon_{\sigma_1(i)}, \dots, \varepsilon_{\sigma_{d/3}(i)}, \varepsilon_{\sigma_1(j)}, \dots, \varepsilon_{\sigma_{d/3}(j)}, \varepsilon_{\sigma_1(k)}, \dots, \varepsilon_{\sigma_{d/3}(k)} \right)$$

for $d/3 = 2r - 1$ and (i, j, k) in \mathfrak{W}_{rn} . We now set about choosing the functions σ_k . For this, let $n \geq 3$ and fix a natural number $r \geq 2$. It is convenient to use an *adjusted* modulo n function $\text{mod}' n$ that takes values in $[n]$, i.e. instead of zero it outputs n . For $i \in [r]$ we consider

$$\begin{aligned} \sigma_i: [rn] &\rightarrow [n], \quad j \mapsto \left\lfloor \frac{j + (i - 1)}{r} \right\rfloor \text{ mod}' n \\ \sigma_{r+i} &:= \sigma_1 \circ (r - i + 1 \ r + 1): [rn] \rightarrow [n] \end{aligned}$$

where $(r - i + 1 \ r + 1)$ denotes the corresponding transposition in the symmetric group of $[rn]$.¹¹ We only need the first $2r - 1$ of these functions and combine them to obtain

$$\sigma: [rn] \rightarrow [n]^{2r-1}, \quad j \mapsto (\sigma_1(j), \sigma_2(j), \dots, \sigma_{2r-1}(j)).$$

▶ **Example 2.9.** For $r = 3$ the functions $\sigma_1, \sigma_2, \dots, \sigma_6$ are sketched by the following table.

j	1	2	3	4	5	6	...	$3n - 5$	$3n - 4$	$3n - 3$	$3n - 2$	$3n - 1$	$3n$
σ_1	1	1	1	2	2	2	...	$n - 1$	$n - 1$	$n - 1$	n	n	n
σ_2	1	1	2	2	2	3	...	$n - 1$	$n - 1$	n	n	n	1
σ_3	1	2	2	2	3	3	...	$n - 1$	n	n	n	1	1
σ_4	1	1	2	1	2	2	...	$n - 1$	$n - 1$	$n - 1$	n	n	n
σ_5	1	2	1	1	2	2	...	$n - 1$	$n - 1$	$n - 1$	n	n	n
σ_6	2	1	1	1	2	2	...	$n - 1$	$n - 1$	$n - 1$	n	n	n

For $r = 3$ and $n = 5$ the functions $\sigma_1, \sigma_2, \dots, \sigma_6$ are given by the following table.

¹¹ We stress that we always take σ_1 (and *not* σ_i) to define σ_{r+i} .

13:18 Barriers in Geodesic Optimization

j	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
σ_1	1	1	1	2	2	2	3	3	3	4	4	4	5	5	5
σ_2	1	1	2	2	2	3	3	3	4	4	4	5	5	5	1
σ_3	1	2	2	2	3	3	3	4	4	4	5	5	5	1	1
σ_4	1	1	2	1	2	2	3	3	3	4	4	4	5	5	5
σ_5	1	2	1	1	2	2	3	3	3	4	4	4	5	5	5
σ_6	2	1	1	1	2	2	3	3	3	4	4	4	5	5	5

► **Remark 2.10.** By construction, each element of $[n]$ is attained exactly r -times by σ_k , $k \in [2r - 1]$. Moreover, the definition of $\sigma_1, \dots, \sigma_r$ yields that σ is injective.

For $i, j, k \in [rn]$ we introduce the short-hand

$$\varepsilon_{\sigma(i)} := (\varepsilon_{\sigma_1(i)}, \varepsilon_{\sigma_2(i)}, \dots, \varepsilon_{\sigma_{2r-1}(i)}) \in (\mathbb{R}^n)^{2r-1}$$

$$\varepsilon_{\sigma(i), \sigma(j), \sigma(k)} := (\varepsilon_{\sigma_1(i)}, \dots, \varepsilon_{\sigma_{2r-1}(i)}, \varepsilon_{\sigma_1(j)}, \dots, \varepsilon_{\sigma_{2r-1}(j)}, \varepsilon_{\sigma_1(k)}, \dots, \varepsilon_{\sigma_{2r-1}(k)}) \in (\mathbb{R}^n)^{6r-3}$$

and we set¹²

$$\mathfrak{J}_r := \{(s, 1, s), (s, s, 1) \mid s = 2, 3, \dots, r\} \subseteq \mathbb{Z}^3.$$

In the following we show that the convex hull of the set

$$\Gamma_{n,6r-3} = \{\varepsilon_{\sigma(i), \sigma(j), \sigma(k)} \mid (i, j, k) \in \mathfrak{W}_{rn} \setminus \mathfrak{J}_r\} \subseteq \Omega_{n,6r-3} \subseteq \left((\mathbb{R}^n)^{2r-1}\right)^3$$

does not contain the zero vector, but is very close to it.

► **Lemma 2.11.** For $n \geq 3$ and $r \geq 2$ it holds that $0 \notin \text{Aff}(\Gamma_{n,6r-3})$.

Below we give the proof in the special case $r = 3$, in which all main ideas of the general proof become apparent and visible. The proof for the general statement is given in Appendix D and certainly looks technical at a first encounter. Therefore, we strongly suggest that the reader first reads the proof for $r = 3$ below.

Proof of Lemma 2.11 for $r = 3$. For the sake of contradiction assume that $0 \in \text{Aff}(\Gamma_{n,15})$. Then there are coefficients $a_s, b_s, c_s \in \mathbb{R}$, where $2 \leq s \leq 3n$, such that $a_2 = a_3 = b_2 = b_3 = 0$, $\sum_s (a_s + b_s + c_s) = 1$ and

$$\sum_{s=2}^{3n} (a_s \varepsilon_{\sigma(s), \sigma(1), \sigma(s)} + b_s \varepsilon_{\sigma(s), \sigma(s), \sigma(1)} + c_s \varepsilon_{\sigma(s-1), \sigma(s), \sigma(s)}) = 0 \in (\mathbb{R}^n)^{15}. \quad (14)$$

The bulk of our work will consist of proving the equations

$$b_2 + c_2 = b_3 + c_3 = \dots = b_{3n} + c_{3n} \quad (15)$$

$$a_2 + c_2 = a_3 + c_3 = \dots = a_{3n} + c_{3n}. \quad (16)$$

¹²One could suggest to consider the set $\{\varepsilon_{\sigma(i), \sigma(j), \sigma(k)} \mid (i, j, k) \in \mathfrak{W}_{rn}\}$, but this still won't ensure that zero is not in the convex hull. The intuition behind is, that $\Gamma_{n,3}$ from the last section is "nearly at the limit", i.e. $0 \notin \text{conv}(\Gamma_{n,3})$ but $0 \in \text{conv}(\Gamma_{n,3} \cup \{(\varepsilon_1, \varepsilon_1, \varepsilon_1)\})$. Now the function σ "introduces $2r - 2$ additional linear relations" as $\varepsilon_{\sigma(i)} \in (\mathbb{1}_n^\perp)^{2r-1}$, since the orthogonal complement $\mathbb{1}_n^\perp \subseteq \mathbb{R}^n$ has codimension one while $(\mathbb{1}_n^\perp)^{2r-1} \subseteq (\mathbb{R}^n)^{2r-1}$ has codimension $2r - 1$. Thus, it is reasonable to remove $2r - 2$ many elements from \mathfrak{W}_{rn} .

From here we will derive a contradiction. We now set about proving Equations (15) and (16). Rewrite the left-hand-side of Equation (14) as the collection for $k \in [5]$ of the following affine linear combinations of $\varepsilon_1, \dots, \varepsilon_n$ in \mathbb{R}^n :

$$\sum_{s=2}^{3n} (a_s \varepsilon_{\sigma_k(s)} + b_s \varepsilon_{\sigma_k(s)} + c_s \varepsilon_{\sigma_k(s-1)}) = 0 \tag{17}$$

$$\sum_{s=2}^{3n} (a_s \varepsilon_{\sigma_k(1)} + b_s \varepsilon_{\sigma_k(s)} + c_s \varepsilon_{\sigma_k(s)}) = 0 \tag{18}$$

$$\sum_{s=2}^{3n} (a_s \varepsilon_{\sigma_k(s)} + b_s \varepsilon_{\sigma_k(1)} + c_s \varepsilon_{\sigma_k(s)}) = 0. \tag{19}$$

If we expand each expression as an affine linear combination of the ε_l , then by Lemma 2.2 the coefficient of ε_l must be n^{-1} for all $l \in [n]$. Translating this for equation (17) with $k = 2, l = 2, \dots, n$ and using Example 2.9 we obtain

$$(a_{m-3} + a_{m-2} + a_{m-1}) + (b_{m-3} + b_{m-2} + b_{m-1}) + (c_{m-2} + c_{m-1} + c_m) = \frac{1}{n} \tag{20}$$

for $m = 6, 9, 12, \dots, 3n$. A similar calculation for $k = 1, 3$ and $l = 2, \dots, n$ shows Equation (20) holds for all $5 \leq m \leq 3n + 1$, where we set $c_{3n+1} := 0$.

Similarly for Equation (18) with $l = 2, \dots, n$ and $k = 1, 2, 3$ we obtain for $4 \leq m \leq 3n$ that

$$(b_{m-2} + c_{m-2}) + (b_{m-1} + c_{m-1}) + (b_m + c_m) = \frac{1}{n} \tag{21}$$

and the same equations with “ b ” replaced by “ a ” when considering Equation (19).

In the following we prove Equation (15). Subtracting (21) from (21) with values of m differing by one, we deduce that

$$\begin{aligned} b_2 + c_2 &= b_5 + c_5 = \dots = b_{3n-1} + c_{3n-1} \\ b_3 + c_3 &= b_6 + c_6 = \dots = b_{3n} + c_{3n}, \\ \text{and } b_4 + c_4 &= b_7 + c_7 = \dots = b_{3n-2} + c_{3n-2}. \end{aligned}$$

Next we deduce Equation (15) by showing $b_2 + c_2 = b_3 + c_3 = b_4 + c_4$.

To do so, we apply Lemma 2.2 to (18) for the coefficient of ε_2 using Example 2.9, which yields for $k = 4, 5$ the equations

$$(b_3 + c_3) + (b_5 + c_5) + (b_6 + c_6) = \frac{1}{n} \tag{22}$$

$$(b_2 + c_2) + (b_5 + c_5) + (b_6 + c_6) = \frac{1}{n} \tag{23}$$

respectively. Subtracting the two shows $b_2 + c_2 = b_3 + c_3$, and we have $b_3 + c_3 = b_4 + c_4$ via subtracting (22) from (21) for $m = 6$. This completes the proof of Equation (15); using Equation (19) we similarly deduce Equation (16).

To get a contradiction we show that $a_s = b_s = c_s = 0$ for all $s = 2, 3, \dots, 3n$. For this, we set $a := \sum_s a_s$ and $b := \sum_s b_s$, and recall that we have defined $a_2 = a_3 = b_2 = b_3 = 0$. This time we use Lemma 2.2 applied to the coefficient of ε_1 in (17), in (18) and in (19) respectively for $k = 1$ to get

$$c_2 + c_3 + c_4 = \frac{1}{n}, \quad a + c_2 + c_3 = \frac{1}{n} \quad \text{and} \quad b + c_2 + c_3 = \frac{1}{n} \tag{24}$$

13:20 Barriers in Geodesic Optimization

respectively. We deduce from these three equations that $a = b = c_4$. Furthermore, $b_2 = b_3 = 0$ shows that (21) for $m = 4$ is $b_4 + (c_2 + c_3 + c_4) = n^{-1}$. Subtracting from the latter the left-hand equation in (24) yields $b_4 = 0$. Similarly, $a_4 = 0$ follows from $a_2 = a_3 = 0$ and the analogous equation of (21) with a 's replaced by b 's.

Now, (20) for $m = 5$ simplifies to $c_3 + c_4 + c_5 = n^{-1}$. Thus, $c_2 = c_5$ with (24) and therefore $a_5 = b_5 = 0$ by (15), (16) and $a_2 = b_2 = 0$. This simplifies (20) for $m = 6$ to $c_4 + c_5 + c_6 = n^{-1}$. Hence, $c_3 = c_6$ as we also have $c_3 + c_4 + c_5 = n^{-1}$ and we get via (15) and (16) that $a_6 = b_6 = 0$. The latter in turn shows that (20) for $m = 7$ becomes $c_5 + c_6 + c_7 = n^{-1}$, so $c_4 = c_7$ and $a_7 = b_7 = 0$ by, again, (15) and (16).

It should have become apparent that we can proceed inductively in the same manner with (20) for $m = 5, \dots, 3n + 1$; thereby using (15) and (16) to deduce $a_s = b_s = 0$ for all $s = 2, 3, \dots, 3n$. In particular, $a = b = c_4 = 0$. Finally, Equation (15) implies $c_4 = c_s$ for all $s = 2, 3, \dots, 3n$, which gives the desired contradiction. \blacktriangleleft

We finish the proof of part (c) of Theorem 2.1 by showing the following Lemma.

► **Lemma 2.12.** *Let $n \geq 3$ and $r \geq 2$. Then*

$$\text{dist}(0, \text{conv}(\Gamma_{n,6r-3})) \leq \frac{\sqrt{6}}{(n-1)\sqrt{r}} 2^{-r(n-1)+1} \leq 2^{-r(n-1)+1}.$$

Proof. We set $N := rn$ and for $i, j, k \in [N]$ we set $\lambda_{i,j,k}$ as in Lemma 2.5 applied for the dimension N . Then Equation (11) of Lemma 2.5 yields

$$\begin{aligned} & \sum_{i,j,k=1}^N \lambda_{i,j,k} (\varepsilon_{\sigma(i)}, \varepsilon_{\sigma(j)}, \varepsilon_{\sigma(k)}) \\ &= \sum_{i,j,k=1}^N \lambda_{i,j,k} (\varepsilon_{\sigma(i)}, 0, 0) + \sum_{i,j,k=1}^N \lambda_{i,j,k} (0, \varepsilon_{\sigma(j)}, 0) + \sum_{i,j,k=1}^N \lambda_{i,j,k} (0, 0, \varepsilon_{\sigma(k)}) \\ &= \sum_{i=1}^N (\varepsilon_{\sigma(i)}, 0, 0) + \sum_{j=1}^N (0, \varepsilon_{\sigma(j)}, 0) + \sum_{k=1}^N (0, 0, \varepsilon_{\sigma(k)}) = \sum_{i=1}^N \varepsilon_{\sigma(i), \sigma(i), \sigma(i)} = 0 \in (\mathbb{R}^n)^{6r-3}, \end{aligned}$$

where we used in the last step equation (5) and Remark 2.10, i.e. that each element of $[n]$ is attained exactly r -many times by all $\sigma_k: [rn] \rightarrow [n]$, $k \in [2r-1]$. Because \mathfrak{W}_N contains the support of λ apart from the element $(1, 1, 1)$, we have

$$\sum_{(i,j,k) \in \mathfrak{W}_N \setminus \mathfrak{J}_r} \lambda_{i,j,k} \varepsilon_{\sigma(i), \sigma(j), \sigma(k)} \tag{25}$$

$$= -\lambda_{1,1,1} \varepsilon_{\sigma(1), \sigma(1), \sigma(1)} - \sum_{(i,j,k) \in \mathfrak{J}_r} \lambda_{i,j,k} \varepsilon_{\sigma(i), \sigma(j), \sigma(k)} =: x \in (\mathbb{R}^n)^{6r-3}, \tag{26}$$

which is an element in the positive cone of $\Gamma_{n,6r-3} = \{\varepsilon_{\sigma(i), \sigma(j), \sigma(k)} \mid (i, j, k) \in \mathfrak{W}_N \setminus \mathfrak{J}_r\}$. Normalizing the latter equation with

$$c := \sum_{(i,j,k) \in \mathfrak{W}_N \setminus \mathfrak{J}_r} \lambda_{i,j,k} = \sum_{i,j,k=1}^N \lambda_{i,j,k} - \left(\lambda_{1,1,1} + \sum_{(i,j,k) \in \mathfrak{J}_r} \lambda_{i,j,k} \right) \geq N - 1$$

shows $c^{-1}x \in \text{conv}(\Gamma_{n,6r-3})$. To bound the norm of $c^{-1}x$ we compute

$$\begin{aligned} \lambda_{1,1,1} + \sum_{(i,j,k) \in \mathfrak{J}_r} \lambda_{i,j,k} &= 2^{-N+1} + \sum_{s=2}^r (\lambda_{s,1,s} + \lambda_{s,s,1}) \\ &= 2^{-N+1} + \sum_{s=2}^r (2^{-N+s-1} + 2^{-N+s-1}) = \sum_{s=1}^r 2^{-N+s} < 2^{-N+r+1}. \end{aligned}$$

Finally, using $\|\varepsilon_{i_1, i_2, \dots, i_{6r-3}}\| \leq \sqrt{6r-3}$ for any $i_1, i_2, \dots, i_{6r-3} \in [n]$ together with the triangle inequality on Equation (26) implies

$$\|c^{-1}x\| \leq \frac{\sqrt{6r-3}}{N-1} 2^{-N+r+1} \leq \frac{\sqrt{6}}{(n-1)\sqrt{r}} 2^{-N+r+1} \leq 2^{-N+r+1} = 2^{-r(n-1)+1},$$

where we used $n \geq 3$ and $r \geq 2$ for $\sqrt{6} \leq (n-1)\sqrt{r}$. ◀

2.4 Polynomial scaling

A simple example of Equation (2) is the minimization of an n -variate homogeneous polynomial of degree d with nonnegative coefficients over the set $x_1, \dots, x_n > 0, \prod x_i = 1$, as studied in [30]. In this case the sets $\text{conv}(S)$ for $S \subseteq \Omega$ are Newton polytopes of homogeneous polynomials, and the minimum of a polynomial is bounded below if and only if the Newton polytope contains $\frac{d}{n} \mathbb{1}_n$. If the polynomials are hyperbolic of degree n , as in [30], their Newton polytope either contains $\mathbb{1}_n$ or is at least $1/\sqrt{n}$ away from it. However, we show that for general homogeneous polynomials the margin can get exponentially small in n even for $d = 3$.

Minimizing a degree d homogeneous polynomial $\sum_{\alpha \in \mathbb{Z}_{\geq 0}^n} p_\alpha x^\alpha$ with nonnegative coefficients over the set $x_1, \dots, x_n > 0, \prod x_i = 1$ is the same as computing Equation (2) for

$$\Omega' := \left\{ -\alpha + \frac{d}{n} \mathbb{1}_n \mid \alpha \in (\mathbb{Z}_{\geq 0})^n \text{ with } |\alpha| = d \right\}. \tag{27}$$

If $n = dm$ for some integer $m \geq 1$, then we have $-\Omega_{m,d} \subseteq \Omega'$. Therefore, Theorem 2.1(b) and (c) and the padding from Appendix C directly yield the following.

► **Corollary 2.13** (Margin for Polynomial scaling). *Fix some $d \geq 3$ and assume $n = dm$ for some $m \geq 3$. Let Ω' be as in Equation (27). Then*

$$\gamma(\Omega') \leq \gamma(\Omega_{m,d}) \leq 2^{-m+1} = 2^{-\frac{n}{d}+1}.$$

and for $d \geq 9$ we even have

$$\gamma(\Omega') \leq \gamma(\Omega_{m,d}) \leq 2^{-\lfloor \frac{(m-1)(d+3)}{6} \rfloor + 1} \approx 2^{-\frac{n}{6}}.$$

Thus, for fixed $d \geq 3$ and $n \rightarrow \infty$ the margin of Ω' can be exponentially small in n . In terms of polynomials, this states that the Newton polytope of a degree $d \geq 3$ homogeneous polynomial can be exponentially close to the origin without containing it.

3 Diameter bounds in the commutative case

In this section we describe an array such that all approximate scalings are very ill conditioned, proving Theorem 1.1. Let us define the diameter bound.

► **Definition 3.1.** *Let $\varepsilon \rightarrow 0$ and $f : \mathbb{R}^m \rightarrow \mathbb{R}$. The diameter bound $D_f(\varepsilon)$ is defined as the infimum over $R > 0$ such that*

$$\inf_{\|x\| \leq R} f(x) \leq \varepsilon + \inf_{x \in \mathbb{R}^m} f(x).$$

Thus, Theorem 1.1 is equivalent to the statement that $D_f(\varepsilon) = \Omega(2^{n/3} \log(1/\varepsilon))$ for $\varepsilon \leq e^{-Cn^2 \log n}$. We now give a proof outline for Theorem 1.1.

3.1 Proof outline

The high-level intuition applies not only to array scaling but to the capacity in general. Recall that the array scaling capacity is

$$\inf_{x \in \mathbb{R}^{3n}} \sum_{\omega \in \Omega} p_{\omega} e^{\omega \cdot x}$$

for $\Omega = \Omega_{n,3} = \{e_i - \frac{1}{n} \mathbb{1}_n : i \in [n]\} \subseteq \mathbb{R}^{3n}$. We build both the support $\Omega' \subseteq \Omega_{n,3}$ and the entries p in the following way. We construct a set $\Omega_0 \subseteq \Omega_{n,3}$, another element $\omega \in \Omega_{n,3}$, and an array q with the following properties.

1. The set $\Omega_0 \subseteq \Omega_{n,3}$ should be the support of a tristochastic array q .
2. The affine hull of Ω_0 , should have codimension one¹³ in \mathbb{R}^{3n} .
3. The origin is in the relative interior of $\text{conv}(\Omega_0)$. Note that the origin is already in $\text{conv}(\Omega_0)$ by the tristochasticity of q .
4. The vector $\omega \in \Omega_{n,3}$ should be at a very small, but positive, distance η from $\text{Aff}(\Omega_0)$. Note that this already implies that the facet gap of $\Omega_0 \cup \omega$ is small.

Finally, we define the entries of p by $p|_{\Omega_0} = \frac{1}{2}q$, $p_{\omega} = \frac{1}{2}$, and $p_{\omega} = 0$ elsewhere. Assuming we have found p according to this process, we now give intuition for the diameter bound.

Let v be the projection of ω to the orthogonal complement of $\text{Aff}(\Omega_0)$. Intuitively, the capacity is only approximately attained by vectors very far in the $-v$ direction. Indeed, first note that $\text{cap}(p) = 1/2$, because $\text{cap}(q) = 1$ by tristochasticity, $\text{cap}(p) \geq \frac{1}{2} \text{cap}(q) = \frac{1}{2}$, and $f_p(-tv/\|v\|) = \frac{1}{2} + e^{-\eta t}$ so $f_p(-tv/\|v\|)$ tends to $\frac{1}{2}$. However, $f_p(-tv/\|v\|)$ tends to $\frac{1}{2}$ slowly if η is small. Indeed, $f_p(-tv/\|v\|) \leq \frac{1}{2}(1 + \varepsilon)$ only if $t \geq \frac{1}{\eta} \log(1/\varepsilon)$.

To conclude rigorously that the capacity is only approached by vectors very far in the $-v$ direction, we must rule out directions with nonzero components in $\text{Aff}(\Omega_0)$. For this, we must use the assumption that 0 is rather deep in the relative interior of $\text{conv}(\Omega_0)$. If this is the case, then any ε -approximate minimizer must have a bounded component in $\text{Aff}(\Omega_0)$, for otherwise the contribution to f_p from the elements of Ω_0 alone will be larger than $\frac{1}{2} + \varepsilon$.

The remainder of the section will be concerned with the construction of a subset Ω_0 , an array q , and an element ω with these properties.


3.2 The construction

We construct the subset Ω_0 from a directed graph D on $[n]$, which we will determine later. If i, j is an edge in D , then Ω_0 includes the elements $(\varepsilon_i, \varepsilon_i, \varepsilon_j)$ as well as the three cyclic permutations of it. That is,

$$\Omega_0 = \{(\varepsilon_j, \varepsilon_i, \varepsilon_i), (\varepsilon_i, \varepsilon_j, \varepsilon_i), (\varepsilon_i, \varepsilon_i, \varepsilon_j) : ij \in E(D)\}.$$

We now describe the graph, as seen in Figure 2.

¹³This will not quite apply in our setting, because $\text{Aff}(\Omega_{n,3})$ is not full-dimensional. Instead, $\text{Aff}(\Omega_0)$ will be codimension one in $\text{Aff}(\Omega_{n,3})$.



■ **Figure 2** The graph D_l from Definition 3.2 with the edge labels proportional to the edge labeling q in Item 1 of Lemma 3.3 (the constant factor $1/6n$ is omitted for readability). We have also omitted the directions, which are all towards the root r .

► **Definition 3.2.** The graph $D_l = (W, E)$ is a directed tree with $l + 1$ levels, where the root is on the 0^{th} level and the leaves are on the l^{th} level. The tree is constructed as follows.

- All the edges are directed towards the root and are between adjacent levels.
- The root has three children, and on the $l - 1$ levels below the root every node has one child.
- Additionally, one of the vertices on level $l - 2$ has an additional child which has its own child.

Explicitly, the vertices W and edges E are given by


$$W = \{u_i, v_i, w_i : i \in [l]\} \cup \{w_0 := u_0 := v_0 := r, \bar{w}_{l-1}, \bar{w}_l\}.$$

$$E = \{u_i u_{i-1}, v_i v_{i-1}, w_i w_{i-1} : i \in [l]\} \cup \{\bar{w}_{l-1} w_{l-2}, \bar{w}_l \bar{w}_{l-1}\}.$$

Note that D_l has $3(l + 1)$ vertices so we set $n = 3(l + 1)$. Thus D_l has $3l + 2$ edges and so $|\Omega_0| = 3(3l + 2) = 3n - 3$. It is helpful to construct the matrix M whose set of rows is Ω_0 . To make the matrix sparser, first replace ε_i by e_i by restricting the minimization to the

$$\begin{bmatrix}
 0 & \begin{array}{cccc} A & I & \dots & 0 \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & I \\ 0 & \dots & 0 & A \end{array} & 0 & 0 & \begin{array}{c} 0 \\ \vdots \\ \vdots \\ I \end{array} \\
 0 & 0 & \begin{array}{cccc} A & I & \dots & 0 \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & I \\ 0 & \dots & 0 & A \end{array} & 0 & \begin{array}{c} 0 \\ \vdots \\ \vdots \\ I \end{array} \\
 0 & 0 & 0 & \begin{array}{cccc} A & I & \dots & 0 \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & I \\ 0 & \dots & 0 & A \end{array} & \begin{array}{c} 0 \\ \vdots \\ \vdots \\ I \end{array} \\
 \begin{array}{cc} A & I \\ 0 & A \end{array} & 0 & 0 & \begin{array}{cccc} 0 & 0 & 0 & \dots \\ 0 & 0 & I & \dots \end{array} & \begin{array}{c} 0 \\ \vdots \\ \vdots \\ 0 \end{array}
 \end{bmatrix}$$

■ **Figure 3** The matrix M written in the reordered basis described before Lemma 3.3. From the left, the five groups of columns correspond to the \bar{w} 's, the u 's, the v 's, the w 's, and r among the vertices of D_l . As such the dimensions of the five column groups, from left, are $3 \cdot 2$, $3(l-1)$, $3(l-1)$, $3(l-1)$, 3 , and the dimensions of the four groups of rows from top are $3(l-1)$, $3(l-1)$, $3(l-1)$, $3 \cdot 2$. A is as in Equation (28) and I is the 3×3 identity matrix.



■ **Figure 4** If v is a vertex of D_l with edges weighted q_1 and q_2 incident to it, then the column v , i of M for $i \in [3]$ sums to $q_1 + 2q_2$. That is, the incoming edge contributes its weight and the outgoing edge contributes twice its weight.

subspace $\sum x_i = \sum y_i = \sum z_i = 0$, which is without loss of generality. We define $\Omega'_0 \subseteq \mathbb{R}^{3n}$ to be Ω_0 but with each $(\varepsilon_i, \varepsilon_j, \varepsilon_k)$ replaced by (e_i, e_j, e_k) ; define $\Omega'_{n,3}$ similarly and define $p(e_i, e_j, e_k) := p(\varepsilon_i, \varepsilon_j, \varepsilon_k)$. Then

$$\inf_{x \in \mathbb{R}^{3n}} \sum_{\omega \in \Omega_{n,3}} p_\omega e^{(\varepsilon_i, \varepsilon_j, \varepsilon_k) \cdot x} = \inf_{\substack{x, y, z \in \mathbb{R}^n \\ \sum x_i = \sum y_i = \sum z_i = 0}} \sum_{\omega \in \Omega'_{n,3}} p_\omega e^{(e_i, e_j, e_k) \cdot (x, y, z)}.$$

Moreover, when we write the matrix M , it is easier to write the vector (x, y, z) in the order $(x_1, y_1, z_1, x_2, y_2, z_2, \dots)$ instead of the order $(x_1, \dots, x_n, y_1, \dots, y_n, z_1, \dots, z_n)$. With this ordering, the matrix M with rows in Ω'_0 is a block matrix M with blocks of size 3, with $n-1$ block rows, and with n block columns. Each block row corresponds to an edge in the directed graph $D_l = (W, E)$ on $n = 3(l+1)$ vertices. If $e \in E$ is an edge from $i \rightarrow j$, then the e^{th} row of M has the matrix

$$A = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix} \quad (28)$$

in the i^{th} block entry and

$$I = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

in the j^{th} block entry and zeroes elsewhere. See Figure 3 for a portrayal of the whole matrix M .

The first three properties for Ω_0 in the proof plan translate to the following three claims about M . The first relates to the tristochasticity of q , the second to the codimension of $\text{Aff}(\Omega'_0)$ in the subspace $(\mathbb{1}_n^\perp)^3$, and the third to the depth of the point $\frac{1}{n}\mathbb{1}_{3n}$ in $\text{conv}(\Omega'_0)$.

► **Lemma 3.3.** *Let $n = 3(l + 1)$.*

1. *The probability distribution q on $E \times [3]$ defined (for $i \in [3]$) by*

$$\begin{aligned} \text{for } j \in [l], \quad q_{u_j u_{j-1}, i} &= q_{v_j v_{j-1}, i} = \frac{1}{6n} \left(2 + (-2)^{-(l-j)} \right) \\ \text{for } j \in [l-2], \quad q_{w_j w_{j-1}, i} &= \frac{1}{6n} \left(2 + (-2)^{-(l-j-1)} \right) \\ q_{w_{l-1} w_{l-2}, i} &= q_{\bar{w}_{l-1} \bar{w}_{l-2}, i} = \frac{1}{2} q_{w_l w_{l-1}, i} = \frac{1}{2} q_{\bar{w}_l \bar{w}_{l-1}, i} = \frac{1}{6n} \left(\frac{3}{2} \right) \end{aligned}$$

on the rows of M has expectation $\frac{1}{n}\mathbb{1}_{3n}$. That is, if the rows of M are scaled by the values of q , each column sums to $1/n$. Note that the entries of q are $\Theta(\frac{1}{n})$. Ignoring the index i in $q_{uv,i}$ allows us to view q as a labeling of the edges of the graph D_l ; see Figures 2 and 4.

2. *$\ker M = \text{span}(\Omega'_0)^\perp$ is spanned by the 2 dimensional space $S \subseteq \mathbb{R}^{W \times [3]}$ given by*

$$S = \{s : s(v, 1) = \alpha, s(v, 2) = \beta, s(v, 3) = \gamma \text{ for all } v \in W, \alpha + \beta + \gamma = 0\}$$

and the function $f \in \mathbb{R}^{W \times [3]}$ which for all $i \in [3]$ assigns

$$\begin{aligned} f(u_j, i) &= f(v_j, i) = f(w_j, i) = (-2)^{-j} \text{ for } j \in [l] \cup \{0\} \\ \text{and } f(\bar{w}_{l-k}, i) &= f(w_{l-k}, i) \text{ for } k \in \{0, 1\}. \end{aligned} \quad (29)$$

Note that $f \in (\mathbb{1}_n^\perp)^3 \subseteq S^\perp$. Thus we have the orthogonal decomposition $\text{span}(\Omega'_0)^\perp = S \oplus \text{span } f$.

3. *Apart from the three zero singular values, all singular values of M are $\Omega(1/n)$.*

Given the lemma, let us prove that the diameter bound holds according to the proof outline at the beginning of the section.

Proof of Theorem 1.1. We first show the claim for n of the form $n = 3(l + 1)$; the bound follows for $3(l + 1) < n < 3(l + 2)$ by applying Proposition 3.5 with $t = 3(l + 1)$, using that the array we construct has capacity $1/2$ and $t/n \geq 2/3$.

We now show the diameter lower bound for $n = 3(l + 1)$. It is enough to exhibit a constant $C > 0$, and a probability distribution p on $\Omega'_{n,3} = \{e_i : i \in [n]\}^3$ such that for all $N \geq Cn^2 \log n$ and all $x, y, z \in \mathbb{1}_n^\perp$,

$$\sum_{\omega \in \Omega'_{n,d}} p_\omega e^{\omega \cdot (x,y,z)} \leq e^{-N} + \inf_{x', y', z' \in \mathbb{1}_n^\perp} \sum_{\omega \in \Omega'_{n,d}} p_\omega e^{\omega \cdot (x,y,z)}$$

only if $\|(x, y, z)\|_2 = \Omega(2^{n/3}N)$. Note that the space $(\mathbb{1}_n^\perp)^3$ over which we are infimizing is a subspace of S^\perp where S is as in Lemma 3.3, and that $\Omega'_{n,d} \subseteq S^\perp$. The proof will follow the outline in Section 3.1; namely, we will consider a subset $\Omega'_0 \subseteq \Omega'_{n,d}$ and an element $\omega' \in \Omega'_{n,d}$ very close to, but outside of, $\text{Aff}(\Omega'_0)$.

Consider the set $\Omega'_0 \subseteq \Omega'_{n,d}$ of rows of M in Lemma 3.3 and the probability distribution q on Ω'_0 from Lemma 3.3. Let $\omega' = (e_{u_l}, e_{v_l}, e_{w_l})$ for the vertices $u_l, v_l, w_l \in D_l$. Let $\Omega = \Omega'_0 \cup \{\omega'\}$, and define the probability distribution p on Ω by $p_{\omega'} = \frac{1}{2}$ and $p_\omega = \frac{1}{2}q_\omega$ for $\omega \in \Omega'_0$. Recall from Lemma 3.3 the orthogonal decomposition $\text{span}(\Omega'_0)^\perp = S \oplus \text{span } f$. As $\mathbb{R}^{3n} = \text{span}(\Omega'_0) \oplus \text{span}(\Omega'_0)^\perp$, we have the orthogonal decomposition $S^\perp = \text{span}(\Omega'_0) \oplus \text{span } f$. Observe that $\omega' \notin \text{span}(\Omega'_0)$, because by Lemma 3.3 we have $\text{span}(\Omega'_0)^\perp = \ker M = S + \text{span } f$ and clearly $f \cdot \omega' \neq 0$.

By Item 1 of Lemma 3.3 we have $\sum_{\omega \in \Omega'_0} q_\omega \omega = \frac{1}{n}(\mathbb{1}_n, \mathbb{1}_n, \mathbb{1}_n)$ and thus $\text{cap}(q) = 1$. Therefore, $\omega' \notin \text{span}(\Omega'_0)$ implies that the infimum is $1/2$ for this choice of Ω and p . We claim that the infimum can only be approximately attained by $h \in (\mathbb{1}_n^\perp)^3$ with a very large component in the one-dimensional space $\text{span } f = \text{span}(\Omega'_0)^\perp \cap (\mathbb{1}_n^\perp)^3$. As in the proof outline, we must bound the components in $\text{span}(\Omega'_0)$ of the approximate minimizer h . For $h \in (\mathbb{1}_n^\perp)^3$ write $h = h_0 + af$ and $\omega' = \omega_0 + bf$ where $h_0, \omega_0 \in \text{span } \Omega'_0$. Note that $|b| = \frac{|f \cdot \omega'|}{\|f\|^2} = O(2^{-l}) = O(2^{-n/3})$ and that $h_0 \in (\mathbb{1}_n^\perp)^3$, because h and f are. Suppose

$$\sum_{\omega \in \Omega} p_\omega e^{\omega \cdot h} \leq \frac{1}{2}e^{-N} + \frac{1}{2}.$$

Equivalently,

$$\sum_{\omega \in \Omega'_0} q_\omega e^{\omega \cdot h_0} + e^{h_0 \cdot \omega_0 + ab\|f\|^2} \leq e^{-N} + 1. \quad (30)$$

Suppose $\|h_0\|$ is bounded by L . If $e^{h_0 \cdot \omega_0 + ab\|f\|^2} \leq e^{-N}$, then $|ab| = \Omega(N - L)$. In particular, $\|h\| \geq \|af\| = |ab|\|f\|/|b| = \Omega((N - L)2^{n/3})$ because of the previous bounds on $|ab|, |b|$, and the fact that $\|f\| = \Theta(1)$. It remains to prove a bound L for $\|h_0\|$. We will do this by showing that if $\|h_0\|$ were too large, then the first term of the left-hand side of Equation (30) would be too large. This amounts to $\frac{1}{n}\mathbb{1}_{3n}$ being in the relative interior of $\text{conv}(\Omega'_0)$, but will be proved using lower bounds on the singular values of M .

Let α denote the least nonzero singular value of M ; by Item 3 of Lemma 3.3 $\alpha = \Omega(1/n)$. As $h_0 \in \text{span}(\Omega'_0) = \text{rowspan}(M)$, we have $\|Mh_0\| \geq \alpha\|h_0\|$ by the singular value bound. We claim that there is some $\omega \in \Omega'_0$ satisfying $\omega \cdot h_0 = \Omega(\alpha\|h_0\|/n)$. To prove this, first note that the $\sum_{\omega \in \Omega'_0} q_\omega \omega \cdot h_0 = \frac{1}{n}(\mathbb{1}_n, \mathbb{1}_n, \mathbb{1}_n) \cdot h_0 = 0$ because $h_0 \in (\mathbb{1}_n^\perp)^3$. Moreover, by Lemma 3.3 we have $q_\omega = \Theta(1/n)$. The claim follows from Lemma 3.4 below applied to the sequence $(\omega \cdot h_0 : \omega \in \Omega'_0)$.

Because $q_\omega = \Theta(\frac{1}{n})$, we must have that $\omega \cdot h_0 = O(\log n)$ for all $\omega \in \Omega'_0$. Else, the contribution from the term $q_\omega e^{\omega \cdot h_0}$ alone is larger than 1, in which case x cannot be an e^{-N} -approximate minimizer. Finally, $\|h_0\| = O(n(\log n)/\alpha) = O(n^2 \log n)$, and so we may take $L = O(n^2 \log n)$ and $N \geq 2L$. \blacktriangleleft

In the above proof, we used the following simple lemma.

► Lemma 3.4. *Let $0 < \beta < \gamma$. Suppose $z \in \mathbb{R}^m$ is such that $\sum_{i=1}^m q_i z_i = 0$ for $q_i \in (\beta/m, \gamma/m)$. Then there exists $i \in [m]$ such that $z_i \geq \frac{\beta}{2\gamma m} \|z\|_2$.*

Proof. Because $\sum q_i z_i = 0$,

$$\sum_{i: z_i < 0} q_i |z_i| = \sum_{i: z_i \geq 0} q_i z_i,$$

and

$$\sum_{i:z_i < 0} q_i |z_i| + \sum_{i:z_i \geq 0} q_i z_i \geq (\beta/m) \|z\|_1 \geq (\beta/m) \|z\|_2.$$

Thus $\sum_{i:z_i \geq 0} q_i |z_i| \geq \frac{\beta}{2m} \|z\|_2$, so there is some i such that $q_i z_i > \frac{1}{m} \frac{\beta}{2m} \|z\|_2$. Thus $z_i > \frac{\beta}{2\gamma m} \|z\|_2$. ◀

To show that our diameter lower bound holds for all values of n , we need the following proposition, which is proved in Appendix E. The idea is to prove diameter bounds for larger arrays from diameter bounds for smaller ones by embedding the smaller array in a “corner” of the larger array.

► **Proposition 3.5.** *Suppose $1 \leq t \leq n$. Let p be a d -dimensional array in $(\mathbb{R}_{\geq 0}^t)^{\otimes d}$ with unit sum; in particular $\text{cap}(p) \leq 1$. Let q be the d -dimensional array in $(\mathbb{R}_{\geq 0}^n)^{\otimes d}$ array such that $q_{i_1, \dots, i_d} = \frac{t}{n} p_{i_1, \dots, i_d}$ for $i_1, \dots, i_d \in [t]$, $q_{iii} = 1/n$ for $t + 1 \leq i \leq n$, and $q_{i_1, \dots, i_d} = 0$ otherwise. For $\varepsilon \leq 1 - \text{cap}(p)$,*

$$D_{f_q}(\varepsilon) \geq D_{f_p} \left(\frac{(1 - \text{cap}(p))\varepsilon}{1 - \text{cap}(p)^{t/n}} \right).$$

In particular, the norm of any ε -approximate minimizer of f_q is at least the norm of some $(\frac{1 - \text{cap}(p)}{1 - \text{cap}(p)^{t/n}})\varepsilon$ -approximate minimizer of f_p .

As a corollary of the proof of Theorem 1.1, we have a bound on the *facet gap* of [14]. The facet gap of a finite set Ω is defined to be the least distance of an element of Ω to the affine hull of a facet of $\text{conv}(\Omega)$. We have shown that the distance between $\text{Aff}(\Omega'_0)$ and ω' is $O(2^{-l})$, or $O(2^{-n/3})$.

► **Corollary 3.6** (Facet gap of array scaling). *There is a subset $\Omega_1 \subseteq \Omega_{n,3}$ with facet gap $O(2^{-n/3})$.*

Analogously to what is done for the margin in Proposition C.1, we may also embed this array inside a larger array to obtain a diameter bound for $d \geq 3$. For $d \geq 3$, take $q(i, j, k, l, \dots, l) = \frac{1}{n} p_{ijk}$ for all $i, j, k, l \in [n]$. Then for $(x_1, \dots, x_d) \in (\mathbb{1}_n^{\frac{1}{n}})^d$ we have

$$f_q(x_1, \dots, x_d) = \frac{1}{n} f_p(x_1, x_2, x_3) \sum_{l=1}^n e^{\sum_{j=4}^d (x_j)_l}.$$

For fixed x_1, x_2, x_3 , by Jensen’s inequality f_q is minimized when $x_j = 0_n$ for $j \geq 4$ and takes value $f_p(x_1, x_2, x_3)$, and thus f_q has the same diameter bound as f_p .

► **Corollary 3.7** (Diameter bound for $d \geq 3$). *There is an absolute constant $C > 0$ such that the following holds. For all $d \geq 3$, there is a family of arrays $q \in (\mathbb{R}_{\geq 0}^n)^{\otimes d}$ with $O(n^2)$ nonzero entries, each of bit-complexity $O(n)$, that satisfies the following property. For all $0 < \varepsilon \leq \exp(-Cn^2 \log n)$ and $x \in \mathbb{R}^{dn}$, if*

$$f_q(x) \leq \text{cap}(p) + \varepsilon$$

then $\|x\|_2 = \Omega(2^{n/3} \log(1/\varepsilon))$.

3.3 Proof of the properties of the construction

We now prove Lemma 3.3.

Proof of Lemma 3.3. It is first helpful to change basis on each copy of \mathbb{R}^3 so that the A blocks are diagonalized. Let $U \in \text{Mat}(3)$ be an orthogonal matrix such that

$$U^\dagger AU = \begin{bmatrix} 2 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix}.$$

This is possible because $2, -1, -1$ are the eigenvalues of the symmetric matrix A . In particular, the first column of U is $(1, 1, 1)/\sqrt{3}$, and the second two columns span the space of vectors with sum zero. Then $M' = (U^{\oplus n})^\dagger MU^{\oplus n}$ is of the form $P \oplus L \oplus L$ where $P_{e,v} = M'_{(e,1),(v,1)}$ for $e, v \in E \times V$ and $L_{e,v} = M'_{(e,2),(v,2)}$. Note that L is the edge-vertex incidence matrix of the directed graph D_l , the row corresponding to the edge (u, v) of D_l has a -1 in the column indexed by the vertex u and a $+1$ in the column indexed by v . Moreover, P is the matrix obtained from L by replacing every -1 entry by a 2 .

To prove Item 2, observe that $\ker M$ is $(U^{\oplus n}) \ker M' = (U^{\oplus n}) \ker P \oplus \ker L \oplus \ker L$. Because D_l is connected, $\ker L = \text{span } \mathbb{1}_n$. As the second two columns of U span the subspace of \mathbb{R}^3 of vectors with sum 0, the two-dimensional space S is given by $(U^{\oplus n})0 \oplus \text{span } \mathbb{1}_n \oplus \text{span } \mathbb{1}_n = (U^{\oplus n})0 \oplus \ker L \oplus \ker L$. We next reason for $\ker P$, the other summand of the orthogonal decomposition of $\ker M'$. The graph D_l is a connected tree, so $\ker P$ is one dimensional. This is because every choice of $g(w_0) \in \mathbb{R}$ determines a unique function $g : V \rightarrow \mathbb{R}$ in $\ker P$. We claim that the function $g(v) = f(v, 1)$ for f as in Equation (29) is in $\ker P$, and hence spans it. To check this, one must check that for every edge $(v, w) \in E$ we have $2g(v) + g(w) = 0$. It is instructive to look at Figure 2. Observe that this property holds for the edges $u_{k,k-1}$ if the sequence $g(u_k)$ obeys the recurrence relation $g(u_{k-1}) = -2g(u_k)$ for $k \in [l]$, which is indeed true by the definition of f . Checking the condition for v and w is similar. As the first column of U is proportional to $\mathbb{1}_3$, $(U^{\oplus n}) \ker P \oplus 0 \oplus 0$ is spanned by the function f . This proves Item 2.

To show Item 3, it is enough to argue that the singular values of P, L, L obey the desired bound. For L this follows straightforwardly from the fact that L is an incidence matrix of a connected, directed tree and so is totally unimodular with linearly independent rows. The singular value bound follows by Lemma 3.8. Rather than arguing spectrally for P , we make an ad-hoc argument using the structure of D_l . We first show that $\|x^t P\|_\infty = \Omega(\|x\|_\infty)$ for all $x \in \mathbb{R}^{n-1}$, which suffices because $\|x^t P\|_2 \geq \|x^t P\|_\infty$ and $\|x\|_\infty \geq \frac{1}{\sqrt{n}}\|x\|_2$.

Let $x \in \mathbb{R}^{n-1} \setminus \{0\}$ and e be an edge in D_l such that $|x(e)| = \|x\|_\infty$. If $e = u_i u_{i-1}$ for $i \in [l]$, then $|x^t P(u_i)| \geq \|x\|_\infty$ because either $i < l$ in which case

$$|x^t P(u_i)| = |2x(u_i u_{i-1}) + x(u_{i+1} u_i)| \geq 2\|x\|_\infty - |x(u_{i+1} u_i)| \geq \|x\|_\infty$$

or $i = l$ and so $|x^t P(u_i)| = |2x(e)| = 2\|x\|_\infty$. The same argument applies to all other edges except $e = w_{l-2} w_{l-3}$. In the latter case we are done if $x^t P(w_{l-2}) \geq 1/3\|x\|_\infty$. Otherwise we necessarily have $|x(w_{l-1} w_{l-2})| + |x(\bar{w}_{l-1} w_{l-2})| \geq 5/3\|x\|_\infty$, since $x^t P(w_{l-2}) = 2x(e) + x(w_{l-1} w_{l-2}) + x(\bar{w}_{l-1} w_{l-2})$. It follows that $|x(w_{l-1} w_{l-2})| \geq 5/3\|x\|_\infty - |x(\bar{w}_{l-1} w_{l-2})| \geq 5/3\|x\|_\infty - \|x\|_\infty \geq 2/3\|x\|_\infty$. As $|x^t P(w_{l-1})| = 2x(w_{l-1} w_{l-2}) + x(w_l w_{l-1})$, we have

$$|x^t P(w_{l-1})| \geq 2|x(w_{l-1} w_{l-2})| - |x(w_l w_{l-1})| \geq \frac{4}{3}\|x\|_\infty - \|x\|_\infty \geq \frac{1}{3}\|x\|_\infty.$$

In any case, there is some value of $x^t P$ with absolute value greater or equal $1/3\|x\|_\infty$.

Finally, for Item 1 we note that the probability distribution q on the rows of M has expectation equal to the all $1/n$ function if and only if the probability distribution q' defined by $q'_e = 3q_{e,1}$ on the rows of P has expectation equal to the all $3/n$ function on the vertices of D_l . Recall that P is obtained from the edge-vertex incidence matrix of D_l by replacing every -1 with a 2 . Thus the expectation of the rows under q' at a vertex v is $\sum_{w:(w,v) \in D_l} q'_{(w,v)} + \sum_{w:(v,w) \in D_l} 2q'_{(v,w)}$; see Figure 4. We now check that this is equal to $3/n$ for each vertex of D_l ; it is helpful to look at Figure 2. The leaves u_l, v_l, w_l , and \bar{w}_l all have outdegree one and indegree zero, and q' takes the value $3 \cdot 3/6n = 3/2n$ on the outgoing edges. The expectation under q' thus takes value $3/n$ on these vertices. On vertices of indegree one and outdegree one, q' takes the value $\frac{1}{2n} (2 + (-2)^{-k})$ on the incoming edge and $\frac{1}{2n} (2 + (-2)^{-(k+1)})$ on the outgoing edge. Thus the expectation takes the value $\frac{1}{2n} (2 + (-2)^{-k}) + \frac{1}{2n} (4 - (-2)^{-k}) = 3/n$. The remaining vertices to check, those of total degree three, are r and w_{l-2} . For r , which has only incoming edges, the expectation under q' is $2 \cdot \frac{1}{2n} (2 + (-2)^{-(l-1)}) + \frac{1}{2n} (2 + (-2)^{-(l-2)})$, which is again $3/n$. For w the expectation is $2 \cdot \frac{1}{2n} (2 - \frac{1}{2}) + 2 \cdot \frac{1}{2n} (2 - \frac{1}{2}) = 3/n$. This completes the proof. ◀

► **Lemma 3.8.** *If A is an $n \times k$ totally unimodular matrix with linearly independent columns, then the eigenvalues of $A^T A$ are all at least $1/n^2$.*

Proof. First note that $k \leq n$ by the linear independence of the columns of A . The least eigenvalue of $A^T A$ is $\min_{x \in \mathbb{R}^k \setminus \{0\}} (x^T A^T A x) / \|x\|^2 = \min_{x \in \mathbb{R}^k \setminus \{0\}} \|Ax\|^2 / \|x\|^2$, so it suffices to show that for all $x \in \mathbb{R}^k$, Ax has norm at least $\|x\|/n$. Indeed, if $Ax = y$, then there is some invertible $k \times k$ submatrix A' of A and $k \times 1$ submatrix y' of y such that $A'x = y'$. By Cramer's rule and unimodularity of A' we have that, for $i \in [k]$,

$$x_i = \frac{\det(B_i)}{\det(A')} = \pm \det(B_i)$$

where B_i is simply the matrix that one obtains by replacing the i^{th} column of A' with the vector y' . By performing the Laplace expansion with respect to the i^{th} column, and by unimodularity of the minors, we have that $x_i \leq \|y\|_1$, and so $\|x\|_2 \leq \sqrt{k} \|y\|_1 \leq n \|Ax\|_2$ (using $k \leq n$). ◀

4 The noncommutative case

In this section we extend the results from the commutative to the noncommutative case. For this, we recall in the first subsection necessary concepts such as moment maps and moment polytopes, and we define the weight margin and the gap of a representation. The second subsection introduces the key concept of a *free* subset of weights, see [24]. This concept dates at least back to [18, Proposition 1.2], where it is called *strong orthogonality*. Freeness will be used to transfer results from the commutative to the noncommutative case.¹⁴ The latter is done in the following three subsections, where we prove bounds on the tensor gap, on the gap for homogeneous polynomials and on the diameter for the natural $\text{SL}(n)^3$ action on 3-tensors. Finally, we show a bound for the weight margin of certain quiver representations. This provides an example, where the constructed set of weights is *not* free, compare Remark 4.28. Still, after adding enough arrows to the considered quiver, we are able to ensure the same bound for the gap.

¹⁴ Actually all presented concepts in the first two subsections work in the very general setting of reductive groups and their rational representations. For the sake of clarity and concreteness we stick to the special case needed in this paper, i.e. the reductive group $\text{SL}(n)^d := \text{SL}(n) \times \dots \times \text{SL}(n)$ with $d \geq 1$ many copies of $\text{SL}(n)$.

4.1 Moment maps and moment polytopes

In the following we introduce the null-cone problem and its dual characterization via moment maps and moment polytopes. This allows us to rigorously introduce the weight margin and the gap of a rational representation. Thereby we establish precise meaning and interpretation of our results regarding these two notions (in view of the null-cone problem). We stick to the notation of [12], where the gap (implicitly) and the weight margin have been introduced. A reader unfamiliar with representation theory is referred to Appendix B.

Let $G = \text{SL}(n)^d$, $K = \text{SU}(n)^d$, $T = \text{ST}(n)^d$ and $T_K = K \cap T$ be matrix Lie subgroups of $\text{GL}(dn)$ via block-diagonal embedding. Then we can think of their Lie algebras $\text{Lie}(G)$ etc. as being block diagonally embedded into $\mathbb{C}^{dn \times dn}$. For a rational representation $\pi: G \rightarrow \text{GL}(V)$ we write $g \cdot v := \pi(g)v$ for the induced action, where $g \in G$ and $v \in V$. Moreover, we denote the set of weights of π by $\Omega(\pi) \subseteq i \text{Lie}(T_K)$ and the induced representation on Lie algebras by $\Pi: \text{Lie}(G) \rightarrow \text{End}(V)$. We remark that we usually identify $i \text{Lie}(T_K) \cong (\mathbb{1}_n^\perp)^d \subseteq (\mathbb{R}^n)^d$, where $\mathbb{1}_n^\perp$ denotes the orthogonal complement of the all-ones vector $\mathbb{1}_n$ in \mathbb{R}^n .

The orbit of $v \in V$ is $G \cdot v := \{g \cdot v \mid g \in G\}$ and we denote its closure¹⁵ by $\overline{G \cdot v}$. A vector v is called G -unstable, if $0 \in \overline{G \cdot v}$, and otherwise v is G -semistable. Equivalently, a vector $v \in V$ is G -unstable if and only if its capacity

$$\text{cap}_G(v) := \inf_{g \in G} \|g \cdot v\|^2$$

equals zero. The G -unstable vectors form an affine subvariety of V - the null-cone (with respect to G). Orbit, stability, and capacity can also be defined for T by replacing G by T in the definitions.

As discussed in Section 1.2, the null-cone problem has many applications in different fields of computer science, mathematics and physics.

Next, we introduce the moment map. Given a rational representation $\pi: G \rightarrow \text{GL}(V)$ there exists an Hermitian inner product $\langle \cdot, \cdot \rangle$ on V , by convention linear in the second argument, such that $\langle k \cdot v, k \cdot w \rangle = \langle v, w \rangle$ holds for all $k \in K$ and all $v, w \in V$.¹⁶

► **Definition 4.1.** For $v \in V \setminus \{0\}$ we define $\mu_G(v) \in i \text{Lie}(K)$ as the unique element of the real vector space $i \text{Lie}(K)$, which satisfies for all $A \in i \text{Lie}(K)$

$$\text{tr}(\mu_G(v)A) = \frac{\langle v, \Pi(A)v \rangle}{\langle v, v \rangle}.$$

This defines the moment map $\mu_G: V \setminus \{0\} \rightarrow i \text{Lie}(T)$ of G . Replacing G by T and K by T_K we derive the moment map $\mu_T: V \setminus \{0\} \rightarrow i \text{Lie}(T_K)$ of T .

The maps μ_G and μ_T are indeed moment maps in the sense of symplectic geometry; namely for the induced action of K and, respectively, T_K on the projective space $\mathbb{P}(V)$. Recall $i \text{Lie}(K) \subseteq \mathbb{C}^{dn \times dn}$ so we can consider $\|\mu_G(v)\|_F$ and $\|\mu_T(v)\|_F$.

An important application of these moment maps is due to the Kempf-Ness theorem [36], which provides a duality for the null-cone membership problem:

$$\text{cap}_G(v) = 0 \quad \Leftrightarrow \quad 0 < \inf_{g \in G} \|\mu_G(g \cdot v)\|_F = \min_{0 \neq w \in G \cdot v} \|\mu_G(w)\|_F \tag{31}$$

and similarly for T , replacing G by T in the above equation. The two moment maps are related as follows.

¹⁵The Euclidean- and the Zariski-closure of $G \cdot v$ coincide.

¹⁶In our concrete representations later on this will be the standard inner product.

► **Proposition 4.2.** *Let $p: i\text{Lie}(K) \rightarrow i\text{Lie}(\mathbb{T}_K)$ be the orthogonal projection. Then $\mu_{\mathbb{T}} = p \circ \mu_G$ and $\|\mu_{\mathbb{T}}(v)\|_F \leq \|\mu_G(v)\|_F$ for all $v \in V \setminus \{0\}$.*

Proof. Since $i\text{Lie}(\mathbb{T}_K) \subseteq i\text{Lie}(K)$ the definition of the moment maps gives $\text{tr}[\mu_{\mathbb{T}}(v)H] = \text{tr}[\mu_G(v)H]$ for all $H \in i\text{Lie}(\mathbb{T}_K)$. But $\mu_{\mathbb{T}}(v) \in i\text{Lie}(\mathbb{T}_K)$ is the unique element with this property, hence $p(\mu_G(v)) = \mu_{\mathbb{T}}(v)$. The inequality $\|\mu_{\mathbb{T}}(v)\|_F \leq \|\mu_G(v)\|_F$ follows directly from the first part. ◀

Now, we explain how the moment maps induce certain polytopes, which can also be used to express the duality in (31). Moreover, the combinatorics of these polytopes captures the important complexity measures (*weight margin* and *gap*). Indeed, one of our main contributions is to analyze parts of this combinatorics, thereby deducing complexity barriers for certain computational problems.

Since the action of \mathbb{T} via π is completely determined by the weight space decomposition $V = \bigoplus_{\omega \in \Omega(\pi)} V_{\omega}$ of V , one can compute $\mu_{\mathbb{T}}(v)$ in terms of this decomposition. For this, write $v = \sum_{\omega} v_{\omega}$ with $v_{\omega} \in V_{\omega}$ and define the support of v with respect to π as

$$\text{supp}(v) := \{\omega \in \Omega(\pi) \mid v_{\omega} \neq 0\}.$$

Using that distinct weight spaces are orthogonal, one computes

$$\mu_{\mathbb{T}}(v) = \sum_{\omega} \frac{\langle v_{\omega}, v_{\omega} \rangle}{\langle v, v \rangle} \omega,$$

which is a convex combination of the weights in $\text{supp}(v)$. Noting that $\text{supp}(v) = \text{supp}(t \cdot v)$ for $t \in \mathbb{T}$ also $\mu_{\mathbb{T}}(t \cdot v) \in \Delta_{\mathbb{T}}(v) := \text{conv}\{\omega \mid \omega \in \text{supp}(v)\}$. In fact,

$$\Delta_{\mathbb{T}}(v) = \overline{\{\mu_{\mathbb{T}}(t \cdot v) \mid t \in \mathbb{T}\}} = \{\mu_{\mathbb{T}}(w) \mid w \in \overline{\mathbb{T} \cdot v}, w \neq 0\} \subseteq i\text{Lie}(\mathbb{T}_K)$$

and $\Delta_{\mathbb{T}}(v)$ is called the *weight polytope* of v .

It is an astonishing result that for fixed $v \in V \setminus \{0\}$, the set $\{\mu_G(g \cdot v) : g \in G\}$ gives rise to a polytope as follows. Let $\text{spec}: \text{Herm}(n) \rightarrow \mathbb{R}^n$ be the function sending a Hermitian matrix to its eigenvalues in decreasing order. Recalling that $i\text{Lie}(K) \subseteq \text{Herm}(n)^d$ is block-diagonally embedded in $\mathbb{C}^{dn \times dn}$, we set

$$s: i\text{Lie}(K) \rightarrow (\mathbb{R}^n)^d, \quad \text{diag}(A_1, \dots, A_d) \mapsto (\text{spec}(A_1), \dots, \text{spec}(A_d)).$$

Then for $v \in V \setminus \{0\}$ the set

$$\Delta_G(v) := \{s(\mu_G(w)) \mid w \in \overline{G \cdot v}, w \neq 0\}$$

is a rational convex polytope, see e.g. [28] or [45, Appendix] by Mumford. We call $\Delta_G(v)$ the *moment polytope* of v . Noting that $\|A\|_F = \|\text{spec}(A)\|_2$ for any $A \in \text{Herm}(n)$ we have $\|\mu_G(v)\|_F = \|s(\mu_G(v))\|_2$ for all $v \in V \setminus \{0\}$. Thus, we can formulate the duality from (31) also as follows:

$$\text{cap}_G(v) = 0 \quad \Leftrightarrow \quad \text{dist}(0, \Delta_G(v)) > 0 \quad \Leftrightarrow \quad 0 \notin \Delta_G(v),$$

and similarly for \mathbb{T} . This motivates the following two definitions.

► **Definition 4.3.** Let $\pi: G \rightarrow \mathrm{GL}(V)$ be a rational representation. We define the gap of π as¹⁷

$$\gamma_G(\pi) := \min \{ \|\mu_G(v)\|_F \mid v \neq 0 \text{ is } G\text{-unstable} \} = \min \{ \mathrm{dist}(0, \Delta_G(v)) \mid v \neq 0 \text{ is } G\text{-unstable} \},$$

and the weight margin of π as

$$\gamma_T(\pi) := \min \{ \|\mu_T(v)\|_F \mid v \neq 0 \text{ is } T\text{-unstable} \} = \min \{ \mathrm{dist}(0, \Delta_T(v)) \mid v \neq 0 \text{ is } T\text{-unstable} \}.$$

Equivalently, $\gamma_T(\pi)$ is the margin of the set of weights $\Omega(\pi)$, i.e. $\gamma_T(\pi) = \gamma(\Omega(\pi))$.

Thus, the gap $\gamma_G(\pi)$ is the largest constant $C > 0$ with the following property: If $\|\mu_G(v)\|_F < C$ for some vector $v \in V$, then v is G -semistable. The same statement holds for the weight margin $\gamma_T(\pi)$ replacing G by T . Therefore, these notions capture how small $\mu_G(g \cdot v)$ (respectively $\mu_T(t \cdot v)$) must be to certify null-cone non-membership. The next remark connects the gap to the classical notion of *instability* due to Mumford [44].

► **Remark 4.4.** The gap is twice the minimum value of all positive instabilities. Indeed, let $M(v)$ denote the instability of a non-zero vector v , see e.g. [45, eq. (9)]. Then $\mathrm{dist}(0, \Delta_G(v)) \geq 2M(v)$ and [45, Theorem 6.1] implies

$$\gamma_G(\pi) = \inf \{ 2M(v) : v \neq 0, v \text{ is } G\text{-unstable} \}.$$

► **Example 4.5.** Recall the tensor scaling action, in which the group $G = \mathrm{SL}(n)^d$ acts on $(\mathbb{C}^n)^{\otimes d}$ via the representation

$$\pi_{n,d}: \mathrm{SL}(n)^d \rightarrow \mathrm{GL}((\mathbb{C}^n)^{\otimes d}), (g_1, \dots, g_d) \mapsto g_1 \otimes \dots \otimes g_d.$$

Similar computations to those in Example B.2 show that the set of weights of $\pi_{n,d}$ is

$$\Omega(\pi_{n,d}) = \Omega_{n,d} = \{ \varepsilon_i \mid i \in [n] \}^d \subseteq (\mathbb{R}^n)^d.$$

Therefore, the weight margin $\gamma_T(\pi_{n,d})$ is the margin $\gamma(\Omega_{n,d})$ for the array scaling problem from Theorem 1.3 and Theorem 2.1. Moreover, the moment map μ_G for $\pi_{n,d}$ can be computed in terms of the quantum marginals as described in the introduction, i.e. $\gamma_G(\pi_{n,d})$ is indeed the tensor gap.

The weight margin and the gap satisfy the following inequality.

► **Proposition 4.6.** It holds that $\gamma_T(\pi) \leq \gamma_G(\pi)$.

Proof. Let $v \neq 0$ be G -unstable. Then there exists $k \in K$ such that $k \cdot v$ is T -unstable; see [51, Theorem 3.25]. By Proposition 4.2 we obtain

$$\|\mu_G(v)\|_F = \|\mu_G(k \cdot v)\|_F \geq \|\mu_T(k \cdot v)\|_F \geq \gamma_T(\pi)$$

where we used in the first equality that $\mu_G(k \cdot v) = k\mu_G(v)k^\dagger$. Therefore $\gamma_G(\pi) \geq \gamma_T(\pi)$. ◀

This inequality motivates the next subsection.

¹⁷ Gap and weight margin are well-defined, i.e. the minimum is attained. Indeed, the moment maps give rise to continuous maps on $\mathbb{P}(V)$ and the non-zero G -unstable (respectively non-zero T -unstable) vectors form a projective subvariety of $\mathbb{P}(V)$; in particular they form a compact set.

4.2 Free sets of weights

Proposition 4.6 from the preceding subsection shows us that an upper bound for the weight margin $\gamma_{\mathbb{T}}(\pi)$ need not necessarily apply to the gap $\gamma_G(\pi)$. Still, many of our bounds in the commutative case (weight margin and diameter) transfer to the noncommutative case (gap and diameter). We use crucially the notion of a *free* subset of weights (or [24]). Freeness is also known as *strong orthogonality* [18].

► **Definition 4.7.** *Let $\pi: G \rightarrow \mathrm{GL}(V)$ be a rational representation with set of weights $\Omega(\pi)$.*

A subset $\Gamma \subseteq \Omega(\pi)$ is called free if no two distinct elements of Γ differ by a root of G . In other words, $\Gamma \cap (\Gamma + \alpha) = \emptyset$ holds for all roots α of G .

Furthermore, a vector $v \in V \setminus \{0\}$ is called free if its support $\mathrm{supp}(v) \subseteq \Omega(\pi)$ is free.

We transfer the results from the commutative to the noncommutative case with the upcoming Proposition 4.8. It is known that for vectors v with free support one has $\mu_G(v) = \mu_{\mathbb{T}}(v)$. This appears implicitly in [49, Lemma 7.1] and [24, Proposition 2.2], but we prove it below for completeness. We thank Visu Makam for pointing out to us that this equality still holds under a weaker condition on v , when the representation decomposes into orthogonal subrepresentations. This can be used to turn our weight margin upper bound for quivers into a gap upper bound (Theorem 4.25). This weaker condition also appears in [21, Theorem 6.5].

► **Proposition 4.8.** *Let $\pi: G \rightarrow \mathrm{GL}(V)$ be a rational representation and suppose $V = \bigoplus_{i=1}^k V_i$ is an orthogonal decomposition into G -subrepresentations with respect to the K -invariant inner product, that is used to define $\mu_{\mathbb{T}}$ and μ_G . Let $v = (v_1, \dots, v_k) \in V \setminus \{0\}$, $v_i \in V_i$ be such that all supports $\Gamma_i := \mathrm{supp}(v_i) \subseteq \Omega(\pi)$ are free. Then for all $t \in \mathbb{T}$ it holds that $\mu_G(t \cdot v) \in i \mathrm{Lie}(\mathbb{T}_K)$ and $\mu_G(t \cdot v) = \mu_{\mathbb{T}}(t \cdot v)$.*

If additionally $0 \notin \Delta_{\mathbb{T}}(v) = \mathrm{conv}(\Gamma)$, where $\Gamma = \bigcup_i \Gamma_i$, then the upper bound $\mathrm{dist}(0, \mathrm{conv}(\Gamma))$ for the weight margin $\gamma_{\mathbb{T}}(\pi)$ also applies to the gap, i.e. $\gamma_G(\pi) \leq \mathrm{dist}(0, \mathrm{conv}(\Gamma))$.

Proof. The action of \mathbb{T} preserves the supports Γ_i , and in particular preserves their freeness. Hence, it suffices to show $\mu_G(v) \in i \mathrm{Lie}(\mathbb{T}_K)$, which immediately yields $\mu_G(v) = \mu_{\mathbb{T}}(v)$ by Proposition 4.2. Moreover, the orthogonality with respect to the K -invariant inner product shows $\mu_G(v) = H_1 \oplus \dots \oplus H_k$, where $H_i = \mu_G^{(i)}(v_i)$ is given by the moment map $\mu_G^{(i)}$ of the G -module V_i if $v_i \neq 0$ and otherwise $H_i = 0$. The latter holds similarly for $\mu_{\mathbb{T}}$.

Therefore, we may assume $k = 1$, i.e. $v \neq 0$ has free support Γ . We write $v = \sum_{\omega \in \Gamma} v_{\omega}$ for $v_{\omega} \in V_{\omega}$. Then, for any root α of G and all $A \in i \mathrm{Lie}(K) \cap \mathrm{Lie}(G)_{\alpha}$ we have $\Pi(A)v_{\omega} = 0$ by $\Gamma \cap (\Gamma + \alpha) = \emptyset$ (i.e., freeness) and Proposition B.4. Thus, $\Pi(A)v = 0$ and $\mathrm{tr}(\mu_G(v)A) = 0$ for all roots α and all $A \in i \mathrm{Lie}(K) \cap \mathrm{Lie}(G)_{\alpha}$. With the root space decomposition $\mathrm{Lie}(G) = \mathrm{Lie}(\mathbb{T}) \oplus \bigoplus_{\alpha} \mathrm{Lie}(G)_{\alpha}$ (see also Example B.3) we conclude $\mu_G(v) \in i \mathrm{Lie}(\mathbb{T}_K)$. The first statement is proven.

For the second claim we note that indeed $\bigcup_i \Gamma_i = \mathrm{supp}(v)$. If additionally $0 \notin \mathrm{conv}(\Gamma) = \Delta_{\mathbb{T}}(v)$, then v is \mathbb{T} -unstable. In particular, v is G -unstable and thus

$$\gamma_G(\pi) \leq \mathrm{dist}(0, \Delta_G(v)).$$

On the other hand, we have

$$\mathrm{dist}(0, \Delta_G(v)) = \inf_{g \in G} \|\mu_G(g \cdot v)\|_F \leq \inf_{t \in \mathbb{T}} \|\mu_G(t \cdot v)\|_F \stackrel{(*)}{=} \mathrm{dist}(0, \mathrm{conv}(\Gamma)),$$

where we used $\mu_G(t \cdot v) = \mu_{\mathbb{T}}(t \cdot v)$ in (*). We conclude by combining the two inequalities. ◀

► **Remark 4.9.** *It is well-known that any rational representation $\pi: G \rightarrow \mathrm{GL}(V)$ can be decomposed into G -irreducible subrepresentations that are pairwise orthogonal with respect to the fixed K -invariant inner product. Proposition 4.8 shows that ensuring freeness on the irreducible subrepresentations suffices.*

We end the section with an interesting connection between the weight margin and the gap.

► **Proposition 4.10.** *Let $\pi: G \rightarrow \mathrm{GL}(V)$ be a rational representation and denote its m -fold direct sum by π^m .*

1. *The weight margin satisfies $\gamma_{\mathrm{T}}(\pi) = \gamma_{\mathrm{T}}(\pi^m)$ for all $m \geq 1$.*
2. *The gap satisfies $\gamma_G(\pi^m) \geq \gamma_G(\pi^{m+1})$ for all $m \geq 1$.*
3. *There exists some $m \leq \dim(V)$ such that $\gamma_G(\pi^m) = \gamma_{\mathrm{T}}(\pi^m) = \gamma_{\mathrm{T}}(\pi)$.*

Proof. We note that π^m is given by the action $g \cdot (v_1, \dots, v_m) = (g \cdot v_1, \dots, g \cdot v_m)$ on V^m . Furthermore, the K -invariant inner product $\langle \cdot, \cdot \rangle$ of V induces naturally a K -invariant product on V^m by

$$\langle (v_1, \dots, v_m), (w_1, \dots, w_m) \rangle_{V^m} := \sum_{i=1}^m \langle v_i, w_i \rangle.$$

For the first claim just note that the weight space decomposition for π^m is $V^m = \bigoplus_{\omega \in \Omega(\pi)} V_{\omega}^m$ and hence $\Omega(\pi^m) = \Omega(\pi)$.

For the second claim, let $(v_1, \dots, v_m) \in V^m \setminus \{0\}$ be G -unstable such that $\|\mu_G(v_1, \dots, v_m)\|_F = \gamma_G(\pi^m)$. Then $(v_1, \dots, v_m, 0) \in V^{m+1} \setminus \{0\}$ is G -unstable as well, so $\|\mu_G(v_1, \dots, v_m, 0)\|_F \geq \gamma_G(\pi^{m+1})$. Moreover, under the inner product $\langle \cdot, \cdot \rangle_{V^{m+1}}$ the first m copies of V are orthogonal to the last copy. Thus, $\mu_G(v_1, \dots, v_m, 0)$ is the 2×2 block-diagonal matrix $\mathrm{diag}(\mu_G(v_1, \dots, v_m), 0)$ and hence $\|\mu_G(v_1, \dots, v_m, 0)\|_F = \|\mu_G(v_1, \dots, v_m)\|_F = \gamma_G(\pi^m)$.

Finally, let $\Gamma = \{\omega_1, \dots, \omega_m\} \subseteq \Omega(\pi)$ be such that $0 \notin \mathrm{conv}(\Gamma)$ and $\mathrm{dist}(0, \mathrm{conv}(\Gamma)) = \gamma_{\mathrm{T}}(\pi)$. We have $m \leq |\Omega(\pi)| \leq \dim(V)$ by the weight space decomposition $V = \bigoplus_{\omega \in \Omega(\pi)} V_{\omega}$. Now, for each $\omega_i \in \Gamma$ fix some weight vector $v_i \in V_{\omega_i} \setminus \{0\}$. Then $v := (v_1, \dots, v_m) \in V^m$ satisfies the assumptions of Proposition 4.8, because $\Gamma_i = \{\omega_i\}$ is free and the distinct copies of V are orthogonal under $\langle \cdot, \cdot \rangle_{V^m}$. Thus, we obtain

$$\gamma_G(\pi^m) \leq \mathrm{dist}(0, \mathrm{conv}(\Gamma)) = \gamma_{\mathrm{T}}(\pi) = \gamma_{\mathrm{T}}(\pi^m),$$

but on the other hand $\gamma_G(\pi^m) \geq \gamma_{\mathrm{T}}(\pi^m)$ by Proposition 4.6. ◀

4.3 Freeness for tensors

We recall from Example 4.5 that $\pi_{n,d}$ denotes the natural representation of $G = \mathrm{SL}(n)^d$ on $(\mathbb{C}^n)^{\otimes d}$ and that the weight margin $\gamma_{\mathrm{T}}(\pi_{n,d})$ is the margin $\gamma(\Omega_{n,d})$ for the array scaling problem from Theorem 1.3 and Theorem 2.1. The purpose of this subsection is to prove the bounds for $\gamma_{\mathrm{T}}(\pi_{n,d})$ from Theorem 2.1 also for the gap $\gamma_G(\pi_{n,d})$.

► **Theorem 4.11.** *Let $\pi_{n,d}$ be the representation induced by the natural action of $G := \mathrm{SL}(n)^d$ on $(\mathbb{C}^n)^{\otimes d}$. Then the weight margin $\gamma_{\mathrm{T}}(\pi_{n,d})$ and the gap $\gamma_G(\pi_{n,d})$ can be bounded as follows:*

- (a) *If $n = 2$ and $d \geq 3$, then $\gamma_{\mathrm{T}}(\pi_{2,d}) \leq \gamma_G(\pi_{2,d}) \leq 2^{-\frac{d}{2}+1}$.*
- (b) *If $n \geq 3$ and $d = 3$, then $\gamma_{\mathrm{T}}(\pi_{n,3}) \leq \gamma_G(\pi_{n,3}) \leq 2^{-n+1}$.*
- (c) *If $n \geq 3$ and $d = 6r - 3$ for some integer $r \geq 2$, then*

$$\gamma_{\mathrm{T}}(\pi_{n,d}) \leq \gamma_G(\pi_{n,d}) \leq \frac{\sqrt{6}}{(n-1)\sqrt{r}} 2^{-r(n-1)+1} \leq 2^{-r(n-1)+1} = 2^{-\frac{(d+3)(n-1)}{6}+1}.$$

Though the above theorem only applies to certain d , we can “pad” the tensors to obtain similar results for all $d \geq 3$. This is because bounds for $\gamma_G(\pi_{n,d})$ via free subsets of weights also hold for $\gamma_G(\pi_{n,d+2})$ and $\gamma_G(\pi_{n,d+3})$, see Proposition C.1. The missing case $n \geq 3$ and $d = 4$ is treated in Proposition C.2. Therefore, we can conclude Theorem 1.6 from the above Theorem 4.11.

Our main method for transferring the bounds from the commutative case (Theorem 2.1) to the noncommutative case is to use the concept of freeness in conjunction with Proposition 4.8. The following definition will be convenient for proving freeness of tensors.

► **Definition 4.12 (Free sets).** A set $M \subseteq [n]^d$ is called free, if $i = (i_1, \dots, i_d), j = (j_1, \dots, j_d) \in M$ with $i \neq j$ always implies $|\{i_l \neq j_l \mid l = 1, \dots, d\}| \geq 2$.

► **Proposition 4.13.** Let $M \subseteq [n]^d$ and denote the induced subset of weights by

$$\Gamma_M := \{(\varepsilon_{i_1}, \dots, \varepsilon_{i_d}) \mid (i_1, \dots, i_d) \in M\} \subseteq (\mathbb{R}^n)^d.$$

Then M is a free set if and only if the set of weights $\Gamma_M \subseteq \Omega(\pi_{n,d})$ is free as in Definition 4.7.

Proof. We recall that Γ_M is free if and only if no two distinct elements of Γ_M differ by a root of $G = \text{SL}(n)^d$, see Definition 4.7. Furthermore, remember that the roots of G are

$$(e_i - e_j, 0_n, \dots, 0_n), (0_n, e_i - e_j, 0_n, \dots, 0_n), \dots, (0_n, \dots, 0_n, e_i - e_j) \in (\mathbb{R}^n)^d$$

for $i, j \in [n]$ with $i \neq j$; see also Example B.3. Now, if $M \subseteq [n]^d$ is not free, then there exist $i = (i_1, \dots, i_d), j = (j_1, \dots, j_d) \in M$ with $i \neq j$ such that they exactly differ one component. Without loss of generality we assume $i_1 \neq j_1$ and $i_l = j_l$ for $l = 2, \dots, n$. But then

$$(\varepsilon_{i_1}, \dots, \varepsilon_{i_d}) = (\varepsilon_{j_1}, \dots, \varepsilon_{j_d}) + (e_{i_1} - e_{j_1}, 0_n, \dots, 0_n),$$

and hence Γ_M is not free. Clearly, the argument can be inverted to show that if Γ_M is not free, then M is not free. ◀

The above proposition shows how the equality $\mu_G(t \cdot v) = \mu_T(t \cdot v)$ of Proposition 4.8 can be verified directly for tensors. For tensors, the moment map components are the quantum marginals, and the equality $\mu_G(t \cdot v) = \mu_T(t \cdot v)$ simply says that the quantum marginals are diagonal. Each off-diagonal entry of a quantum marginal is the inner product between distinct $d - 1$ -dimensional slices of a tensor, and if the support of the tensor is free then the supports of such slices are entirely disjoint - thus the quantum marginals are diagonal.

In the following two Propositions we show, that the subsets of weights, which witness the upper bounds for the (weight) margin in Theorem 2.1, are all free. Thereby, we will implicitly use Proposition 4.13.

► **Proposition 4.14.** For $r \geq 2$ the rows of A_{2r} form a free subset of $[2]^{2r}$, i.e. $\Gamma_{2,2r}$ is free. Moreover, for $r \geq 1$ the set of weights $\Gamma_{2,2r+1}$ is free.

Proof. Clearly, $\Gamma_{2,3} = \{\varepsilon_{1,1,1}, \varepsilon_{2,1,2}\}$ is free. Recall the constructions of $\Gamma_{2,2r}$ and $\Gamma_{2,2r+1}$ from Section 2.1. If $\Gamma_{2,2r}$ is free, then $\Gamma_{2,2r+1}$ is clearly also free. Thus, we are left to prove the former.

Consider A_{2r} as defined in Equation (6). We must show that distinct rows of A_{2r} differ in at least two entries for all $r \geq 2$. The claim is proven by induction on $r \geq 3$. For $r = 3$, we verify the claim by inspection of A_6 . Let a_i be the i^{th} row of A_6 ; its definition is recalled in the left-hand table below. The right-hand table lists for each pair a_i, a_j with $i < j$ two distinct entries in which a_i and a_j differ, which shows the claim for $r = 3$.

13:36 Barriers in Geodesic Optimization

entry	1	2	3	4	5	6
a_1	1	1	1	1	1	1
a_2	2	1	2	2	2	2
a_3	1	2	2	1	1	1
a_4	2	2	1	1	2	2
a_5	1	2	1	2	2	1
a_6	2	2	2	2	1	1

	a_2	a_3	a_4	a_5	a_6
a_1	1,3	2,3	1, 2	2,4	1,2
a_2		1,2	2,3	1,2	5,6
a_3			1,3	3,4	1, 4
a_4				1,4	3,4
a_5					1,3

In fact, the table also proves the claim for $r = 2$, since a_1, \dots, a_4 already pairwise differ in at least two of the first four entries.

Now assume that the claim holds for some fixed $r \geq 3$. Let a_i, a_j be distinct rows of A_{2r+2} ; we will show they differ in at least two entries. If $1 \leq i < j \leq 2r$, then by our inductive hypothesis there is nothing to prove because the first $2r$ rows of A_{2r+2} contain A_{2r} as a submatrix.

To complete the proof, it is enough to show that the $4 \times (2r + 2)$ submatrix formed by restricting to the m^{th} block row, $m \in [r]$, and the last block row of A_{2r+2} satisfies the hypothesis, i.e. any two distinct rows of this submatrix differ in at least two entries. This is the case as restricting to its 1^{st} , m^{th} and last block columns yields a 4×6 submatrix of A_6 if $m \neq 1$, namely

$$\begin{pmatrix} B_2 & B_3 & B_1 \\ B_2 & B_2 & B_3 \end{pmatrix},$$

and a 4×4 submatrix equal to A_4 if $m = 1$. ◀

► **Proposition 4.15.** *For $n \geq 3$ the set $\mathfrak{W}_n \subseteq [n]^3$ is free, i.e. $\Gamma_{n,3} \subseteq \Omega(\pi_{n,3})$ is free. Furthermore, for $n \geq 3$ and $r \geq 2$ the set of weights $\Gamma_{n,6r-3} \subseteq \Omega(\pi_{n,6r-3})$ is free.*

Proof. We remind the reader that

$$\mathfrak{W}_n = \{(s, 1, s), (s, s, 1), (s - 1, s, s) \mid s = 2, 3, \dots, n\}.$$

Let $x = (x_1, x_2, x_3), y = (y_1, y_2, y_3) \in \mathfrak{W}_n$ be such that $x \neq y$. We prove by a distinction of cases that x and y differ in at least two entries. First, we assume $x_1 = y_1$. Then $a := x_1 = y_1 \geq 2$, otherwise $x = (1, 2, 2) = y$ contradicts $x \neq y$. Thus $x, y \in \{(a, 1, a), (a, a, 1), (a, a + 1, a + 1)\}$ and we conclude that x and y differ in at least two entries as $x \neq y$. Second, we assume $x_1 \neq y_1$. There is nothing to show if $x_2 \neq y_2$, so we additionally assume $b := x_2 = y_2$. If $b = 1$, then we are done by $x = (x_1, 1, x_1)$ and $y = (y_1, 1, y_1)$. On the other hand, $b \geq 2$ yields $x, y \in \{(b, b, 1), (b - 1, b, b)\}$ and as $x \neq y$ they differ in the first and third entry. This proves the first statement.

For the second claim, recall that

$$\Gamma_{n,6r-3} = \{\varepsilon_{\sigma(i),\sigma(j),\sigma(k)} \mid (i, j, k) \in \mathfrak{W}_{rn} \setminus \mathfrak{J}_r\},$$

where $\sigma: [rn] \rightarrow [n]^{2r-1}$ is injective, compare Remark 2.10. By the first part \mathfrak{W}_{rn} is free and so is its subset $\mathfrak{W}_{rn} \setminus \mathfrak{J}_r$. Hence $\Gamma_{n,6r-3}$ is free as σ is injective. ◀

We are now ready to deduce Theorem 4.11.

Proof of Theorem 4.11. Recall that all the bounds in Theorem 4.11 hold for the weight margin $\gamma_T(\pi)$ by Theorem 2.1. This was proven by exhibiting witness sets $\Gamma_{n,d} \subseteq \Omega(\pi_{n,d})$ such that $0 \notin \text{conv}(\Gamma_{n,d})$, which gives the bound $\gamma_T(\pi_{n,d}) \leq \text{dist}(0, \text{conv}(\Gamma_{n,d}))$. But if $\Gamma_{n,d}$ is free, then we even have

$$\gamma_G(\pi_{n,d}) \leq \text{dist}(0, \text{conv}(\Gamma_{n,d}))$$

by Proposition 4.8. By Proposition 4.14 the witness sets $\Gamma_{2,3}$ and $\Gamma_{2,2r}, \Gamma_{2,2r+1}, r \geq 2$ for Theorem 2.1(a) are free, which proves Theorem 4.11(a). Similarly, we conclude parts (b) and (c) with Proposition 4.15, which shows that for $n \geq 3$ and $r \geq 2$ the witness sets $\Gamma_{n,3}$ and $\Gamma_{n,6r-3}$ are free. \blacktriangleleft

4.4 Freeness for homogeneous polynomials

In the following we transfer the result from d -tensors to the natural $\mathrm{SL}(n)$ action on homogeneous d -forms in n variables. This representation is given by

$$\varrho_{n,d}: \mathrm{SL}(n) \rightarrow \mathrm{GL}(\mathbb{C}[x_1, \dots, x_n]_d), \quad g \mapsto (p(x) \mapsto p(g^{-1}x)).$$

Each monomial $x^\alpha = x_1^{\alpha_1} \cdots x_n^{\alpha_n}$, given by a multi-index $\alpha = (\alpha_1, \dots, \alpha_n) \in (\mathbb{Z}_{\geq 0})^n$ with $|\alpha| := \sum_i \alpha_i = d$, is a weight vector for $\varrho_{n,d}$ with weight $-\alpha + \frac{d}{n} \mathbb{1}_n$. Therefore

$$\Omega(\varrho_{n,d}) = \left\{ -\alpha + \frac{d}{n} \mathbb{1}_n \mid \alpha \in (\mathbb{Z}_{\geq 0})^n \text{ with } |\alpha| = d \right\},$$

i.e. $\Omega(\varrho_{n,d}) = \Omega'$ from Equation (27) and the bounds from Corollary 2.13 apply to $\gamma_{\mathrm{ST}(n)}(\varrho_{n,d}) = \gamma(\Omega')$. If $n = dm$ for some integer $m \geq 1$, then we have $-\Omega(\pi_{m,d}) \subseteq \Omega(\varrho_{n,d})$.

► Proposition 4.16. *Let $n = dm$ for some integer $m \geq 1$. If $\Gamma \subseteq \Omega(\pi_{m,d})$ is free, then $-\Gamma \subseteq \Omega(\varrho_{n,d})$ is free.*

Proof. We prove the statement by contraposition. Assume that $-\Gamma \subseteq \Omega(\varrho_{n,d})$ is not free. Then there exists a root $\alpha = e_i - e_j \in \mathbb{R}^n$ of $\mathrm{SL}(n)$, where $i, j \in [n]$ with $i \neq j$, and two distinct weights $\omega, \omega' \in -\Gamma$ such that $\omega = \omega' + e_i - e_j$, equivalently $-\omega = -\omega' - e_i + e_j$. The latter equation enforces $-\alpha$ to be of the form

$$(0_m, \dots, 0_m, e_k - e_l, 0_m, \dots, 0_m) \in (\mathbb{R}^m)^d \cong \mathbb{R}^n \quad \text{for some } k, l \in [m] \text{ with } k \neq l,$$

because $-\omega, -\omega' \in \Omega(\pi_{m,d})$. Thus, $-\alpha$ is a root of $\mathrm{SL}(m)^d$ and hence $\Gamma \subseteq \Omega(\pi_{m,d})$ is not free. \blacktriangleleft

As a consequence of the preceding Proposition we obtain bounds for the gap $\gamma_{\mathrm{SL}(n)}(\varrho_{n,d})$.

► Theorem 4.17 (Gap for Polynomial scaling). *Let $d \geq 3$ and let $n = dm$ for some integer $m \geq 2$. Then there exists a constant $C > 0$, independent of n and d such that*

$$\gamma_{\mathrm{SL}(n)}(\varrho_{n,d}) \leq 2^{-Cdm} = 2^{-Cn}.$$

More concretely, for $d = 3$ and $m \geq 3$ it holds that

$$\gamma_{\mathrm{SL}(n)}(\varrho_{n,d}) \leq \mathrm{dist}(0, \Gamma_{m,3}) \leq 2^{-m+1} = 2^{-\frac{n}{3}+1},$$

and if $m \geq 3$ and $d = 6r - 3$ for some $r \geq 2$, we have

$$\gamma_{\mathrm{SL}(n)}(\varrho_{n,d}) \leq \mathrm{dist}(0, \Gamma_{m,6r-3}) \leq 2^{-r(m-1)+1} = 2^{-\frac{(d+3)(m-1)}{6}+1} \approx 2^{-\frac{n}{6}}.$$

Proof. We recall that Theorem 1.6 was proven by padding the results from Theorem 4.11. Thus, for each $m \geq 2$ and $d \geq 3$ the bound $\gamma_{\mathrm{SL}(m)^d}(\pi_{m,d}) \leq 2^{-Cm^d}$ from Theorem 1.6 is witnessed by a free set of weights $\Gamma_{m,d} \subseteq \Omega(\pi_{m,d})$, i.e. $0 < \mathrm{dist}(0, \mathrm{conv}(\Gamma_{m,d})) \leq 2^{-Cdm}$. But then $0 \notin \mathrm{conv}(-\Gamma_{m,d})$ and $-\Gamma_{m,d} \subseteq \Omega(\varrho_{n,d})$ is free by Proposition 4.16. Therefore, Proposition 4.8 yields

$$\gamma_{\mathrm{SL}(n)}(\varrho_{n,d}) \leq \mathrm{dist}(0, \mathrm{conv}(-\Gamma_{m,d})) = \mathrm{dist}(0, \mathrm{conv}(\Gamma_{m,d})) \leq 2^{-Cdm}.$$

Similarly, we get the other bounds by using freeness of $\Gamma_{m,3}$ and, respectively, $\Gamma_{m,6r-3}$ (see Proposition 4.15) combined with the distance bounds Lemma 2.7 and Lemma 2.12, respectively. \blacktriangleleft

4.5 Freeness and diameter bound

In this section we show that the diameter lower bound of Theorem 1.1 generalizes to diameter bounds for the capacity Equation (4) over the noncommutative group $G = \mathrm{SL}(n)^d$. Many algorithms for computing the capacity have resorted to geodesically convex optimization - G can be viewed as a manifold on which $g \mapsto \|g \cdot v\|^2$ is geodesically convex. The distance between an element of g and the identity in this geometry is closely related to the condition number of the matrix g . The diameter bound question is the following: *given an input v and $\varepsilon > 0$, how large a ball in G about the identity must we optimize over to find an approximate minimizer $g \in G$ such that $\|g \cdot v\|^2 - \mathrm{cap}(v) \leq \varepsilon$?* In other words, how well-conditioned can we expect approximate minimizers to Equation (4) to be? This matters because all the algorithms we know start at the origin and take small steps in the manifold, and if all the high-precision solutions are far from the origin then such algorithms cannot reach any of them quickly.

Before tackling this question we must make our notions of distance more precise. The manifold we use is actually not G but rather the manifold P of Hermitian, positive-definite matrices in G . Indeed, we can write

$$\inf_{g \in G} \|g \cdot v\|^2 = \inf_{g \in G} \langle v, g^\dagger g \cdot v \rangle = \inf_{x \in P} \langle v, x \cdot v \rangle.$$

Thus we may instead optimize the function $f_v : g \mapsto \langle v, g \cdot v \rangle$ over P . The manifold P is a prototypical example of a *Hadamard manifold*, a complete, simply connected Riemannian manifold of non-positive sectional curvature [8]. For us, $G = \mathrm{SL}(n)^d$ for some d , and so P is just the set of d -tuples of positive-definite matrices of determinant 1. Even for $d = 1$, P contains a totally geodesic submanifold isometric to the hyperbolic plane; as such the volumes of balls grow exponentially in their radius.¹⁸ The function $f_v : g \mapsto \|g \cdot v\|^2$ is convex along geodesics in this manifold [12]¹⁹. The geodesics through a point $X \in P$ are given by $\gamma(t) = \sqrt{X} e^{Ht} \sqrt{X}$ for Hermitian H . The Riemannian gradient $\nabla \log f_v(g)$ of $\log f_v$ at $g \in P$ is given by the moment map $\mu_G(g \cdot v)$. The geodesic ball of radius R in P about the identity is given by

$$B_R := \{e^A : A \text{ traceless, Hermitian, } \|A\|_F \leq R\} \subseteq P.$$

In a slight abuse of notation, we define the geodesic ball in G (rather than P) to be KB_R , as in the introduction. The values taken by f_v over B_{2R} are the same as the values taken by $g \mapsto \|g \cdot v\|^2$ on KB_R . We now define diameter bounds.

► **Definition 4.18.** *The diameter bound $D_f(\varepsilon)$ for a function f on P and a real number $\varepsilon > 0$ is defined as the infimum over $R > 0$ such that*

$$\inf_{g \in B_R} f(g) \leq \varepsilon + \inf_{g \in P} f(g).$$

We will show that the diameter bound for the norm-squared function can grow faster than $\mathrm{poly}(n, \log(1/\varepsilon))$ for $d = 3$. Firstly, we need to review how diameter bounds for tensors in $(\mathbb{R}_{\geq 0}^n)^d$ like that in Theorem 1.1 relate to diameter bounds for tensors in $(\mathbb{C}^{\otimes n})^d$ over $\mathrm{SL}(n)^d$

¹⁸The volume of a ball can be computed exactly [27], but the very crude bound of volume $\Omega(e^{\Theta(r)} - O(n \log n))$ for the geodesic ball of radius r can be proved elementarily. The manifold $\mathrm{PD}(n) \cap \mathrm{SL}(n)$ contains the hyperbolic plane as a totally geodesic submanifold, in which the ball of radius r has area $e^{\Theta(r)}$ [15]. This shows the ball of radius r in $\mathrm{PD}(n) \cap \mathrm{SL}(n)$ contains $\Omega(e^{\Theta(r)})$ balls of radius 1, which themselves have volume at least $e^{-O(n \log n)}$ by comparison with the Euclidean ball.

¹⁹This was implicitly shown much earlier in [36].

and $\text{ST}(n)^d$. Infimizing $f_v(g)$ over the subset $P \cap \text{ST}(n)^d \subseteq P$, or the tuples of positive-definite diagonal matrices within $\text{SL}(n)^d$, results in a program of the form Equation (2). For $d = 3$, for example,

$$\inf_{g \in P \cap \text{ST}(n)^3} \langle v, g \cdot v \rangle = \text{cap}(p) = \inf_{x \in (\mathbb{R}^n)^3} \sum_{\omega \in \Omega_{n,3}} p_\omega e^{\omega \cdot x} = \inf_{x \in (\mathbb{1}_n^\perp)^3} \sum_{\omega \in \Omega_{n,3}} p_\omega e^{\omega \cdot x} \quad (32)$$

where $\Omega_{n,3} = \{(\varepsilon_i, \varepsilon_j, \varepsilon_k) : i, j, k \in [n]\}$ and $p_{(\varepsilon_i, \varepsilon_j, \varepsilon_k)} = |v_{ijk}|^2$. The correspondence is exactly $g = e^{\text{diag}(x)}$ for $x \in (\mathbb{1}_n^\perp)^3$, which implies the following.

► **Lemma 4.19.** *For all $\varepsilon > 0$, the diameter bound $D_f(\varepsilon)$ for the function $f_v : g \mapsto \langle v, g \cdot v \rangle$ on $\text{ST}(n)^3$ is equal to the diameter bound $D_h(\varepsilon)$ of the function f_p where $p_{ijk} = |v_{ijk}|^2$, or*

$$f_p : (\mathbb{R}^n)^3 \rightarrow \mathbb{R}, \quad x \mapsto \sum_{i,j,k \in [n]} |v_{ijk}|^2 e^{(\varepsilon_i, \varepsilon_j, \varepsilon_k) \cdot x}.$$

Of course, there's nothing special about $d = 3$ here, and the lemma generalizes straightforwardly to other d . For instance, applying Lemma 4.19 for $d = 2$ shows that restricting operator scaling to diagonal matrices yields an instance of matrix scaling. We have shown how diameter bounds over $\text{ST}(n)^d$ relate to those over $(\mathbb{R}^n)^d$. Now we complete the chain by showing how to relate diameter bounds over $\text{SL}(n)^d$ to those over $\text{ST}(n)^d$. We will show that tensors with free support (defined in Definition 4.12) have the same diameter bound over $\text{SL}(n)^d$ as they do over $\text{ST}(n)^d$, which by Theorem 1.1 and Lemma 4.19 we have shown can be superpolynomial. We then show that the construction from Section 3.2 is free.

► **Theorem 4.20.** *Let G denote $\text{SL}(n)^d$, and let T denote $\text{ST}(n)^d$. Suppose $\mu_T(t \cdot v) = \mu_G(t \cdot v)$ for all $t \in T$ (which holds if v has free support). Then for any $R > 0$ we have*

$$\inf_{g \in B_R} f_v(g) = \inf_{g \in T \cap B_R} f_v(g),$$

where B_R denotes the geodesic ball of radius R about the identity in G .

Proof. Define $B := B_R$ and recall that P denotes the positive-definite matrices in G . Let $f : P \rightarrow \mathbb{R}$ be given by $f : g \rightarrow \langle v, g \cdot v \rangle$. Clearly $\inf_{g \in B} f(g) \leq \inf_{g \in T \cap B} f(g)$. We must show the converse inequality. Let $g^* := \arg \min_{g \in B} f(g)$. Recall that P is a Hadamard manifold. Define T_+ to be $T \cap P$. Let πg^* denote the projection of g^* to T_+ , that is, the closest point in T_+ to g^* . As T_+ is a geodesically convex set, projections to T_+ are unique and distances decrease under the projection [8, Theorem 2.1.12]. Thus, $\pi g^* \in B$. If we can show that $f(\pi g^*) \leq f(g^*)$ then the proof is complete.

Let $g^* = \exp_{\pi g^*}(x)$ for some x in the tangent space $T_{\pi g^*}P$ to P at πg^* . That is, $\gamma : [0, 1] \rightarrow P$, $t \mapsto \exp_{\pi g^*}(tx)$ is the geodesic between πg^* and g^* . Then, in the local inner product $\langle \cdot, \cdot \rangle_{\pi g^*}$ at πg^* , x is orthogonal to the tangent space $T_{\pi g^*}T_+ \subseteq T_{\pi g^*}P$ of T_+ at πg^* , because πg^* is a local minimum of the geodesically convex function $d(g^*, \cdot)^2$ on T_+ and x is proportional to the gradient of $d(g^*, \cdot)^2$ at πg^* .

The function f is geodesically convex, and its gradient $\nabla f(\pi g^*)$ is proportional to the moment map $\mu_G(\pi g^* \cdot v)$. By the assumption that $\mu_T(t \cdot v) = \mu_G(t \cdot v)$ for all $t \in T$, $\mu_G(\pi g^* \cdot v)$ is in $i \text{Lie}(T_K)$, which is precisely the tangent space of T_+ at πg^* . Thus

$$f(g^*) = f(\exp_{\pi g^*}(x)) \geq f(\pi g^*) + \langle x, \nabla f(\pi g^*) \rangle_{\pi g^*} = f(\pi g^*),$$

which completes the proof. ◀

► **Lemma 4.21.** *The support of the tensor p from Theorem 1.1 is free.*

13:40 Barriers in Geodesic Optimization

Proof. Recall that a tensor in $(\mathbb{C}^n)^{\otimes 3}$ is free if and only if the supports of distinct rows of its weight matrix intersect in at most one element. The construction in Proposition 3.5 preserves freeness, so we can consider the case $n = 3(l+1)$ treated in the proof of Theorem 1.1. Recall that, in this case, the support of p is $\Omega'_0 \cup \omega'$ where Ω'_0 is the rows of a matrix M defined from the directed graph D_l . Each row in the matrix M corresponds to some edge D_l . Let us first verify that Ω'_0 is free. Assuming the rows correspond to the same edge, they can be verified to have intersection in at most one element, because the nonzero entries of the three rows corresponding to an edge are contained in a 3×6 submatrix with the following form:

$$[A \ I] = \begin{bmatrix} 0 & 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Here the cells containing 1 are colored for readability. Now consider the case that the rows belong to two different edges. If the two edges share no vertices, then clearly the corresponding edges do not intersect. Because the graph is a directed tree, edges may only share a vertex which is the sink of at least one of the edges. If the vertex is a sink for both edges, then the nonzero entries in the 6 rows belonging to either edge (after permutation) take the form

$$\begin{bmatrix} 0 & A & I \\ A & 0 & I \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

If the shared vertex is a sink for only one edge, then the rows are

$$\begin{bmatrix} 0 & A & I \\ A & I & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}.$$

In all these cases it can be verified that supports of distinct rows intersect in at most one element. Lastly, we need to make sure that the intersection of the support of ω' with the support of any element of Ω'_0 is at most one. Recall that ω' is defined to have entry one in each block corresponding to the leaves u_l, v_l, w_l in D_l . However, there are no edges between the leaves, so the support of no row can intersect that of ω' in more than one element. ◀

We are now nearly ready to prove Theorem 1.4. We would simply use the array p from the proof of Theorem 1.1, but setting $|v_{ijk}|^2 = p_{ijk}$ would not be solvable over the rationals. Therefore we must round $\sqrt{p_{ijk}}$, which requires some additional technical lemmas proven in Appendix E.

► **Lemma 4.22** (Rounding and diameter bounds). *Let $p, q : \Omega \rightarrow \mathbb{R}_{\geq 0}$ be positive functions on a finite set $\Omega \subseteq \mathbb{R}^m$. Suppose there is a set B such that*

$$\inf_{x \in B} f_p(x) \geq (1 + \varepsilon) \text{cap } p,$$

and let $M = \max\{1/q_\omega, 1/p_\omega : \omega \in \Omega\}$. Then

$$\inf_{x \in B} f_q(x) \geq ((1 + \varepsilon)(1 - M\|p - q\|_\infty) - M\|p - q\|_1) \text{cap } q.$$

► **Lemma 4.23** (Rounding and capacity). *Let $\Omega \subseteq \mathbb{R}^m$ be finite and let $p, q : \Omega \rightarrow \mathbb{R}_{\geq 0}$ be positive functions on Ω . Let $M_0 = \max_{\omega \in \Omega} 1/q_\omega$. Then*

$$\log \text{cap } q \geq \log \text{cap } p - M_0 \|p - q\|_\infty.$$

Proof of Theorem 1.4. First recall that the values taken by $g \mapsto \|g \cdot v\|^2$ on the geodesic ball KB_R in G are the same as the values taken by $f_v : g \mapsto \langle v, g \cdot v \rangle$ on B_{2R} in P . Thus it is enough to show that $f := f_v$ has diameter bound $D_f(\varepsilon) = \Omega(2^{n/3} \log(1/\varepsilon))$ for $\varepsilon \leq e^{-Cn^2 \log n}$.

We will apply Lemma 4.22 with p as in the proof of Theorem 1.1 and $q_{ijk} = |v_{ijk}|^2$, with v_{ijk} chosen so that v has the same support as p and $p_{ijk} - \delta < |v_{ijk}|^2 \leq p_{ijk}$ for δ small. Because v is free, by Theorem 4.20 the diameter bound for f_v is the same as the diameter bound for f_v over $\text{ST}(n)^3$. By Lemma 4.19, this is the same as the diameter bound for f_q . It remains to show that $D_{f_q}(\varepsilon) = \Omega(2^{n/3} \log(1/\varepsilon))$. We will do this by relating $D_{f_q}(\varepsilon)$ to $D_{f_p}(\varepsilon)$; in particular we will show $D_{f_q}(\Omega(\varepsilon)) \geq D_{f_p}(\varepsilon)$.

Let $R = D_{f_p}(\varepsilon)$. We have $\inf_{x \in (\mathbb{R}^n)^3, \|x\| \leq R} f_p(x) \geq \text{cap}(p) + \varepsilon = (1 + 2\varepsilon) \text{cap}(p)$, recalling that $\text{cap}(p) = 1/2$. By Lemma 4.22,

$$\inf_{x \in (\mathbb{R}^n)^3, \|x\| \leq R} f_q(x) \geq ((1 + 2\varepsilon)(1 - M\|p - q\|_\infty) - M\|p - q\|_1) \text{cap}(q).$$

As $\text{cap } q \leq 1/2$, if $M\|p - q\|_\infty \leq M\|p - q\|_1 \leq c\varepsilon$ for c a small enough constant, then we have $((1 + 2\varepsilon)(1 - M\|p - q\|_\infty) - M\|p - q\|_1) \text{cap } q = \text{cap } q + \Omega(\varepsilon)$, so

$$\inf_{x \in (\mathbb{R}^n)^3, \|x\| \leq R} f_q(x) \geq \text{cap } q + \Omega(\varepsilon).$$

Thus $D_{f_q}(\Omega(\varepsilon)) \geq D_{f_p}(\varepsilon)$ assuming $M\|p - q\|_1 \leq c\varepsilon$. To ensure that this constraint is satisfied, choose v of bit complexity $O(\log n + \log(1/\varepsilon))$ such that $\|p - q\|_1 = \frac{c}{n}\varepsilon$. Because $p_{ijk} = \Omega(1/n)$ for i, j, k in the support of p by construction, we have $q_{ijk} = \Omega(1/n)$ for i, j, k in the support of q and hence $M = O(n)$. Thus $M\|p - q\|_1 \leq c\varepsilon$. Applying Lemma 4.23 together with our assumptions about the size of $p - q$ and the fact that $\text{cap}(q) = \text{cap}(v)$ implies the final claim that $\text{cap}(v) \geq 1/4$ and that $1 \geq \|v\| \geq 1/2$. ◀

Finally, we remark that the same diameter bound holds for $d \geq 3$ for tuples of tensors. We note that if $v \in (\mathbb{C}^n)^{\otimes 3}$ has free support, then so does the tensor $v \otimes e_l \otimes \dots \otimes e_l \in (\mathbb{C}^n)^{\otimes d}$ for $d \geq 3$. By Proposition 4.8, the tuple $w \in ((\mathbb{C}^n)^{\otimes d})^n$ given by

$$w_l = \frac{1}{n} v \otimes e_l \otimes \dots \otimes e_l \text{ for } l \in [n]$$

has $\mu_T(t \cdot v) = \mu_G(t \cdot v)$ for all $t \in \text{ST}(n)^d$. The commutative problem obtained by restricting to $\text{SL}(n)^d$ as in Lemma 4.19 is precisely f_q as in Corollary 3.7. As in the proof of Theorem 1.4, by Theorem 4.20, Lemma 4.19 and Corollary 3.7, we have the following.

► **Corollary 4.24.** *There is a constant $C > 0$ such that the following holds for all $d \geq 3$. For all $\varepsilon \leq \exp(-Cn^2 \log n)$, there is a tuple of tensors $w = w(\varepsilon) \in ((\mathbb{C}^n)^{\otimes d})^n$ with $O(n^2)$ nonzero entries of bit complexity $O(\log n + \log(1/\varepsilon))$, and a geodesic ball $B = B(\varepsilon)$ of radius $\Omega(2^{n/3} \log(1/\varepsilon))$ about the identity in $\text{SL}(n)^d$, such that*

$$\inf_{g \in B} \|g \cdot w\|^2 \geq \text{cap}(v) + \varepsilon.$$

Moreover, it holds that $1/4 \leq \text{cap}(w) \leq 1$ and $1/2 \leq \|w\| \leq 1$.

4.6 A bound on weight margin and gap for quivers

For $d \geq 2$ let Q_d be the quiver

$$\begin{aligned} 1 \longleftarrow 2 \longrightarrow 3 \cdots \cdots d-2 \longrightarrow d-1 \longleftarrow d & \quad \text{if } d \text{ even} \\ 1 \longrightarrow 2 \longleftarrow 3 \cdots \cdots d-2 \longrightarrow d-1 \longleftarrow d & \quad \text{if } d \text{ odd.} \end{aligned}$$

and let $Q_d^{(k)}$ be the quiver one obtains from Q_d by adding $k - 1$ additional copies of each arrow in Q_d . As before, let $G = \text{SL}(n)^d$ and $T = \text{ST}(n)^d$. Then G acts on the quiver Q_d with dimension vector (n, \dots, n) as described in the introduction. We denote the corresponding representation by π_d . Note that the action of G on $Q_d^{(k)}$ with dimension vector (n, \dots, n) is given by π_d^k . In this subsection we prove a bound on the weight margin of π_d and on the gap of π_d^n . The bound on $\gamma_G(\pi_d^n)$ is thanks to the refinement of freeness in Proposition 4.8 pointed out by Visu Makam.

► **Theorem 4.25.** *Let $n, d \geq 2$ and denote the natural action of $G = \text{SL}(n)^d$ on the quiver Q_d with dimension vector (n, \dots, n) by $\pi_d: \text{SL}(n)^d \rightarrow \text{GL}(V_d)$, where $V_d = (\mathbb{C}^{n \times n})^{d-1}$. The representation π_d^n corresponds to the G -action on the quiver $Q_d^{(n)}$ with dimension vector (n, \dots, n) . It holds that*

$$\gamma_T(\pi_d) \leq (n - 1)^{-d+1} \quad \text{and} \quad \gamma_G(\pi_d^n) \leq (n - 1)^{-d+1}.$$

► **Remark 4.26.** *Before proving the theorem, we point out a few consequences.*

1. *Theorem 4.25 shows that $\gamma_T(\pi_d)^{-1}$ and $\gamma_G(\pi_d^n)^{-1}$ are not polynomially bounded with respect to $\dim V_d = (d - 1)n^2$ and $\dim \text{SL}(n)^d = d(n^2 - 1)$. Instead we see for fixed n and $d \rightarrow \infty$ an exponential behaviour in the number of vertices d . Thus, our bound shows that the exponential behaviour in d cannot be avoided in general lower bounds for quiver actions like [12, Theorem 6.21 Item 4]. The latter applied to π_d shows $\gamma_T(\pi_d) \geq n^{-d^2 - (3/2)d} (dn + 1)^{-d}$.*
2. *The proof of Theorem 4.25 below shows that for the bound on the gap it is enough to consider the quiver $Q_d^{(n-1)}$ with an additional n^{th} arrow from d to $d - 1$.*
3. *The ideas presented below can be adjusted to prove similar bounds for other dimension vectors. For example, one can show that the gap for the SL -action on $Q_d^{(2)}$ with dimension vector $(1, 3, 3, \dots, 3, 2)$ is inverse exponential in d . This aligns with an algebraic barrier for this action; the invariants that cut out the null cone for this action have exponential degree [19, Proposition 1.5].*
4. *The quiver Q_d is of finite representation type and has no oriented cycles. Therefore, the null-cone membership problem for π_d can be solved in polynomial-time by algebraic algorithms.²⁰ This means Q_d is an example where the weight margin is very small but there still exist efficient algorithms. Can the existence of efficient algorithms still be explained by a large gap in this case? This leads to the following interesting open question.*

► **Problem 4.27.** *Is the gap $\gamma_G(\pi_d)$ inverse polynomial in n and d ?*

A positive answer would provide an interesting example, since in this case the weight margin of π_d would be *significantly* smaller than the gap of π_d .

We now introduce several lemmas needed to prove Theorem 4.25. Note that the set of weights of π_d viewed as a subset of $(\mathbb{R}^n)^d$ is

$$\left\{ ((-1)^d \varepsilon_i, (-1)^{d-1} \varepsilon_j, 0, \dots, 0), (0, (-1)^{d-1} \varepsilon_i, (-1)^{d-2} \varepsilon_j, 0, \dots, 0), \dots, (0, \dots, 0, \varepsilon_i, -\varepsilon_j) \mid i, j \in [n] \right\}.$$

²⁰ Personal communication with Visu Makam. There does not seem to be an explicit reference in the literature.

We define recursively the subsets of weights

$$\Gamma_2 := \{(\varepsilon_i, -\varepsilon_j) \mid i \in [n-1], j \in [n]\} \subseteq \Omega(\pi_2) \subseteq \mathbb{R}^{2n}$$

$$\text{for } d \geq 3, \Gamma_d := \left\{((-1)^d \varepsilon_i, (-1)^{d-1} \varepsilon_n, 0_n, \dots, 0_n) \mid i \in [n-1]\right\} \cup (\{0_n\} \times \Gamma_{d-1}) \subseteq \Omega(\pi_d) \subseteq \mathbb{R}^{dn}.$$

► **Remark 4.28.** We note that for $d \geq 2$, Γ_d is not free. For instance, we can always write

$$(0_n, \dots, 0_n, \varepsilon_1, -\varepsilon_1) = (0_n, \dots, 0_n, \varepsilon_1, -\varepsilon_2) + (0_n, \dots, 0_n, 0_n, e_2 - e_1),$$

i.e. the weights $(0_n, \dots, 0_n, \varepsilon_1, -\varepsilon_1), (0_n, \dots, 0_n, \varepsilon_1, -\varepsilon_2) \in \Gamma_d$ differ by the root $(0_n, \dots, 0_n, 0_n, e_2 - e_1)$ of $\text{SL}(n)^d$. Therefore, we cannot deduce a bound on the gap $\gamma_G(\pi_d)$ via Proposition 4.8. However, the latter allows us to deduce at least a bound on the gap of π_d^n .

In the next two lemmas we show that Γ_d witnesses the bound on $\gamma_{\text{T}}(\pi_d)$ and afterwards we use Proposition 4.8 to transfer this bound to $\gamma_G(\pi_d^n)$.

► **Lemma 4.29.** For all $d \geq 2$ it holds that $0 \notin \text{conv}(\Gamma_d)$.

Proof. We prove the statement by induction on $d \geq 2$. For $d = 2$, just note that any element in $\text{conv}(\Gamma_2) \subseteq \mathbb{R}^{2n}$ has value $-1/n$ in the n -th entry. In particular, $0 \notin \text{conv}(\Gamma_2)$. For $d \geq 3$ let

$$x = \sum_{\omega \in \Gamma_d} \lambda_{\omega} \omega, \quad \lambda_{\omega} \geq 0$$

be a convex combination of the elements in Γ_d . Assume there is an $i \in [n-1]$ such that for

$$\omega_i := ((-1)^d \varepsilon_i, (-1)^{d-1} \varepsilon_n, 0_n, \dots, 0_n)$$

one has $\lambda_{\omega_i} > 0$. Then the n -th entry of x is non-zero, since ω_i has n -th entry $(-1)^{d+1}/n$ and all (other) $\omega \in \Gamma_d$ have $(-1)^{d+1}/n$ or zero as n -th entry. On the other hand, if $\lambda_{\omega_i} = 0$ for all $i \in [n-1]$, then $x \in \{0_n\} \times \text{conv}(\Gamma_{d-1})$. By induction hypothesis on $d-1$ we necessarily have $x \neq 0$. ◀

► **Lemma 4.30.** For $d \geq 2$ it holds that $x_d := \lambda_d((-1)^{d-1} \varepsilon_n, 0_n, \dots, 0_n) \in \text{conv}(\Gamma_d)$, where

$$\lambda_d := \left(\sum_{i=1}^{d-1} (n-1)^i \right)^{-1}.$$

In particular, $\|x_d\|_2 < |\lambda_d| \leq (n-1)^{-d+1}$.

Proof. We proceed by induction on $d \geq 2$. In the case $d = 2$, consider the convex combination

$$\sum_{i=1}^{n-1} \sum_{j=1}^n \frac{1}{(n-1)n} (\varepsilon_i, -\varepsilon_j) = \frac{1}{n-1} (-\varepsilon_n, 0_n) = x_2,$$

where we used (5). Now assume the claim is proven for some $d \geq 2$, hence

$$\lambda_d(0_n, (-1)^{d-1} \varepsilon_n, 0_n, \dots, 0_n) \in \{0_n\} \times \text{conv}(\Gamma_d) \subseteq \text{conv}(\Gamma_{d+1}). \tag{33}$$

Setting $\mu := (n-1)\lambda_{d+1}\lambda_d^{-1}$ we have $\mu\lambda_d = (n-1)\lambda_{d+1}$ and $\mu + (n-1)\lambda_{d+1} = 1$. Together with (5) and (33) we deduce $x_{d+1} \in \text{conv}(\Gamma_{d+1})$ via

$$\mu \lambda_d(0_n, (-1)^{d-1} \varepsilon_n, 0_n, \dots, 0_n) + \lambda_{d+1} \sum_{i=1}^{n-1} ((-1)^{d+1} \varepsilon_i, (-1)^d \varepsilon_n, 0_n, \dots, 0_n) = x_{d+1}.$$

This ends the induction. Finally, $\|x_d\|_2 < |\lambda_d|$ follows from $\|\varepsilon_n\|_2 < 1$. ◀

Proof of Theorem 4.25. By Lemma 4.29 and Lemma 4.30 we have

$$\gamma_{\mathbb{T}}(\pi_d) \leq (n-1)^{-d+1}.$$

With the fact $\Omega(\pi_d) = \Omega(\pi_d^n)$ and with Proposition 4.8 we transfer this bound to the gap of π_d^n . To do so, we note that the natural inner product on $V_d^n = (\mathbb{C}^{n \times n})^{n(d-1)}$, given by the trace inner product on each $\mathbb{C}^{n \times n}$ copy, is invariant under the action of $K = \text{SU}(n)^d$. Clearly, distinct $\mathbb{C}^{n \times n}$ copies are orthogonal under this inner product. Thus, to be able to apply Proposition 4.8 it is enough to assign to each $\mathbb{C}^{n \times n}$ copy, i.e. to each arrow of $Q_d^{(n)}$, a matrix M_i such that $\text{supp}(M_i)$ is free and $\Gamma_d = \bigcup_i \text{supp}(M_i)$.

For this, we consider the $n \times n$ matrices

$$M := \begin{pmatrix} I_{n-1} & 0 \\ 0 & 0 \end{pmatrix} \quad \text{and} \quad P := \begin{pmatrix} 0 & I_{n-1} \\ 1 & 0 \end{pmatrix},$$

and $E_{i,j}$ is the matrix with (i,j) -entry one and all other entries zero. Then $E_{i,i}P = E_{i,\sigma(i)}$, where $\sigma: [n] \rightarrow [n]$ is the cycle $(1\ 2\ \dots\ n)$. Therefore, for $k \in [n]$ we have

$$\text{supp}(MP^{k-1}) = \{(0_{n(d-2)}, \varepsilon_i, -\varepsilon_{\sigma^{k-1}(i)}) \mid i \in [n-1]\} \text{ and } \{0_{n(d-2)}\} \times \Gamma_2 = \bigcup_{k \in [n]} \text{supp}(MP^{k-1}).$$

For fixed k , $i_1 \neq i_2$ implies $\sigma^{k-1}(i_1) \neq \sigma^{k-1}(i_2)$, so any distinct elements of $\text{supp}(MP^{k-1})$ differ in the last two \mathbb{R}^n -components. Hence, each $\text{supp}(MP^{k-1})$ is free and we assign M, MP, \dots, MP^{n-1} to the n arrows that go from vertex d to vertex $d-1$. For $l \in [d-2]$, we assign to the n arrows between the vertices l and $l+1$ each of the matrices $E_{1,n}, E_{2,n}, \dots, E_{n-1,n}$ at least once. (Exactly one of the latter matrices is assigned to two of these arrows.) Clearly, the support of $E_{i,n}$, $i \in [n-1]$ is free as it contains just one weight. By construction, this assignment does the job. Moreover, the argument shows that $n-1$ arrows between the vertices l and $l+1$, $l \in [d-2]$, suffice. \blacktriangleleft

References

- 1 P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, Princeton, NJ, 2008. With a foreword by Paul Van Dooren. doi:10.1515/9781400830244.
- 2 Zeyuan Allen-Zhu, Ankit Garg, Yuanzhi Li, Rafael Oliveira, and Avi Wigderson. Operator scaling via geodesically convex optimization, invariant theory and polynomial identity testing. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 172–181, 2018.
- 3 Zeyuan Allen-Zhu, Yuanzhi Li, Rafael Oliveira, and Avi Wigderson. Much faster algorithms for matrix scaling. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 890–901. IEEE, 2017.
- 4 Noga Alon and Van H. Vu. Anti-Hadamard matrices, coin weighing, threshold gates, and indecomposable hypergraphs. *Journal of Combinatorial Theory, Series A*, 79(1):133–160, 1997.
- 5 Jason M. Altschuler and Enric Boix-Adsera. Polynomial-time algorithms for Multimarginal Optimal Transport problems with structure, 2020. arXiv:2008.03006.
- 6 Carlos Amendola, Kathlen Kohn, Philipp Reichenbach, and Anna Seigal. Invariant theory and scaling algorithms for maximum likelihood estimation, 2020. arXiv:2003.13662.
- 7 Nima Anari, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials, entropy, and a deterministic approximation algorithm for counting bases of matroids. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 35–46. IEEE, 2018.

- 8 Miroslav Bacák. *Convex analysis and optimization in Hadamard spaces*, volume 22. Walter de Gruyter GmbH & Co KG, 2014.
- 9 Rajendra Bhatia. *Positive definite matrices*. Princeton Series in Applied Mathematics. Princeton University Press, Princeton, NJ, 2007.
- 10 Peter Bürgisser, Matthias Christandl, Ketan D. Mulmuley, and Michael Walter. Membership in moment polytopes is in NP and coNP. *SIAM J. Comput.*, 46(3):972–991, 2017. doi: 10.1137/15M1048859.
- 11 Peter Bürgisser, Cole Franks, Ankit Garg, Rafael Oliveira, Michael Walter, and Avi Wigderson. Efficient algorithms for tensor scaling, quantum marginals, and moment polytopes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 883–897. IEEE, 2018.
- 12 Peter Bürgisser, Cole Franks, Ankit Garg, Rafael Oliveira, Michael Walter, and Avi Wigderson. Towards a theory of non-commutative optimization: geodesic first and second order methods for moment maps and polytopes, 2019. arXiv:1910.12375.
- 13 Peter Bürgisser, Ankit Garg, Rafael Oliveira, Michael Walter, and Avi Wigderson. Alternating Minimization, Scaling Algorithms, and the Null-Cone Problem from Invariant Theory. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*, volume 94 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 24:1–24:20, 2018. doi: 10.4230/LIPIcs.ITCS.2018.24.
- 14 Peter Bürgisser, Yinan Li, Harold Nieuwboer, and Michael Walter. Interior-point methods for unconstrained geometric programming and scaling problems, 2020. arXiv:2008.12110.
- 15 James W. Cannon, William J. Floyd, Richard Kenyon, Walter R. Parry, et al. Hyperbolic geometry. *Flavors of geometry*, 31:59–115, 1997.
- 16 Michael B. Cohen, Aleksander Madry, Dimitris Tsipras, and Adrian Vladu. Matrix scaling and balancing via box constrained Newton’s method and interior point methods. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 902–913. IEEE, 2017.
- 17 Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pages 2292–2300, 2013.
- 18 Jiri Dadok and Victor Kac. Polar representations. *J. Algebra*, 92(2):504–524, 1985. doi: 10.1016/0021-8693(85)90136-X.
- 19 Harm Derksen and Visu Makam. Degree bounds for semi-invariant rings of quivers. *J. Pure Appl. Algebra*, 222(10):3282–3292, 2018. doi:10.1016/j.jpaa.2017.12.007.
- 20 Harm Derksen and Visu Makam. Algorithms for orbit closure separation for invariants and semi-invariants of matrices. *Algebra Number Theory*, 14(10):2791–2813, 2020. doi: 10.2140/ant.2020.14.2791.
- 21 Harm Derksen and Visu Makam. An exponential lower bound for the degrees of invariants of cubic forms and tensor actions. *Adv. Math.*, 368:107136, 25, 2020. doi:10.1016/j.aim.2020.107136.
- 22 Michael A Forbes and Amir Shpilka. Explicit noether normalization for simultaneous conjugation via polynomial identity testing. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 527–542. Springer, 2013.
- 23 Cole Franks and Ankur Moitra. Rigorous Guarantees for Tyler’s M-estimator via quantum expansion, 2020. arXiv:2002.00071.
- 24 Matthias Franz. Moment polytopes of projective G -varieties and tensor products of symmetric group representations. *J. Lie Theory*, 12(2):539–549, 2002.
- 25 Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. A deterministic polynomial time algorithm for non-commutative rational identity testing. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 109–117. IEEE, 2016.
- 26 Ankit Garg, Christian Ikenmeyer, Visu Makam, Rafael Oliveira, Michael Walter, and Avi Wigderson. Search Problems in Algebraic Complexity, GCT, and Hardness of Generators for Invariant Rings. In *35th Computational Complexity Conference (CCC 2020)*, volume 169 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 12:1–12:17, 2020. doi:10.4230/LIPIcs.CCC.2020.12.

- 27 X. Gual-Arnau and A. M. Naveira. Volume of tubes in noncompact symmetric spaces. *Publ. Math. Debrecen*, 54(3-4):313–320, 1999.
- 28 V. Guillemin and S. Sternberg. Convexity properties of the moment mapping. II. *Invent. Math.*, 77(3):533–546, 1984. doi:10.1007/BF01388837.
- 29 Leonid Gurvits. Classical complexity and quantum entanglement. *Journal of Computer and System Sciences*, 69(3):448–484, 2004.
- 30 Leonid Gurvits. Combinatorial and algorithmic aspects of hyperbolic polynomials, 2004. arXiv:math/0404474.
- 31 Brian C. Hall. *Lie groups, Lie algebras, and representations*, volume 222 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2003. An elementary introduction. doi:10.1007/978-0-387-21554-9.
- 32 Linus Hamilton and Ankur Moitra. No-go Theorem for Acceleration in the Hyperbolic Plane, 2021. arXiv:2101.05657.
- 33 Moritz Hardt and Ankur Moitra. Algorithms and hardness for robust subspace recovery. In *Conference on Learning Theory*, pages 354–375, 2013.
- 34 Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Constructive non-commutative rank computation is in deterministic polynomial time. *Comput. Complexity*, 27(4):561–593, 2018. doi:10.1007/s00037-018-0165-7.
- 35 Bahman Kalantari and Leonid Khachiyan. On the complexity of nonnegative-matrix scaling. *Linear Algebra and its applications*, 240:87–103, 1996.
- 36 George Kempf and Linda Ness. The length of vectors in representation spaces. In *Algebraic geometry (Proc. Summer Meeting, Univ. Copenhagen, Copenhagen, 1978)*, volume 732 of *Lecture Notes in Math.*, pages 233–243. Springer, Berlin, 1979.
- 37 M. K. Kravtsov and V. E. Lukshin. On some properties of noninteger vertices of a three-index axial transportation polytope. *Tr. Inst. Matematiki NAN Belarusi*, 13(2):31–36, 2005.
- 38 V. M. Kravtsov. Combinatorial properties of noninteger vertices of a polytope in a three-index axial assignment problem. *Kibernet. Sistem. Anal.*, 43(1):33–44, 189, 2007. doi:10.1007/s10559-007-0023-0.
- 39 Tianyi Lin, Nhat Ho, Marco Cuturi, and Michael I. Jordan. On the complexity of approximating multimarginal optimal transport, 2019. arXiv:1910.00152.
- 40 Nathan Linial and Zur Luria. On the vertices of the d-dimensional Birkhoff polytope. *Discrete & Computational Geometry*, 51(1):161–170, 2014.
- 41 Tomasz Maciążek and Adam Sawicki. Critical points of the linear entropy for pure L-qubit states. *Journal of Physics A: Mathematical and Theoretical*, 48(4):045305, January 2015. doi:10.1088/1751-8113/48/4/045305.
- 42 Tomasz Maciążek and Adam Sawicki. Asymptotic properties of entanglement polytopes for large number of qubits. *Journal of Physics A: Mathematical and Theoretical*, 51(7):07LT01, January 2018. doi:10.1088/1751-8121/aaa4d7.
- 43 Ketan Mulmuley. Geometric complexity theory V: Efficient algorithms for Noether normalization. *Journal of the American Mathematical Society*, 30(1):225–309, 2017.
- 44 David Mumford. *Geometric Invariant Theory*. Ergebnisse der Mathematik und ihrer Grenzgebiete, Neue Folge, Band 34. Springer-Verlag, Berlin-New York, 1965.
- 45 Linda Ness. A stratification of the null cone via the moment map. *Amer. J. Math.*, 106(6):1281–1329, 1984. With an appendix by David Mumford. doi:10.2307/2374395.
- 46 Beresford N. Parlett and Christian Reinsch. Balancing a matrix for calculation of eigenvalues and eigenvectors. In *Handbook for Automatic Computation*, pages 315–326. Springer, 1971.
- 47 Alexander Rusciano. A Riemannian Corollary of Helly’s theorem. *J. Convex Anal.*, 27(4):1261–1275, 2020.
- 48 Mohit Singh and Nisheeth K. Vishnoi. Entropy, optimization and counting. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 50–59, 2014.
- 49 Reyer Sjamaar. Convexity properties of the moment mapping re-examined. *Adv. Math.*, 138(1):46–91, 1998. doi:10.1006/aima.1998.1739.

- 50 Damian Straszak and Nisheeth K. Vishnoi. Maximum entropy distributions: Bit complexity and stability. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 2861–2891. PMLR, 25–28 June 2019. [arXiv:1711.02036](https://arxiv.org/abs/1711.02036).
- 51 Nolan R. Wallach. *Geometric Invariant Theory: Over the real and complex numbers*. Universitext. Springer, Cham, 2017. doi:10.1007/978-3-319-65907-7.
- 52 Hermann Weyl. *The classical groups: their invariants and representations*, volume 45. Princeton university press, 1946.
- 53 Hongyi Zhang and Suvrit Sra. First-order methods for geodesically convex optimization. In *Conference on Learning Theory*, pages 1617–1638. PMLR, 2016.

A Notation

f_p	the function $\mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}, x \mapsto \sum_{\omega \in \Omega} p_\omega e^{\omega \cdot x}$, see Equation (2)
$\text{cap}(p)$	the capacity of a non-negative function p on a finite set $\Omega \subseteq \mathbb{R}^m$, see Equation (2)
$\text{cap}(v)$	the capacity of a vector v under a group action, see Equation (4)
$[n]$	the set $\{1, 2, \dots, n\}$
0_n	the zero vector in \mathbb{R}^n
e_i	the i^{th} canonical unit vector in \mathbb{R}^n
$\mathbb{1}_n$	the all-ones vector in \mathbb{R}^n
$\mathbb{1}_n^\perp$	the orthogonal complement of $\mathbb{1}_n$ in \mathbb{R}^n , i.e. $\{(v_1, \dots, v_n) \in \mathbb{R}^n : \sum_i v_i = 0\}$
ε_i	the vector $e_i - \frac{1}{n} \mathbb{1}_n$
I_n	the $n \times n$ identity matrix
$\text{dist}(0, S)$	the distance from the origin to the set S
$\text{conv}(S)$	the convex hull of S in \mathbb{R}^n
$\text{Aff}(S)$	the affine hull of S in \mathbb{R}^n
$\pi_{n,d}$	the representation for d -dimensional tensor scaling
$\Omega(\pi)$	the set of weights of a representation π
$\Omega_{n,d} = \Omega(\pi_{n,d})$	the set $\{\varepsilon_i : i \in [n]\}^d$ corresponding to d -dimensional array scaling; equal to the set of weights of the tensor scaling representation $\pi_{n,d}$, see Example 4.5
$\gamma(\Omega)$	the <i>margin</i> of the finite set $\Omega \subseteq \mathbb{R}^m$, see Definition 1.2
$\gamma_T(\pi)$	the <i>weight margin</i> of a representation π , i.e. $\gamma(\Omega(\pi))$, see Definition 4.3
$\gamma_G(\pi)$	the <i>gap</i> of a representation π , see Definition 4.3
$\text{tr}(A)$	the trace of a square matrix A
$D_f(\varepsilon)$	the diameter bound of a function f for $\varepsilon > 0$, see Definition 3.1 respectively Definition 4.18
$\ A\ _F$	the Frobenius norm of a square matrix A
e^A	the exponential of a square matrix A
$\text{Lie}(G)$	the Lie algebra of a matrix Lie group G
$\text{GL}(n)$	the group of invertible <i>complex</i> $n \times n$ matrices
$\text{SL}(n)$	the group of invertible <i>complex</i> $n \times n$ matrices with determinant one
$\text{ST}(n)$	the group of diagonal invertible <i>complex</i> $n \times n$ matrices with determinant one
$\text{SU}(n)$	the group of unitary matrices of size $n \times n$ and determinant one
$\text{Herm}(n)$	the set of complex Hermitian $n \times n$ matrices
$\text{GL}(V)$	the group of \mathbb{C} -linear, bijective maps $V \rightarrow V$, where V is a \mathbb{C} -vector space

B Representation theory background

In this section we briefly recall some representation theory. All the concepts we present here actually work in the very general setting of reductive groups and their rational representations, see e.g. [12, section 2]. For the sake of clarity and concreteness we stick to the special case needed in this paper, i.e. the reductive group $\text{SL}(n)^d := \text{SL}(n) \times \dots \times \text{SL}(n)$ with $d \geq 1$ many copies of $\text{SL}(n)$.

We call a Euclidean-closed subgroup $H \subseteq \text{GL}(n)$ a *matrix Lie group*. Indeed, such an H is naturally a Lie group (c.f. [31, Theorem 1.19]) with real *Lie algebra*

$$\text{Lie}(H) := \{A \in \mathbb{C}^{n \times n} \mid \forall t \in \mathbb{R}: e^{tA} \in H\}.$$

13:48 Barriers in Geodesic Optimization

The Lie bracket for $\text{Lie}(H)$ is the commutator $[A, B] := AB - BA$. Moreover, for $d \geq 1$ the product $H^d := H \times \cdots \times H$ becomes a matrix Lie group via block-diagonal embedding into $\text{GL}(dn)$, i.e.

$$H^d \hookrightarrow \text{GL}(dn), \quad (h_1, \dots, h_d) \mapsto \begin{pmatrix} h_1 & & \\ & \ddots & \\ & & h_d \end{pmatrix}$$

Then the Lie algebra of H^d is $\text{Lie}(H^d) = \text{Lie}(H) \times \cdots \times \text{Lie}(H)$ block-diagonally embedded into $\mathbb{C}^{dn \times dn}$. If $G \subseteq \text{GL}(n)$ is another matrix Lie group, then $G \times H$ is again a matrix Lie group with Lie algebra $\text{Lie}(G \times H) = \text{Lie}(G) \times \text{Lie}(H)$.

► **Example B.1.** The groups $\text{GL}(n)$, $\text{SL}(n)$, $\text{U}(n)$ and $\text{GT}(n)$ are matrix Lie groups with Lie algebras

$$\begin{aligned} \text{Lie}(\text{GL}(n)) &= \mathbb{C}^{n \times n} & \text{Lie}(\text{U}(n)) &= \{A \in \mathbb{C}^{n \times n} \mid A^\dagger = -A\} = i \text{Herm}(n) \\ \text{Lie}(\text{SL}(n)) &= \{A \in \mathbb{C}^{n \times n} \mid \text{tr}(A) = 0\} & \text{Lie}(\text{GT}(n)) &= \{A \in \mathbb{C}^{n \times n} \mid A \text{ diagonal matrix}\}. \end{aligned}$$

Therefore, also $\text{SU}(n)$, $\text{ST}(n)$ and $\text{U}(n) \cap \text{ST}(n)$ are matrix Lie groups and their Lie algebras are obtained by corresponding intersections of the above Lie algebras. In particular, we have

$$\text{Lie}(\text{U}(n) \cap \text{ST}(n)) = \{i \text{diag}(x_1, \dots, x_n) \mid x_j \in \mathbb{R}, x_1 + \dots + x_n = 0\}.$$

Thus, we can identify $i \text{Lie}(\text{U}(n) \cap \text{ST}(n))$ with the orthogonal complement $(\mathbb{1}_n)^\perp \subseteq \mathbb{R}^n$ of the all-ones vector $\mathbb{1}_n$.

In the following, let $G := \text{SL}(n)^d$ for some $d \geq 1$. Then $K := \text{SU}(n)^d$ is a maximal compact subgroup of G , and $\text{T} := \text{ST}(n)^d$ and $\text{T}_K := K \cap \text{T}$ are maximal tori of G and K , respectively. As explained above, we think of all these groups as matrix Lie subgroups of $\text{GL}(dn)$, and hence of their Lie algebras as subsets of $\mathbb{C}^{dn \times dn}$.

A rational representation of $G = \text{SL}(n)^d$ is a group morphism $\pi: G \rightarrow \text{GL}(V)$, such that in some basis of V the matrix entries of $\pi(g) \in \text{GL}(V)$ are polynomials in the matrix entries of g .²¹ Such a rational representation of G induces a representation of the Lie algebras by

$$\Pi: \text{Lie}(G) \rightarrow \text{End}(V), \quad A \mapsto \left. \frac{d}{dt} \right|_{t=0} \pi(e^{tA})$$

with the property $\pi(e^A) = e^{\Pi(A)}$ for all $A \in \text{Lie}(G)$. Restricting π to the commutative subgroup T induces a so-called *weight space decomposition* of V . That is, there is some finite set $\Omega(\pi) \subseteq i \text{Lie}(\text{T}_K)$ and a decomposition $V = \bigoplus_{\omega \in \Omega(\pi)} V_\omega$ into non-zero subspaces such that each $\omega \in \Omega(\pi)$ and any $v_\omega \in V_\omega$ satisfy

$$\forall A \in \text{Lie}(\text{T}): \quad \pi(e^A) v_\omega = e^{\text{tr}(A\omega)} v_\omega$$

or, equivalently,

$$\forall A \in \text{Lie}(\text{T}): \quad \Pi(A) v_\omega = \text{tr}(A\omega) v_\omega.$$

The elements $\omega \in \Omega(\pi)$ are called *weights* of π and the $v_\omega \in V_\omega$ are called *weight vectors*. Considering Example B.1 we frequently use the identification $i \text{Lie}(\text{T}_K) \cong (\mathbb{1}_n)^\perp$, where $(\mathbb{1}_n)^\perp$ is the orthogonal complement of $\mathbb{1}_n$ in \mathbb{R}^n . We note that for $\omega \in i \text{Lie}(\text{T}_K) \subseteq \mathbb{C}^{dn \times dn}$ the Frobenius norm $\|\omega\|_F$ becomes under this identification the 2-norm $\|\omega\|_2$ in $(\mathbb{R}^n)^d$.

²¹ In other words, π is a morphism of affine algebraic groups.

► **Example B.2.** Let $d = 1$. The group $G = \mathrm{SL}(n)$ acts on \mathbb{C}^n by left-multiplication, which induces the rational representation $\pi: \mathrm{SL}(n) \rightarrow \mathrm{GL}(n), g \mapsto g$ with corresponding Lie algebra representation $\Pi: \mathrm{Lie}(\mathrm{SL}(n)) \rightarrow \mathbb{C}^{n \times n}, A \mapsto A$. For $i \in [n]$ we set

$$\varepsilon_i := e_i - \frac{1}{n} \mathbb{1}_n \in \mathbb{1}_n^\perp \subseteq \mathbb{R}^n.$$

For all $A = \mathrm{diag}(a_1, \dots, a_n) \in \mathrm{Lie}(\mathrm{T})$ and all $i \in [n]$

$$\pi(e^A)e_i = \mathrm{diag}(e^{a_1}, \dots, e^{a_n})e_i = e^{a_i}e_i \stackrel{(*)}{=} e^{\mathrm{tr}(A \mathrm{diag}(\varepsilon_i))}e_i$$

where we used $a_1 + \dots + a_n = 0$ in $(*)$. Thus, $\varepsilon_i \in \mathbb{1}_n^\perp \cong i \mathrm{Lie}(\mathrm{T}_K)$ is a weight of π with weight vector e_i . Since $\mathbb{C}^n = \bigoplus_i \mathbb{C}e_i$, we deduce $\Omega(\pi) = \{\varepsilon_i \mid i \in [n]\}$.

► **Example B.3.** Of particular importance in representation theory is the adjoint representation. That is, $G = \mathrm{SL}(n)^d$ acts on its Lie algebra by conjugation $\mathrm{Ad}: G \rightarrow \mathrm{GL}(\mathrm{Lie}(G)), g \mapsto (A \mapsto gAg^{-1})$, which induces the representation of Lie algebras $\mathrm{ad}: \mathrm{Lie}(G) \mapsto \mathrm{End}(\mathrm{Lie}(G)), A \mapsto (B \mapsto [A, B])$. The non-zero weights $\alpha \in \Omega(\mathrm{Ad})$ are called roots of G and the weight spaces $\mathrm{Lie}(G)_\alpha$ are called root spaces.

Let $d = 1$ and for $i, j \in [n]$ denote by $E_{i,j}$ the matrix with entry one at position i, j and all other entries being zero. Then for $i, j \in [n]$ with $i \neq j$ and for all $A = \mathrm{diag}(a_1, \dots, a_n), B \in \mathrm{Lie}(\mathrm{T})$ we compute

$$\begin{aligned} \mathrm{ad}(A)E_{i,j} &= [A, E_{i,j}] = (a_i - a_j)E_{i,j} = \mathrm{tr}(A \mathrm{diag}(e_i - e_j))E_{i,j}, \\ \mathrm{ad}(A)(B) &= [A, B] = 0. \end{aligned}$$

Since $0_n, e_i - e_j \in \mathbb{1}_n^\perp \cong i \mathrm{Lie}(\mathrm{T}_K)$, we deduce $e_i - e_j \in \Omega(\mathrm{Ad})$ with weight vector $E_{i,j}$ and $0_n \in \Omega(\mathrm{Ad})$ with weight vector $B \in \mathrm{Lie}(\mathrm{T})$. Therefore, the set of roots of $\mathrm{SL}(n)$ is $\{e_i - e_j \mid i, j \in [n], i \neq j\}$, because $\mathrm{Lie}(G) = \mathrm{Lie}(\mathrm{T}) \oplus \bigoplus_{i \neq j} \mathbb{C}E_{i,j}$.

More generally, one can deduce that the roots of $G = \mathrm{SL}(n)^d$ are the

$$(e_i - e_j, 0_n, \dots, 0_n), (0_n, e_i - e_j, 0_n, \dots, 0_n), \dots, (0_n, \dots, 0_n, e_i - e_j) \in (\mathbb{R}^n)^d$$

for $i, j \in [n]$ with $i \neq j$ and that $\mathrm{Lie}(G) = \mathrm{Lie}(\mathrm{T}) \oplus \bigoplus_\alpha \mathrm{Lie}(G)_\alpha$.

We need the following property of roots, see e.g. [31, Lemma 7.11].

► **Proposition B.4.** Let α be a root of $G = \mathrm{SL}(n)^d$ and let $\pi: G \rightarrow \mathrm{GL}(V)$ be a rational representation of G . If V_ω is the weight space of some weight $\omega \in \Omega(\pi)$, then

$$\Pi(\mathrm{Lie}(G)_\alpha)(V_\omega) \subseteq V_{\omega+\alpha},$$

where $V_{\omega+\alpha} := \{0\}$, if $\omega + \alpha \notin \Omega(\pi)$.

C Padding for tensor margin and tensor gap

The Theorems 2.1 and 4.11 only give for all $n \geq 2$ bounds for certain sub-families of $\{(n, d) \mid d \geq 3\}$. Still, we can deduce Theorems 1.3 and 1.6 via some padding on the number of tensor factors d ; that padding is provided in Proposition C.1 below. Recall the representation for tensor scaling

$$\pi_{n,d}: \mathrm{SL}(n)^d \rightarrow \mathrm{GL}((\mathbb{C}^n)^{\otimes d}), (g_1, \dots, g_d) \mapsto g_1 \otimes \dots \otimes g_d,$$

which set of weights is $\Omega(\pi_{n,d}) = \Omega_{n,d} = \{\varepsilon_i \mid i \in [n]\}^d \subseteq (\mathbb{R}^n)^d$.

► **Proposition C.1.** *Let $G := \text{SL}(n)^d$ and $n, d \geq 1$. Consider a set of weights $\Gamma_{n,d} \subseteq \Omega_{n,d}$ such that $0 \notin \text{conv}(\Gamma_{n,d})$, i.e. $\Gamma_{n,d}$ witnesses the inequality $\gamma(\Omega_{n,d}) = \gamma_{\text{T}}(\pi_{n,d}) \leq \text{dist}(0, \text{conv}(\Gamma_{n,d}))$.*

1. *Then $\gamma(\Omega_{n,d+1}) \leq \text{dist}(0, \text{conv}(\Gamma_{n,d}))$. Consequently, $\gamma(\Omega_{n,d+1}) \leq \gamma(\Omega_{n,d})$.*
2. *If additionally $\Gamma_{n,d}$ is free, then $\gamma_G(\pi_{n,d+r}) \leq \text{dist}(0, \text{conv}(\Gamma_{n,d}))$ for all $r \geq 2$.*

Proof. To prove the statement we set for $r \geq 1$

$$\Delta_r := \{(\varepsilon_i, \dots, \varepsilon_i) \mid i \in [n]\} \subseteq (\mathbb{R}^n)^r \quad \text{and} \quad \Gamma_{n,d+r} := \Gamma_{n,d} \times \Delta_r \subseteq \Omega(\pi_{n,d+r}).$$

By Equation (5) we have $0 \in \text{conv}(\Delta_r)$ and therefore

$$\text{conv}(\Gamma_{n,d+r}) = \text{conv}(\Gamma_{n,d}) \times \text{conv}(\Delta_r) \supseteq \text{conv}(\Gamma_{n,d}) \times \{0\}.$$

The latter implies

$$\text{dist}(0, \text{conv}(\Gamma_{n,d+r})) \leq \text{dist}(0, \text{conv}(\Gamma_{n,d})). \quad (34)$$

Clearly, $0 \in \text{conv}(\Gamma_{n,d+r})$ implies $0 \in \text{conv}(\Gamma_{n,d})$ or, by contraposition, the assumption $0 \notin \text{conv}(\Gamma_{n,d})$ yields $0 \notin \text{conv}(\Gamma_{n,d+r})$. The latter for $r = 1$ shows $\gamma_{\text{T}}(\pi_{n,d+1}) \leq \text{dist}(0, \text{conv}(\Gamma_{n,d+1}))$ and we conclude the first assertion with Equation (34).

Assume in addition that $\Gamma_{n,d}$ is free and let $r \geq 2$. Considering Definition 4.12 and Proposition 4.13 we prove that also $\Gamma_{n,d+r}$ is free. For this, let $M \subseteq [n]^d$ be such that $\Gamma_M = \Gamma_{n,d}$ and consider $(x, i, \dots, i), (y, j, \dots, j) \in M \times [n]^r$ with $(x, i, \dots, i) \neq (y, j, \dots, j)$. If $x \neq y$, then x and y differ in at least two components by freeness of M . If $x = y$, then we have $i \neq j$ and so (x, i, \dots, i) and (y, j, \dots, j) differ in at least two components, using $r \geq 2$. This shows that $\Gamma_{n,d+r}$ is free for $r \geq 2$. Since also $0 \notin \text{conv}(\Gamma_{n,d+r})$ we obtain with Proposition 4.8 that $\gamma_G(\pi_{n,d+r}) \leq \text{dist}(0, \text{conv}(\Gamma_{n,d+r}))$ holds for all $r \geq 2$. Finally, we deduce the second statement using Equation (34). ◀

► **Proposition C.2.** *For $n \geq 3$ it holds that $\gamma_{\text{T}}(\pi_{n,4}) \leq \gamma_G(\pi_{n,4}) \leq 2^{-n+1}$.*

Proof. This result can be obtained by imitating the proof of Theorem 2.1(b) in subsection 2.2 by using

$$\Gamma_{n,4} := \{(\varepsilon_i, \varepsilon_j, \varepsilon_k, \varepsilon_i) \mid (i, j, k) \in \mathfrak{W}_n\} \subseteq \Omega(\pi_{n,4}).$$

Clearly, $0 \notin \text{conv}(\Gamma_{n,4})$ as $0 \notin \text{conv}(\Gamma_{n,3})$ by Lemma 2.8. Moreover, one can show with Lemma 2.5 (similar to the proof of Lemma 2.7) that

$$x := -\frac{1}{c 2^{n-1}}(\varepsilon_1, \varepsilon_1, \varepsilon_1, \varepsilon_1) \in \text{conv}(\Gamma_{n,4}), \quad \text{where} \quad c = n - 2^{-n+1} \geq 2.$$

Thus, $\|(\varepsilon_1, \varepsilon_1, \varepsilon_1, \varepsilon_1)\| \leq \sqrt{4}$ implies $\|x\| \leq c^{-1} 2^{-n+1} \sqrt{4} \leq 2^{-n+1}$. This proves $\gamma_{\text{T}}(\pi_{n,4}) \leq 2^{-n+1}$.

Since \mathfrak{W}_n is free by Proposition 4.15, the set $\{(i, j, k, i) \mid (i, j, k) \in \mathfrak{W}_n\}$ is free. Hence, we conclude $\gamma_G(\pi_{n,4}) \leq 2^{-n+1}$ with Proposition 4.13 and Proposition 4.8. ◀

D Proof of Lemma 2.11

Proof. For the sake of contradiction assume that $0 \in \text{Aff}(\Gamma_{n,6r-3})$. Then there are coefficients $a_s, b_s, c_s \in \mathbb{R}$, where $2 \leq s \leq rn$, such that $a_2 = \dots = a_r = b_2 = \dots = b_r = 0$, $\sum_s (a_s + b_s + c_s) = 1$ and

$$\sum_{s=2}^{rn} (a_s \varepsilon_{\sigma(s), \sigma(1), \sigma(s)} + b_s \varepsilon_{\sigma(s), \sigma(s), \sigma(1)} + c_s \varepsilon_{\sigma(s-1), \sigma(s), \sigma(s)}) = 0 \in (\mathbb{R}^n)^{6r-3}. \quad (35)$$

The bulk of our work will consist of proving the equations

$$b_2 + c_2 = b_3 + c_3 = \dots = b_{rn} + c_{rn} \tag{36}$$

$$a_2 + c_2 = a_3 + c_3 = \dots = a_{rn} + c_{rn}. \tag{37}$$

From here we will derive a contradiction. We now set about proving Equations (36) and (37). Rewrite the left-hand-side of Equation (35) as the collection for $k \in [2r - 1]$ of the following affine linear combinations of $\varepsilon_1, \dots, \varepsilon_n$ in \mathbb{R}^n :

$$\sum_{s=2}^{rn} (a_s \varepsilon_{\sigma_k(s)} + b_s \varepsilon_{\sigma_k(s)} + c_s \varepsilon_{\sigma_k(s-1)}) = 0 \tag{38}$$

$$\sum_{s=2}^{rn} (a_s \varepsilon_{\sigma_k(1)} + b_s \varepsilon_{\sigma_k(s)} + c_s \varepsilon_{\sigma_k(s)}) = 0 \tag{39}$$

$$\sum_{s=2}^{rn} (a_s \varepsilon_{\sigma_k(s)} + b_s \varepsilon_{\sigma_k(1)} + c_s \varepsilon_{\sigma_k(s)}) = 0. \tag{40}$$

If we expand this expressions as affine linear combinations of the ε_l , then by Lemma 2.2 the coefficient of ε_l must be n^{-1} for all $l \in [n]$. Translating this for equations (38), (39) and (40) respectively with $2 \leq l \leq n$ and $k \in [r]$, and using for $j \in [r]$ that

$$\sigma_k(r(l-1) + j - k + 1) = \left\lceil \frac{(r(l-1) + j - k + 1) + (k-1)}{r} \right\rceil = l \tag{41}$$

we get

$$\forall k \in [r], l \in \{2, 3, \dots, n\}: \quad \sum_{j=1}^r (a_{r(l-1)+j-k+1} + b_{r(l-1)+j-k+1} + c_{r(l-1)+j-k+2}) = \frac{1}{n} \tag{42}$$

$$\forall k \in [r], l \in \{2, 3, \dots, n\}: \quad \sum_{j=1}^r (b_{r(l-1)+j-k+1} + c_{r(l-1)+j-k+1}) = \frac{1}{n} \tag{43}$$

$$\forall k \in [r], l \in \{2, 3, \dots, n\}: \quad \sum_{j=1}^r (a_{r(l-1)+j-k+1} + c_{r(l-1)+j-k+1}) = \frac{1}{n} \tag{44}$$

respectively, where we set $c_{rn+1} := 0$. Fixing some $l \geq 2$ and subtracting Equation (43) with $k = 1$ from Equation (43) for $k = 2$, we find a telescoping sum that reduces to $b_{r(l-1)} + c_{r(l-1)} = b_{rl} + c_{rl}$. Indeed, subtracting the two yields

$$\begin{aligned} 0 &= \sum_{j=1}^r (b_{r(l-1)+j-1} + c_{r(l-1)+j-1}) - \sum_{j=1}^r (b_{r(l-1)+j} + c_{r(l-1)+j}) \\ &= \sum_{j=0}^{r-1} (b_{r(l-1)+j} + c_{r(l-1)+j}) - \sum_{j=1}^r (b_{r(l-1)+j} + c_{r(l-1)+j}) \\ &= (b_{r(l-1)} + c_{r(l-1)}) - (b_{rl} + c_{rl}). \end{aligned}$$

More generally, for $k \in [r-1]$ combining (43) for k and $k \leftarrow k+1$, implies $b_{rl-k+1} + c_{rl-k+1} = b_{r(l-1)-k+1} + c_{r(l-1)-k+1}$ for all $l = 2, \dots, n$, i.e. for every $k \in [r-1]$ we have

$$c_{r-k+1} = b_{r-k+1} + c_{r-k+1} = b_{2r-k+1} + c_{2r-k+1} = \dots = b_{rn-k+1} + c_{rn-k+1}. \tag{45}$$

13:52 Barriers in Geodesic Optimization

We are still missing the value $k = 0$, or the equations

$$b_{r+1} + c_{r+1} = b_{2r+1} + c_{2r+1} = \dots = b_{r(n-1)+1} + c_{r(n-1)+1}. \quad (46)$$

We obtain this by subtracting, for $l = 2, \dots, n$, (43) for $k = 1$ and l from (43) with $k = r$ and $l \leftarrow l + 1$. Indeed,

$$\begin{aligned} 0 &= \sum_{j=1}^r (b_{rl+j-r+1} + c_{rl+j-r+1}) - \sum_{j=1}^r (b_{r(l-1)+j} + c_{r(l-1)+j}) \\ &= \sum_{j=2}^{r+1} (b_{r(l-1)+j} + c_{r(l-1)+j}) - \sum_{j=1}^r (b_{r(l-1)+j} + c_{r(l-1)+j}) \\ &= (b_{rl+1} + c_{rl+1}) - (b_{r(l-1)+1} + c_{r(l-1)+1}). \end{aligned}$$

Lastly, we are missing the equations $b_2 + c_2 = b_3 + c_3 = \dots = b_{r+1} + c_{r+1}$ for Equation (36). We have not yet used in Equation (39) the values $k = r + m$ with $m \in [r - 1]$. For this we note that

$$\sigma_{r+m}(j) = 2 \quad \text{for } j \in \{r - m + 1\} \cup \{r + 2, r + 3, \dots, 2r\}.$$

We use this equation to apply Lemma 2.2 to (39) for ε_2 and $k = r + m$ with $m \in [r - 1]$ to obtain

$$b_{r-m+1} + c_{r-m+1} + \sum_{j=2}^r (b_{r+j} + c_{r+j}) = \frac{1}{n}.$$

We need one more equation to eliminate the right-hand term, so we use the following. Lemma 2.2 applied to equation (43) for $k = 1$ and $l = 2$ yields

$$\sum_{j=1}^r (b_{r+j} + c_{r+j}) = \frac{1}{n}.$$

Subtracting this equation from the previous one yields, $b_{r-m+1} + c_{r-m+1} = b_{r+1} + c_{r+1}$ for all $m = 1, \dots, r - 1$. Together with the equations (45) and (46) we conclude Equation (36). Analogously, (40) and (44) can be used to obtain Equation (37).

To get a contradiction we show that $a_s = b_s = c_s = 0$ for all $s = 2, 3, \dots, rn$. For this, we set $a := \sum_s a_s$ and $b := \sum_s b_s$. Equation (41) still applies for $l = 1, k = 1$, so Lemma 2.2 applied to the coefficient of ε_1 in (38), in (39) and in (40) respectively for $k = 1$ gives

$$\sum_{j=1}^r c_{j+1} = \frac{1}{n}, \quad a + \sum_{j=1}^{r-1} c_{j+1} = \frac{1}{n} \quad \text{and} \quad b + \sum_{j=1}^{r-1} c_{j+1} = \frac{1}{n}$$

respectively. Subtracting the second equation from the first gives $a = c_{r+1}$, and reasoning analogously for the third yields $a = b = c_{r+1}$. Moreover, (43) with $k = r$ and $l = 2$ is $\sum_{j=1}^r (b_{j+1} + c_{j+1}) = n^{-1}$. Using the latter together with $b_2 = \dots = b_r = 0$ and $\sum_{j=1}^r c_{j+1} = n^{-1}$ yields $b_{r+1} = 0$ and similarly $a_{r+1} = 0$ via (44) with $k = r$ and $l = 2$.

Since now also $a_{r+1} = b_{r+1} = 0$, the equation (42) with $k = r$ and $l = 2$ simplifies to $\sum_{j=1}^r c_{j+2} = n^{-1}$. In conjunction with $\sum_{j=1}^r c_{j+1} = n^{-1}$ we deduce $c_2 = c_{r+2}$ and hence $b_{r+2} = 0 = a_{r+2}$ by (36) and (37). But now (42) with $k = r - 1$ and $l = 2$ is $\sum_{j=1}^r c_{j+3} = n^{-1}$ and together with $\sum_{j=1}^r c_{j+2} = n^{-1}$ we get $c_3 = c_{r+3}$. Continuing inductively we obtain

$$\forall j \in [r]: \quad c_{j+1} = c_{r+j+1} \quad \text{and} \quad a_{r+j+1} = b_{r+j+1} = 0$$

via (42) with $l = 2$, $k \in [r]$ and via (36), (37). Then (42) with $k = r$ and $l = 3$ simplifies to $\sum_{j=1}^r c_{r+j+2} = n^{-1}$ and together with $n^{-1} = \sum_{j=1}^r c_{j+1} = \sum_{j=1}^r c_{r+j+1}$ we have $c_{r+2} = c_{2r+2}$. Hence, $b_{2r+2} = 0 = a_{2r+2}$ via (36) respectively (37). Continuing inductively in the outlined manner with equation (42) for $k \in [r]$, $l = 3, \dots, n$ and with the equations (36) and (37) we conclude $a_s = b_s = 0$ for all $s = 2, 3, \dots, rn$, so $a = b = 0$. Finally, (36) implies $c_{r+1} = c_s$ for all $s = 2, \dots, rn$, but $c_{r+1} = b = 0$ giving the desired contradiction. ◀

E Padding and rounding for diameter bounds

We begin with the proof of Proposition 3.5. We prove it only for $d = 3$, but the proof goes through mutatis mutandis for all $d \geq 1$.

Proof of Proposition 3.5. Recall that q is the $n \times n \times n$ array such that $q_{ijk} = \frac{t}{n} p_{ijk}$ for $i, j, k \in [t]$, $q_{iii} = 1/n$ for $t + 1 \leq i \leq n$, and $q_{ijk} = 0$ otherwise. We may split the inputs $x, y, z \in \mathbb{1}_n^\perp$ into

$$\begin{aligned} x &= \left(x' + \alpha_1 \mathbb{1}_t, x'' - \frac{t}{n-t} \alpha_1 \mathbb{1}_{n-t} \right), \\ y &= \left(y' + \alpha_2 \mathbb{1}_t, y'' - \frac{t}{n-t} \alpha_2 \mathbb{1}_{n-t} \right), \\ z &= \left(z' + \alpha_3 \mathbb{1}_t, z'' - \frac{t}{n-t} \alpha_3 \mathbb{1}_{n-t} \right) \end{aligned}$$

where $x', y', z' \in \mathbb{R}^t$, $x'', y'', z'' \in \mathbb{R}^{n-t}$ each sum to zero; write $w = (x', y', z')$. As $\|(x, y, z)\|_2 \geq \|w\|_2$, it is enough to prove that $\|w\|_2$ is large for any approximate minimizer. By optimizing over α_i and x'', y'', z'' for fixed w , one computes that the optimum value for f_q for any fixed w is $f_p(w)^{t/n}$. To see this, write

$$f_q(x, y, z) = \frac{te^{\alpha_1 + \alpha_2 + \alpha_3}}{n} f_p(w) + \frac{e^{-\frac{t}{n-t}(\alpha_1 + \alpha_2 + \alpha_3)}}{n} \sum_{i=t+1}^n e^{x''_i + y''_i + z''_i}.$$

First note that for fixed α_i 's, the second term is minimized at $x'' = y'' = z'' = 0$ by Jensen's inequality. Furthermore, the value only depends on $\alpha := \alpha_1 + \alpha_2 + \alpha_3$. With $x'', y'', z'' = 0$, we have

$$f_q(x, y, z) = g(w, \alpha) := \frac{te^\alpha}{n} f_p(w) + \frac{(n-t)}{n} e^{-\frac{t}{n-t}\alpha}.$$

Taking the derivative in α , we see that this is minimized when $f_p(w)e^\alpha = e^{-\frac{t}{n-t}\alpha}$, or $e^\alpha = f_p(w)^{-1/(1+\frac{t}{n-t})} = f_p(w)^{-\frac{n-t}{n}}$. Plugging this value in proves that the optimum is $f_p(w)^{t/n}$. By concavity of $x^{t/n}$, provided $f_p(w) \leq 1$ we have

$$f_p(w)^{t/n} - \text{cap}(p)^{t/n} \geq \frac{1 - \text{cap}(p)^{t/n}}{1 - \text{cap}(p)} (f_p(w) - \text{cap}(p)).$$

The first factor in the second term is the slope of the line from $(\text{cap}(p), \text{cap}(p)^{t/n})$ to $(1, 1)$. Thus for any $\varepsilon \leq 1 - \text{cap}(p)$, any ε -approximate minimizer for f_q has norm at least that of some $(\frac{1 - \text{cap}(p)}{1 - \text{cap}(p)^{t/n}}) \varepsilon$ -approximate minimizer for f_p . ◀

Proof of Lemma 4.23. We use the dual expression: $\log \text{cap } q = -\inf_{\mathbb{E}_r, \omega=0} D_{KL}(r||q)$ where r ranges over probability distributions on Ω . In particular,

$$\log \text{cap } q \geq -D_{KL}(r||q)$$

13:54 Barriers in Geodesic Optimization

for any distribution r on Ω with $\mathbb{E}_r \omega = 0$. Let r be a probability distribution; calculate

$$\begin{aligned} \log \text{cap } q &\geq -D_{KL}(r||q) = -D_{KL}(r||p) + D_{KL}(r||p) - D_{KL}(r||q) \\ &= -D_{KL}(r||p) + \sum_{\omega \in \Omega} r_\omega \log(r_\omega/p_\omega) - \sum_{\omega \in \Omega} r_\omega \log(r_\omega/q_\omega) \\ &= -D_{KL}(r||p) + \sum_{\omega \in \Omega} r_\omega (\log q_\omega - \log p_\omega). \end{aligned}$$

We lower bound $\log q_\omega - \log p_\omega \geq \frac{1}{q_\omega}(q_\omega - p_\omega)$ by applying the inequality $\log x \leq x - 1$ to $x = p_\omega/q_\omega$. Hence

$$\begin{aligned} \log \text{cap } q &\geq -D_{KL}(r||p) + \sum_{\omega \in \Omega} r_\omega \frac{1}{q_\omega}(q_\omega - p_\omega) \\ &\geq -D_{KL}(r||p) - M\|p - q\|_\infty. \end{aligned}$$

Allowing $-D_{KL}(r||p)$ to tend to $\log \text{cap } p$ completes the proof. \blacktriangleleft

Proof of Lemma 4.22. Applying Lemma 4.23 with the roles of p and q switched yields

$$\log \text{cap } p \geq \log \text{cap } q - M\|p - q\|_\infty.$$

Exponentiating both sides and applying the inequality $e^x \geq 1 + x$ yields $\text{cap } p \geq (1 - M\|p - q\|_\infty) \text{cap } q$. Thus

$$\inf_{x \in B} f_q(x) = \inf_{x \in S} f_q(x) \geq -\sup_{x \in S} |f_q(x) - f_p(x)| + \inf_{x \in S} f_p(x).$$

Note that the minimizer for f_q over B lies in the set $S := B \cap \{x : \forall \omega, q_\omega e^{x \cdot \omega} \leq f_q(0) = \|q\|_1\}$. For all $x \in S$, we have $e^{x \cdot \omega} \leq \text{cap } q/p_\omega$ for all $\omega \in \Omega$, so

$$\begin{aligned} f_q(x) - f_p(x) &\leq \sum_{\omega \in \Omega} |p_\omega - q_\omega| e^{x \cdot \omega} \\ &\leq \sum_{\omega \in \Omega} |p_\omega - q_\omega| \|q\|_1 / q_\omega \\ &\leq \|p - q\|_1 M \|q\|_1. \end{aligned}$$

Combining the above inequality with the lower bound for $\text{cap}(p)$,

$$\begin{aligned} \inf_{x \in B} f_q(x) &\geq -M\|q\|_1\|p - q\|_1 + (1 + \varepsilon) \text{cap } p \\ &\geq (1 + \varepsilon)(1 - M\|p - q\|_\infty) \text{cap } q - M\|p - q\|_1\|q\|_1. \end{aligned} \quad \blacktriangleleft$$