

Quantum Query Algorithms are Completely Bounded Forms

Srinivasan Arunachalam*

Jop Briët†

Carlos Palazuelos‡

Abstract

We prove a characterization of t -query quantum algorithms in terms of the unit ball of a space of degree- $2t$ polynomials. Based on this, we obtain a refined notion of approximate polynomial degree that equals the quantum query complexity, answering a question of Aaronson et al. (CCC'16). Our proof is based on a fundamental result of Christensen and Sinclair (*J. Funct. Anal.*, 1987) that generalizes the well-known Stinespring representation for quantum channels to multilinear forms. Using our characterization, we show that many polynomials of degree four are far from those coming from two-query quantum algorithms. We also give a simple and short proof of one of the results of Aaronson et al. showing an equivalence between one-query quantum algorithms and bounded quadratic polynomials.

1 Introduction

In the black-box model of quantum computation one is given access to a unitary operation, usually referred to as an oracle, that allows one to probe the bits of an unknown binary string $x \in \{-1, 1\}^n$ in superposition. Promised that x lies in a subset $D \subseteq \{-1, 1\}^n$, the goal in this model is to learn some property of x given by a Boolean function $f : D \rightarrow \{-1, 1\}$, when only given access to x through the oracle. An application of the oracle is usually referred to as a *query*. The bounded-error quantum query complexity of f , denoted $Q_\varepsilon(f)$, is the minimal number of queries a quantum algorithm must make on the worst-case input $x \in D$ to compute $f(x)$ with probability at least $1 - \varepsilon$, where $\varepsilon \in (0, 1/2)$ is usually some fixed but arbitrary positive constant.

Many of the best-known quantum algorithms are naturally captured by this model. Famous partial functions whose quantum query complexity is exponentially smaller than their classical counterpart (the decision-tree complexity) are period finding [Sho97], Simon's problem [Sim97] and Forrelation [AA15]. Famous problems related to total functions that admit polynomial quantum speed-ups include unstructured search [Gro96], element distinctness [Amb07] and NAND-tree evaluation [FGG08]. It is well-known that for all total functions, the quantum and classical query complexities are polynomially related [BBC⁺01]; see Ambainis et al. [ABB⁺16] and Aaronson et al. [ABK16] for recent progress on the largest possible separations.

*QuSoft, CWI and University of Amsterdam, the Netherlands. Supported by ERC Consolidator Grant QPROGRESS. E-mail: arunacha@cwi.nl

†CWI, QuSoft. Supported by a VENI grant and the Gravitation-grant NETWORKS-024.002.003 from the Netherlands Organisation for Scientific Research (NWO). E-mail: j.briet@cwi.nl

‡Facultad de C.C. Matemáticas, UCM. Instituto de Ciencias Matemáticas, Madrid Spain. Supported by the Ramon y Cajal program (RYC-2012-10449), the Spanish MINECO MTM2014-54240-P, Comunidad de Madrid (QUITEMAD+ Project S2013/ICE-2801) and ICMAT Severo Ochoa Grant No. SEV-2015-0554. E-mail: carlospalazuelos@mat.ucm.es

Despite the simplicity of the query model, determining the quantum query complexity of a given function f appears to be highly non-trivial. Several methods were introduced to tackle this problem. For constructing quantum query algorithms, there are general methods based on quantum walks [Amb07, MNRS11], span programs [Rei09] and learning graphs [Bel12]. For proving lower bounds there are two main methods, known as the *polynomial method* [BBC⁺01] and the *adversary method* [Amb02]. The latter was eventually generalized to the “negative weight” adversary method [HLŠ07] and was shown to *characterize* quantum query complexity [HLŠ07, Rei09, Rei11, LMR⁺11], but proving lower bounds using this method appears to be hard in general. This paper will focus on the polynomial method.

1.1 The polynomial method

The polynomial method is based on a connection between quantum query algorithms and polynomials discovered by Beals et al. [BBC⁺01]. They observed that for every t -query quantum algorithm \mathcal{A} that on input $x \in \{-1, 1\}^n$ returns a random sign $\mathcal{A}(x)$, there exists a degree- $(2t)$ polynomial p such that $p(x) = \mathbb{E}[\mathcal{A}(x)]$ for every x (where the expectation is taken over the randomness of the output). It follows that if \mathcal{A} computes $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ with probability at least $1 - \varepsilon$, then p satisfies $|p(x) - f(x)| \leq 2\varepsilon$ for every x . The polynomial method thus converts the problem of lower bounding quantum query complexity to the problem of proving lower bounds on the minimum degree of a polynomial p such that $|p(x) - f(x)| \leq 2\varepsilon$ holds for all inputs x . The minimal degree of such a polynomial is called the *approximate (polynomial) degree* and is denoted by $\deg_\varepsilon(f)$. Notable applications of this approach showed optimality for Grover’s search algorithm [BBC⁺01]¹ and the above-mentioned algorithms for collision-finding and element distinctness [AS04]. In a recent work, Bun et al. [BKT17] use the polynomial method to resolve the quantum query complexity of several other well-studied Boolean functions.

Converses to the polynomial method. A natural question is whether the polynomial method admits a converse. If so, this would imply a succinct characterization of quantum algorithms in terms of basic mathematical objects. However, Ambainis [Amb06] answered this question in the negative, showing that for infinitely many n , there is a function f with $\deg_{1/3}(f) \leq n^\alpha$ and $Q_{1/3}(f) \geq n^\beta$ for some positive constants $\beta > \alpha$ (recently larger separations were obtained in [ABK16]). The approximate degree thus turns out to be an imprecise measure for quantum query complexity in general. These negative results leave open the following two possibilities:

1. There is a (simple) refinement of approximate polynomial degree that approximates $Q_\varepsilon(f)$ up to a constant factor.
2. Constant-degree polynomials characterize constant-query quantum algorithms.

These avenues were recently explored by Aaronson et al. [AA15, AAI⁺16]. The first work strengthened the polynomial method by observing that quantum algorithms give rise to polynomials with a so-called *block-multilinear* structure. Based on this observation, they introduced a refined degree measure, $\text{bm-deg}_\varepsilon(f)$ which lies between $\deg_\varepsilon(f)$ and $2Q_\varepsilon(f)$, prompting the immediate question of how well that approximates $Q_\varepsilon(f)$. The subsequent work showed, among other things, that

¹The first quantum lower bound for the search problem was proven by Bennett et al. [BBBV97] using the so-called hybrid method. Beals et al. [BBC⁺01] reproved their result using the polynomial method.

for infinitely many n , there is a function f with $\text{bm-deg}_{1/3}(f) = O(\sqrt{n})$ and $Q_{1/3}(f) = \Omega(n)$, thereby also ruling out the possibility that this degree measure validates possibility 1. The natural next question then asks if there is another refined notion of polynomial degree that approximates quantum query complexity [AAI⁺16, Open problem 3].

In the direction of the second avenue, [AAI⁺16] showed a surprising converse to the polynomial method for quadratic polynomials. Say that a polynomial $p \in \mathbb{R}[x_1, \dots, x_n]$ is *bounded* if it satisfies $p(x) \in [-1, 1]$ for all $x \in \{-1, 1\}^n$.

Theorem 1.1 (Aaronson et al.). *There exists an absolute constant $C \in (0, 1]$ such that the following holds. For every bounded quadratic polynomial p , there exists a one-query quantum algorithm that, on input $x \in \{-1, 1\}^n$, returns a random sign with expectation $Cp(x)$.*

This implies that possibility 2 holds true for quadratic polynomials. It also leads to the problem of finding a similar converse for higher-degree polynomials, asking for instance whether two-query quantum algorithms are equivalent to quartic polynomials [AAI⁺16, Open problem 1].

1.2 Our results

This paper addresses the above-mentioned two problems. Our first result is a new notion of polynomial degree that gives a tight characterization of quantum query complexity (Definition 1.4 and Corollary 1.5 below), giving an answer to [AAI⁺16, Open problem 3]. Using this characterization, we show that there is no generalization of Theorem 1.1 to higher-degree polynomials, in the sense that there is no absolute constant $C \in (0, 1]$ for which the analogous statement holds true. This gives a partial answer to [AAI⁺16, Open problem 1], ruling out a strong kind of equivalence. Finally, we give a simplified shorter proof of Theorem 1.1. Below we explain our results in more detail.

Quantum algorithms are completely bounded forms. For the rest of the discussion, all polynomials will be assumed to be bounded, real and $(2n)$ -variate if not specified otherwise. We refer to a homogeneous polynomial as a *form*. For $\alpha \in \{0, 1, 2, \dots\}^{2n}$ and $x \in \mathbb{R}^{2n}$, we write $|\alpha| = \alpha_1 + \dots + \alpha_{2n}$ and $x^\alpha = x_1^{\alpha_1} \dots x_{2n}^{\alpha_{2n}}$. Then, any form p of degree t can be written as

$$p(x) = \sum_{\alpha \in \{0, 1, \dots, t\}^{2n} : |\alpha| = t} c_\alpha x^\alpha, \quad (1)$$

where c_α are some real coefficients. Our new notion of polynomial degree is based on a characterization of quantum query algorithms in terms of forms satisfying a certain norm constraint. The norm we assign to a form as in (1) is given by a norm of the symmetric t -tensor $T_p \in \mathbb{R}^{2n \times \dots \times 2n}$ with (i_1, \dots, i_t) -coordinate

$$(T_p)_{i_1, \dots, i_t} = \frac{c_{e_{i_1} + \dots + e_{i_t}}}{|\{i_1, \dots, i_t\}|!}, \quad (2)$$

where e_i is the i th standard basis vector for \mathbb{R}^{2n} and $|\{i_1, \dots, i_t\}|$ denotes the number of distinct elements in the set $\{i_1, \dots, i_t\}$. Note that p can then also be written as

$$p(x) = \sum_{i_1, \dots, i_t=1}^{2n} (T_p)_{i_1, \dots, i_t} x_{i_1} \dots x_{i_t}. \quad (3)$$

The relevant norm of T_p is in turn given in terms of an infimum over decompositions of the form $T_p = \sum_{\sigma \in S_t} T^\sigma \circ \sigma$, where the sum is over permutations of $\{1, \dots, t\}$, each T^σ is a t -tensor, and $T^\sigma \circ \sigma$ is the permuted version of T^σ given by

$$(T^\sigma \circ \sigma)_{i_1, \dots, i_t} = T_{i_{\sigma(1)}, \dots, i_{\sigma(t)}}^\sigma.$$

Finally, the actual norm is based on the *completely bounded norm* of each of the T^σ . Given a t -tensor $T \in \mathbb{R}^{2^n \times \dots \times 2^n}$, its completely bounded norm $\|T\|_{\text{cb}}$ is given by the supremum over positive integers k and collections of $k \times k$ unitary matrices $U_1(i), \dots, U_t(i)$, for $i \in [2n]$, of the operator norm

$$\left\| \sum_{i_1, \dots, i_t=1}^{2n} T_{i_1, \dots, i_t} U_1(i_1) \cdots U_t(i_t) \right\|. \quad (4)$$

Definition 1.2 (Completely bounded norm of a form). *Let p be a form of degree t and let T_p be the symmetric t -tensor as in (2). Then, the completely bounded norm of p is defined by*

$$\|p\|_{\text{cb}} = \inf \left\{ \sum_{\sigma \in S_t} \|T^\sigma\|_{\text{cb}} : T_p = \sum_{\sigma \in S_t} T^\sigma \circ \sigma \right\}. \quad (5)$$

This norm was originally introduced in the general context of tensor products of operator spaces in [OP99]. In that framework, the definition considered here corresponds to a particular operator space based on ℓ_1^n , but we shall not use this fact here. Our characterization of quantum query algorithms is as follows.

Theorem 1.3 (Characterization of quantum algorithms). *Let $\beta : \{-1, 1\}^n \rightarrow [-1, 1]$ and let t be a positive integer. Then, the following are equivalent.*

1. *There exists a form p of degree $2t$ such that $\|p\|_{\text{cb}} \leq 1$ and $p((x, \mathbf{1})) = \beta(x)$ for every $x \in \{-1, 1\}^n$, where $\mathbf{1} \in \mathbb{R}^n$ is the all-ones vector.*
2. *There exists a t -query quantum algorithm that, on input $x \in \{-1, 1\}^n$, returns a random sign with expected value $\beta(x)$.*

It may be observed that the content of the polynomial method is contained in the above statement, since any $(2n)$ -variate form p defines an n -variate polynomial given by $q(x) = p((x, \mathbf{1}))$. The above theorem refines the polynomial method in the sense that quantum algorithms can only yield polynomials of the form $q(x) = p((x, \mathbf{1}))$ where p has completely bounded norm at most one. Our proof is based on a fundamental result of Christensen and Sinclair [CS87] concerning multilinear forms on C^* -algebras that generalizes the well-known Stinespring representation theorem for quantum channels (see also [PS87] and [Pis03, Chapter 5]).

Completely bounded approximate degree. Theorem 1.3 motivates the following new notion of approximate degree for partial Boolean functions.

Definition 1.4 (Completely bounded approximate degree). *For $D \subseteq \{-1, 1\}^n$, let $f : D \rightarrow \{-1, 1\}$ be a (possibly partial) Boolean function and let $\varepsilon > 0$. Then, the ε -completely bounded approximate degree of f , denoted $\text{cb-deg}_\varepsilon(f)$, is the smallest positive integer t for which there exists a form p of degree $2t$ such that $\|p\|_{\text{cb}} \leq 1$ as in Eq. (5) and we have $|p((x, \mathbf{1})) - f(x)| \leq 2\varepsilon$ for every $x \in D$.*

As a corollary of Theorem 1.3, we get the following characterization of quantum query complexity.

Corollary 1.5. *For every $D \subseteq \{-1, 1\}^n$, $f : D \rightarrow \{-1, 1\}$ and $\varepsilon > 0$, we have $\text{cb-deg}_\varepsilon(f) = Q_\varepsilon(f)$.*

Separations for higher-degree forms. Theorem 1.1 follows from our Theorem 1.3 and the fact that for every bounded quadratic form $p(x) = x^T A x$, the matrix A has completely bounded norm bounded from above by an absolute constant (independent on n); this is discussed in more detail below. If the same were true for the tensors T_p corresponding to higher-degree forms p then Theorem 1.3 would give higher-degree extensions of Theorem 1.1. Unfortunately, this is false. Bounded forms whose associated tensors have unbounded completely bounded norm appeared before in the work of Smith [Smi88], who gave an explicit example with completely bounded norm $\sqrt{\log n}$. Since $\|p\|_{\text{cb}}$ involves an infimum over decompositions of T_p , this does not yet imply a counterexample to higher-degree versions of Theorem 1.1. However, such counterexamples are implied by recent work on Bell inequalities, multiplayer XOR games in particular. It is not difficult to see that $\|p\|_{\text{cb}}$ is bounded from below by the so-called *jointly completely bounded norm* of the tensor T_p , a quantity that in quantum information theory is better known as the entangled bias of the XOR game whose (unnormalized) game tensor is given by T_p . One obtains this quantity by inserting tensor products between the unitaries appearing in (4). Pérez-García et al. [PGWP⁺08] and Vidick and the second author [BV13] gave examples of bounded cubic forms with unbounded jointly completely bounded norm. Both constructions are non-explicit, the first giving a completely bounded norm of order $\Omega((\log n)^{1/4})$ and the latter of order $\tilde{\Omega}(n^{1/4})$. Here, we explain how to get a larger separation by means of a much simpler (although still non-explicit) construction and show that a bounded cubic form p given by a suitably normalized random sign tensor has completely bounded norm $\|p\|_{\text{cb}} = \Omega(\sqrt{n})$ with high probability (Theorem 4.1). The result presented here is not new, but it follows from the existence of commutative operator algebras which are not \mathcal{Q} -algebras. Here, we present a self-contained proof which follows the same lines as in [DJT95, Theorem 18.16] and, in addition, we prove the result with high probability (rather than just the existence of such trilinear forms). We also explain how to obtain from this result quartic examples by embedding into 3-dimensional “tensor slices”, which in turn imply counterexamples to a quartic versus two-query version of Theorem 1.1.

Short proof of Theorem 1.1. As shown in [AAI⁺16], Theorem 1.1 is yet another surprising consequence of the ubiquitous Grothendieck inequality [Gro53] (Theorem 5.2 below), well known for its relevance to Bell inequalities [Tsi87, CHTW04] and combinatorial optimization [AN06, KN12], not to mention its fundamental importance to Banach spaces [Pis12]. An equivalent formulation of Grothendieck’s inequality again recovers Theorem 1.1 for quadratic forms $p(x) = x^T A x$ given by a matrix $A \in \mathbb{R}^{n \times n}$ satisfying a certain norm constraint $\|A\|_{\ell_\infty \rightarrow \ell_1} \leq 1$, which in particular implies that p is bounded (see Section 2 for more on this norm). Indeed, in that case Grothendieck’s inequality implies that $\|A\|_{\text{cb}} \leq K_G$ for some absolute constant $K_G \in (0, \infty)$ (independent of n and A). Normalizing by K_G^{-1} , one obtains Theorem 1.1 with $C = K_G^{-1}$ for such quadratic forms from Theorem 1.3. The general version of Theorem 1.1 for quadratic polynomials follows from this via a so-called decoupling argument (see Section 5). This arguably does not simplify the original proof of Theorem 1.1, as Theorem 1.3 relies on deep results itself. However, in Section 5 we give a short simplified proof, showing that Theorem 1.1 follows almost directly from a “factorization version” of Grothendieck’s inequality (Theorem 5.3) that follows from the more standard version (Theorem 5.2). The factorization version was used in the original proof as well, but only as a lemma in a more intricate argument. In computer science, this factorization version already found applications in an algorithmic version of the Bourgain–Tzafriri Column Subset Theorem [Tro09] and algorithms for community detection in the stochastic block model [LLV15]. This appears to

be its first occurrence in quantum computing.

1.3 Related work

Although there is no converse to the polynomial method for arbitrary polynomials, equivalences between quantum algorithms and polynomials have been studied before in certain models of computation. For example, we do know of such characterization in the model of non-deterministic query complexity [Wol03], the unbounded-error query complexity [MNR11] and quantum query complexity in expectation [KLW15]. We remark here that in all these settings, the quantum algorithms constructed from polynomials were *non-adaptive* algorithms, i.e., the quantum algorithm begins with a quantum state, repeatedly applies the oracle some fixed number of times and then performs a projective measurement. Crucially, these algorithms do not contain interlacing unitaries that are present in the standard model of query complexity, hence are known to be a much weaker class of algorithms (see Montanaro [Mon10] for more details).

Our main result is yet another demonstration of the expressive power of C^* -algebras and operator space theory in quantum information theory; for a survey on applications of these areas to two-prover one-round games, see [PV16]. The appearance of Q -algebras (mentioned in the above paragraph on separations) is also not a first in quantum information theory, see for instance [PGWP⁺08, BBLV12, BBLV13].

1.4 Organization

In Section 2, we give a brief introduction to normed vector spaces, C^* -algebras and define the model of quantum query complexity. In Section 3, we prove our main theorem characterizing quantum query algorithms. In Section 4, we explain the separation obtained for higher-degree forms. In Section 5, we give a short proof of the main theorem in Aaronson et al. [AAI⁺16].

2 Preliminaries

Notation. For a positive integer t denote $[t] = \{1, \dots, t\}$. For $x \in \mathbb{C}^n$, let $\text{Diag}(x)$ be the $n \times n$ diagonal matrix whose diagonal forms x . Given a matrix $X \in \mathbb{C}^{n \times n}$, let $\text{diag}(X) \in \mathbb{C}^n$ denote its diagonal vector. For $x \in \{0, 1\}^n$, denote $(-1)^x = ((-1)^{x_1}, \dots, (-1)^{x_n})$. Let $e_1, e_2, \dots, e_n \in \mathbb{C}^n$ be the standard basis vectors and let $E_{ij} = e_i e_j^*$. For $i, j \in [n]$, let $\delta_{i,j}$ be the indicator for the event $[i = j]$. Let $\mathbf{1} = (1, \dots, 1)$ and $\mathbf{0} = (0, \dots, 0)$ denote the n -dimensional all-ones (resp. all-zeros) vector.

Normed vector spaces. For parameter $p \in [1, \infty)$, the p -norm of a vector $x \in \mathbb{R}^n$ is defined by $\|x\|_p = (|x_1|^p + \dots + |x_n|^p)^{1/p}$ and for $p = \infty$ by $\|x\|_\infty = \max\{|x_i| : i \in [n]\}$. Denote the n -dimensional Euclidean unit ball by $B_2^n = \{x \in \mathbb{R}^n : \|x\|_2 \leq 1\}$. For a matrix $A \in \mathbb{R}^{n \times n}$, denote the standard operator norm by $\|A\|$ and define

$$\|A\|_{\ell_\infty \rightarrow \ell_1} = \sup \{ \|Ax\|_{\ell_1} : \|x\|_{\ell_\infty} \leq 1 \} = \max_{x, y \in \{-1, 1\}^n} x^\top A y.$$

We denote the norm of a general normed vector space X by $\|\cdot\|_X$, if there is a danger of ambiguity. Denote by $\mathbb{1}_X$ the identity map on X and by $\mathbb{1}_d$ the identity map on \mathbb{C}^d . For normed vector spaces X, Y , let $L(X, Y)$ be the collection of all linear maps $T : X \rightarrow Y$. We will use the notation $L(X)$

as a shorthand for $L(X, X)$. The (operator) norm of a linear map $T \in L(X, Y)$ is given by $\|T\| = \sup\{\|T(x)\|_Y : \|x\|_X \leq 1\}$. Such a map is an *isometry* if $\|T(x)\|_Y = \|x\|_X$ for every $x \in X$ and a *contraction* if $\|T(x)\|_Y \leq \|x\|_X$ for every $x \in X$. Throughout we endow \mathbb{C}^d with the standard Euclidean norm. Note that the space $L(\mathbb{C}^d)$ is naturally identified with the set of $d \times d$ matrices, sometimes denoted $M_d(\mathbb{C})$, and we use the two notations interchangeably. For a Hilbert space \mathcal{H} , we endow $\mathcal{H} \otimes \mathbb{C}^d$ with the norm given by the inner product $\langle f \otimes a, g \otimes b \rangle = \langle f, g \rangle_{\mathcal{H}} \langle a, b \rangle$, making this space isometric to $\mathcal{H} \oplus \dots \oplus \mathcal{H}$ (d times). Similarly, we endow $L(\mathcal{H}) \otimes L(\mathbb{C}^d)$ with the operator norm of the space $L(\mathcal{H} \otimes \mathbb{C}^d)$ of linear operators on the Hilbert space $\mathcal{H} \otimes \mathbb{C}^d$; with some abuse of notation, we shall identify the two spaces of operators.

C^* -algebras. We collect a few basic facts of C^* -algebras that we use later and refer to [Arv12] for an extensive introduction. A C^* -algebra $\mathcal{X} = (X, \cdot, *)$ is a normed complex vector space X , complete with respect to its norm (i.e., a Banach space), that is endowed with two operations in addition to the standard vector-space addition and scalar multiplication operations:

1. an associative multiplication $\cdot : X \times X \rightarrow X$, denoted $x \cdot y$ for $x, y \in X$, that is distributive with respect to the vector space addition and continuous with respect to the norm of X , which is to say that $\|x \cdot y\|_X \leq \|x\|_X \|y\|_X$ for all $x, y \in X$;
2. an involution $*$: $X \rightarrow X$, that is, a conjugate linear map that sends $x \in X$ to (a unique) $x^* \in X$ satisfying $(x^*)^* = x$ and $(xy)^* = y^* x^*$ for any $x, y \in X$, and such that $\|x \cdot x^*\|_X = \|x\|_X^2$.

Any finite-dimensional normed vector space is a Banach space. A C^* -algebra \mathcal{X} is *unital* if it has a multiplicative identity, denoted $\mathbb{1}_{\mathcal{X}}$. The most important example of a unital C^* -algebra is $M_n(\mathbb{C})$, where the involution operator is the conjugate-transpose and the norm is the operator norm. A linear map $\pi : \mathcal{X} \rightarrow \mathcal{Y}$ from one C^* -algebra \mathcal{X} to another \mathcal{Y} is a *$*$ -homomorphism* if it preserves the multiplication operation, $\pi(xy) = \pi(x)\pi(y)$, and satisfies $\pi(x)^* = \pi(x^*)$ for all $x, y \in \mathcal{X}$. For a complex Hilbert space \mathcal{H} , a mapping $\pi : \mathcal{X} \rightarrow L(\mathcal{H})$ is a *$*$ -representation* if it is a $*$ -homomorphism. An important fact is the Gelfand–Naimark Theorem [Mur14, Theorem 3.4.1] asserting that any C^* -algebra admits an isometric (that is, norm-preserving) $*$ -representation for some complex Hilbert space.

Completely bounded norms. We also collect a few basic facts about completely bounded norms that we use later and refer to [Pau02] for an extensive introduction. For a C^* -algebra \mathcal{X} and positive integer d , we denote by $M_d(\mathcal{X})$ the set of d -by- d matrices with entries in \mathcal{X} . Note that this set can naturally be identified with the algebraic tensor product $\mathcal{X} \otimes L(\mathbb{C}^d)$, that is, the linear span of all elements of the form $x \otimes M$, where $x \in \mathcal{X}$ and $M \in L(\mathbb{C}^d)$. We shall endow $M_d(\mathcal{X})$ with a norm induced by an isometric $*$ -representation π of \mathcal{X} into $L(\mathcal{H})$ for a Hilbert space \mathcal{H} . The linear map $\pi \otimes \mathbb{1}_{L(\mathbb{C}^d)}$ sends elements in $M_d(\mathcal{X})$ (or $\mathcal{X} \otimes L(\mathbb{C}^d)$) to elements (operators) in $L(\mathcal{H} \otimes \mathbb{C}^d)$. The norm of an element $A \in M_d(\mathcal{X})$ is then defined to be $\|A\| = \|(\pi \otimes \mathbb{1}_{L(\mathbb{C}^d)})(A)\|$. The notation $\|A\|$ reflects the fact that this norm is in fact independent of the particular $*$ -representation π . Based on this, we can define a norm on linear maps $\sigma : \mathcal{X} \rightarrow L(\mathcal{H})$ as follows:

$$\|\sigma\|_{\text{cb}} = \sup \left\{ \frac{\|(\sigma \otimes \mathbb{1}_{L(\mathbb{C}^d)})(A)\|}{\|A\|} : d \in \mathbb{N}, A \in \mathcal{X} \otimes L(\mathbb{C}^d), A \neq 0 \right\}$$

Tensors and multilinear forms. For vector spaces X, Y over the same field and positive integer t , recall that a mapping

$$T : \underbrace{X \times \cdots \times X}_{t \text{ times}} \rightarrow Y$$

is *t-linear* if for every $x_1, \dots, x_t \in X$ and $i \in [t]$, the map $y \mapsto T(x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_t)$ is linear. A *t-tensor* of dimension n is a map $T : [n] \times \cdots \times [n] \rightarrow \mathbb{C}$, which can alternatively be identified by $T = (T_{i_1, \dots, i_t})_{i_1, \dots, i_t=1}^n \in \mathbb{C}^{n \times \cdots \times n}$. With abuse of notation we identify a *t-tensor* $T \in \mathbb{C}^{n \times \cdots \times n}$ with the *t-linear* form $T : \mathbb{C}^n \times \cdots \times \mathbb{C}^n \rightarrow \mathbb{C}$ given by

$$T(x_1, \dots, x_t) = \sum_{i_1, \dots, i_t=1}^n T_{i_1, \dots, i_t} x_1(i_1) \cdots x_t(i_t).$$

Next, we introduce the completely bounded norm of a *t-linear* form $T : \mathcal{X} \times \cdots \times \mathcal{X} \rightarrow \mathbb{C}$ on a C^* -algebra \mathcal{X} . First, we use the standard identification of such forms with the linear form on the tensor product $\mathcal{X} \otimes \cdots \otimes \mathcal{X}$ given by $T(x_1 \otimes \cdots \otimes x_t) = T(x_1, \dots, x_t)$. We consider a bilinear map $\odot : (\mathcal{X} \otimes L(\mathbb{C}^d), \mathcal{X} \otimes L(\mathbb{C}^d)) \rightarrow \mathcal{X} \otimes \mathcal{X} \otimes L(\mathbb{C}^d)$ for any positive integer d defined as follows. For $x, y \in \mathcal{X}$ and $M_x, M_y \in L(\mathbb{C}^d)$, let

$$(x \otimes M_x) \odot (y \otimes M_y) = (x \otimes y) \otimes (M_x M_y).$$

Observe that this operation changes the order of the tensor factors and *multiplies* M_x with M_y . This operation is associative but *not* commutative. Extend the definition of the \odot operation bi-linearly to its entire domain. Define the *t-linear* map $T_d : M_d(\mathcal{X}) \times \cdots \times M_d(\mathcal{X}) \rightarrow L(\mathbb{C}^d)$ by

$$T_d(A_1, \dots, A_t) = (T \otimes \mathbb{1}_{L(\mathbb{C}^d)})(A_1 \odot \cdots \odot A_t).$$

The completely bounded norm of T is now defined by

$$\|T\|_{\text{cb}} = \sup \left\{ \|T_d(A_1, \dots, A_t)\| : d \in \mathbb{N}, A_j \in M_d(\mathcal{X}), \|A_j\| \leq 1 \right\}.$$

Note that the definition given in (4) corresponds to the particular case where the C^* -algebra \mathcal{X} is formed by the $n \times n$ diagonal matrices. Since any square matrix with operator norm at most 1 is a convex combination of unitary matrices (by the Russo-Dye Theorem)², the completely bounded norm can also be defined by taking the supremum over unitaries $A_j \in M_d(\mathcal{X})$. The completely bounded norm can be defined more generally for multilinear maps into $L(\mathcal{H})$, for some Hilbert space \mathcal{H} , to yield the definition of this norm for linear maps given above, but we will not use this here.

Quantum query complexity. The quantum query model was formally defined by Beals et al. in [BBC⁺01]. In this model, we are given black-box access to a unitary operator, often called an oracle O_x , whose description depends in a simple way on some binary input string $x \in \{0, 1\}^n$. An application of the oracle on a quantum register is referred to as a quantum *query* to x . In the standard form of the model, a query acts on a pair of registers on (Q, A) , where Q is an n -dimensional query register and A is a one-qubit auxiliary register. A query to the oracle effects the unitary transformation given by

$$O_x : |i, b\rangle \rightarrow |i, b \oplus x_i\rangle$$

²A precise statement and short proof of the Russo-Dye theorem can be found in [Gar84].

where $i \in [n]$, $b \in \{0,1\}$. (These oracles are also commonly called *bit oracles*.)

A quantum query algorithm consists of a fixed sequence of unitary operations acting on (Q,A) in addition to a *workspace* register W . A t -query quantum algorithm begins by initializing the joint register (Q,A,W) in the all-zero state and continues by interleaving a sequence of unitaries U_0, \dots, U_t on (Q,A,W) with oracles O_x on (Q,A) . Finally, the algorithm performs a 2-outcome measurement on A and returns the measurement outcome.




Figure 1: A t -query quantum algorithm that starts with the all-zero state and concludes by measuring the register A .

For a Boolean function $f : \{0,1\}^n \rightarrow \{0,1\}$, the algorithm is said to compute f with error $\varepsilon > 0$ if for every x , the measurement outcome of register A equals $f(x)$ with probability at least $1 - \varepsilon$. The *bounded-error query complexity* of f , denoted $Q_\varepsilon(f)$, is the smallest t for which such an algorithm exists. Note that in this model, we are not concerned with the amount of time (i.e., the number of gates) it takes to implement the interlacing unitaries, which could be much bigger than the query complexity itself.

Here we will work with a slightly less standard oracle sometime referred to as a *phase oracle*, in which the standard oracle is preceded and followed by a Hadamard on A . Since the Hadamards can be undone by the unitaries surrounding the queries in a quantum query algorithm, using the phase oracle does not reduce generality. A query to this oracle, sometimes denoted $O_{x,\pm}$, applies the (controlled) unitary $\text{Diag}((1, (-1)^x))$ to joint register (A,Q) . To avoid having to write $(-1)^x$ later on, we shall work in the equivalent setting where Boolean functions send $\{-1,1\}^n$ to $\{-1,1\}$.

3 Characterization of quantum query algorithms

In this section we prove Theorem 1.3. The main ingredient of the proof is the following celebrated representation theorem by Christensen and Sinclair [CS87] showing that completely-boundedness of a multilinear form is equivalent to the existence of an exceedingly nice factorization.

Theorem 3.1 (Christensen–Sinclair). *Let t be a positive integer and let \mathcal{X} be a C^* -algebra. Then, for any t -linear form $T : \mathcal{X} \times \dots \times \mathcal{X} \rightarrow \mathbb{C}$, we have $\|T\|_{\text{cb}} \leq 1$ if and only if there exist Hilbert spaces $\mathcal{H}_0, \dots, \mathcal{H}_{t+1}$ where $\mathcal{H}_0 = \mathcal{H}_{t+1} = \mathbb{C}$, $*$ -representations $\pi_i : \mathcal{X} \rightarrow L(\mathcal{H}_i)$ for each $i \in [t]$ and contractions $V_i \in L(\mathcal{H}_i, \mathcal{H}_{i-1})$, for each $i \in [t+1]$ such that for any $x_1, \dots, x_t \in \mathcal{X}$, we have*

$$T(x_1, \dots, x_t) = V_1 \pi_1(x_1) V_2 \pi_2(x_2) V_3 \dots V_t \pi_t(x_t) V_{t+1}. \quad (6)$$

We first show how the above result simplifies when restricting to the special case in which the C^* -algebra \mathcal{X} is formed by the set of diagonal n -by- n matrices.

Corollary 3.2. *Let m, n, t be positive integers such that $t \geq 2$ and $m = n^t$. Let $T \in \mathbb{C}^{n \times \dots \times n}$ be a t -tensor. Then, $\|T\|_{\text{cb}} \leq 1$ if and only if there exist a positive integer d , unit vectors $u, v \in \mathbb{C}^m$ and contractions $U_i, V_i \in L(\mathbb{C}^m, \mathbb{C}^{dn})$ such that for any $x_1, \dots, x_t \in \mathbb{C}^n$, we have*

$$T(x_1, \dots, x_t) = u^* U_1^* (\text{Diag}(x_1) \otimes \mathbb{1}_d) V_1 \cdots U_t^* (\text{Diag}(x_t) \otimes \mathbb{1}_d) V_t v. \quad (7)$$

The proof of the above corollary uses the following fact about the completely bounded norm of $*$ -representations of C^* -algebras [Pis03, Theorem 1.6].

Lemma 3.3. *Let \mathcal{X} be a finite-dimensional C^* -algebra, $\mathcal{H}, \mathcal{H}'$ be Hilbert spaces, $\pi : \mathcal{X} \rightarrow L(\mathcal{H})$ be a $*$ -representation and $U \in L(\mathcal{H}, \mathcal{H}')$ and $V \in L(\mathcal{H}', \mathcal{H})$ be linear maps. Then, the map $\sigma : \mathcal{X} \rightarrow L(\mathcal{H}')$, defined as $\sigma(x) = U\pi(x)V$, satisfies that $\|\sigma\|_{\text{cb}} \leq \|U\| \|V\|$.*

In addition, we use the famous Fundamental Factorization Theorem [Pau02, Theorem 8.4]. Below we state the theorem when restricted to finite-dimensional spaces (see also the remark after [JKP09, Theorem 16]).

Theorem 3.4 (Fundamental factorization theorem). *Let $\sigma : L(\mathbb{C}^n) \rightarrow L(\mathbb{C}^m)$ be a linear map and let $d = nm$. Then, there exist $U, V \in L(\mathbb{C}^m, \mathbb{C}^{dn})$ such that $\|U\| \|V\| \leq \|\sigma\|_{\text{cb}}$ and for any $M \in L(\mathbb{C}^n)$, we have $\sigma(M) = U^*(M \otimes \mathbb{1}_d)V$.*

Proof of Corollary 3.2. The set $\mathcal{X} = \text{Diag}(\mathbb{C}^n)$ of diagonal matrices is a (finite-dimensional) C^* -algebra (endowed with the standard matrix product and conjugate-transpose involution). Define the t -linear form $R : \mathcal{X} \times \dots \times \mathcal{X} \rightarrow \mathbb{C}$ by $R(X_1, \dots, X_t) = T(\text{diag}(X_1), \dots, \text{diag}(X_t))$. We claim that $\|R\|_{\text{cb}} = \|T\|_{\text{cb}}$. Observe that for every positive integer d , the set $\{B \in M_d(\mathcal{X}) : \|B\| \leq 1\}$ can be identified with the set of block-diagonal matrices $B = \sum_{i=1}^n E_{i,i} \otimes B(i)$ of size $nd \times nd$ and blocks $B(1), \dots, B(n)$ of size $d \times d$ satisfying $\|B(i)\| \leq 1$ for all $i \in [n]$. It follows that

$$R_d(B_1, \dots, B_t) = \sum_{i_1, \dots, i_t=1}^n R(E_{i_1, i_1}, \dots, E_{i_t, i_t}) B_1(i_1) \cdots B_t(i_t) = \sum_{i_1, \dots, i_t=1}^n T_{i_1, \dots, i_t} B_1(i_1) \cdots B_t(i_t),$$

which shows the claim.

Next, we show that (6) is equivalent to (7). The fact that (7) implies (6) follows immediately from the fact that the map $\text{Diag}(x) \mapsto \text{Diag}(x) \otimes \mathbb{1}_d$ is a $*$ -representation. Now assume (6). Without loss of generality, we may assume that each of the Hilbert spaces $\mathcal{H}_1, \dots, \mathcal{H}_t$ has dimension at least m . If not, we can expand the dimensions of the ranges and domains of the representations π_i and contractions V_i by dilating with appropriate isometries into larger Hilbert spaces (“padding with zeros”). For each $i \in [t]$, let $S_i \subseteq \mathcal{H}_i$ be the subspace

$$S_i = \text{Span} \{ \pi_i(x_i) V_{i+1} \cdots V_t \pi_t(x_t) V_{t+1} : x_i, \dots, x_t \in \mathcal{X} \}.$$

Since $\dim(\mathcal{X}) = n$, we have that $\dim(S_i) \leq m$. For each $i \in [t]$, let $Q_i \in L(\mathbb{C}^m, \mathcal{H}_i)$ be an isometry such that $S_i \subseteq \text{Im}(Q_i)$. Note that V_{i+1} is a vector in the unit ball of \mathcal{H}_i . Let $Q_{t+1} \in L(\mathbb{C}^m, \mathcal{H}_t)$ be an isometry such that $V_{t+1} \in \text{Im}(Q_{t+1})$. Note that for each $i \in [t+1]$, the map $Q_i Q_i^*$ acts as the identity on $\text{Im}(Q_i)$. For each $i \in \{2, \dots, t\}$ define the map $\sigma_i : \mathcal{X} \rightarrow L(\mathbb{C}^m)$ by $\sigma_i(x) = Q_i^* V_i \pi_i(x) Q_{i+1}$ and $\sigma_1(x) = Q_1^* \pi_1(x) Q_2$. Finally define $u = Q_1^* V_1^*$ and $v = Q_{t+1}^* V_{t+1}$. Then, the right-hand side of (6) can be written as

$$u^* \sigma_1(x_1) \cdots \sigma_t(x_t) v.$$

It follows from Lemma 3.3 that $\|\sigma_i\|_{\text{cb}} \leq 1$. Let $\sigma'_i : L(\mathbb{C}^n) \rightarrow L(\mathbb{C}^m)$ be the linear map given by $\sigma'_i(M) = \sigma_i(\text{Diag}(M_{11}, \dots, M_{nn}))$ for any $M \in L(\mathbb{C}^m)$. Then, for any diagonal matrix $x \in \mathcal{X}$, we have $\sigma_i(x) = \sigma'_i(x)$ and $\|\sigma'_i\|_{\text{cb}} = \|\sigma_i\|_{\text{cb}}$. It follows from Theorem 3.4 that there exists a positive integer d_i and contractions $U_i, V_i : L(\mathbb{C}^m, \mathbb{C}^{d_i})$ such that $\sigma_i(x) = U_i^*(x \otimes \mathbb{1}_{d_i})V_i$ for any $x \in \mathcal{X}$. We can take all d_i equal to $d = \max_i \{d_i\}$ by suitably dilating the contractions U_i, V_i . Setting $u' = u/\|u\|_{\ell_2}$ and $U'_1 = \|u\|_{\ell_2} U_1$, and similarly defining v', V'_{i+1} gives the remaining implication. \square

Corollary 3.2 implies the following lemma, from which Theorem 1.3 easily follows.

Lemma 3.5. *Let $\beta : \{-1, 1\}^n \rightarrow [-1, 1]$ be some map and let t be a positive integer. Then, the following are equivalent.*

1. *There exists a $(2t)$ -tensor $T \in \mathbb{R}^{2n \times \dots \times 2n}$ such that $\|T\|_{\text{cb}} \leq 1$ and for every $x \in \{-1, 1\}^n$ and $y = (x, \mathbf{1})$, we have*

$$\sum_{i_1, \dots, i_{2t}=1}^{2n} T_{i_1, \dots, i_{2t}} y_{i_1} \cdots y_{i_{2t}} = \beta(x).$$

2. *There exists a t -query quantum algorithm that, on input $x \in \{-1, 1\}^n$, returns a random sign with expected value $\beta(x)$.*

Proof. We first prove that (2) implies (1). As discussed in Section 2, a t -query quantum algorithm with phase oracles initializes the joint register (A, Q, W) in the all-zero state on which it then performs a sequence of unitaries U_1, \dots, U_t interlaced with queries $D(x) = \text{Diag}((\mathbf{1}, x)) \otimes \mathbb{1}_W$. Let $\{P_0, P_1\}$ be the two-outcome measurement done at the end of the algorithm and assume that it returns $+1$ on measurement outcome zero and -1 otherwise. Let $Q = P_0 - P_1$ and note that Q is a contraction since P_0, P_1 are positive semi-definite and satisfy $P_0 + P_1 = \mathbb{1}$. The expected value of the measurement outcome is then given by

$$e_0^* U_1^* D(x) U_2^* \cdots D(x) U_t^* Q U_t D(x) \cdots U_2 D(x) U_1 e_0. \quad (8)$$

By assumption, this expected value equals $\beta(x)$ for every $x \in \{-1, 1\}^n$. For $z \in \mathbb{C}^{2n}$, denote $D'(z) = \text{Diag}((z_{n+1}, \dots, z_{2n}, z_1, \dots, z_n)) \otimes \mathbb{1}_W$ and $\tilde{U}_t = U_t^* Q U_t$. Define the $(2t)$ -linear form T by

$$T(y_1, \dots, y_{2t}) = u^* U_1^* D'(y_1) U_2^* \cdots D'(y_t) \tilde{U}_t D'(y_{t+1}) \cdots U_2 D'(y_{2t}) U_1 u.$$

Clearly $T((x, \mathbf{1}), \dots, (x, \mathbf{1})) = \beta(x)$ for every $x \in \{-1, 1\}^n$. Moreover, by definition T admits a factorization as in (7). It thus follows from Corollary 3.2 that $\|T\|_{\text{cb}} \leq 1$. We turn T into a real tensor by taking its real part $T' = (T + \bar{T})/2$, where \bar{T} is the coordinate-wise complex conjugate of T . Since for any $x \in \{-1, 1\}^n$ and $y = (x, \mathbf{1})$, the value $T(y, \dots, y)$ is real, we have $T'(y, \dots, y) = \beta(x)$. We need to show that $\|T'\|_{\text{cb}} \leq 1$. To this end, consider an arbitrary positive integer d , unit vectors $v, w \in \mathbb{C}^d$ and sequences of unitary matrices $V_1(i), \dots, V_{2t}(i)$ for $i \in [n]$ such that

$$\left\| \sum_{i_1, \dots, i_{2t}=1}^{2n} \overline{T_{i_1, \dots, i_{2t}}} V_1(i_1) \cdots V_{2t}(i_{2t}) \right\| = \left| \sum_{i_1, \dots, i_{2t}=1}^{2n} \overline{T_{i_1, \dots, i_{2t}}} v^* V_1(i_1) \cdots V_{2t}(i_{2t}) w \right|.$$

Note that $\|\bar{T}\|_{\text{cb}}$ is given by the supremum over d and $V_j(i)$. Taking the complex conjugate of the above summands on the right-hand side allows us to express the above absolute value as

$$\left| \sum_{i_1, \dots, i_{2t}=1}^{2n} T_{i_1, \dots, i_{2t}} \bar{v}^* \overline{V_1(i_1)} \cdots \overline{V_{2t}(i_{2t})} \bar{w} \right|, \quad (9)$$

where $\bar{v}, \bar{w}, \overline{V_j(i)}$ denote the coordinate-wise complex conjugates. Since each $\overline{V_j(i)}$ is still unitary, it follows that (9) is at most $\|T\|_{\text{cb}}$ and so $\|\bar{T}\|_{\text{cb}} \leq \|T\|_{\text{cb}} \leq 1$. Hence, by the triangle inequality, $\|T'\|_{\text{cb}} \leq (\|T\|_{\text{cb}} + \|\bar{T}\|_{\text{cb}})/2 \leq 1$ as desired.

Next, we show that (1) implies (2). Let T be a $(2t)$ -tensor as in item 1. Since any matrix with operator norm at most 1 is a convex combination of unitary matrices (by the Russo-Dye Theorem), it follows from Corollary 3.2 that T admits a factorization as in (7). Let $V_0, U_{2t+1} \in L(\mathbb{C}^m, \mathbb{C}^{2dn})$ be isometries. For each $i \in [2t+1]$, define the map $W_i \in L(\mathbb{C}^{2dn})$ by $W_i = V_{i-1}U_i^*$. Observe that each W_i is a contraction and recall that unitaries are contractions. For the moment, assume for simplicity that each W_i is in fact unitary. Define two vectors $\tilde{u} = V_0u$ and $\tilde{v} = U_{2t+1}v$ and observe that these are unit vectors in \mathbb{C}^{2dn} . The right-hand side of (7) then gives us

$$T(y_1, \dots, y_{2t}) = \tilde{u}^* W_1 (\text{Diag}(y_1) \otimes \mathbb{1}_d) W_2 (\text{Diag}(y_2) \otimes \mathbb{1}_d) W_3 \cdots W_{2t} (\text{Diag}(y_{2t}) \otimes \mathbb{1}_d) W_{2t+1} \tilde{v}. \quad (10)$$

Based on this, we obtain the quantum query algorithm described in Figure 2.




Figure 2: The registers C, Q, W denote the control, query and workspace registers. Let U, V be unitaries with $W_1\tilde{u}$ and $W_{2t+1}\tilde{v}$ as their first rows, respectively and for $x \in \{-1, 1\}^n$ and $y = (x, \mathbf{1})$, let $\text{Diag}(y)$ be the query operator. The algorithm begins by initializing the joint register (C, Q, W) in the all-zero state and proceeds by performing the displayed operations. The algorithm returns $+1$ if the outcome of the measurement on C equals zero and -1 otherwise.

To see why this algorithm satisfies the requirements, first note that the algorithm makes t queries to the input x . For the correctness of the algorithm, we begin by observing that before the application of the first query, the state of the joint register (C, Q, W) is

$$\frac{1}{\sqrt{2}}(e_0 \otimes W_1\tilde{u} + e_1 \otimes W_{t+1}\tilde{v}).$$

Before the final Hadamard gate, the state of the joint register is given by

$$\begin{aligned} & \frac{1}{\sqrt{2}}e_0 \otimes ((\text{Diag}(y) \otimes \mathbb{1}_d)W_t \cdots W_2(\text{Diag}(y) \otimes \mathbb{1}_d)W_1\tilde{u}) \\ & + \frac{1}{\sqrt{2}}e_1 \otimes (W_{t+1}^*(\text{Diag}(y) \otimes \mathbb{1}_d)W_{t+2}^* \cdots W_{2t}^*(\text{Diag}(y) \otimes \mathbb{1}_d)W_{2t+1}^*\tilde{v}). \end{aligned}$$

A standard calculation and (10) then show that after the final Hadamard gate, the expected output of the algorithm is precisely $T((x, \mathbf{1}), \dots, (x, \mathbf{1})) = \beta(x)$. In the general case where the W_i s are not necessarily unitary, we can use the fact that, by the Russo-Dye Theorem and Carathéodory's

Theorem, each W_i is a convex combination of at most $(dn)^2 + 1$ unitaries. The algorithm can thus use randomness to effect each W_i on expectation. Alternatively, by linear algebra there exists a unitary matrix $W'_i \in \mathbb{C}^{2dn \times 2dn}$ that has W_i as its upper-left corner (see [AAI⁺16, Lemma 7]), through which the algorithm could implement W_i by working on a larger quantum register. \square

Using Lemma 3.5, we now prove our main Theorem 1.3.

Proof of Theorem 1.3. We first show that (2) implies (1). Using the equivalence in Lemma 3.5, it follows that there exists a $(2t)$ -tensor $T \in \mathbb{R}^{2n \times \dots \times 2n}$ such that $\|T\|_{\text{cb}} \leq 1$ and for every $x \in \{-1, 1\}^n$ and $y = (x, \mathbf{1})$, we have

$$\sum_{i_1, \dots, i_{2t}=1}^{2n} T_{i_1, \dots, i_{2t}} y_{i_1} \cdots y_{i_{2t}} = \beta(x).$$

Define the symmetric $2t$ -tensor $T_p = \frac{1}{(2t)!} \sum_{\sigma \in S_{2t}} T \circ \sigma$. Let $p \in \mathbb{R}[x_1, \dots, x_{2n}]$ be the form of degree $2t$ associated with T_p by (2) (note that there is a unique polynomial associated with the symmetric tensor T_p). Then, $p((x, \mathbf{1})) = \beta(x)$ for every $x \in \{-1, 1\}^n$. Moreover, if we set $T^\sigma = T$ for each $\sigma \in S_{2t}$, it follows from the above decomposition of T_p and Definition 1.2 that $\|p\|_{\text{cb}} \leq \|T\|_{\text{cb}} \leq 1$.

Next, we show that (1) implies (2). Let p be a degree- $(2t)$ form satisfying $\|p\|_{\text{cb}} \leq 1$. Suppose T_p as defined in Eq. (2) can be written as $T_p = \sum_{\sigma \in S_{2t}} T^\sigma \circ \sigma$ and $\sum_{\sigma \in S_{2t}} \|T^\sigma\|_{\text{cb}} = \|p\|_{\text{cb}} \leq 1$. Define $T = \sum_{\sigma \in S_{2t}} T^\sigma$. Then, using the triangle inequality, it follows that $\|T\|_{\text{cb}} \leq \sum_{\sigma \in S_{2t}} \|T^\sigma\|_{\text{cb}} \leq 1$. Also note that for any $y \in \mathbb{R}^{2n}$,

$$T(y, \dots, y) = \sum_{\sigma \in S_{2t}} T^\sigma(y, \dots, y) = \sum_{\sigma \in S_{2t}} (T^\sigma \circ \sigma)(y, \dots, y) = T_p(y, \dots, y) = p(y).$$

Using Lemma 3.5 (in particular (1) \implies (2)) for the tensor T , the theorem follows. \square

We now prove Corollary 1.5, which is an immediate consequence of our main theorem.

Proof of Corollary 1.5. We first show $\text{cb-deg}_\varepsilon(f) \geq Q_\varepsilon(f)$: Suppose $\text{cb-deg}_\varepsilon(f) = d$, then there exists a degree- $(2d)$ form p satisfying: $|p(x) - f(x)| \leq 2\varepsilon$ for every $x \in D$ and $\|p\|_{\text{cb}} \leq 1$. Using our characterization in Theorem 1.3, it follows that there exists a d -query quantum algorithm \mathcal{A} , that on input $x \in D$, returns a random sign with expected value $p(x)$. So, our ε -error quantum algorithm for f simply runs \mathcal{A} and outputs the random sign.

We next show $\text{cb-deg}_\varepsilon(f) \leq Q_\varepsilon(f)$. Suppose $Q_\varepsilon(f) = t$. Then, there exists a t -query quantum algorithm that, on input $x \in D$, outputs a random sign with expected value $\beta(x)$ satisfying $|\beta(x) - f(x)| \leq 2\varepsilon$. Note that we could also run the quantum algorithm for $x \notin D$ and let $\beta(x)$ be the expected value of the quantum algorithm for such x s. Using Theorem 1.3, we know that there exists a degree- $(2t)$ form p satisfying $\beta(x) = p(x)$ for every $x \in \{-1, 1\}^n$ and $\|p\|_{\text{cb}} \leq 1$. Clearly p satisfies the conditions of Definition 1.4, hence $\text{cb-deg}_\varepsilon(f) \leq t$. \square

4 Separations for quartic polynomials

In this section we show the existence of a quartic polynomial p that is bounded but for which any two-query quantum algorithm \mathcal{A} satisfying $\mathbb{E}[\mathcal{A}(x)] = Cp(x)$ for every $x \in \{-1, 1\}^n$ must necessarily have $C = O(n^{-1/2})$. We show this using a (random) *cubic* form that is bounded, but whose completely bounded norm is $\text{poly}(n)$, following a construction of [DJT95, Theorem 18.16].

Given a form $p : \mathbb{R}^n \rightarrow \mathbb{R}$, we define its norm as

$$\|p\| = \sup\{|p(x)| : x \in \{-1, 1\}^n\}.$$

Note that the condition $\|p\| \leq 1$ is equivalent to p being bounded.

Theorem 4.1. *There exist absolute constants $C, c \in (0, \infty)$ such that the following holds. Let*

$$p(x) = \sum_{\alpha \in \{0,1,2,3\}^n : |\alpha|=3} c_\alpha x^\alpha$$

be the random cubic form such the coefficients c_α are independent uniformly distributed $\{-1, 1\}$ -valued random variables. Then, with probability at least $1 - Cne^{-cn}$, we have $\|p\|_{\text{cb}} \geq c\sqrt{n}\|p\|$.

We shall use the following standard concentration-of-measure results. The first is the Hoeffding bound [Pol12, Corollary 3 (Appendix B)].

Lemma 4.2 (Hoeffding bound). *Let X_1, \dots, X_m be independent uniformly distributed $\{-1, 1\}$ -random variables and let $a \in \mathbb{R}^m$. Then, for any $\tau > 0$, we have*

$$\Pr\left[\left|\sum_{i=1}^m a_i X_i\right| > \tau\right] \leq 2e^{-\frac{\tau^2}{2(a_1^2 + \dots + a_m^2)}}$$

The second result is one from random matrix theory concerning upper tail estimates for Wigner ensembles (see [Tao12, Corollary 2.3.6]).

Lemma 4.3. *There exist absolute constants $C, c \in (0, \infty)$ such that the following holds. Let n be a positive integer and let M be a random $n \times n$ symmetric random matrix such that for $j \geq i$, the entries M_{ij} are independent random variables with mean zero and absolute value at most 1. Then, for any $\tau \geq C$, we have*

$$\Pr[\|M\| > \tau\sqrt{n}] \leq Ce^{-c\tau n}.$$

We also use the following proposition.

Proposition 4.4. *Let m, n, t be positive integers, let $p \in \mathbb{R}[x_1, \dots, x_n]$ be a t -linear form, let $T_p \in \mathbb{R}^{n \times \dots \times n}$ be as in (2) and let $A_1, \dots, A_n \in L(\mathbb{R}^m)$ be pairwise commuting contractions. Then,*

$$\|p\|_{\text{cb}} \geq \left\| \sum_{i_1, \dots, i_t=1}^n (T_p)_{i_1, \dots, i_t} A_{i_1} \cdots A_{i_t} \right\|.$$

Proof. Consider an arbitrary decomposition $T_p = \sum_{\sigma \in S_t} T_p^\sigma \circ \sigma$. Then, the definition of the completely bounded norm and triangle inequality show that

$$\sum_{\sigma \in S_t} \|T_p^\sigma\|_{\text{cb}} \geq \sum_{\sigma \in S_t} \left\| \sum_{i_1, \dots, i_t=1}^n T_{i_1, \dots, i_t}^\sigma A_{i_1} \cdots A_{i_t} \right\| \geq \left\| \sum_{\sigma \in S_t} \sum_{i_1, \dots, i_t=1}^n T_{i_1, \dots, i_t}^\sigma A_{i_1} \cdots A_{i_t} \right\|.$$

Since the A_i commute, the above reduces to

$$\begin{aligned} \left\| \sum_{\sigma \in S_t} \sum_{i_1, \dots, i_t=1}^n T_{i_1, \dots, i_t}^\sigma A_{\sigma^{-1}(i_1)} \cdots A_{\sigma^{-1}(i_t)} \right\| &= \left\| \sum_{\sigma \in S_t} \sum_{i_1, \dots, i_t=1}^n (T_p^\sigma \circ \sigma)_{i_1, \dots, i_t} A_{i_1} \cdots A_{i_t} \right\| \\ &= \left\| \sum_{i_1, \dots, i_t=1}^n (T_p)_{i_1, \dots, i_t} A_{i_1} \cdots A_{i_t} \right\|. \end{aligned}$$

The claim now follows from the definition of $\|p\|_{\text{cb}}$ and since the decomposition of T_p was arbitrary. \square

Proof of Theorem 4.1. We begin by showing that with high probability, $\|p\| \leq O(n^2)$. To this end, let us fix an arbitrary $x \in \{-1, 1\}^n$. Then, $p(x)$ is a sum of at most n^3 independent uniformly distributed random $\{-1, 1\}$ -random variables. It therefore follows from Lemma 4.2 that

$$\Pr[|p(x)| > 2n^2] \leq 2e^{-2n},$$

By the union bound over $x \in \{-1, 1\}^n$, it follows that $\|p\| > 2n^2$ with probability at most $2e^{-n}$, which gives the claim.

We now lower bound $\|p\|_{\text{cb}}$. Let $\tau > 0$ be a parameter to be set later. Let $T \in \mathbb{R}^{n \times n \times n}$ be the random symmetric 3-tensor associated with p as in (2). For every $i \in [n]$, we define the linear map $A_i : \mathbb{R}^{2n+2} \rightarrow \mathbb{R}^{2n+2}$ by

$$\begin{cases} A_i e_0 = e_i \\ A_i e_j = \frac{1}{\tau\sqrt{n}} \sum_{k=1}^n T_{i,j,k} e_{k+n} \\ A_i e_{j+n} = \delta_{i,j} e_{2n+1} \\ A_i e_{2n+1} = 0. \end{cases}$$

Observe that for every $i, j, k \in [n]$, we have

$$e_{2n+1}^* A_i A_j A_k e_0 = \frac{1}{\tau\sqrt{n}} T_{i,j,k}. \quad (11)$$

Since T is symmetric, it follows easily that these maps commute, which is to say that $A_i A_j = A_j A_i$ for every $i, j \in [n]$. In addition, we claim that with high probability, these maps are contractions (i.e., the associated matrices have operator norm at most 1). To see this, for each $i \in [n]$, let M_i be the random matrix given by $M_i = (T_{i,j,k})_{j,k=1}^n$. Observe that M_i is symmetric and its entries have mean zero and absolute value at most 1. By Lemma 4.3 and a union bound, we get that

$$\Pr\left[\max_{i \in [n]} \|M_i\| > \tau\sqrt{n}\right] \leq Cne^{-c\tau n}. \quad (12)$$

for absolute constants c, C and provided $\tau \geq C$. Now, for any Euclidean unit vector $u \in \mathbb{R}^{2n+2}$, we have

$$\begin{aligned} \|A_i u\|^2 &= |u_0|^2 + \frac{1}{\tau^2 n} \sum_{k=1}^n \left| \sum_{j=1}^n u_j T_{i,j,k} \right|^2 + |u_{i+n}|^2 \\ &\leq |u_0|^2 + \frac{\|M_i\|^2}{\tau^2 n} \sum_{j=1}^n |u_j|^2 + |u_{i+n}|^2. \end{aligned}$$

It follows from (12) that $\max_i \|M_i\| \leq \tau\sqrt{n}$ with probability at least $1 - Cne^{-c\tau n}$, which in turn implies the above is at most $\|u\|^2 \leq 1$ and therefore that all A_i have operator norm at most 1.

By Proposition 4.4,

$$\|p\|_{\text{cb}} \geq \left\| \sum_{i,j,k=1}^n T_{i,j,k} A_i A_j A_k \right\|,$$

provided that the A_i s are contractions. By (11), and since $|T_{i,j,k}| \geq 1/6$ for every $i, j, k \in [n]$, the above is at least $n^{5/2}/(36\tau)$ with probability at least $1 - Cne^{-c\tau n}$. Letting τ be a sufficiently large constant then gives the result. \square

To demonstrate the failure of Theorem 1.1 for quartic polynomials, we embed As mentioned in the introduction, one can easily extend this result to the case of 4-linear forms.

Corollary 4.5. *There exists a bounded quartic form*

$$q(x_1, \dots, x_n) = \sum_{\alpha \in \{0,1\}^n: |\alpha|=4} d_\alpha x^\alpha, \quad (13)$$

and pairwise commuting contractions $A_1, \dots, A_n \in L(\mathbb{R}^{2n+2})$ such that

$$\left\| \sum_{i,j,k,\ell=1}^n (T_q)_{i,j,k,\ell} A_i A_j A_k A_\ell \right\| \geq c\sqrt{n}$$

where $c \in (0,1]$ is some absolute constant.

Proof. Let p be a bounded multi-linear cubic form such that $\|p\|_{\text{cb}} \geq C\sqrt{n}$, the existence of which is guaranteed by Theorem 4.1. Let $T_p \in \mathbb{R}^{n \times n \times n}$ be the random symmetric 3-tensor associated to p . Consider the symmetric 4-tensor $S \in \mathbb{R}^{(n+1) \times (n+1) \times (n+1) \times (n+1)}$ defined by $S_{0,j,k,\ell} = T_{j,k,\ell}$, $S_{i,0,k,\ell} = T_{i,k,\ell}$, $S_{i,j,0,\ell} = T_{i,j,\ell}$, $S_{i,j,k,0} = T_{i,j,k}$ for every $i, j, k, \ell \in [n]$ and $S_{i,j,k,\ell} = 0$ otherwise. Since S is symmetric, there exists a unique multi-linear quartic form q associated to S . It follows easily that $\|q\| = 4\|p\|$. Moreover, by considering the contractions A_i used in the proof of Theorem 4.1 and defining $A_0 = \mathbb{1}_{n+2}$, it follows that $\|q\|_{\text{cb}} \geq 4\|p\|_{\text{cb}}$. The form $q/4$ is thus as desired. \square

We claim that a form q as in Corollary (4.5) gives a counterexample to possible quartic extensions of Theorem 1.1. To see this, suppose there exists a two-query quantum algorithm \mathcal{A} and a $C \in (0, \infty)$ such that $\mathbb{E}[\mathcal{A}(x)] = Cq(x)$ for each $x \in \{-1, 1\}^n$. By Theorem 1.3 that there exists a $(2n)$ -variate quartic form h such that $h(x, \mathbf{1}) = Cq(x)$ for each $x \in \{-1, 1\}^n$ and $\|h\|_{\text{cb}} \leq 1$. We now show that the degree-4 coefficients in $h(x, y)$ are completely determined by $q(x)$. Indeed, if we expand

$$h(x, y) = \sum_{\alpha, \beta \in \{0,1,2,3,4\}^n: |\alpha|+|\beta|=4} d'_{\alpha,\beta} x^\alpha y^\beta,$$

then

$$h(x, \mathbf{1}) = \sum_{\alpha, \beta \in \{0,1,2,3,4\}^n: |\alpha|+|\beta|=4} d'_{\alpha,\beta} x^\alpha = C \sum_{\alpha \in \{0,1\}^n: |\alpha|=4} d_\alpha x^\alpha = Cq(x). \quad (14)$$

It follows from the above that $d'_{\alpha,0} = Cd_\alpha$ for all $\alpha \in \{0,1\}^n$ such that $|\alpha| = 4$.

In order to lower bound $\|h\|_{\text{cb}}$, let $T_h \in \mathbb{R}^{(2n) \times (2n) \times (2n) \times (2n)}$ be the symmetric 4-tensor associated to h . By Proposition (4.4), we have

$$\|h\|_{\text{cb}} \geq \left\| \sum_{i,j,k,\ell=1}^{2n} (T_h)_{i,j,k,\ell} B_i B_j B_k B_\ell \right\|,$$

for every set of pairwise commuting contractions B_1, \dots, B_{2n} . In particular, set $B_i = A_i$ as in Corollary (4.5) for $i \in [n]$ and let B_i be the all-zero matrix for $i \in \{n+1, \dots, 2n\}$. Since the A_i s were pairwise commuting in Corollary (4.5) (which ofcourse commute with the all-zero matrix), the B_i s are pairwise commuting. Finally, observe that for all $i, j, k, \ell \in [n]$, we have $(T_h)_{i,j,k,\ell} =$

$d'_{\alpha,0}/(|\{i,j,k,\ell\}|!)$, which is equal to $Cd_\alpha/(|\{i,j,k,\ell\}|!)$ (by Eq. (14)). In particular, using Corollary (4.5), we have

$$\|h\|_{\text{cb}} \geq \left\| \sum_{i,j,k,\ell=1}^{2n} (T_h)_{i,j,k,\ell} B_i B_j B_k B_\ell \right\| = C \left\| \sum_{i,j,k,\ell=1}^n (T_q)_{i,j,k,\ell} A_i A_j A_k A_\ell \right\| \geq Cc\sqrt{n}.$$

This implies that $1 \geq \|h\|_{\text{cb}} = C\|q\|_{\text{cb}} \geq Cc\sqrt{n}$, and so $C \leq 1/(c\sqrt{n})$.

5 Short proof of Theorem 1.1.

In this section, we give a short proof of Theorem 1.1, restated below for convenience.

Theorem 1.1 (Aaronson et al.). *There exists an absolute constant $C \in (0,1]$ such that the following holds. For any bounded quadratic polynomial p , there exists a one-query quantum algorithm that, on input $x \in \{-1,1\}^n$, returns a random sign with expectation $Cp(x)$.*

We begin by giving a brief sketch of the original proof.

Proof sketch of Theorem 1.1. The first step is to show that without loss of generality, we may assume that the polynomial p is a quadratic form. This is the content of the decoupling argument mentioned in the introduction, proved for polynomials of arbitrary degree in [AAI⁺16], but stated here only for the quadratic case.

Lemma 5.1. *There exists an absolute constant $C \in (0,1]$ such that the following holds. For any bounded quadratic polynomial p , there exists a matrix $A \in \mathbb{R}^{(n+1) \times (n+1)}$ with $\|A\|_{\ell_\infty \rightarrow \ell_1} \leq 1$, such that the quadratic form $q(y) = y^\top A y$ satisfies $q((x,1)) = Cp(x)$ for all $x \in \{-1,1\}^n$.*

To prove the theorem, we may thus restrict to a quadratic form $p(x) = x^\top A x$ given by some matrix $A \in \mathbb{R}^{n \times n}$ such that $\|A\|_{\ell_\infty \rightarrow \ell_1} \leq 1$. The next step is to massage the matrix A into a unitary matrix (that can be applied by a quantum algorithm). To obtain this unitary, the authors use an argument based on two versions of Grothendieck's inequality and a technique known as *variable splitting*, developed in earlier work of Aaronson and Ambainis [AA15]. The first version of Grothendieck's inequality is the one most commonly used in applications [Gro53].

Theorem 5.2 (Grothendieck). *There exists a universal constant $K_G \in (0,\infty)$ such that the following holds. For every positive integer n and matrix $A \in \mathbb{R}^{n \times n}$, we have*

$$\sup \left\{ \sum_{i,j=1}^n A_{ij} \langle u_i, v_j \rangle : d \in \mathbb{N}, u_i, v_j \in B_2^n \right\} \leq K_G \|A\|_{\ell_\infty \rightarrow \ell_1}.$$

Elementary proofs of this theorem can be found for instance in [AN06]. The *Grothendieck constant* K_G is the smallest real number for which Theorem 5.2 holds true. The problem of determining its exact value, posed in [Gro53], remains open. The best lower and upper bounds $1.6769 \dots \leq K_G < 1.7822 \dots$ were proved by Davie and Reeds [Dav84, Ree91], and Braverman et al. [BMMN13], resp. The second version of Grothendieck's inequality is as follows.

Theorem 5.3 (Grothendieck). *For every positive integer n and matrix $A \in \mathbb{R}^{n \times n}$, there exist $u, v \in (0, 1]^n$ such that $\|u\|_{\ell_2} = \|v\|_{\ell_2} = 1$ and such that the matrix*

$$B = \frac{1}{K_G} \text{Diag}(u)^{-1} A \text{Diag}(v)^{-1} \quad (15)$$

satisfies $\|B\| \leq \|A\|_{\ell_\infty \rightarrow \ell_1}$, where $\text{Diag}(w)$ denotes the square diagonal matrix whose diagonal is w .

Our contribution. The first (standard) version of Grothendieck's inequality (Theorem 5.2) easily implies that any matrix A such that $\|A\|_{\ell_\infty \rightarrow \ell_1} \leq 1$ has completely bounded norm at most K_G . Combing this fact with our Theorem 1.3 and Lemma 5.1, one quickly retrieves Theorem 1.1. However, Theorem 1.3 is based on the rather deep Theorem 3.1. We observe that Theorem 1.1 also follows readily from the much simpler Theorem 5.3 alone (proved below for completeness), after one assumes that p is a quadratic form as above. Indeed, Theorem 5.3 gives unit vectors u, v such that the matrix B as in (15) has (operator) norm at most 1. Unitary matrices have norm exactly 1 and of course represent the type of operation a quantum algorithm can implement. Moreover, since u, v are unit vectors, they represent $(\log n)$ -qubit quantum states. Using the fact that for $w, z \in \mathbb{R}^n$, we have $\text{Diag}(w)z = \text{Diag}(z)w$, we get the following *factorization* formula (not unlike the one of Corollary 3.2, which is of course no coincidence):

$$\frac{x^\top A x}{K_G} = x^\top \text{Diag}(u) B \text{Diag}(v) x = u^\top \text{Diag}(x) B \text{Diag}(x) v. \quad (16)$$

If we assume for the moment that the matrix B actually is unitary, then the right-hand side of (16) suggests the simple one-query quantum algorithm described in Figure 3.




Figure 3: Let U_u, U_v be unitaries that have u, v as their first rows, respectively. The algorithm initializes a $(1 + \log n)$ -qubit register in the all-zero state, transforms this state into the superposition $\frac{1}{\sqrt{2}}(e_0 \otimes u + e_1 \otimes v)$, queries the input x via the unitary $\text{Diag}(x)$ applied to the $(\log n)$ -qubit register, applies a controlled- B , and finishes by measuring the first qubit in the Hadamard basis.

Using (16), we observe that the algorithm returns zero with probability

$$\frac{1}{2} + \frac{1}{2} \langle \text{Diag}(u)x, B \text{Diag}(v)x \rangle = \frac{1}{2} + \frac{x^\top A x}{2K_G},$$

Now, it is clear that the expected value of the measurement result is precisely $p(x)/K_G$, giving Theorem 1.1 with $C = 1/K_G$. In case B is not unitary, one can use the same argument used in the final step of the proof of Theorem 1.3.

5.1 Factorization version of Grothendieck's inequality

For completeness and because of its relevance to Theorem 1.1, we give a proof of Theorem 5.3 here. The proof relies on the standard version of Grothendieck's inequality (Theorem 5.2). In addition, the proof makes use of the following version of the Hahn–Banach theorem [Rud91, Theorem 3.4].

Theorem 5.4 (Hahn–Banach separation theorem). *Let $C, D \subseteq \mathbb{R}^n$ be convex sets and let C be algebraically open. Then the following are equivalent:*

- *The sets C and D are disjoint.*
- *There exists a vector $\lambda \in \mathbb{R}^n$ and a constant $\alpha \in \mathbb{R}$ such that $\langle \lambda, c \rangle < \alpha$ for every $c \in C$ and $\langle \lambda, d \rangle \geq \alpha$ for every $d \in D$.*

Moreover, if C and D are convex cones, we may take $\alpha = 0$.

Proof of Theorem 5.3. Let $M = A / (K_G \|A\|_{\ell_\infty \rightarrow \ell_1})$. By Theorem 5.2 (the standard Grothendieck inequality), we have that

$$\sum_{i,j=1}^n M_{ij} \langle x_i, y_j \rangle \leq 1$$

for all vectors x_i, y_j with Euclidean norm at most 1. Then, for arbitrary vectors x_i, y_j , we have

$$\sum_{i,j=1}^n M_{ij} \langle x_i, y_j \rangle \leq \max_{i,j \in [n]} \|x_i\| \|y_j\| \leq \frac{1}{2} \max_{i,j \in [n]} (\|x_i\|^2 + \|y_j\|^2), \quad (17)$$

where the second inequality is by AM-GM inequality. Define the set $K \subseteq \mathbb{R}^{n \times n}$ by

$$K = \left\{ \left(\|x_i\|^2 + \|y_j\|^2 - 2 \sum_{k,\ell=1}^n M_{k\ell} \langle x_k, y_\ell \rangle \right)_{i,j=1}^n : d \in \mathbb{N}, x_i, y_j \in \mathbb{R}^d \right\}.$$

We claim that K is a convex cone. Observe that for every $t \in \mathbb{R}_+$ and matrix $Q \in K$ given by vectors x_i, y_j , the vectors $x'_i = \sqrt{t}x_i$ and $y'_j = \sqrt{t}y_j$ similarly define tQ , and so K is a cone. We now show K is a convex set. Let $Q, Q' \in K$ be specified by x_i, y_j and x'_i, y'_j respectively. Then, for any $\lambda \in [0, 1]$, the convex combination $\lambda Q + (1 - \lambda)Q'$ also belongs to K , as it can be specified by the vectors $(\sqrt{\lambda}x_i, \sqrt{1 - \lambda}x'_i), (\sqrt{\lambda}y_j, \sqrt{1 - \lambda}y'_j)$.

Additionally, it follows from Eq. (17) that K is disjoint from the open convex cone $\mathbb{R}_{<0}^{n \times n}$ of matrices with strictly negative entries. By Theorem 5.4 (the Hahn–Banach separation theorem), we conclude that there exists a nonzero matrix $L \in \mathbb{R}^{n \times n}$ such that $\langle L, Q \rangle > 0$ for every $Q \in K$ and $\langle L, N \rangle \leq 0$ for every $N \in \mathbb{R}_{<0}^{n \times n}$. In particular, the second inequality implies that $L \in \mathbb{R}_+^{n \times n}$. Let $P = L / \sum_{ij} L_{ij}$, so that $\{P_{ij}\}_{i,j=1}^n$ defines a probability distribution over $[n]^2$. Then, for any $Q \in K$,

$$\begin{aligned} 0 &\leq \langle P, Q \rangle \\ &= \sum_{i,j=1}^n P_{ij} (\|x_i\|^2 + \|y_j\|^2) - 2 \sum_{k,\ell=1}^n M_{k\ell} \langle x_k, y_\ell \rangle \\ &= \sum_{i=1}^n \sigma_i \|x_i\|^2 + \sum_{j=1}^n \mu_j \|y_j\|^2 - 2 \sum_{k,\ell=1}^n M_{k\ell} \langle x_k, y_\ell \rangle, \end{aligned}$$

where $\sigma_i = P_{i1} + \dots + P_{in}$ and $\mu_j = P_{1j} + \dots + P_{nj}$. Rearranging the inequality above and using bi-linearity, it follows that for every $\lambda > 0$, we have

$$2 \sum_{k,\ell=1}^n M_{k\ell} \langle x_k, y_\ell \rangle = 2 \sum_{k,\ell=1}^n M_{k\ell} \langle \lambda x_k, \lambda^{-1} y_\ell \rangle \leq \lambda^2 \sum_{i=1}^n \sigma_i \|x_i\|^2 + \lambda^{-2} \sum_{j=1}^n \mu_j \|y_j\|^2.$$

Setting

$$\lambda = \left(\frac{\sum_{j=1}^n \mu_j \|y_j\|^2}{\sum_{i=1}^n \sigma_i \|x_i\|^2} \right)^{1/4}$$

in the inequality above, we find that

$$\sum_{k,\ell=1}^n M_{k\ell} \langle x_k, y_\ell \rangle \leq \left(\sum_{i=1}^n \sigma_i \|x_i\|^2 \right)^{1/2} \left(\sum_{j=1}^n \mu_j \|y_j\|^2 \right)^{1/2}.$$

In particular, for the case where $x_k, y_\ell \in \mathbb{R}$, i.e., the scalar case, we have

$$x^T M y \leq \|\text{diag}(\sigma)^{1/2} x\| \|\text{diag}(\mu)^{1/2} y\|.$$

Since the σ_i, μ_j are strictly positive, the inequality above implies

$$x^T \left(\text{Diag}(\sigma)^{-1/2} M \text{Diag}(\mu)^{-1/2} \right) y \leq \|x\| \cdot \|y\|,$$

which in particular implies that $\|\text{Diag}(\sigma)^{-1/2} M \text{Diag}(\mu)^{-1/2}\| \leq 1$. Using the definition of M , we have

$$\|\text{Diag}(\sigma)^{-1/2} A \text{Diag}(\mu)^{-1/2}\| \leq K_G \|A\|_{\ell_\infty \rightarrow \ell_1}.$$

The theorem follows by letting $u_i = \sqrt{\sigma_i}$, $v_i = \sqrt{\mu_i}$ for every $i \in [n]$. □

Acknowledgments.

J.B. thanks Farrokh Labib for useful discussions. S.A. thanks Tom Bannink, Nathaniel Johnston for useful discussions and Ronald de Wolf for helpful discussions and encouragement.

References

- [AA15] S. Aaronson and A. Ambainis. Forrelation: A problem that optimally separates quantum from classical computing. In *Proceedings of 47th ACM STOC*, pages 307–316, 2015. arXiv:1411.5729v1. ^{1,2,17}
- [AAI⁺16] S. Aaronson, A. Ambainis, J. Iraids, M. Kokainis, and J. Smotrovs. Polynomials, quantum query complexity, and Grothendieck’s inequality. In *31st Conference on Computational Complexity, CCC 2016*, pages 25:1–25:19, 2016. arXiv:1511.08682. ^{2,3,5,6,13,17}
- [ABB⁺16] A. Ambainis, K. Balodis, A. Belovs, T. Lee, M. Santha, and J. Smotrovs. Separations in query complexity based on pointer functions. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016*, pages 800–813, 2016. arXiv:1506.04719. ²

- [ABK16] S. Aaronson, S. Ben-David, and R. Kothari. Separations in query complexity using cheat sheets. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC*, pages 863–876, 2016. arXiv:1511.01937. ²
- [Amb02] A. Ambainis. Quantum lower bounds by quantum arguments. *J. Comput. Syst. Sci.*, 64(4):750–767, 2002. Earlier version in STOC’00. arXiv:quant-ph/0002066. ²
- [Amb06] A. Ambainis. Polynomial degree vs. quantum query complexity. *J. Comput. System Sci.*, 72(2):220–238, 2006. Earlier version in FOCS’03. quant-ph/0305028. ²
- [Amb07] A. Ambainis. Quantum walk algorithm for element distinctness. *SIAM J. Comput.*, 37(1):210–239, 2007. Earlier version in FOCS’04. arXiv:quant-ph/0311001. ^{1,2}
- [AN06] N. Alon and A. Naor. Approximating the cut-norm via Grothendieck’s inequality. *SIAM Journal of Computing*, 35(4):787–803, 2006. Earlier version in STOC’04. ^{5,17}
- [Arv12] W. Arveson. *An invitation to C^* -algebras*, volume 39 of *Graduate Texts in Mathematics*. Springer, 2012. ⁷
- [AS04] S. Aaronson and Y. Shi. Quantum lower bounds for the collision and the element distinctness problems. *J. ACM*, 51(4):595–605, 2004. ²
- [BBBV97] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal of Computing*, 26(5):1510–1523, 1997. quant-ph/9701001. ²
- [BBC⁺01] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, 2001. Earlier version in FOCS’98. quant-ph/9802049. ^{2,8}
- [BBLV12] J. Briët, H. Buhrman, T. Lee, and T. Vidick. All Schatten spaces endowed with the Schur product are Q -algebras. *Journal of Functional Analysis*, 262(1):1–9, 2012. ⁶
- [BBLV13] J. Briët, H. Buhrman, T. Lee, and T. Vidick. Multipartite entanglement in XOR games. *Quantum Information & Computation*, 13(3-4):334–360, 2013. arXiv:0911.4007. ⁶
- [Bel12] A. Belovs. Span programs for functions with constant-sized 1-certificates. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012*, pages 77–84, 2012. arXiv:1105.4024. ²
- [BKT17] M. Bun, R. Kothari, and J. Thaler. The polynomial method strikes back: Tight quantum query bounds via dual polynomials. arXiv:1710.09079, 2017. ²
- [BMMN13] M. Braverman, K. Makarychev, Y. Makarychev, and A. Naor. The Grothendieck constant is strictly smaller than Krivine’s bound. *Forum Math. Pi*, 1:453–462, 2013. Preliminary version in FOCS’11. arXiv:1103.6161. ¹⁷
- [BV13] J. Briët and T. Vidick. Explicit lower and upper bounds on the entangled value of multiplayer XOR games. *Communications in Mathematical Physics*, 321(1):181–207, 2013. arXiv:1108.5647. ⁵

- [CHTW04] R. Cleve, P. Høyer, B. Toner, and J. Watrous. Consequences and limits of nonlocal strategies. In *Computational Complexity, 2004. Proceedings. 19th IEEE Annual Conference on*, pages 236–249. IEEE, 2004. arXiv:quant-ph/0404076. ⁵
- [CS87] E. Christensen and A. M. Sinclair. Representations of completely bounded multilinear operators. *Journal of Functional analysis*, 72(1):151–181, 1987. ^{4,9}
- [Dav84] A. Davie. Lower bound for K_G . Unpublished, 1984. ¹⁷
- [DJT95] J. Diestel, H. Jarchow, and A. Tonge. *Absolutely summing operators*, volume 43 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1995. ^{5,13}
- [FGG08] E. Farhi, J. Goldstone, and S. Gutmann. A quantum algorithm for the Hamiltonian NAND tree. *Theory of Computation*, 4(8):169–190, 2008. arXiv:quant-ph/0702144. ¹
- [Gar84] L. T. Gardner. An elementary proof of the Russo-Dye theorem. *Proceedings of the American Mathematical Society*, 90(1):171, 1984. ⁸
- [Gro53] A. Grothendieck. Résumé de la théorie métrique des produits tensoriels topologiques (French). *Bol. Soc. Mat. São Paulo*, 8:1–79, 1953. ^{5,17}
- [Gro96] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219. ACM, 1996. ¹
- [HLŠ07] P. Høyer, T. Lee, and R. Špalek. Negative weights make adversaries stronger. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, 2007*, pages 526–535, 2007. arXiv:quant-ph/0611054. ²
- [JKP09] N. Johnston, D. W. Kribs, and V. I. Paulsen. Computing stabilized norms for quantum operations via the theory of completely bounded maps. *Quantum Information & Computation*, 9(1):16–35, 2009. arXiv:0711.3636. ¹⁰
- [KLW15] J. Kaniewski, T. Lee, and R. de Wolf. Query complexity in expectation. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP*, pages 761–772, 2015. arXiv:1411.7280. ⁶
- [KN12] S. Khot and A. Naor. Grothendieck-type inequalities in combinatorial optimization. *Communications on Pure and Applied Mathematics*, 65(7):992–1035, 2012. arXiv:1108.2464. ⁵
- [LLV15] C. M. Le, E. Levina, and R. Vershynin. Sparse random graphs: regularization and concentration of the Laplacian. 2015. arXiv:1502.03049. ⁵
- [LMR⁺11] T. Lee, R. Mittal, B. W. Reichardt, R. Špalek, and M. Szegedy. Quantum query complexity of state conversion. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011*, pages 344–353, 2011. arXiv:1011.3020. ²
- [MNR11] A. Montanaro, H. Nishimura, and R. Raymond. Unbounded-error quantum query complexity. *Theor. Comput. Sci.*, 412(35):4619–4628, 2011. arXiv:0712.1446. ⁶

- [MNRS11] F. Magniez, A. Nayak, J. Roland, and M. Santha. Search via quantum walk. *SIAM J. Comput.*, 40(1):142–164, 2011. Earlier version in STOC’07. arXiv:quant-ph/0608026. ²
- [Mon10] A. Montanaro. Nonadaptive quantum query complexity. *Inf. Process. Lett.*, 110(24), 2010. arXiv:1001.0018. ⁶
- [Mur14] G. J. Murphy. *C*-algebras and operator theory*. Academic press, 2014. ⁷
- [OP99] T. Oikhberg and G. Pisier. The “maximal” tensor product of operator spaces. *Proceedings of the Edinburgh Mathematical Society*, 42(2):267–284, 1999. ⁴
- [Pau02] V. Paulsen. *Completely bounded maps and operator algebras*, volume 78. Cambridge University Press, Cambridge, 2002. ^{7,10}
- [PGWP⁺08] D. Pérez-García, M. Wolf, C. Palazuelos, I. Villanueva, and M. Junge. Unbounded violation of tripartite Bell inequalities. *Communications in Mathematical Physics*, 279:455, 2008. arXiv:quant-ph/0702189. ^{5,6}
- [Pis03] G. Pisier. *Introduction to operator space theory*, volume 294 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 2003. ^{4,10}
- [Pis12] G. Pisier. Grothendieck’s theorem, past and present. *Bull. Amer. Math. Soc.*, 49(2):237–323, 2012. also available at arXiv:1101.4195. ⁵
- [Pol12] D. Pollard. *Convergence of stochastic processes*. Science & Business Media. Springer, 2012. ¹⁴
- [PS87] V. I. Paulsen and R. R. Smith. Multilinear maps and tensor norms on operator systems. *Journal of functional analysis*, 73(2):258–276, 1987. ⁴
- [PV16] C. Palazuelos and T. Vidick. Survey on nonlocal games and operator space theory. *Journal of Mathematical Physics*, 57(1):015220, 2016. ⁶
- [Ree91] J. Reeds. A new lower bound on the real Grothendieck constant. Manuscript (<http://www.dtc.umn.edu/~reedsj/bound2.dvi>), 1991. ¹⁷
- [Rei09] B. Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every boolean function. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009*, pages 544–551, 2009. arXiv:0904.2759. ²
- [Rei11] B. Reichardt. Reflections for quantum query algorithms. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011*, pages 560–569, 2011. arXiv:1005.1601. ²
- [Rud91] W. Rudin. *Functional analysis*. McGraw-Hill Science, 1991. ¹⁹
- [Sho97] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal of Computing*, 26(5):1484–1509, 1997. Earlier version in FOCS’94. ¹
- [Sim97] D. Simon. On the power of quantum computation. *Siam journal of computing*, 26(5):1474–1483, 1997. Earlier version in FOCS’94. ¹

- [Smi88] R.R. Smith. Completely bounded multilinear maps and Grothendieck's inequality. *Bulletin of the London Mathematical Society*, 20(6):606–612, 1988. ⁵
- [Tao12] T. Tao. *Topics in random matrix theory*, volume 132. American Mathematical Society, 2012. ¹⁴
- [Tro09] J. A. Tropp. Column subset selection, matrix factorization, and eigenvalue optimization. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 978–986, 2009. ⁵
- [Tsi87] B. S. Tsirelson. Quantum analogues of the Bell inequalities. The case of two spatially separated domains. *J. Soviet Math.*, 36:557–570, 1987. ⁵
- [Wol03] R. de Wolf. Nondeterministic quantum query and communication complexities. *SIAM J. Comput.*, 32(3):681–699, 2003. cs.CC/0001014. ⁶