

VANISHING-ERROR APPROXIMATE DEGREE AND QMA COMPLEXITY

ALEXANDER A. SHERSTOV* AND JUSTIN THALER†

ABSTRACT. The ϵ -approximate degree of a function $f: X \rightarrow \{0, 1\}$ is the least degree of a multivariate real polynomial p such that $|p(x) - f(x)| \leq \epsilon$ for all $x \in X$. We determine the ϵ -approximate degree of the element distinctness function, the surjectivity function, and the permutation testing problem, showing they are $\Theta(n^{2/3} \log^{1/3}(1/\epsilon))$, $\tilde{\Theta}(n^{3/4} \log^{1/4}(1/\epsilon))$, and $\Theta(n^{1/3} \log^{2/3}(1/\epsilon))$, respectively. Previously, these bounds were known only for constant ϵ .

We also derive a connection between vanishing-error approximate degree and quantum Merlin–Arthur (QMA) query complexity. We use this connection to show that the QMA complexity of permutation testing is $\Omega(n^{1/4})$. This improves on the previous best lower bound of $\Omega(n^{1/6})$ due to Aaronson (*Quantum Information & Computation*, 2012), and comes somewhat close to matching a known upper bound of $O(n^{1/3})$.

1. INTRODUCTION

The ϵ -approximate degree of a function $f: X \rightarrow \{0, 1\}$, denoted $\deg_\epsilon(f)$, is the least degree of a multivariate real-valued polynomial p such that $|p(x) - f(x)| \leq \epsilon$ for all inputs $x \in X$. Lower bounds on approximate degree have many applications in theoretical computer science, ranging from quantum query and communication lower bounds, to oracle separations and cryptographic secret sharing schemes. Upper bounds on approximate degree have important algorithmic implications in learning theory and differential privacy, and underlie state-of-the-art circuit and formula size lower bounds. The interested reader can find a bibliographic overview of these applications in [BKT18, She18].

This paper focuses on three well-studied functions whose approximation by polynomials has applications to quantum computing and beyond. The first function is *element distinctness* ED_n , where the input is a list of n numbers from $\{1, 2, \dots, n\}$ and the objective is to determine if the numbers are pairwise distinct. The second function is *surjectivity* $\text{SURJ}_{n,r}$, where the input is a list of n numbers from the range $\{1, 2, \dots, r\}$ and the goal is to check whether every range element appears on the list. The canonical setting is $r = \lfloor cn \rfloor$ for some constant $0 < c < 1$. The third problem that we study is *permutation testing* $\text{PTP}_{n,\alpha}$, parameterized by a constant $0 < \alpha < 1$. Here, the input is a list of n numbers from $\{1, 2, \dots, n\}$, and the objective is to distinguish the case when the list contains every range element from the case when the list contains at most αn range elements. In the context of polynomial approximation, it is customary to represent the input to these functions

* Computer Science Department, UCLA, Los Angeles, CA 90095. Supported by NSF grant CCF-1814947. Email: sherstov@cs.ucla.edu

† Department of Computer Science, Georgetown University, Washington, DC 20057. Supported by NSF CAREER award CCF-1845125. Email: justin.thaler@georgetown.edu.

as a Boolean matrix $x = [x_{i,j}]$, where $x_{i,j} = 1$ if and only if the i th element on the list equals j .

Vanishing-error approximate degree. Much work in the area has focused on *bounded-error* approximate degree, defined for a Boolean function f as the quantity $\deg_{1/3}(f)$. The choice of constant $1/3$ here is arbitrary, as $\deg_\epsilon(f) = \Theta(\deg_{1/3}(f))$ for all constants $0 < \epsilon < 1/2$. In particular, the bounded-error approximate degrees of element distinctness, surjectivity, and permutation testing are known to be $\Theta(n^{2/3})$, $\tilde{\Theta}(n^{3/4})$, and $\Theta(n^{1/3})$, respectively [AS04, Amb05, Kut05, Aar12, She18, BKT18]. Our understanding of approximate degree with vanishing error, $\epsilon = o(1)$, is far less complete. Among the very few functions whose vanishing-error approximate degree has been determined is the n -bit AND function, with the asymptotic bound $\deg_\epsilon(\text{AND}_n) = \Theta(n^{1/2} \log^{1/2}(1/\epsilon))$ due to Buhrman et al. [BCWZ99]. We give a new and entirely different proof of their result. Our technique further allows us to settle the vanishing-error approximate degrees of the much more complicated functions of element distinctness, surjectivity, and permutation testing:

Theorem 1.1. *Let $0 < c < 1$ and $0 < \alpha < 1$ be arbitrary constants. Then*

$$\begin{aligned} \deg_\epsilon(\text{ED}_n) &= \Omega\left(n^{2/3} \left(\log \frac{1}{\epsilon}\right)^{1/3}\right), \\ \deg_\epsilon(\text{SURJ}_{n, \lfloor cn \rfloor}) &= \tilde{\Omega}\left(n^{3/4} \left(\log \frac{1}{\epsilon}\right)^{1/4}\right), \\ \deg_\epsilon(\text{PTP}_{n, \alpha}) &= \Omega\left(n^{1/3} \left(\log \frac{1}{\epsilon}\right)^{2/3}\right) \end{aligned}$$

for all $1/3^n \leq \epsilon \leq 1/3$.

This theorem is optimal with respect to all parameters. The lower bounds for element distinctness and surjectivity match the vanishing-error constructions in [She18], whereas the lower bound for permutation testing is tight by a quantum query argument which we include as Theorem 3.8. A comment is in order on ϵ -approximate degree in the complementary range, $\epsilon < 1/3^n$. Routine interpolation gives an exact representation for each of the functions in Theorem 1.1 as a polynomial of degree at most n . Theorem 1.1 shows that this upper bound is asymptotically tight, settling the ϵ -approximate degree for $\epsilon < 1/3^n$ as well.

We prove a result analogous to Theorem 1.1 for k -element distinctness ED_n^k , a well-studied generalization of ED_n . Specifically, we prove that if ED_n^k has bounded-error approximate degree $\Omega(n^\ell)$, then it has ϵ -approximate degree $\Omega(n^\ell \log^{1-\ell}(1/\epsilon))$. The best known lower bound on the bounded-error approximate degree of ED_n^k is $\tilde{\Omega}(n^{3/4-1/(2k)})$ [BKT18], so this yields

$$\deg_\epsilon(\text{ED}_n^k) = \tilde{\Omega}\left(n^{\frac{3}{4}-\frac{1}{2k}} \left(\log \frac{1}{\epsilon}\right)^{\frac{1}{4}+\frac{1}{2k}}\right).$$

For large k , this comes close to the best known upper bound [She18]:

$$\deg_\epsilon(\text{ED}_n^k) = O\left(n^{\frac{3}{4} - \frac{1}{4(2^k-1)}} \left(\log \frac{1}{\epsilon}\right)^{\frac{1}{4} + \frac{1}{4(2^k-1)}}\right).$$

Our techniques are quite general, and we are confident that they will find other applications in the area. The technical core of our results establishes that for any function f_n that contains $\text{AND}_k \circ f_{\lfloor n/k \rfloor}$ as a subfunction for each $k \leq n$, any bounded-error approximate degree lower bound for f_n automatically implies a strong lower bound for the ϵ -approximate degree of f_n . This allows us to prove tight lower bounds on the vanishing-error approximate degrees of AND_n , ED_n , ED_n^k , and $\text{SURJ}_{n,r}$. To handle $\text{PTP}_{n,\alpha}$, we generalize our technique to other outer functions. Our analysis is based on the so-called method of dual polynomials, whereby one proves approximate degree lower bounds by constructing explicit dual solutions to a certain linear program capturing the approximate degree of the given function.

In the remainder of the introduction, we focus on an application of Theorem 1.1 to quantum Merlin–Arthur complexity.

The Merlin–Arthur model. The Merlin–Arthur (MA) model of query complexity features a function f and two asymmetric players, Merlin and Arthur. Arthur’s goal is to compute f on some unknown input x while querying as few bits of x as possible. Merlin, who knows x , can help Arthur compute $f(x)$ by sending him a single witness, i.e., an arbitrary message of some bit length m . However, Merlin is untrusted. The model requires that, for any $x \in f^{-1}(1)$, there is some Merlin message causing Arthur to output 1 with probability at least $2/3$, and for any $x \in f^{-1}(0)$, no Merlin message can cause Arthur to output 1 with probability more than $1/3$. The cost of the protocol is the sum of the witness length m and the number of bits of x queried by Arthur. In *quantum* Merlin–Arthur (QMA) query complexity, the witness sent by Merlin is allowed to be an arbitrary m -qubit quantum message, and Arthur is permitted to query bits of the input x in superposition. The MA and QMA query models have important analogues in communication complexity and Turing machine complexity. In the former setting, Arthur is replaced by two parties Alice and Bob, and the input x is split between them.

The complexity class QMA is a quantum analog of NP and accordingly has received considerable attention. It is well known that any QMA protocol can be simulated by an SBQP \subseteq PP protocol with at most a quadratic blowup in cost, i.e., $\text{QMA}(f) \geq \Omega(\text{SBQP}(f)^{1/2})$ [Vya03].¹ In turn, the existence of an SBQP query protocol that makes at most c queries implies that the *one-sided* approximate degree of f is at most $O(c)$. Here, the one-sided ϵ -approximate degree of f is the least degree of a real polynomial p such that $|p(x)| \leq \epsilon$ for all $x \in f^{-1}(0)$, and $p(x) \geq 1 - \epsilon$ for all $x \in f^{-1}(1)$ [BT15]. As a consequence, one can prove QMA query lower bounds for f by lower bounding the one-sided approximate degree of f .

Only a handful of additional results are known about QMA query and communication complexity. Raz and Shpilka [RS04] showed that AND_n has QMA query

¹An SBQP protocol \mathcal{A} is a quantum protocol for which there is some α such that \mathcal{A} accepts every input in $f^{-1}(1)$ with probability at least α , and every input in $f^{-1}(0)$ with probability at most $\alpha/2$ [Kup15].

complexity $\Theta(\sqrt{n})$. Klauck [Kla11] showed that the QMA communication complexity of the disjointness problem is $\Omega(n^{1/3})$. Neither of these results follows from a naïve application of the bound $\text{QMA}(f) \geq \Omega(\sqrt{\text{SBQP}(f)})$.

QMA complexity of permutation testing. The permutation testing problem $\text{PTP}_{n,\alpha}$ has played an important role in the study of interactive proof systems because it possesses a simple non-interactive perfect zero knowledge (NIPZK) protocol of logarithmic cost, yet is a hard problem in many other models. Hence, it has been used to prove a variety of complexity class separations. In particular, Aaronson [Aar12] showed that the QMA query complexity of $\text{PTP}_{n,\alpha}$ is $\Omega(n^{1/6})$, and thereby gave an oracle separating NIPZK from QMA. Bouland et al. [BCH⁺17] built on Aaronson’s result to give an oracle separating non-interactive *statistical* zero knowledge (NISZK) from the complexity class UPP, answering a question of Watrous from 2002. Gur, Liu, and Rothblum [GLR18] showed that the MA query complexity of $\text{PTP}_{n,\alpha}$ is $\Omega(n^{1/4})$. Despite this progress, the precise QMA complexity of $\text{PTP}_{n,\alpha}$ has remained open, with the best upper bound being $O(n^{1/3})$ [BHT16, Aar12] and the best lower bound being Aaronson’s $\Omega(n^{1/6})$. We obtain a polynomially stronger lower bound.

Theorem 1.2. *Let $0 < \alpha < 1$ be an arbitrary constant. Then any QMA query protocol for $\text{PTP}_{n,\alpha}$ with witness length m has query cost $\Omega(n/m)^{1/3}$. In particular, $\text{PTP}_{n,\alpha}$ has QMA complexity $\Omega(n^{1/4})$.*

This result quantitatively matches the MA lower bound of Gur et al. [GLR18] but holds in the more powerful quantum setting. Theorem 1.2 comes reasonably close to matching the known QMA query upper bound of $O(n^{1/3})$, which holds even if Merlin does not send any message to Arthur; see Theorem 3.8.

To prove Theorem 1.2, we derive a connection between QMA query complexity and vanishing-error approximate degree for a class of functions that includes AND_n , ED_n , and $\text{PTP}_{n,\alpha}$. This connection amounts to the observation that, for these particular functions, the one-sided ϵ -approximate degree is *equal* to the ϵ -approximate degree. Prior work on QMA complexity (e.g., [Kla11]) has implicitly exploited a similar observation in the special case of AND_n . Our analysis substantially generalizes the insights of prior work, and makes explicit the key phenomenon at play, namely the equivalence of one-sided vs. standard approximate degree for these functions. Combining this connection with our new vanishing-error approximate degree lower bounds in Theorem 1.1 establishes Theorem 1.2.

2. PRELIMINARIES

For a function f , we let $\text{dom } f$ and $\text{im } f$ stand for the domain and image of f , respectively. We view Boolean functions as mappings $f: X \rightarrow \{0, 1\}$ for a finite set X . For functions $g: X \rightarrow Y$ and $f: Y^n \rightarrow Z$, we let $f \circ g$ denote the block-composition of f and g . In more detail, $f \circ g: X^n \rightarrow Z$ is the function that maps $(x_1, \dots, x_n) \in X^n$ to $f(g(x_1), \dots, g(x_n))$. We generalize block-composition to the case when the domain of f is properly contained in Y^n by defining the domain of $f \circ g$ as the set of $(x_1, \dots, x_n) \in X^n$ such that $(g(x_1), \dots, g(x_n)) \in \text{dom } f$.

2.1. Polynomial approximation. For a multivariate real polynomial $p: \mathbb{R}^n \rightarrow \mathbb{R}$, we let $\text{deg } p$ denote the total degree of p , i.e., the largest degree of any monomial of p . It will be convenient to define the degree of the zero polynomial by $\text{deg } 0 = -\infty$. For

a real-valued function ϕ supported on a finite subset of \mathbb{R}^n , we define the *orthogonal content* of ϕ , denoted $\text{orth } \phi$, to be the minimum degree of a real polynomial p for which $\langle \phi, p \rangle \neq 0$. We adopt the convention that $\text{orth } \phi = \infty$ if no such polynomial exists. For two functions $f, \psi: X \rightarrow \mathbb{R}$, let $\langle f, \psi \rangle = \sum_{x \in X} f(x)\psi(x)$ denote the correlation of f and ψ , and let $\|\psi\|_1 = \sum_{x \in X} |\psi(x)|$. For any real-valued function $\psi: X \rightarrow \mathbb{R}$, its k -th tensor power $\psi^{\otimes k}: X^k \rightarrow \mathbb{R}$ is given by $\psi^{\otimes k}(x_1, \dots, x_k) = \psi(x_1) \cdots \psi(x_k)$.

The ϵ -approximate degree of a function $f: X \rightarrow \mathbb{R}$, denoted $\text{deg}_\epsilon(f)$, is the least degree of a polynomial $p: X \rightarrow \mathbb{R}$ such that $|p(x) - f(x)| \leq \epsilon$ for all $x \in X$. We emphasize that no restriction is placed on the behavior of p at inputs outside f 's domain of definition, X . For most functions of interest to us, the domain X is a proper subset of $\{0, 1\}^n$ and thus their approximating polynomials may take on arbitrary values on $\{0, 1\}^n \setminus X$. The following dual characterization of approximate degree is well known and can be verified using linear programming duality.

Fact 2.1. *Fix $d > 0$ and a function $f: X \rightarrow \mathbb{R}$. Then $\text{deg}_\epsilon(f) \geq d$ if and only if there exists a function $\psi: X \rightarrow \mathbb{R}$ such that*

$$\begin{aligned} \langle f, \psi \rangle &> \epsilon \|\psi\|_1, \\ \text{orth } \psi &\geq d. \end{aligned}$$

The simplest function of interest to us is $\text{AND}_n: \{0, 1\}^n \rightarrow \{0, 1\}$, given as usual by $\text{AND}_n(x) = \bigwedge_{i=1}^n x_i$. Its bounded-error approximate degree was determined by Nisan and Szegedy [NS94], as follows.

Theorem 2.2. *For all $n \geq 1$,*

$$\text{deg}_{1/3}(\text{AND}_n) = \Theta(\sqrt{n}).$$

2.2. Surjectivity. Let $\mathcal{D}_{n,r}$ stand for the set of Boolean matrices of size $n \times r$ in which every row has exactly one 1. Every matrix $x \in \mathcal{D}_{n,r}$ has a natural interpretation as specifying a mapping $\phi: \{1, 2, \dots, n\} \rightarrow \{1, 2, \dots, r\}$, where $\phi(i) = j$ if and only if $x_{i,j} = 1$. Our next three functions are defined on $\mathcal{D}_{n,r}$ and can thus be regarded as ‘‘function properties.’’ To start with, the *surjectivity problem* with n elements and range size r is defined as $\text{SURJ}_{n,r}: \mathcal{D}_{n,r} \rightarrow \{0, 1\}$, where

$$\text{SURJ}_{n,r}(x) = \bigwedge_{j=1}^r \bigvee_{i=1}^n x_{i,j}.$$

Thus, $\text{SURJ}_{n,r}$ takes as input an $n \times r$ Boolean matrix in which every row contains exactly one 1, and evaluates to 1 if and only if every column of the input contains at least one 1. Interpreting the input matrix as a mapping, $\text{SURJ}_{n,r}$ evaluates to 1 if and only if that mapping is surjective. This surjectivity property is trivially false for $r > n$, and the standard setting of parameters is $r = \lfloor cn \rfloor$ for some constant $0 < c < 1$. The choice of constant c is unimportant because it affects $\text{deg}_{1/3}(\text{SURJ}_{n, \lfloor cn \rfloor})$ by at most a multiplicative constant. It was shown in [She18] that the surjectivity function has bounded-error approximate degree $O(n^{3/4})$. Bun et al. [BKT18] gave an alternate proof of this upper bound and additionally proved that it is tight up to a polylogarithmic factor. We thus have:

Theorem 2.3. *Let $0 < c < 1$ be an arbitrary constant. Then*

$$\deg_{1/3}(\text{SURJ}_{n, \lfloor cn \rfloor}) = \tilde{\Theta}(n^{3/4}).$$

2.3. Element distinctness. Another well-studied function is *element distinctness* $\text{ED}_{n,r}: \mathcal{D}_{n,r} \rightarrow \{0,1\}$, defined by $\text{ED}_{n,r}(x) = 1$ if and only if every column of the input matrix x has at most one 1. Switching to the interpretation of x as a mapping, $\text{ED}_{n,r}(x)$ evaluates to true if and only if the mapping is one-to-one. This property is trivially false for $r < n$. In the complementary case, Ambainis [Amb05] proved that for any given ϵ , the ϵ -approximate degree of $\text{ED}_{n,r}$ is the same for all $r \geq n$. This means that one may without loss of generality focus on the special case $r = n$, with the shorthand notation $\text{ED}_n = \text{ED}_{n,n}$. Aaronson and Shi [AS04], Ambainis [Amb05], and Kutin [Kut05] showed that ED_n has bounded-error approximate degree $\Theta(n^{2/3})$.

Theorem 2.4. *For all $n \geq 1$,*

$$\deg_{1/3}(\text{ED}_n) = \Theta(n^{2/3}).$$

Element distinctness generalizes in a natural way to a function called *k-element distinctness*, denoted $\text{ED}_n^k: \mathcal{D}_{n,n} \rightarrow \{0,1\}$. This new function evaluates to true if and only if the input matrix has no column with k or more 1s. Viewing the input as a mapping, ED_n^k evaluates to true if and only if no range element occurs k or more times. With these definitions, we have $\text{ED}_n = \text{ED}_n^2$.

2.4. Permutation testing. The final problem of interest to us is a restriction of element distinctness ED_n . In more detail, fix an integer $n \geq 1$ and a real number $0 < \alpha < 1$. The domain of the *permutation testing problem* $\text{PTP}_{n,\alpha}$ is the set of all matrices $x \in \mathcal{D}_{n,n}$ in which the number of columns containing a 1 is either exactly n or at most αn . The function evaluates to true in the former case and to false in the latter. Equivalently, $\text{PTP}_{n,\alpha}(x) = 1$ if and only if x is a permutation matrix. In the regime of interest to us, $0 < \alpha < 1$ is a constant independent of n .

The permutation testing problem was introduced by Aaronson [Aar12], who defined it somewhat differently. In his variant of permutation testing, which we denote by $\text{PTP}_{n,\delta}^*$, one is given a matrix $x \in \mathcal{D}_{n,n}$ that is either (i) a permutation matrix, or (ii) disagrees from every permutation matrix in at least δn rows. The function evaluates to true in case (i) and to false in case (ii). As the following proposition shows, Aaronson's $\text{PTP}_{n,\delta}^*$ is precisely the same function as our $\text{PTP}_{n,1-\delta}$.

Proposition 2.5. *Let $0 < \delta < 1$ and $n \geq 1$ be given. Then as functions,*

$$\text{PTP}_{n,\delta}^* = \text{PTP}_{n,1-\delta}.$$

Specifically, the l.h.s. and r.h.s. have the same domain and agree at every point thereof.

Proof. This claim is easiest to verify by interpreting an input $x \in \mathcal{D}_{n,n}$ as a mapping $\phi: \{1, 2, \dots, n\} \rightarrow \{1, 2, \dots, n\}$. A moment's reflection shows that ϕ disagrees from every permutation $\{1, 2, \dots, n\} \rightarrow \{1, 2, \dots, n\}$ in at least $n - |\text{im } \phi|$ points, and there is a permutation that achieves this lower bound. Restating this in matrix terminology, a matrix $x \in \mathcal{D}_{n,n}$ disagrees from every permutation matrix in at

least δn rows if and only if the number of columns of x containing a 1 is at most $n - \delta n$. \square

By adapting earlier analyses of element distinctness, Aaronson [Aar12] obtained the following result.

Theorem 2.6. *Let $0 < \delta < 1$ be an arbitrary constant. Then*

$$\deg_{1/3}(\text{PTP}_{n,\delta}^*) = \Omega(n^{1/3}).$$

This result is stated in [Aar12] specifically for $\delta = 1/8$, but the proof actually allows any $0 < \delta < 1$. Combining this theorem with Proposition 2.5 gives the following corollary.

Corollary 2.7. *Let $0 < \alpha < 1$ be an arbitrary constant. Then*

$$\deg_{1/3}(\text{PTP}_{n,\alpha}) = \Omega(n^{1/3}).$$

We close this section with a remark on input encoding. In this work, functions like $\text{SURJ}_{n,r}$ take as input a Boolean matrix x in which every row has exactly one 1. Some other works [BM12, BKT18] represent the input as a list $y_1, \dots, y_n \in \{0, 1\}^{\lceil \log r \rceil}$, where y_i encodes the location of the unique 1 in the i -th row of the matrix representation x . Switching to this alternate representation affects the approximate degree by at most a logarithmic factor. See [She18] for a detailed treatment of the relationship between these representations.

3. APPROXIMATE DEGREE LOWER BOUNDS

In this section, we study the vanishing-error approximate degree of element distinctness, surjectivity, and permutation testing, and in particular settle Theorem 1.1 from the introduction. The core of our technique is the following auxiliary result.

Proposition 3.1. *For any $\epsilon \geq 0$ and any function $f: X \rightarrow \mathbb{R}$ on a finite subset X of Euclidean space,*

$$\deg_{\epsilon^k}(f^{\otimes k}) \geq k \deg_{\epsilon}(f), \quad k = 1, 2, 3, \dots$$

In particular, every function $f: X \rightarrow \{0, 1\}$ satisfies

$$\deg_{\epsilon^k}(\text{AND}_k \circ f) \geq k \deg_{\epsilon}(f), \quad k = 1, 2, 3, \dots$$

Proof. We may assume that $\deg_{\epsilon}(f) \neq 0$ since the proposition is trivial otherwise. Let ψ be an ϵ -error dual polynomial for f , as guaranteed by Fact 2.1:

$$\begin{aligned} \langle f, \psi \rangle &> \epsilon \|\psi\|_1, \\ \text{orth } \psi &= \deg_{\epsilon}(f). \end{aligned}$$

Then

$$\begin{aligned} \langle f^{\otimes k}, \psi^{\otimes k} \rangle &= \langle f, \psi \rangle^k \\ &> (\epsilon \|\psi\|_1)^k \\ &= \epsilon^k \|\psi^{\otimes k}\|_1. \end{aligned}$$

Applying Fact 2.1 once again,

$$\begin{aligned} \deg_{\epsilon^k}(f^{\otimes k}) &\geq \text{orth } \psi^{\otimes k} \\ &= k \text{ orth } \psi \\ &= k \deg_{\epsilon}(f). \end{aligned} \quad \square$$

The proof of Proposition 3.1 applies more generally to the conjunction of k distinct functions, but we will not need this generalization.

3.1. Warmup. To illustrate our technique in the simplest possible setting, we consider the well-studied AND_n function. Buhrman et al. [BCWZ99] proved that its ϵ -error approximate degree is $\Theta(\sqrt{n \log(1/\epsilon)})$. We give a new and simple proof of their lower bound.

Theorem 3.2. *For all $1/3^n \leq \epsilon \leq 1/3$,*

$$\deg_{\epsilon}(\text{AND}_n) = \Omega\left(\sqrt{n \log \frac{1}{\epsilon}}\right). \quad (3.1)$$

Proof. For $k = 1, 2, \dots, n$, we have

$$\begin{aligned} \deg_{3^{-k}}(\text{AND}_n) &\geq \deg_{3^{-k}}(\text{AND}_k \circ \text{AND}_{\lfloor n/k \rfloor}) \\ &\geq k \deg_{1/3}(\text{AND}_{\lfloor n/k \rfloor}) \\ &= k \cdot \Omega\left(\sqrt{\frac{n}{k}}\right) \\ &= \Omega(\sqrt{nk}), \end{aligned}$$

where the first, second, and third steps use the identity $\text{AND}_{n_1 n_2} = \text{AND}_{n_1} \circ \text{AND}_{n_2}$, Proposition 3.1, and Theorem 2.2, respectively. This directly implies (3.1). \square

3.2. Element distinctness. Our next result is a tight lower bound on the vanishing error approximate degree of element distinctness, matching the upper bound from [She18].

Theorem 3.3. *For all $1/3^n \leq \epsilon \leq 1/3$,*

$$\deg_{\epsilon}(\text{ED}_n) = \Omega\left(n^{2/3} \left(\log \frac{1}{\epsilon}\right)^{1/3}\right). \quad (3.2)$$

Proof. For any $k = 1, 2, 3, \dots, n$, we claim that $\text{AND}_k \circ \text{ED}_{\lfloor n/k \rfloor}$ is a subproblem of ED_n . To see why, recall that the input to ED_n is an $n \times n$ Boolean matrix in which every row i contains exactly one 1, corresponding to the value of the i th element. Now, fix $k \in \{1, 2, \dots, n\}$ and consider the restriction of ED_n to input matrices that are *block-diagonal*, with k blocks of size $\lfloor n/k \rfloor$ each and an additional block of $n - k\lfloor n/k \rfloor$ ones on the diagonal. Each of the first k blocks corresponds to an instance of $\text{ED}_{\lfloor n/k \rfloor}$, and the overall problem amounts to computing the AND of these k instances. Therefore, $\text{AND}_k \circ \text{ED}_{\lfloor n/k \rfloor}$ is a subproblem of ED_n , and

$$\deg_{\epsilon}(\text{ED}_n) \geq \deg_{\epsilon}(\text{AND}_k \circ \text{ED}_{\lfloor n/k \rfloor}) \quad (3.3)$$

for all ϵ and all $k = 1, 2, 3, \dots, n$.

The rest of the proof is closely analogous to that for AND_n . For $k = 1, 2, \dots, n$,

$$\begin{aligned} \deg_{3^{-k}}(\text{ED}_n) &\geq \deg_{3^{-k}}(\text{AND}_k \circ \text{ED}_{\lfloor n/k \rfloor}) \\ &\geq k \deg_{1/3}(\text{ED}_{\lfloor n/k \rfloor}) \\ &\geq k \cdot \Omega\left(\frac{n}{k}\right)^{2/3} \\ &= \Omega(n^{2/3} k^{1/3}) \end{aligned}$$

where the first three steps use (3.3), Proposition 3.1, and Theorem 2.4, respectively. This directly implies (3.2). \square

The previous proof shows more generally that $\text{AND}_k \circ \text{ED}_{\lfloor n/k \rfloor}^r$ is a subfunction of ED_n^r for any $k = 1, 2, \dots, n$. As a result, our analysis of element distinctness proves the following statement.

Theorem 3.4. *Fix constants $r \geq 2$ and $\ell \in [0, 1]$ such that*

$$\deg_{1/3}(\text{ED}_n^r) = \Omega(n^\ell).$$

Then

$$\deg_\epsilon(\text{ED}_n^r) = \Omega\left(n^\ell \left(\log \frac{1}{\epsilon}\right)^{1-\ell}\right), \quad \frac{1}{3^n} \leq \epsilon \leq \frac{1}{3}.$$

Combining Theorem 3.4 with the known lower bound

$$\deg_{1/3}(\text{ED}_n^r) = \tilde{\Omega}\left(n^{\frac{3}{4} - \frac{1}{2r}}\right)$$

due to [BKT18], we conclude that

$$\deg_\epsilon(\text{ED}_n^r) = \tilde{\Omega}\left(n^{\frac{3}{4} - \frac{1}{2r}} \left(\log \frac{1}{\epsilon}\right)^{\frac{1}{4} + \frac{1}{2r}}\right)$$

for $1/3^n \leq \epsilon \leq 1/3$. Moreover, Theorem 3.4 will, in a black-box manner, translate any future improvement in the bounded-error lower bound for ED_n^r into an improved vanishing-error lower bound.

3.3. Surjectivity. An instance x of the surjectivity problem $\text{SURJ}_{n,r}$ can be embedded inside a larger instance of surjectivity in many ways, e.g., by duplicating a row of x or by forming a block-diagonal matrix with blocks x and 1. These two transformations yield

$$\deg_\epsilon(\text{SURJ}_{n,r}) \leq \deg_\epsilon(\text{SURJ}_{n+1,r}), \quad (3.4)$$

$$\deg_\epsilon(\text{SURJ}_{n,r}) \leq \deg_\epsilon(\text{SURJ}_{n+1,r+1}), \quad (3.5)$$

respectively. We will now prove an essentially tight lower bound on the vanishing-error approximate degree of surjectivity, matching the upper bound from [She18] up to a logarithmic factor.

Theorem 3.5. *Let $0 < c < 1$ be an arbitrary constant. Then*

$$\deg_\epsilon(\text{SURJ}_{n, \lfloor cn \rfloor}) = \tilde{\Omega} \left(n^{3/4} \left(\log \frac{1}{\epsilon} \right)^{1/4} \right), \quad \frac{1}{3^n} \leq \epsilon \leq \frac{1}{3}.$$

Proof. The proof is a cosmetic adaptation of the analysis of element distinctness. To start with, we claim that for any positive integers n, r, k such that $k \mid n$ and $k \mid r$, the composition $\text{AND}_k \circ \text{SURJ}_{n/k, r/k}$ is a subproblem of $\text{SURJ}_{n, r}$. Indeed, the input to $\text{SURJ}_{n, r}$ is an $n \times r$ Boolean matrix in which every row i contains exactly one 1. Consider the restriction of $\text{SURJ}_{n, r}$ to input matrices that are block-diagonal, with k blocks of size $n/k \times r/k$ each. Each of these blocks corresponds to an instance of $\text{SURJ}_{n/k, r/k}$, and the overall problem amounts to computing the AND of these k instances. This settles the claim.

Now let n be arbitrary. Then for all positive integers $k \leq \min\{cn, (1-c)n\}$,

$$\begin{aligned} \deg_{3^{-k}}(\text{SURJ}_{n, \lfloor cn \rfloor}) &\geq \deg_{3^{-k}}(\text{SURJ}_{n - (\lfloor cn \rfloor - k \lfloor cn/k \rfloor), k \lfloor cn/k \rfloor}) \\ &\geq \deg_{3^{-k}}(\text{SURJ}_{n-k, k \lfloor cn/k \rfloor}) \\ &\geq \deg_{3^{-k}}(\text{SURJ}_{k(\lfloor n/k \rfloor - 1), k \lfloor cn/k \rfloor}) \\ &\geq \deg_{3^{-k}}(\text{AND}_k \circ \text{SURJ}_{\lfloor n/k \rfloor - 1, \lfloor cn/k \rfloor}) \\ &\geq k \deg_{1/3}(\text{SURJ}_{\lfloor n/k \rfloor - 1, \lfloor cn/k \rfloor}) \\ &\geq k \cdot \tilde{\Omega} \left(\frac{n}{k} \right)^{3/4} \\ &= \tilde{\Omega}(n^{3/4} k^{1/4}), \end{aligned}$$

where the first step uses (3.5); the second and third steps use (3.4); the fourth step applies the claim from the opening paragraph of the proof; the fifth step is valid by Proposition 3.1; and the sixth step invokes Theorem 2.3. This settles the theorem. \square

3.4. Permutation testing. We now turn to the permutation testing problem, which requires a more subtle analysis than the functions that we have examined so far. The difficulty is that permutation testing does not admit a self-reduction with AND as an outer function. To address this, we will need to generalize Proposition 3.1 appropriately. For a real $0 \leq \alpha < 1$ and an integer $k \geq 1$, we define $\text{AND}_{k, \alpha}$ to be the restriction of AND_k to inputs whose Hamming weight is either k or at most αk . The following result subsumes Proposition 3.1 as the special case $\alpha = (k-1)/k$.

Proposition 3.6. *Fix a real number $0 \leq \alpha < 1$ and an integer $k \geq 1$. Then for any $\epsilon \geq 0$ and any function $f: X \rightarrow \{0, 1\}$ on a finite subset X of Euclidean space,*

$$\deg_{\epsilon^{k/\binom{k-1}{\lfloor \alpha k \rfloor}}}(\text{AND}_{k, \alpha} \circ f) \geq (\lfloor \alpha k \rfloor + 1) \deg_\epsilon(f).$$

In particular,

$$\deg_{(\epsilon/2)^k}(\text{AND}_{k, \alpha} \circ f) \geq \alpha k \deg_\epsilon(f).$$

Proof. We may assume that $\deg_\epsilon(f) \neq 0$ since the proposition is trivial otherwise. Let ψ be an ϵ -error dual polynomial for f , as guaranteed by Fact 2.1:

$$\begin{aligned} \langle f, \psi \rangle &> \epsilon \|\psi\|_1, \\ \text{orth } \psi &= \deg_\epsilon(f). \end{aligned}$$

Abbreviate $\ell = \lfloor \alpha k \rfloor$ and define $\Psi: X^k \rightarrow \mathbb{R}$ by

$$\Psi(x_1, x_2, \dots, x_k) = \prod_{i=1}^{\ell} \psi(x_i) \cdot \prod_{i=\ell+1}^{k-1} (f(x_1) + f(x_2) + \dots + f(x_k) - i).$$

Observe that Ψ is supported on the domain of $\text{AND}_{k,\alpha} \circ f$. Moreover, we have the pointwise inequality

$$\begin{aligned} |\Psi| &\leq |\psi^{\otimes k}| \prod_{i=\ell+1}^{k-1} i \\ &= |\psi^{\otimes k}| \cdot \frac{(k-1)!}{\ell!}. \end{aligned} \tag{3.6}$$

Now

$$\begin{aligned} \langle \Psi, \text{AND}_{k,\alpha} \circ f \rangle &= \langle \Psi, f^{\otimes k} \rangle \\ &= (k-\ell-1)! \langle \psi^{\otimes k}, f^{\otimes k} \rangle \\ &> (k-\ell-1)! \epsilon^k \|\psi\|_1^k \\ &= (k-\ell-1)! \epsilon^k \|\psi^{\otimes k}\|_1 \\ &\geq \epsilon^k \cdot \frac{(k-\ell-1)! \ell!}{(k-1)!} \|\Psi\|_1 \\ &= \epsilon^k \binom{k-1}{\ell}^{-1} \|\Psi\|_1, \end{aligned}$$

where the next-to-last step uses (3.6). Applying Fact 2.1 once again,

$$\begin{aligned} \deg_{\epsilon^k / \binom{k-1}{\ell}}(\text{AND}_{k,\alpha} \circ f) &\geq \text{orth } \Psi \\ &\geq (\ell+1) \text{orth } \psi \\ &= (\ell+1) \deg_\epsilon(f). \end{aligned} \quad \square$$

For $m \leq n$, a permutation testing instance $\phi: \{1, 2, \dots, m\} \rightarrow \{1, 2, \dots, m\}$ can be extended in a natural way to a larger instance $\Phi: \{1, 2, \dots, n\} \rightarrow \{1, 2, \dots, n\}$ by letting $\Phi(i) = i$ for $i = m+1, m+2, \dots, n$. This gives

$$\deg_\epsilon(\text{PTP}_{m,\alpha}) \leq \deg_\epsilon\left(\text{PTP}_{n, \frac{m}{n} \cdot \alpha + \frac{n-m}{n}}\right), \quad m \leq n. \tag{3.7}$$

We are now in a position to prove our lower bound on the ϵ -approximate degree of permutation testing.

Theorem 3.7. *Let $0 < \alpha < 1$ be a given constant. Then*

$$\deg_\epsilon(\text{PTP}_{n,\alpha}) = \Omega\left(n^{1/3} \left(\log \frac{1}{\epsilon}\right)^{2/3}\right), \quad \frac{1}{3^n} \leq \epsilon \leq \frac{1}{3}. \quad (3.8)$$

Proof. Let $0 < \beta < 1$ be arbitrary. We claim that for any positive integers n and k with $k \mid n$, the permutation testing function $\text{PTP}_{n,\beta}$ contains

$$\text{AND}_{k,\beta/2} \circ \text{PTP}_{n/k,\beta/2} \quad (3.9)$$

as a subfunction. The proof is similar to that for element distinctness. Specifically, view instances of (3.9) as block-diagonal matrices with k blocks of size n/k each. Then a positive instance of (3.9) is a permutation matrix and therefore a positive instance of $\text{PTP}_{n,\beta}$. A negative instance of (3.9), on the other hand, features at least $k - \frac{\beta}{2}k$ blocks from $(\text{PTP}_{n/k,\beta/2})^{-1}(0)$ and therefore corresponds to a mapping $\{1, 2, \dots, n\} \rightarrow \{1, 2, \dots, n\}$ with a range of size at most

$$n - \left(k - \frac{\beta k}{2}\right) \cdot \left(\frac{n}{k} - \frac{\beta n}{2k}\right) \leq \beta n.$$

In particular, any negative instance of (3.9) is also a negative instance of $\text{PTP}_{n,\beta}$. This completes the proof of the claim.

Now for any $\epsilon \geq 0$ and any $k \in \{1, 2, \dots, \lceil \alpha n/2 \rceil\}$, we have

$$\begin{aligned} \deg_\epsilon(\text{PTP}_{n,\alpha}) &\geq \deg_\epsilon(\text{PTP}_{k\lfloor n/k \rfloor, \alpha/2}) \\ &\geq \deg_\epsilon(\text{AND}_{k,\alpha/4} \circ \text{PTP}_{\lfloor n/k \rfloor, \alpha/4}), \end{aligned} \quad (3.10)$$

where the first inequality uses (3.7), and the second inequality follows from the claim established in the previous paragraph. The rest of the proof is analogous to those for AND_n and ED_n . For $k = 1, 2, \dots, \lceil \alpha n/2 \rceil$,

$$\begin{aligned} \deg_{6^{-k}}(\text{PTP}_{n,\alpha}) &\geq \deg_{6^{-k}}(\text{AND}_{k,\alpha/4} \circ \text{PTP}_{\lfloor n/k \rfloor, \alpha/4}) \\ &\geq \frac{\alpha k}{4} \deg_{1/3}(\text{PTP}_{\lfloor n/k \rfloor, \alpha/4}) \\ &= \frac{\alpha k}{4} \cdot \Omega\left(\frac{n}{k}\right)^{1/3} \\ &= \Omega(n^{1/3} k^{2/3}), \end{aligned}$$

where the first three steps are valid by (3.10), Proposition 3.6, and Corollary 2.7, respectively. This directly implies (3.8). \square

We will now show that Theorem 3.7 is optimal with respect to all parameters. In fact, we will prove the stronger result that permutation testing has an ϵ -error quantum query algorithm with cost $O(n^{1/3} \log^{2/3}(1/\epsilon))$. Our quantum algorithm is inspired by the well-known algorithm for the collision problem due to Brassard et al. [BHT16].

Theorem 3.8. *Let $0 < \alpha < 1$ be a given constant. Then for all $n \geq 1$ and $1/3^n \leq \epsilon \leq 1/3$, the permutation testing problem $\text{PTP}_{n,\alpha}$ has an ϵ -error quantum query algorithm with cost $O(n^{1/3} \log^{2/3}(1/\epsilon))$. In particular,*

$$\deg_\epsilon(\text{PTP}_{n,\alpha}) = O\left(n^{1/3} \left(\log \frac{1}{\epsilon}\right)^{2/3}\right). \quad (3.11)$$

Proof. We give an algorithm whose only quantum component is Grover search. Specifically, we will only use the fact that, given query access to N items of which M are marked, Grover search finds a marked item with probability $2/3$ using $O(\sqrt{N/M})$ queries (see, e.g., [BHT16, BHT98]). We will follow the convention in the quantum query literature and view the input to $\text{PTP}_{n,\alpha}$ as a function $\phi: \{1, 2, \dots, n\} \rightarrow \{1, 2, \dots, n\}$, where the algorithm has query access to ϕ .

Let s be an integer parameter to be determined later. Our algorithm starts by choosing a uniformly random subset $S \subseteq \{1, 2, \dots, n\}$ of cardinality $|S| = s$. Next, we query ϕ at every point of S . If ϕ is not one-to-one on S , we output “false.” In the complementary case, we execute Grover search $\log(1/\epsilon)$ times independently, each time looking for a point $i \in \bar{S}$ with the property that $\phi(i) \in \phi(S)$. We output “false” if such a point is found, and “true” otherwise.

If ϕ is a permutation, the described algorithm is always correct. In the complementary case when $|\text{im } \phi| \leq \alpha n$, there are at least $(1 - \alpha)n$ points $i \in \{1, 2, \dots, n\}$ such that $|\phi^{-1}(\phi(i))| \geq 2$. Call such points *special*. We will henceforth assume that S contains at least $(1 - \alpha)s/2$ special points, which happens with probability at least $1 - \exp(-\Theta_\alpha(s))$. If ϕ is not one-to-one on S , the algorithm correctly outputs “false.” If ϕ is one-to-one on S and S contains at least $(1 - \alpha)s/2$ special points, then each of the Grover executions has $\geq (1 - \alpha)s/2$ eligible points to output from among a total of $|\bar{S}| = n - s$ possibilities; this means that each Grover execution finds an eligible point with probability at least $2/3$ using $O(\sqrt{n/((1 - \alpha)s)})$ queries, thereby forcing the correct output. In summary, the described algorithm has error probability at most $\exp(-\Theta_\alpha(s)) + (1/3)^{\log(1/\epsilon)}$ and query cost $s + O(\sqrt{n/((1 - \alpha)s)} \cdot \log(1/\epsilon))$. In particular, error ϵ can be achieved with query cost $O(n^{1/3} \log^{2/3}(1/\epsilon))$. This query bound in turn implies (3.11) using the standard transformation of a quantum query algorithm to a polynomial; see, e.g., Ambainis [Amb05]. \square

4. QMA LOWER BOUNDS

The objective of this section is to “lift” the approximate degree lower bound of Theorem 3.7 to QMA query complexity. As our first step, we generalize our lower bound to one-sided approximation. The *one-sided* ϵ -approximate degree of a function $f: X \rightarrow \mathbb{R}$, denoted $\deg_\epsilon^+(f)$, is the least degree of a polynomial $p: X \rightarrow \mathbb{R}$ such that $|p(x)| \leq \epsilon$ for all $x \in f^{-1}(0)$, and $p(x) \geq 1 - \epsilon$ for all $x \in f^{-1}(1)$. Thus, p approximates f uniformly on $f^{-1}(0)$ but may take on arbitrarily large values on $f^{-1}(1)$. It is clear from the definition that $\deg_\epsilon^+(f) \leq \deg_\epsilon(f)$. The gap between these quantities can be large in general, such as 1 versus $\Omega(\sqrt{n})$ for the bounded-error approximation of OR_n . However, we will show that these two notions of approximation are equivalent for the permutation testing function.

Proposition 4.1. *For all α, ϵ , and n ,*

$$\deg_{\epsilon}^{+}(\text{PTP}_{n,\alpha}) = \deg_{\epsilon}(\text{PTP}_{n,\alpha}). \quad (4.1)$$

This equality of approximate degree and one-sided approximate degree for permutation testing has the important consequence that the lower bound of Theorem 3.7 applies to the one-sided setting as well. The proof of Proposition 4.1 is based on the observation that any one-sided approximant for permutation testing can be symmetrized to be constant on $f^{-1}(1)$, effectively making it a two-sided approximant. This technique was used previously in [BT15, Theorem 2] to argue that $\deg_{\epsilon}^{+}(\text{ED}_n) = \deg_{\epsilon}(\text{ED}_n)$.

Proof of Proposition 4.1. Let p be a one-sided approximant for $\text{PTP}_{n,\alpha}$ with error ϵ , so that $|p| \leq \epsilon$ on $\text{PTP}_{n,\alpha}^{-1}(0)$ and $p \geq 1 - \epsilon$ on $\text{PTP}_{n,\alpha}^{-1}(1)$. Define

$$p^{*}(x) = \mathbf{E} p(\sigma x \tau), \quad (4.2)$$

where σ, τ are uniformly random permutations on $\{1, 2, \dots, n\}$, and $\sigma x \tau$ denotes the matrix obtained by permuting the rows of x according to σ and the columns according to τ . Then p^{*} is also a one-sided approximant for $\text{PTP}_{n,\alpha}$ because $\text{PTP}_{n,\alpha}^{-1}(0)$ and $\text{PTP}_{n,\alpha}^{-1}(1)$ are closed under permutations of rows and columns. Moreover, p^{*} takes on the same value, call it M , at all $x \in \text{PTP}_{n,\alpha}^{-1}(1)$ because $\sigma x \tau$ in (4.2) is a uniformly random permutation matrix in that case. As a result, the normalized polynomial $p^{*} / \max\{1, M\}$ approximates $\text{PTP}_{n,\alpha}$ pointwise within ϵ . Finally, $\deg p^{*} \leq \deg p$ because p^{*} is an average of polynomials, each obtained from p by permuting the input variables. \square

We will also need the following proposition, implicit in Marriott and Watrous's proof [MW05] of Vyalys's result [Vya03] on QMA and SBQP. For completeness, we include its short proof.

Proposition 4.2. *Suppose that $f: X \rightarrow \{0, 1\}$ has a QMA query protocol with witness length m and query cost q . Then there is a polynomial $p: X \rightarrow \mathbb{R}$ such that*

$$\deg p = O(mq), \quad (4.3)$$

$$|p(x)| \leq 2^{-2m} \quad \text{for all } x \in f^{-1}(0), \quad (4.4)$$

$$p(x) \geq 2^{-m-1} \quad \text{for all } x \in f^{-1}(1). \quad (4.5)$$

Proof. Marriott and Watrous [MW05] showed that the soundness and completeness errors of the QMA query protocol for f can be driven down to 2^{-2m} without an increase in witness length, and with only a factor of $O(m)$ increase in query cost. This yields a QMA protocol \mathcal{Q} for f that has witness length m , query cost $O(mq)$, and soundness and completeness errors 2^{-2m} . That is, on any input in $f^{-1}(1)$, there exists a witness that causes Arthur to accept with probability at least $1 - 2^{-2m}$, and on any input in $f^{-1}(0)$, for every witness that might be sent by Merlin, Arthur accepts with probability at most 2^{-2m} .

Now run \mathcal{Q} with the witness fixed to the totally mixed state. This yields a quantum query algorithm \mathcal{A} . On inputs in $f^{-1}(0)$, the acceptance probability of \mathcal{A} is at most the soundness error of \mathcal{Q} , which is at most 2^{-2m} . On inputs in $f^{-1}(1)$,

the acceptance probability of \mathcal{A} is at least $(1 - 2^{-2m}) \cdot 2^{-m} \geq 2^{-m-1}$. Now (4.3)–(4.5) follow from the well-known result of Beals et al. [BBC⁺01] that the acceptance probability of any T -query quantum algorithm on input x is a polynomial $p(x)$ of degree at most $2T$. \square

We have reached our main result on the QMA complexity of permutation testing, stated as Theorem 1.2 in the introduction. For the reader’s convenience, we restate the theorem here.

Theorem. *Let $0 < \alpha < 1$ be an arbitrary constant. Then any QMA query protocol for $\text{PTP}_{n,\alpha}$ with witness length m has query cost $\Omega(n/m)^{1/3}$. In particular, $\text{PTP}_{n,\alpha}$ has QMA complexity $\Omega(n^{1/4})$.*

Proof. Fix a QMA query protocol for $\text{PTP}_{n,\alpha}$ with witness length $m \in [3, n]$ and query cost q . Then Proposition 4.2 gives a polynomial p satisfying (4.3)–(4.5). It follows that $2^{m+1}p$ approximates $\text{PTP}_{n,\alpha}$ in a one-sided manner to error 2^{-m+1} , forcing $\deg_{2^{-m+1}}^+(\text{PTP}_{n,\alpha}) = O(mq)$. On the other hand, taking $\epsilon = 2^{-m+1}$ in Theorem 3.7 and Proposition 4.1 shows that $\deg_{2^{-m+1}}^+(\text{PTP}_{n,\alpha}) = \Omega(n^{1/3}m^{2/3})$. Comparing these complementary bounds on the one-sided approximate degree of permutation testing gives $q = \Omega(n/m)^{1/3}$ and thus $\max\{m, q\} = \Omega(n^{1/4})$. \square

We remind the reader that by virtue of Proposition 2.5, the variant of permutation testing studied in this paper is equivalent to Aaronson’s permutation testing problem [Aar12]. As a result, Theorems 1.2, 3.8, and 3.8 and Proposition 4.1 remain valid with $\text{PTP}_{n,\alpha}$ replaced by $\text{PTP}_{n,\alpha}^*$.

5. OPEN PROBLEMS

A natural next step would be to close the gap between our $\Omega(n^{1/4})$ QMA lower bound for permutation testing and the known upper bound of $O(n^{1/3})$. In addition, we highlight the well-known open question of resolving the QMA communication complexity of set disjointness. The best known lower bound here is $\Omega(n^{1/3})$ [Kla11], while the best upper bound is $O(n^{1/2})$. We believe that both questions highlight significant gaps in our understanding of QMA. Another natural open question is whether the naïve error-reduction method for approximate degree is optimal. Namely, it is well known that $\deg_\epsilon(f) \leq O(\min\{\deg_{1/3}(f) \log(1/\epsilon), n\})$ for every $f: \{0, 1\}^n \rightarrow \{0, 1\}$, yet this bound is not known to be tight for any such f with sublinear approximate degree. It *is* tight for some f whose domain is a proper subset of $\{0, 1\}^n$, based for example on approximate counting.

REFERENCES

- [Aar12] Scott Aaronson. Impossibility of succinct quantum proofs for collision-freeness. *Quantum Information & Computation*, 12(1-2):21–28, 2012.
- [Amb05] Andris Ambainis. Polynomial degree and lower bounds in quantum complexity: Collision and element distinctness with small range. *Theory of Computing*, 1(1):37–46, 2005.
- [AS04] Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *J. ACM*, 51(4):595–605, 2004.
- [BBC⁺01] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001.

- [BCH⁺17] Adam Bouland, Lijie Chen, Dhiraj Holden, Justin Thaler, and Prashant Nalini Vasudevan. On the power of statistical zero knowledge. In *Proceedings of the Fifty-Eighth Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 708–719, 2017.
- [BCWZ99] Harry Buhrman, Richard Cleve, Ronald de Wolf, and Christof Zalka. Bounds for small-error and zero-error quantum algorithms. In *Proceedings of the Fortieth Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 358–368, 1999.
- [BHT98] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum counting. In *Proceedings of the Twenty-Fifth International Colloquium on Automata, Languages and Programming (ICALP)*, pages 820–831, 1998.
- [BHT16] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum algorithm for the collision problem. In *Encyclopedia of Algorithms*, pages 1662–1664. Springer, 2016.
- [BKT18] Mark Bun, Robin Kothari, and Justin Thaler. The polynomial method strikes back: Tight quantum query bounds via dual polynomials. In *Proceedings of the Fiftieth Annual ACM Symposium on Theory of Computing (STOC)*, pages 297–310, 2018.
- [BM12] Paul Beame and Widad Machmouchi. The quantum query complexity of AC^0 . *Quantum Information & Computation*, 12(7-8):670–676, 2012.
- [BT15] Mark Bun and Justin Thaler. Hardness amplification and the approximate degree of constant-depth circuits. In *Proceedings of the Forty-Second International Colloquium on Automata, Languages and Programming (ICALP)*, pages 268–280, 2015.
- [GLR18] Tom Gur, Yang P. Liu, and Ron D. Rothblum. An exponential separation between MA and AM proofs of proximity. In *Proceedings of the Forty-Fifth International Colloquium on Automata, Languages and Programming (ICALP)*, pages 73:1–73:15, 2018.
- [Kla11] Hartmut Klauck. On Arthur Merlin games in communication complexity. In *Proceedings of the Twenty-Sixth Annual IEEE Conference on Computational Complexity (CCC)*, 2011.
- [Kup15] Greg Kuperberg. How hard is it to approximate the Jones polynomial? *Theory of Computing*, 11:183–219, 2015.
- [Kut05] Samuel Kutin. Quantum lower bound for the collision problem with small range. *Theory of Computing*, 1(1):29–36, 2005.
- [MW05] Chris Marriott and John Watrous. Quantum Arthur–Merlin games. *Computational Complexity*, 14(2):122–152, 2005.
- [NS94] Noam Nisan and Mario Szegedy. On the degree of Boolean functions as real polynomials. *Computational Complexity*, 4:301–313, 1994.
- [RS04] Ran Raz and Amir Shpilka. On the power of quantum proofs. In *Proceedings of the Nineteenth Annual IEEE Conference on Computational Complexity (CCC)*, pages 260–274, 2004.
- [She18] Alexander A. Sherstov. Algorithmic polynomials. In *Proceedings of the Fiftieth Annual ACM Symposium on Theory of Computing (STOC)*, pages 311–324, 2018.
- [Vya03] Mikhail N. Vyalyi. QMA=PP implies that PP contains PH. *Electronic Colloquium on Computational Complexity*, 10(021), 2003.