AN OPTIMAL SEPARATION OF RANDOMIZED AND QUANTUM QUERY COMPLEXITY (EXTENDED ABSTRACT)

ALEXANDER A. SHERSTOV*, ANDREY A. STOROZHENKO*, AND PEI WU*

Understanding the relative power of quantum and classical computing is of basic importance in theoretical computer science. This question has been studied most actively in the *query model*, which is tractable enough to allow unconditional lower bounds yet rich enough to capture most of the known quantum algorithms. Illustrative examples include the quantum algorithms of Deutsch and Jozsa [9], Bernstein and Vazirani [5], Grover [12], and Shor's period-finding [17]. In the query model, the task is to evaluate a fixed function f on an unknown n-bit input x. In the classical setting, query algorithms are commonly referred to as *decision trees.* A decision tree accesses the input one bit at a time, choosing the bits to query in adaptive fashion. The objective is to determine $f(x)$ by querying as few bits as possible. The minimum number of queries needed to determine $f(x)$ in the worst case is called the *query complexity of* f. The quantum model is a far-reaching generalization of the classical decision tree whereby all bits can be queried in superposition with a single query. The catch is that the outcomes of those queries are then also in superposition, and it is not clear a priori whether quantum query algorithms are more powerful than decision trees. The focus of our paper is on the *bounded-error* regime, where the query algorithm (quantum or classical) is allowed to err with small constant probability on any given input.

The comparative power of randomized and quantum query algorithms has been studied for more than two decades. In pioneering work, Deutsch and Jozsa [9] gave a quantum query algorithm that solves, with a single query, a problem on n bits that any deterministic decision tree needs at least $n/2$ queries to solve. Unfortunately, this separation does not apply to the more subtle, bounded-error setting. This was addressed in follow-up work by Simon [18], who exhibited a problem with bounded-error quantum query complexity $O(\log^2 n)$ and randomized query complexity $\Omega(\sqrt{n})$. These are striking examples of the computational advantages afforded by the quantum model.

This leaves us with a fundamental question: what is the largest possible separation between bounded-error quantum and randomized query complexity, for a problem with n-bit input? This question was popularized by Buhrman et al. [6] and, a decade later, by Aaronson and Ambainis [1], who presented it as being essential to understanding the phenomenon of quantum speedups. Toward this goal, the authors of [1] obtained both positive and negative results. They showed, for every constant t, that every quantum algorithm with t queries can be converted to a randomized decision tree of cost $O(n^{1-1/2t})$. In particular, this rules out an $O(1)$ versus $\Omega(n)$ separation. In the opposite direction, Aaronson and Ambainis exhibited a problem that can be solved to bounded error with a single quantum query but has randomized query complexity $\tilde{\Omega}(\sqrt{n})$. They left open the challenge of obtaining a separation of $O(1)$ versus $\Omega(n^\alpha)$ for some $\alpha > 1/2$.

In more detail, Aaronson and Ambainis [1] introduced and studied the *k-fold forrelation problem.* The input to the problem is a k-tuple of vectors $x_1, x_2, \ldots, x_k \in \{-1, 1\}^n$, where n is a power of 2. Define

$$\phi_{n,k}(x_1, x_2, \ldots, x_k) = \frac{1}{n} \mathbf{1}^\intercal D_{x_1} H D_{x_2} H D_{x_3} H \cdots H D_{x_k} \mathbf{1}, \tag{1}$$

where $\mathbf{1}$ is the all-ones vector, H is the Hadamard transform matrix of order n, and D_{x_i} is the diagonal matrix with the vector x_i on the diagonal. Since each of the linear transformations $H, D_{x_1}, D_{x_2}, \ldots, D_{x_n}$ preserves Euclidean length, it follows that $|\phi_{n,k}(x_1, x_2, \ldots, x_k)| \leqslant 1$. Given $x_1, x_2, \ldots, x_k$, the forrelation problem is to distinguish between the cases $|\phi_{n,k}(x_1, x_2, \ldots, x_k)| \leqslant \alpha$ and $\phi_{n,k}(x_1, x_2, \ldots, x_k) \geqslant \beta$, where the problem parameters $0 < \alpha < \beta < 1$ are suitably chosen constants. Equation (1) directly gives a quantum algorithm that solves the forrelation problem with bounded error and query cost k, and in fact [1] the cost can be further reduced to $\lceil k/2 \rceil$. Aaronson and Ambainis complemented this with an $\tilde{\Omega}(\sqrt{n})$ lower bound on the randomized query complexity of the forrelation problem for $k = 2$, hence the 1 versus $\tilde{\Omega}(\sqrt{n})$ separation mentioned above.

Building on the work of Aaronson and Ambainis [1], last year Tal [21] gave an improved separation of $O(1)$ versus $\Omega(n^{2/3-\varepsilon})$ for bounded-error quantum and randomized query complexities, for any constant $\varepsilon > 0$.

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

For this, Tal replaced (1) with the more general quantity

$$\phi_{n,k,U}(x_1, x_2, \ldots, x_k) = \frac{1}{n}\mathbf{1}^\intercal D_{x_1} U D_{x_2} U D_{x_3} U \cdots U D_{x_k} \mathbf{1}, \tag{2}$$

where U is an arbitrary but fixed orthogonal matrix. On input $x_1, x_2, \ldots, x_k \in \{-1,1\}^n$, the author of [21] considered the problem of distinguishing between the cases $|\phi_{n,k,U}(x_1, x_2, \ldots, x_k)| \leqslant 2^{-k-1}$ and $\phi_{n,k,U}(x_1, x_2, \ldots, x_k) \geqslant 2^{-k}$. This problem is referred to in [21] as the *k-fold rorrelation problem with respect to U*. The quantum algorithm of Aaronson and Ambainis, adapted to the arbitrary choice of U, solves this new problem with $\lceil k/2 \rceil$ queries and advantage $\Omega(2^{-k})$ over random guessing, which counts as a bounded-error algorithm for any constant k. On the other hand, Tal [21] proved that the randomized query complexity of the k-fold rorrelation problem for uniformly random U is $\Omega(n^{2(k-1)/(3k-1)}/k \log n)$ with high probability. Setting k to a large constant gives a separation of $O(1)$ versus $\Omega(n^{2/3-\varepsilon})$ for bounded-error quantum and randomized query complexity for any constant $\varepsilon > 0$.

Our Results

Separations for partial functions. Prior to our paper, Tal's separation of $O(1)$ versus $\Omega(n^{2/3-\varepsilon})$ was the strongest known, and Aaronson and Ambainis's challenge of obtaining an $O(1)$ versus $\Omega(n^{1-\varepsilon})$ separation remained open. The main contribution of our work is to resolve this question. In what follows, we let $f_{n,k,U}$ denote the k-fold rorrelation problem with respect to U. We prove:

Theorem 1. *Let n and k be positive integers, with $k \leqslant \frac{1}{3}\log n - 1$. Let $U \in \mathbb{R}^{n \times n}$ be a uniformly random orthogonal matrix. Then with probability $1 - o(1)$,*

$$R_{\frac{1}{2}-\gamma}(f_{n,k,U}) = \Omega\left(\frac{\gamma^2}{k} \cdot \frac{n^{1-\frac{1}{k}}}{(\log n)^{2-\frac{1}{k}}}\right), \qquad \forall \gamma \in [0, 1/2]. \tag{3}$$

For $k = 2$, this lower bound is the same as Aaronson and Ambainis's lower bound for the forrelation problem (which is $f_{n,2,H}$ in our notation). For $k = 3$ already, Theorem 1 is a polynomial improvement on all previous work, including Tal's recent result [21]. Theorem 1 is essentially tight for all k, both even and odd, due to the matching upper bound $O_k(n^{1-1/k})$ of Aaronson and Ambainis [1] for bounded block-multilinear polynomials of degree k. Since $f_{n,k,U}$ has an efficient quantum protocol for every U, we obtain the following corollary:

Corollary 2. *Let $\varepsilon > 0$ be given. Then there is a partial Boolean function f on $\{-1,1\}^n$ with $Q_{1/3}(f) = O(1)$ and $R_{1/3}(f) = \Omega(n^{1-\varepsilon})$.*

This separation of bounded-error quantum and randomized query complexities is best possible for all f due to Aaronson and Ambainis's result that every quantum protocol with k queries can be simulated by a randomized query algorithm of cost $O(n^{1-1/2k})$. In particular, Corollary 2 shows that the rorrelation problem separates quantum and randomized query complexity optimally, of all problems f.

Separation for total functions. Our results so far pertain to *partial* Boolean functions, whose domain of definition is a proper subset of the Boolean hypercube. For total Boolean functions, such large quantum-classical gaps are not possible. In a seminal paper, Beals et al. [4] prove that the bounded-error quantum query complexity of a total function f is always polynomially related to the randomized query complexity of f. A natural question to ask is how large this polynomial gap can be. Grover's search [12] shows that the n-bit OR function has bounded-error quantum query complexity $\Theta(\sqrt{n})$ and randomized complexity $\Theta(n)$. For a long time, this quadratic separation was believed to be the largest possible. In a surprising result, Aaronson et al. [2] proved the existence of a total function f with $R_{1/3}(f) = \tilde{\Omega}(Q_{1/3}(f)^{2.5})$. This was improved by Tal [21] to $R_{1/3}(f) \geqslant Q_{1/3}(f)^{8/3-o(1)}$. We give a polynomially stronger separation:

Theorem 3. *There is a function $f\colon \{-1,1\}^n \to \{0,1\}$ with $R_{1/3}(f) \geqslant Q_{1/3}(f)^{3-o(1)}$.*

Theorem 3 follows by combining our quantum-classical query separations (see the full version) with the "cheatsheet" framework of Aaronson et al. [2]. A recent paper of Aaronson et al. [3] conjectures that $R_{1/3}(f) = O(Q_{1/3}(f)^3)$ for every total function f, which would mean that our Theorem 3 is essentially optimal.

Separations for communication complexity. Using standard reductions, our quantum-classical query separations imply analogous separations for communication complexity. We prove:

THEOREM 4 (Partial functions). *Let $\varepsilon > 0$ be given. Then there is a partial Boolean function F on $\{-1,1\}^N \times \{-1,1\}^N$ with $Q^{\mathrm{cc}}_{1/3}(F) = O(\log N)$ and $R^{\mathrm{cc}}_{1/3}(F) = \Omega(N^{1-\varepsilon})$.*

THEOREM 5 (Total functions). *There is $F\colon \{-1,1\}^N \times \{-1,1\}^N \to \{0,1\}$ with $R^{\mathrm{cc}}_{1/3}(F) \geqslant Q^{\mathrm{cc}}_{1/3}(F)^{3-o(1)}$.*

Theorem 4 is near-optimal and a polynomial improvement on previous work. The best previous quantum-classical separation for communication complexity was $O(\log N)$ versus $\Omega(N^{2/3-\varepsilon})$, implicit in Tal [21] and preceded in turn by other exponential separations [15, 16, 10]. Similarly, Theorem 5 is a polynomial improvement on previous work, the best previous result being a power of $8/3$ separation implicit in [21].

Fourier weight of decision trees. It is straightforward to verify that a uniformly random input $x \in (\{-1,1\}^n)^k$ is with high probability a *negative* instance of the rorrelation problem $f_{n,k,U}$. With this in mind, Tal [21] proves his lower bound for rorrelation by constructing a probability distribution $\mathscr{D}_{n,k,U}$ that generates *positive* instances of $f_{n,k,U}$ with nontrivial probability yet is indistinguishable from the uniform distribution by a decision tree T of cost $n^{2/3-O(1/k)}$. His notion of indistinguishability is based on the Fourier spectrum. Specifically, Tal [21] shows that: (i) the *sum* of the absolute values of the Fourier coefficients of T of given order ℓ does not grow too fast with ℓ; and (ii) the *maximum* Fourier coefficient of $\mathscr{D}_{n,k,U}$ of order ℓ decays exponentially fast with ℓ. In Tal's paper, the bound for (ii) is essentially optimal, whereas the bound for (i) is far from tight. The sum of the absolute values of the order-ℓ Fourier coefficients of a decision tree T, which we refer to as the *ℓ-Fourier weight of T*, is shown in [21] to be at most

$$c^\ell \sqrt{d^\ell (1 + \log kn)^{\ell-1}}, \tag{4}$$

where d is the depth of the tree and $c \geqslant 1$ is an absolute constant. This bound is strong for any constant ℓ but degrades rapidly as ℓ grows. In particular, for $\ell = \sqrt{d}$ already, (4) is weaker than the trivial bound $\binom{d}{\ell}$. This is a major obstacle since the indistinguishability proof requires strong bounds for every ℓ. This obstacle is the reason why Tal's analysis gives the randomized query lower bound $n^{2/3-O(1/k)}$ as opposed to the optimal $\tilde{\Omega}(n^{1-1/k})$. Tal conjectured that the ℓ-Fourier weight of a depth-d decision tree is in fact bounded by $c^\ell \sqrt{\binom{d}{\ell}(1 + \log kn)^{\ell-1}}$, which is a factor of $\sqrt{\ell!}$ improvement on (4) and essentially optimal. We prove his conjecture:

THEOREM 6. *Let $T\colon \{-1,1\}^n \to \{0,1\}$ be a function computable by a decision tree of depth d. Then*

$$\sum_{\substack{S \subseteq \{1,2,\ldots,n\}: \\ |S| = \ell}} |\hat{T}(S)| \leqslant c^\ell \sqrt{\binom{d}{\ell}(1 + \log n)^{\ell-1}}, \qquad \ell = 1, 2, \ldots, n,$$

where $c \geqslant 1$ is an absolute constant.

It is well known and easy to show that Theorem 6 is essentially tight, even for *nonadaptive* decision trees [13, Theorem 5.19]. The actual statement that we prove is more precise and takes into account the density parameter $\mathbf{P}[T(x) \neq 0]$; see the full version for details. With Theorem 6 in hand, all our main results (Theorem 1 and its corollaries) follow immediately by combining the new bound on the Fourier weight of decision trees with Tal's near-optimal bounds on the individual Fourier coefficients of $\mathscr{D}_{n,k,U}$.

Theorem 6 is of interest in its own right, independent of its use in this paper to obtain optimal quantum-classical separations. The study of the Fourier spectrum has a variety of applications in theoretical computer science, including circuit complexity, learning theory, pseudorandom generators, and quantum computing. Even prior to Tal's work, the ℓ-Fourier weight of decision trees was studied for $\ell = 1$ by O'Donnell and Servedio [14], who proved the tight $O(\sqrt{d})$ bound and used it to give a polynomial-time learning algorithm for monotone decision trees. Fourier weight has been studied for various other classes of Boolean functions, including bounded-depth circuits, branching programs, low-degree polynomials over finite fields, and functions with bounded sensitivity; see the recent papers [11, 19, 20, 8, 7] and the references therein.

References

[1] S. Aaronson and A. Ambainis, *Forrelation: A problem that optimally separates quantum from classical computing*, SIAM J. Comput., 47 (2018), pp. 982–1038, doi:10.1137/15M1050902.

[2] S. Aaronson, S. Ben-David, and R. Kothari, *Separations in query complexity using cheat sheets*, in *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing* (STOC), 2016, pp. 863–876, doi:10.1145/2897518.2897644.

[3] S. Aaronson, S. Ben-David, R. Kothari, and A. Tal, *Quantum implications of Huang's sensitivity theorem*. Available at `https://arxiv.org/abs/2004.13231`, 2020.

[4] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf, *Quantum lower bounds by polynomials*, J. ACM, 48 (2001), pp. 778–797, doi:10.1145/502090.502097.

[5] E. Bernstein and U. V. Vazirani, *Quantum complexity theory*, SIAM J. Comput., 26 (1997), pp. 1411–1473, doi:10.1137/S0097539796300921.

[6] H. Buhrman, L. Fortnow, I. Newman, and H. Röhrig, *Quantum property testing*, SIAM J. Comput., 37 (2008), pp. 1387–1400, doi:10.1137/S0097539704442416.

[7] E. Chattopadhyay, P. Hatami, K. Hosseini, and S. Lovett, *Pseudorandom generators from polarizing random walks*, in *Proceedings of the Thirty-Third Annual IEEE Conference on Computational Complexity* (CCC), vol. 102, 2018, pp. 1:1–1:21, doi:10.4230/LIPIcs.CCC.2018.1.

[8] E. Chattopadhyay, P. Hatami, O. Reingold, and A. Tal, *Improved pseudorandomness for unordered branching programs through local monotonicity*, in *Proceedings of the Fiftieth Annual ACM Symposium on Theory of Computing* (STOC), 2018, pp. 363–375, doi:10.1145/3188745.3188800.

[9] D. Deutsch and R. Jozsa, *Rapid solution of problems by quantum computation*, Proc. R. Soc. Lond. A, 439 (1992), pp. 553–558, doi:10.1098/rspa.1992.0167.

[10] D. Gavinsky, *Entangled simultaneity versus classical interactivity in communication complexity*, IEEE Trans. Inf. Theory, 66 (2020), pp. 4641–4651, doi:10.1109/TIT.2020.2976074.

[11] P. Gopalan, R. A. Servedio, and A. Wigderson, *Degree and sensitivity: Tails of two distributions*, in *Proceedings of the Thirty-First Annual IEEE Conference on Computational Complexity* (CCC), vol. 50, 2016, pp. 13:1–13:23, doi:10.4230/LIPIcs.CCC.2016.13.

[12] L. K. Grover, *A fast quantum mechanical algorithm for database search*, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (STOC), 1996, pp. 212–219, doi:10.1145/237814.237866.

[13] R. O'Donnell, *Analysis of Boolean Functions*, Cambridge University Press, 2014.

[14] R. O'Donnell and R. A. Servedio, *Learning monotone decision trees in polynomial time*, SIAM J. Comput., 37 (2007), pp. 827–844, doi:10.1137/060669309.

[15] R. Raz, *Exponential separation of quantum and classical communication complexity*, in *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing* (STOC), 1999, pp. 358–367, doi:10.1145/301250.301343.

[16] O. Regev and B. Klartag, *Quantum one-way communication can be exponentially stronger than classical communication*, in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (STOC), 2011, pp. 31–40, doi:10.1145/1993636.1993642.

[17] P. W. Shor, *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*, SIAM J. Comput., 26 (1997), pp. 1484–1509, doi:10.1137/S0097539795293172.

[18] D. R. Simon, *On the power of quantum computation*, SIAM J. Comput., 26 (1997), pp. 1474–1483, doi:10.1137/S0097539796298637.

[19] T. Steinke, S. P. Vadhan, and A. Wan, *Pseudorandomness and Fourier-growth bounds for width-3 branching programs*, Theory Comput., 13 (2017), pp. 1–50, doi:10.4086/toc.2017.v013a012.

[20] A. Tal, *Tight bounds on the fourier spectrum of AC0*, in *Proceedings of the Thirty-Second Annual IEEE Conference on Computational Complexity* (CCC), vol. 79, 2017, pp. 15:1–15:31, doi:10.4230/LIPIcs.CCC.2017.15.

[21] A. Tal, *Towards optimal separations between quantum and randomized query complexities*, in *Proceedings of the Sixty-First Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 2020.