Research Statement

Pei Wu

November 28, 2022

I have broad research interests in theoretical computer science. I am particularly passionate about computational complexity theory for its combination of mathematical depth, extensive and surprising connections to other disciplines, and broad relevance to real-world computing. My work has focused on complexity theory and analytic aspects of Boolean functions. In this research statement, I will present key highlights of my doctoral and postdoctoral work, including the following research directions.

(i) The relative power of quantum versus classical computing in the query, communication, and proof models. Among my contributions is obtaining the largest possible separation of quantum versus classical computing in the extensively studied *query model* (also called the *blackbox model*).

(ii) Analytic properties of Boolean functions and their applications to query complexity, communication complexity, and learning theory. A key contribution of my doctoral work is the determination of the threshold degree of constant-depth circuits, which resolves a 50-year-old question from Minsky and Papert's seminar work on neural networks.

(iii) Interactive coding, a vast generalization of classical error correction to the interactive setting. My work here has produced optimal error-correcting codes for interactive communication in the adversarial model with substitutions, insertions, and deletions.

In the concluding section of this document, I will describe my vision for future work.

1 Background

Communication complexity. The classical model of two-party communication features two geographically separated parties, Alice and Bob, who have private inputs x and y, respectively, and need to communicate back and forth to compute a given function $f(x, y)$. This model can be viewed as a far-reaching generalization of classical information theory, in the sense that information theory studies the *one-way* transmission of information as opposed to interactive communication. Analogous to Shannon's noiseless coding theorem and capacity theorem, the natural questions here are: what is the minimum communication cost for Alice and Bob to compute f, and how does one handle noise if it is present?

Communication complexity theory is of great intrinsic importance because communication is a key resource in computing. Moreover, communication complexity is a powerful tool in studying various other computational models because virtually any computational process involves information flow among two or more components. Indeed, there is a vast body of research applying communication complexity to study computational phenomena as diverse as circuits, streaming algorithms, and computational learning.

Query complexity. In the query model, the task is to evaluate a given function f on an unknown n-bit input x. To access the input, we query an index i of our choice and receive x_i. The goal is to minimize the worst-case number of queries by choosing the query indices strategically.

Even though query complexity is among the simplest computational models, it remains the focus of a large body of research. This is because the query model captures the hardness of many important problems. To illustrate, the vast majority of known quantum algorithms, including Grover's search algorithm and Shor's period finding algorithm, are captured by the query model. In addition, query complexity sheds light on more sophisticated models. For example, it is well known that the query model is closely related to Turing machines with oracles. As another example, there are a variety of *lifting theorems* that make it possible to instantly transfer lower bounds for the query model to the vastly more powerful model of communication complexity. This lifting approach has recently enabled the resolution of several important open problems.

Proof systems. The notion of proof is central to complexity theory. For example, the class NP (nondeterministic polynomial time) has the following proof interpretation. For any language $L \in \text{NP}$, there is a deterministic polynomial-time verifier such that: (1) given any input $x \in L$, there is a polynomial-size proof π such that the verifier accepts the pair (x, π) (the completeness requirement); and (2) given any input $x \notin L$, the verifier rejects the pair (x, π) for every proof π (the soundness requirement). The Merlin-Arthur proof system generalizes the above notion by allowing the verifier to be probabilistic. Roughly speaking, the completeness requirement now states that the verifier should accept every input $x \in L$ with high probability (say, at least $2/3$); whereas the soundness requirement states that the verifier should reject every input $x \notin L$ with high probability (say, at least $2/3$). Quantum computation is inherently probabilistic, and therefore the Merlin-Arthur model generalizes naturally to quantum computation.

2 Quantum versus Classical Computing

2.1 Quantum Supremacy in Query and Communication Models

Quantum query complexity has been extensively studied and can be justly considered to be among the biggest achievements of quantum computing to date. Of particular prominence in this line of research are results demonstrating the superiority of quantum algorithms over their classical counterparts. I am particularly interested in the bounded-error regime, where the query algorithm is allowed to err with a small constant probability. In groundbreaking work, Simon [41] exhibited a partial Boolean function whose bounded-error quantum query complexity is exponentially smaller than its randomized (i.e., classical) query complexity. This raises the question: what is the largest possible separation between quantum and randomized query complexity? This question was first explicitly stated in 2002 by Buhrman et al. [12], and has since been popularized by Aaronson and Ambainis [2].

We settled this 18-year-old problem completely in [38]. Specifically, we proved that for any constant k, there is a partial function f with quantum query complexity at most k and randomized query complexity $\tilde{\Omega}(n^{1-1/2k})$. This gives a quantum-classical separation of $O(1)$ versus $\Omega(n^{1-\epsilon})$ for every $\epsilon > 0$, which is a polynomial improvement on the best previous separation of $O(1)$ versus $\Omega(n^{2/3-\epsilon})$ due to Tal [42]. Our k versus $\tilde{\Omega}(n^{1-1/2k})$ separation is optimal due to Aaronson and Ambainis' result [2] that any k-query quantum algorithm can be simulated by $O(n^{1-1/2k})$ randomized classical queries, for an arbitrary constant k. By the well-known framework of "cheatsheets" due to Aaronson et al. [4], our result also implies a cubic separation between quantum and randomized query complexity for *total* functions. This separation is the largest known and has been conjectured to be tight by other researchers [5]. As a technical centerpiece of our work [38], we prove a tight bound on the ℓ_1 norm of any given level of the Fourier spectrum of decision trees. This bound on Fourier weight settles a conjecture of Tal [42] and is of substantial interest in its own right, considering the central role of the Fourier spectrum in many recent breakthroughs in the area [14, 15, 34].

The analogous question has been extensively studied in the communication model [33, 35, 20, 36]. Via query-to-communication lifting, we obtain near-optimal separations for quantum versus randomized communication complexity. In particular, we obtain an $O(\log n)$ versus $\Omega(n^{1-\epsilon})$ separation for the quantum versus randomized communication complexity of partial functions, for any $\epsilon > 0$. Our separation is essentially optimal and a polynomial improvement on previous work.

2.2 Quantum Proof Systems with Exponentially Short Proofs

The notion of proof plays a central role in classical complexity theory. The *quantum* Merlin-Arthur proof system (QMA) is one of the most basic quantum proof systems and represents the quantum analog of NP. In a QMA proof system associated with some language L, the prover provides a quantum proof for any input x that consists of polynomially many quantum bits, and a quantum polynomial-time verifier checks whether the proof certifies $x \in L$. Due to the mysterious features of quantum mechanics, new and interesting phenomena manifest themselves. For example, the number of provers does not play a role classically (as long as the proof length is restricted to polynomial size), and the resulting complexity class is NP. However, in quantum computation, this is probably not the case—there is evidence that two unentangled quantum provers may have more power than one prover [21]. To make this distinction, the proof system with two unentangled provers is denoted QMA(2).

Since its introduction in 2003 [30], there has been little progress in understanding QMA(2) as a complexity class, despite substantial efforts over the past two decades, motivated by complexity theory and the connection to polynomial optimization problems and quantum entanglement [3, 10, 7, 19, 16, 22]. It remains a major open problem to even slightly improve either side of the trivial bounds (i) that $\text{QMA} \subseteq \text{QMA}(2)$, since the verifier can always ignore one of the two provers; and (ii) $\text{QMA}(2) \subseteq \text{NEXP}$, where the latter denotes the complexity class of nondeterministic exponential time, since a NEXP algorithm can guess all the amplitudes of a quantum proof of polynomial size.

In joint work with F. G. Jeronimo, we studied QMA(2) by restricting the provers to send proofs where the quantum amplitudes are of the same phase. We call this class $\text{QMA}^+(2)$. Under this restriction, it turns out that using only *logarithmic* size quantum proofs, a $\text{QMA}^+(2)$ proof system can certify any NP language L with a constant soundness and completeness gap. In contrast, for the classical NP class, polynomial-size proofs are necessary. We further strengthen this result and show that $\text{QMA}^+(2) = \text{NEXP}$. Namely, with two unentangled proofs of only *polynomial* size, a *quantum polynomial-time* verifier will be able to check whether a given instance is in a NEXP language. Apart from providing a new characterization of NEXP, our work may shed light on directions towards the problem of whether $\text{QMA}(2) = \text{NEXP}$.

3 The Study of Boolean Functions and Applications

3.1 Sign-Representation of Boolean Functions

Representations of Boolean functions by real polynomials are of great importance in a variety of contexts, from communication complexity and quantum computing to machine learning theory. For example, the notion of *approximate degree* has played an essential role in quantum query complexity for decades. *Threshold degree* has an even broader range of applications, including various models of computational learning. The notion of threshold degree originates in the pioneering work of Minsky and Papert [31] and is defined, for a Boolean function $f : \{0,1\}^n \to \{-1,1\}$, as the minimum degree of a real polynomial p that represents f in sign: $f(x) = \operatorname{sgn} p(x)$ for all x.

The threshold degree of polynomial-size constant-depth circuits (AC^0) has been the focus of 50+ years of work. Together with my Ph.D. advisor, we were able to essentially settle this longstanding problem in [39]. More specifically, we proved that for any $\epsilon > 0$, there is an AC^0 circuit with threshold degree $\Omega(n^{1-\epsilon})$. This lower bound essentially matches the trivial upper bound of $O(n)$ and is a polynomial improvement on the best previous lower bound, $\Omega(\sqrt{n})$. We further strengthened this result to handle not only threshold degree but also *sign-rank*—a vastly more general notion than threshold degree.

Our results have far-reaching applications in communication complexity and learning theory. In communication complexity, our results give the strongest known lower bounds for AC^0, showing the optimality of the trivial protocol where Alice sends her entire input to Bob. Our lower bound holds even if Alice and Bob only need to compute f with an arbitrary nonzero advantage over random guessing. In learning theory, our results rule out the possibility of the distribution-free PAC learning of AC^0 based on the powerful *dimension complexity* paradigm. This framework captures nearly all known algorithmic results for distribution-free PAC learning.

This work has been invited to appear in a special issue of *SIAM Journal of Computing* for STOC 2019.

3.2 Random Restrictions of Boolean Functions

For any Boolean function $f : \{-1,1\}^n \to \{0,1\}$, the individual *influence* of the ith coordinate x_i is the probability of flipping the value of f when flipping x_i on a random input $x \in \{-1,1\}^n$. This notion of influences was first introduced by Ben-Or and Linial [9] in the context of *collective coin flipping.* It coincides with the "Banzaf index" studied in game theory. The class of Boolean functions with individual influences bounded by $\tau = o(1)$ is a central topic in the analysis of Boolean functions. There are several motivations to study such functions. First, they arise naturally in social choice theory [27, 28]. For example, in a voting system of two candidates and n voters, each bit x_i represents the individual preference of each voter between the two candidates. When aggregating the social preference, it is natural to use a function f where the potential of any given individual to determine the final outcome is limited. Second, from an algorithmic perspective, suppose that we have access to the input via a limited number of queries. Then, it is natural to query a variable when its individual influence is large. This observation has been applied in different settings [18, 1]. In computational complexity, for example, distinguishing dictatorship functions from functions with small individual influences is a key component of proving optimal NP-hardness for approximation [8, 24, 25, 29].

Let $\mathcal{R}_p$ denote the *random restrictions* of f, where each variable is fixed with probability p to a random value in $\{-1,1\}$, and is left undetermined (*alive*) with the complementary probability $1-p$. Applying a random restriction and studying the properties of the restricted function turn out to be a very useful technique. This approach has led to breakthroughs in a variety of areas. For example, it underlies exponential lower bounds in circuit complexity [23] and the recent dramatic improvements of the sunflower lemma in combinatorics [6].

In order to gain a deeper understanding of query complexity and block sensitivity, together with Ronen Eldan and Avi Wigderson, we studied [17] random restrictions of Boolean functions with small individual influences. We showed that, when the alive probability is $\tilde{\Omega}((\log 1/\tau)^{-1})$, the restricted function remains nonconstant with overwhelming probability. This parameter is the best possible. The related problem was first formulated by Friedgut and Kalai as the "it ain't over till it's over" conjecture in the context of social choice theory. The original conjecture was solved by Mossel, O'Donnell, and Oleszkiewicz [32]. Their approach falls short of obtaining the optimal parameters. Our result has applications in social choice theory and complexity theory. In complexity theory, our result implies, among other things, that a random input will have *block sensitivity* $\tilde{\Omega}(\log 1/\tau)$ with overwhelming probability. It is worth mentioning that, by the well-known Kahn-Kalai-Linial inequality [26], the average block sensitivity is at least $\Omega(\log 1/\tau)$. Our result establishes the following stronger statement: not only is the average block sensitivity large, but in fact essentially all inputs have large block sensitivity.

4 Interactive Coding

Noise is omnipresent in communication. In the classical setting of one-way communication, the study of information transmission under noise forms a large part of classical information theory. In pioneering work, Schulman [37] considered noise in the setting of interactive communication. This area of research, called *interactive coding*, is a fascinating and highly active discipline at the crossroads of information theory and communication complexity. More concretely, consider the following scenario. Alice and Bob would like to execute a communication protocol π defined for a noiseless environment. However, the communication channel is controlled by an adversary who can corrupt any fraction ρ of symbols transmitted through the channel. The question is, can Alice and Bob use some interactive analogue of error-correcting codes to ensure that they are both able to recover, from their noisy communication, the transcript that π would have produced without noise? A far-reaching generalization of this model, proposed by Braverman et al. [11], allows *arbitrary* corruptions: insertions, deletions, and substitutions. For any constant $\epsilon > 0$, the authors of [11] showed how to faithfully simulate any protocol in this generalized model with corruption rate up to $\frac{1}{18} - \epsilon$, using a constant-size alphabet and a constant-factor overhead in communication.

Braverman et al. posed the following natural and fundamental question: what is the maximum corruption rate that can be tolerated in this generalized model of substitutions, insertions, and deletions? We gave a complete answer to this question in [40]. We showed that for any $\epsilon > 0$, there is an interactive coding scheme that uses a constant-size alphabet and achieves noise tolerance $\frac{1}{4} - \epsilon$, at the expense of a constant-factor overhead in communication complexity compared to π. This rate is easily seen to be optimal, even in the presence of substitution errors alone.

5 Future Directions

Theoretical computer science is a young and exceptionally active discipline. I look forward to pursuing my current areas of expertise as well as branching out into new areas of theoretical computer science. In what follows, I mention several of my favorite problems that are closely related to my past research.

Sign-representation of shallow circuits

Our aforementioned threshold degree result for AC^0 uses circuits whose depth is a large constant. If we turn to extremely shallow circuits, there are many unsettled problems. The only case we understand fully is the trivial case of depth-1 circuits, which are just the AND and OR functions. Analyzing the sign-representation and pointwise approximation of circuits of depth as small as 2 is already very challenging—and very rewarding from the standpoint of applications. For example, a major open problem is to determine the quantum query complexity of *triangle detection.* As a function, triangle detection is easily computable by a depth-2 circuit. Establishing a tight approximate degree lower bound is currently the most promising approach to this fundamental problem.

Quantum query/communication complexity

In the query complexity world, our understanding is now more or less complete: we know that the quantum and randomized query complexity can be arbitrarily separated for partial functions. For total functions, my Ph.D. thesis gives a cubic separation, whereas Aaronson et al. [5] prove that any separation is at most quartic. Closing the gap for total functions is the main unresolved question in this line of work. In the communication world, things are more open, especially for total functions. It is a major open problem to decide if quantum protocols can be super-polynomially more efficient than randomized protocols for total functions. There are several reasons why this problem is so challenging. First, the lifting technique does not apply since quantum and randomized query complexity are polynomially related. Moreover, many canonical lower bound techniques for randomized communication complexity (e.g., approximate rank and the discrepancy method) also lower bound quantum communication complexity.

The discrepancy between the exponential and polynomial quantum speedups for partial and total functions begs for a deeper understanding that would allow us to interpolate between these phenomena. A folklore conjecture is that for any q-query quantum algorithm Q, there is a $q^{O(1)}$-query classical algorithm A that approximates Q in the ℓ_2 sense. More concretely, for all but a small fraction of the inputs x, $A(x)$ approximates $\mathbb{E}[Q(x)]$ within a tiny constant additive error. This conjecture generalizes polynomial quantum speedups for total functions and explains why, to achieve exponential speedups for a partial function f, the domain of f must be restricted to a subset of the Boolean cube so structured that either $f^{-1}(0)$ or $f^{-1}(1)$ has only $o(1)$ measure over the uniform distribution. Aaronson and Ambainis reduced this conjecture to a purely analytical problem. The conjecture states that for any polynomial P of degree d, there is a variable whose influence is at least $(\mathrm{Var}[f]/d)^{O(1)}$ [1].

Two-party communication complexity

Analogous to the structural complexity theory of Turing machines, two-party communication has its own complexity classes: P, BPP, NP, PH. A large number of questions regarding the relationships among these complexity classes remain open. For example, in the Turing machine world, we know that BPP lies within the second level of PH, and it is conjectured to be equal to P. In the communication world, BPP is also contained in the second level of PH, but the conjecture is that BPP $\not\subseteq$ P^{NP}. Recently, Chattopadhyay et al. [13] showed that BPP $\not\subseteq$ P^{GT}, where GT is the greater-than function. An ambitious open problem is to prove that BPP $\not\subseteq$ P^{NP}, which would give a highly accurate placement of BPP in the polynomial hierarchy. Looking out further, if we consider higher levels of the polynomial hierarchy, we reach the frontier of research in communication complexity. A central and notoriously hard problem in communication complexity is to exhibit functions that are not contained in PH. Currently, it remains open even to prove explicit lower bounds against AM, a subclass of PH.

References

[1] Scott Aaronson and Andris Ambainis. The need for structure in quantum speedups. *Theory of Computing*, 10(6):133–166, 2014. 3.2, 5

[2] Scott Aaronson and Andris Ambainis. Forrelation: A problem that optimally separates quantum from classical computing. *SIAM J. Comput.*, 47(3):982–1038, 2018. 2.1

[3] Scott Aaronson, Salman Beigi, Andrew Drucker, Bill Fefferman, and Peter Shor. The power of unentanglement. In *Proceedings of the Twenty-Third Annual IEEE Conference on Computational Complexity* (CCC), pages 223–236, 2008. 2.2

[4] Scott Aaronson, Shalev Ben-David, and Robin Kothari. Separations in query complexity using cheat sheets. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing* (STOC), pages 863–876, 2016. 2.1

[5] Scott Aaronson, Shalev Ben-David, Robin Kothari, Shravas Rao, and Avishay Tal. Degree vs. approximate degree and quantum implications of Huang's sensitivity theorem. In *Proceedings of the Fifty-Third Annual ACM Symposium on Theory of Computing* (STOC), pages 1330–1342, 2021. 2.1, 5

[6] Ryan Alweiss, Shachar Lovett, Kewen Wu, and Jiapeng Zhang. Improved bounds for the sunflower lemma. *Annals of Mathematics*, 194(3):795 – 815, 2021. 3.2

[7] Salman Beigi. NP vs QMAlog(2). *Quantum Info. Comput.*, 2010. 2.2

[8] Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, PCPs, and nonapproximability – towards tight results. *SIAM J. Comput.*, 27(3):804–915, 1998. 3.2

[9] Michael Ben-Or and Nathan Linial. Collective coin flipping, robust voting schemes and minima of banzhaf values. In *26th Annual Symposium on Foundations of Computer Science (sfcs 1985)*, pages 408–416, 1985. 3.2

[10] Hugue Blier and Alain Tapp. All languages in NP have very short quantum proofs. In *2009 Third International Conference on Quantum, Nano and Micro Technologies*, pages 34–37, 2009. 2.2

[11] Mark Braverman, Ran Gelles, Jieming Mao, and Rafail Ostrovsky. Coding for interactive communication correcting insertions and deletions. *IEEE Trans. Information Theory*, 63(10):6256–6270, 2017. 4

[12] Harry Buhrman, Lance Fortnow, Ilan Newman, and Hein Röhrig. Quantum property testing. *SIAM J. Comput.*, 37(5):1387–1400, 2008. 2.1

[13] Arkadev Chattopadhyay, Shachar Lovett, and Marc Vinyals. Equality alone does not simulate randomness. In *Proceedings of the Thirty-Fourth Annual IEEE Conference on Computational Complexity* (CCC), volume 137 of *LIPIcs*, pages 14:1–14:11, 2019. 5

[14] Eshan Chattopadhyay, Pooya Hatami, Kaave Hosseini, and Shachar Lovett. Pseudorandom generators from polarizing random walks. In *Proceedings of the Thirty-Third Annual IEEE Conference on Computational Complexity* (CCC), volume 102, pages 1:1–1:21, 2018. 2.1

[15] Eshan Chattopadhyay, Pooya Hatami, Omer Reingold, and Avishay Tal. Improved pseudorandomness for unordered branching programs through local monotonicity. In *Proceedings of the Fiftieth Annual ACM Symposium on Theory of Computing* (STOC), pages 363–375, 2018. 2.1

[16] Alessandro Chiesa and Michael A. Forbes. Improved soundness for QMA with multiple provers. *Chic. J. Theor. Comput. Sci.*, 2013. 2.2

[17] Ronen Eldan, Avi Wigderson, and Pei Wu. An optimal "it ain't over till it's over" theorem. Available at `http://arxiv.org/abs/2208.03450`, 2022. 3.2

[18] E. Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, 1(18):27–35, 1998. 3.2

[19] François Le Gall, Shota Nakagawa, and Harumichi Nishimura. On QMA protocols with two short quantum proofs. *Quantum Info. Comput.*, 2012. 2.2

[20] Dmitry Gavinsky, Julia Kempe, Oded Regev, and R. de Wolf. Bounded-error quantum state identification and exponential separations in communication complexity. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing* (STOC), pages 594–603, 2006. 2.1

[21] Aram W. Harrow and Ashley Montanaro. Testing product states, quantum merlin-arthur games and tensor optimization. *J. ACM*, 60(1), feb 2013. 2.2

[22] Aram W. Harrow, Anand Natarajan, and Xiaodi Wu. An improved semidefinite programming hierarchy for testing entanglement. *Communications in Mathematical Physics*, 2017. 2.2

[23] Johan Håstad. Computational limitations of small-depth circuits. 1987. 3.2

[24] Johan Håstad. Testing of the long code and hardness for clique. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, pages 11–19, New York, NY, USA, 1996. Association for Computing Machinery. 3.2

[25] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, jul 2001. 3.2

[26] J. Kahn, G. Kalai, and N. Linial. The influence of variables on boolean functions. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, pages 68–80, 1988. 3.2

[27] Gil Kalai. A Fourier-theoretic perspective on the Condorcet paradox and Arrow's theorem. *Advances in Applied Mathematics*, 29(3):412–426, 2002. 3.2

[28] Gil Kalai. Social indeterminacy. *Econometrica*, 72(5):1565–1581, 2004. 3.2

[29] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? *SIAM Journal on Computing*, 37(1):319–357, 2007. 3.2

[30] Hirotada Kobayashi, Keiji Matsumoto, and Tomoyuki Yamakami. Quantum merlin-arthur proof systems: Are multiple Merlins more helpful to Arthur? In *Algorithms and Computation*, 2003. 2.2

[31] Marvin L. Minsky and Seymour A. Papert. *Perceptrons: An Introduction to Computational Geometry.* MIT Press, Cambridge, Mass., 1969. 3.1

[32] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. *Annals of Mathematics*, pages 295–341, 2010. 3.2

[33] Ran Raz. Exponential separation of quantum and classical communication complexity. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing* (STOC), pages 358–367, 1999. 2.1

[34] Ran Raz and Avishay Tal. Oracle separation of BQP and PH. In *Proceedings of the Fifty-First Annual ACM Symposium on Theory of Computing* (STOC), pages 13–23, 2019. 2.1

[35] Alexander A. Razborov. Quantum communication complexity of symmetric predicates. *Izvestiya: Mathematics*, 67(1):145–159, 2003. 2.1

[36] Oded Regev and Bo'az 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), pages 31–40, 2011. 2.1

[37] Leonard J. Schulman. Coding for interactive communication. *IEEE Trans. Information Theory*, 42(6):1745–1756, 1996. 4

[38] Alexander A Sherstov, Andrey A Storozhenko, and Pei Wu. An optimal separation of randomized and quantum query complexity. In *Proceedings of the Fifty-Third Annual ACM Symposium on Theory of Computing* (STOC), pages 1289–1302, 2021. 2.1

[39] Alexander A. Sherstov and Pei Wu. Near-optimal lower bounds on the threshold degree and sign-rank of AC^0. In *Proceedings of the Fifty-First Annual ACM Symposium on Theory of Computing* (STOC), pages 401–412, 2019. 3.1

[40] Alexander A. Sherstov and Pei Wu. Optimal interactive coding for insertions, deletions, and substitutions. *IEEE Trans. Inf. Theory*, 65(10):5971–6000, 2019. Preliminary version in *Proceedings of the Fifty-Eighth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 2017. 4

[41] Daniel R. Simon. On the power of quantum computation. *SIAM J. Comput.*, 26(5):1474–1483, 1997. 2.1

[42] Avishay 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. 2.1