

# Open Problems Related to Quantum Query Complexity

Scott Aaronson\*

## Abstract

I offer a case that quantum query complexity still has loads of enticing and fundamental open problems—from relativized QMA versus QCMA and BQP versus IP, to time/space tradeoffs for collision and element distinctness, to polynomial degree versus quantum query complexity for partial functions, to the Unitary Synthesis Problem and more.

## 1 Introduction

Quantum query complexity (see [24] for a classic survey) is the study of how many queries a quantum computer needs to make to an input string  $X$  to learn various properties of  $X$ . The key here is that a single query can access multiple bits of  $X$ , one in each branch of a superposition state. For over thirty years, this subject has been a central source of what we know about both the capabilities and the limitations of quantum computers.

In my view, there are two reasons why query complexity has played such an important role in quantum computing theory as a whole. First, it so happens that most of the famous quantum algorithms—including Deutsch-Jozsa [26], Bernstein-Vazirani [21], Simon [48], Shor [47], and Grover [33]—fit naturally into the query complexity framework, or (in the case of Shor’s algorithm) have a central component that does. Second, query complexity lets us prove not only upper bounds, but also nontrivial and informative lower bounds—as illustrated by the seminal 1994 theorem of Bennett, Bernstein, and Vazirani [20] that a quantum computer needs  $\Omega(\sqrt{N})$  queries to search an unordered list of size  $N$  for a single “marked item.” This both demonstrated the optimality of Grover’s algorithm, two years before that algorithm had been discovered to exist (!), and showed the existence of an oracle relative to which  $\text{NP} \not\subseteq \text{BQP}$ .

Of course, oracle separations sometimes mislead us about the “real world,” where no oracles are present—a famous example being the 1990  $\text{IP} = \text{PSPACE}$  theorem [45]. Even after a quarter century, though, non-relativizing techniques (i.e., techniques that transcend query complexity) have made only minor inroads into quantum complexity theory, at least outside the usual place where those techniques have shined: namely, the study of interactive proof systems.

Yet today, some in our field seem to have the impression that quantum query complexity is more-or-less a closed subject. Certainly, some of quantum computing theorists’ attention has understandably shifted to other topics, from the theoretical foundations of quantum supremacy experiments [39], to potential near-term or “NISQ” quantum algorithms [42], to the quest to prove a quantum PCP Theorem [10]. And certainly, many of the great open problems of quantum query complexity from circa 2000 *were* ultimately solved: to give some well-known examples, the quantum

---

\*University of Texas at Austin. Email: aaronson@utexas.edu. Supported by a Vannevar Bush Fellowship from the US Department of Defense, a Simons Investigator Award, and the Simons “It from Qubit” collaboration.

query complexities of the collision and element distinctness problems [9] and of evaluating read-once formulas [44]; the optimal separation between deterministic and quantum query complexities of total Boolean functions [16, 7]; and an oracle separation between BQP and the polynomial hierarchy [43].

Nevertheless, in this article I’d like to make the case that open problems abound in quantum query complexity—and I don’t mean detail problems, of tightening some bound to remove a logarithmic factor, but big, juicy, important problems. Some of my problems are old and well-known; others are obscure; still others, as far as I know, appear here in writing for the first time.

## 2 QMA and Oracle Separations

Recall that QMA, Quantum Merlin Arthur, is the class of languages  $L$  for which membership in  $L$  can be proven via a polynomial-size quantum witness state  $|\varphi\rangle$  that’s verified in quantum polynomial time. In 2002, Aharonov and Naveh [12] defined QCMA, or Quantum Classical Merlin Arthur, to be the subclass of QMA where the witness  $|\varphi\rangle = |w\rangle$  is required to be a classical basis state (i.e., a string). Ever since, one of the fundamental problems of quantum complexity theory has been whether  $\text{QMA} = \text{QCMA}$ . One distinctive feature of this question is that even its query complexity analogue remains open:

**Problem 1** *Is there an oracle relative to which  $\text{QMA} \neq \text{QCMA}$ ?*

In 2007, Greg Kuperberg and I [8] at least showed that there’s a *quantum* oracle  $U$ —that is, a collection of unitary transformations provided as black boxes—such that  $\text{QMA}^U \neq \text{QCMA}^U$ . This was the first use of quantum oracles in complexity theory; their use has since become standard. But the question of whether QMA and QCMA can be separated by a “standard” oracle remained wide open. In 2015, Fefferman and Kimmel [30] showed that there’s a “randomized, in-place” classical oracle relative to which  $\text{QMA} \neq \text{QCMA}$ , but for proving a conventional classical oracle separation, I believe the best candidate we have remains the “component mixers problem” introduced in 2011 by Lutomirski [40].

Here is another fundamental problem that’s remained open about QMA and oracle separations:

**Problem 2** *Is there an oracle—even a quantum oracle—relative to which  $\text{QMA} \neq \text{QMA}(2)$ ?*

Here,  $\text{QMA}(2)$  is the analogue of QMA to allow *two* unentangled Merlins, so that Arthur can always assume that the witness state he receives is a tensor product  $|\psi\rangle \otimes |\varphi\rangle$  across two polynomial-size registers. Despite 18 years of work on this class, the only inclusions known are still the obvious ones,  $\text{QMA} \subseteq \text{QMA}(2) \subseteq \text{NEXP}$ . Furthermore, unlike with the QMA vs. QCMA problem, here we do not even have a *quantum* oracle relative to which  $\text{QMA} \neq \text{QMA}(2)$ . Watrous (see [5]) conjectured that there is no quantum channel that takes polynomially many qubits as input, produces polynomially many qubits as output, always produces an approximately separable state on two registers as its output, and can approximately produce *any* separable state. Proving Watrous’s “no disentanglers conjecture” is a prerequisite to separating QMA from  $\text{QMA}(2)$  query complexity, since were his conjecture false, we could always use a QMA witness to simulate a  $\text{QMA}(2)$  witness. See Harrow, Natarajan, and Wu [35] for the best current progress toward proving Watrous’s conjecture.

### 3 Query/Space Tradeoffs

In the *collision* problem, we're given black-box access to a function  $f : \{1, \dots, n\} \rightarrow \{1, \dots, m\}$  (where  $n$  is even and  $m \geq n$ ), and are asked to decide whether  $f$  is 1-to-1 or 2-to-1, promised that one of those is the case. In the *element distinctness* problem, we're given black-box access to a function  $f : \{1, \dots, n\} \rightarrow \{1, \dots, m\}$ , with no promise, and are simply asked whether  $f$  is 1-to-1.

**Problem 3** *What are the optimal tradeoffs between the number of queries used by a quantum algorithm to solve the collision or the element distinctness problems, and the number of qubits or classical bits of memory?*

Brassard, Høyer, Tapp [22] gave a quantum algorithm for the collision problem that uses  $O(n^{1/3})$  quantum queries, as well as  $O(n^{1/3})$  bits of classical memory and  $O(\log n)$  qubits. Six years later, Ambainis [15] gave a quantum algorithm for element distinctness that uses  $O(n^{2/3})$  quantum queries and  $O(n^{2/3})$  qubits. The 2002 collision lower bound by me and Yaoyun Shi [9] showed that both of these algorithms were optimal in terms of queries, thereby settling the problems' quantum query complexity.

Here, though, we're asking whether a quantum algorithm for these problems could achieve near-optimal query complexity while also using a small *memory*. Note that, by using Grover's algorithm, we could solve the collision problem using only  $O(\log n)$  qubits in total, but then we'd need  $O(\sqrt{n})$  queries rather than  $O(n^{1/3})$ . For element distinctness, even supposing that we need a large memory, it would also be interesting to know whether the memory needs to be made of qubits, or whether coherently-queryable classical bits (a so-called "qRAM") would suffice.

At present, unfortunately, the only techniques that we have for proving quantum lower bounds that trade off space with query complexity, seem to apply only to problems with many bits of output, such as sorting a list [37]. Proving such lower bounds for decision problems, like collision or element distinctness, will probably require the invention of new techniques.

### 4 Maximal Separations

Given a total Boolean function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ , we denote by  $D(f)$ ,  $R(f)$ , and  $Q(f)$  the deterministic, (bounded-error) randomized, and (bounded-error) quantum query complexities of  $f$  respectively. In their seminal 1998 paper, Beals et al. [18] showed that  $D(f) = O(Q(f)^6)$  for all  $f$ . This stood as the best known relationship between  $D(f)$  and  $Q(f)$  until extremely recently, when, building on Huang's breakthrough proof of the Sensitivity Conjecture [36], some of us [7] showed that  $D(f) = O(Q(f)^4)$  for all total Boolean functions  $f$ .

In the other direction, until 2015 it was widely believed that the largest possible gap between classical and quantum query complexities for total Boolean functions was quadratic, and was achieved by Grover's algorithm applied to the  $n$ -bit OR function. But Ambainis et al. [16] then refuted that conjecture, by giving an example of a Boolean function  $f$  for which  $D(f) \approx Q(f)^4$ . Not long afterward, Ben-David [19] (see also Aaronson, Ben-David, and Kothari [6]) gave an example of an  $f$  for which  $R(f) \approx Q(f)^{2.5}$ , thereby showing that Ambainis et al.'s separation was not just an artifact of ignoring classical randomized algorithms. Ben-David's result was recently improved to give functions  $f$  for which  $R(f) \approx Q(f)^{8/3}$  [49] and even  $R(f) \approx Q(f)^3$  [17, 46].

Yet all this progress, as dramatic as it’s been, still leaves a gap between 3 and 4 in the exponent of the optimal separation between  $R(f)$  and  $Q(f)$ .

**Problem 4** *What is the largest possible gap between  $R(f)$  and  $Q(f)$ , for a total Boolean function  $f$ ?*

## 5 Degree of Partial Functions

Let  $f : S \rightarrow \{0, 1\}$  be a partial Boolean function, where  $S \subseteq \{0, 1\}^n$ . Define the *approximate degree* of  $f$ , or  $\widetilde{\deg}(f)$ , to be the minimum degree of a real polynomial  $p : \mathbb{R}^n \rightarrow \mathbb{R}$  such that

- (i)  $|p(x) - f(x)| \leq \frac{1}{3}$  for all  $x \in S$ , and
- (ii)  $p(x) \in [0, 1]$  for all  $x \in \{0, 1\}^n$ .

The seminal 1998 work of Beals et al. [18] showed that  $\widetilde{\deg}(f) \leq 2Q(f)$  for all  $f$ , where  $Q(f)$  is bounded-error quantum query complexity. This is simply because the acceptance probability of a  $T$ -query quantum algorithm is a real polynomial of degree at most  $2T$ . Beals et al.’s result was the beginning of the wildly-successful *polynomial method* in quantum complexity theory, whose central idea is that to lower-bound quantum query complexity, it suffices to lower-bound approximate degree.

In 2003, Ambainis [14] showed that there can be small polynomial gaps between  $\widetilde{\deg}(f)$  and  $Q(f)$  for total Boolean functions  $f$ —and thus, “the polynomial method is not tight.” Ben-David, Kothari, and I [6] later improved this to get a fourth-power gap between  $\widetilde{\deg}(f)$  and  $Q(f)$ , which is tight by the recent work of Ben-David, Kothari, Rao, Tal, and me [7].

I ask about the situation for *partial* Boolean functions:

**Problem 5** *What is the largest possible gap between  $Q(f)$  and  $\widetilde{\deg}(f)$  for partial  $f$ ? Can the gap even be exponential?*

Note that, if it weren’t for requirement (ii)—namely, that the polynomial must be bounded in  $[0, 1]$  even on inputs that violate the promise—we’d have a degree-1 polynomial representing the  $n$ -bit OR function, whose quantum query complexity is  $\Theta(\sqrt{n})$ .

## 6 Unitary Synthesis Problem

In 2007, Greg Kuperberg and I [8] raised the following question:

**Problem 6** *For every  $n$ -qubit unitary transformation  $U$ , does there exist an oracle  $A : \{0, 1\}^* \rightarrow \{0, 1\}$  such that a  $\text{BQP}^A$  machine can implement  $U$ ?<sup>1</sup>*

In my 2016 Barbados lecture notes [4], I took to calling this the “Unitary Synthesis Problem.” It remains wide open.

---

<sup>1</sup>Technically, Kuperberg and I only asked whether this is true for a Haar-random  $U$ , but I expect the Haar-random case to be essentially the hardest case.

For comparison, it’s not hard to show that, for every  $n$ -qubit state  $|\psi\rangle$ , there exists an oracle  $A$  such that a  $\text{BQP}^A$  machine can prepare  $|\psi\rangle$ . Indeed, for every  $n$ -qubit unitary  $U$  and every polynomial  $p$ , there exists an oracle  $A$  such that a  $\text{BQP}^A$  machine can simulate the behavior of  $U$  on any chosen  $p(n)$  basis states. However, extending this construction to simulate  $U$  on *all* states seems to entail exponentially many queries to  $A$ .

While it might sound esoteric, the Unitary Synthesis Problem has turned up again and again—for example, in the study of the nonabelian hidden subgroup problem [29], of decoding Hawking radiation from a black hole [34, 4], and of schemes for quantum copy-protection and quantum money [2, 4]. In each of those topics, one is interested in certain complicated  $n$ -qubit unitary transformations  $U$ —and especially, whether or not those  $U$ ’s have polynomial-size quantum circuits. The question arises: could we at least show that small quantum circuits would exist *if* (say)  $\text{P} = \text{PSPACE}$ , or some other classical complexity classes dramatically collapsed? While the implication isn’t immediate, a positive answer to the Unitary Synthesis Problem would strongly suggest that the answer was yes. For what it’s worth, though, my conjecture is that the answer is negative—in which case, the study of quantum circuit complexity *cannot* be so easily related to classical complexity theory.

## 7 Verifiability of Quantum Computing

Let  $\text{IP}$  be the class of languages that admit classical interactive proofs. In the unrelativized world,  $\text{IP} = \text{PSPACE}$  [45], but it’s well-known that the situation relative to oracles can be dramatically different [32]. Thus we ask:

**Problem 7** *Does there exist an oracle  $A$  such that  $\text{BQP}^A \not\subseteq \text{IP}^A$ ?*

In my view, the Forrelation problem, which I [3] introduced in 2009, and which Raz and Tal [43] used in 2018 to give an oracle relative to which  $\text{BQP} \not\subseteq \text{PH}$ , provides a compelling candidate for an oracle relative to which  $\text{BQP} \not\subseteq \text{IP}$  as well. However, showing that Forrelation is not in  $\text{IP}$  will require a new circuit lower bound—one that talks about circuits with “expectation” and “maximization” gates, rather than  $\text{AC}^0$  circuits with AND, OR, and NOT gates. As far as I know, Aiello, Goldwasser, and Håstad [13] proved what’s still the best known lower bound against expectation/maximization circuits in 1989, when they gave an oracle relative to which more rounds give interactive protocols more power.

Now let  $\text{IP}_{\text{BQP}}$  be the subclass of  $\text{BQP}$  consisting of all languages for which a  $\text{BQP}$  prover can convince a  $\text{BPP}$  verifier of a “yes” answer, through polynomially many rounds of classical interaction.

**Problem 8** *Is there at least an oracle relative to which  $\text{BQP} \neq \text{IP}_{\text{BQP}}$ ?*

Besides Forrelation, even the complement of Simon’s Problem (i.e., output “yes” if  $f$  is a 1-to-1 function, or “no” if  $f$  satisfies the Simon promise, promised that one of these is the case) seems like a good candidate for an oracle problem in  $\text{BQP}$  but not in  $\text{IP}_{\text{BQP}}$ . In the Simon example, note that there *is* an interactive protocol, based on AM approximate counting—it just doesn’t seem to be a protocol for which a  $\text{BQP}$  machine could implement the prover’s strategy.


Finally, a question that was implicit in our previous one:

**Problem 9** *Are the celebrated protocols for blind and verified quantum computation, due to Broadbent et al. [23], Aharonov et al. [11], and Mahadev [41], inherently non-relativizing?*

Certainly these protocols don't *manifestly* work relative to arbitrary oracles, but are there variants of the protocols that do?

## 8 Glued Trees

In 2002, Childs et al. [25] gave a celebrated quantum walk algorithm that, informally, gets from the leftmost to the rightmost vertex in the following  $\exp(n)$ -sized graph, in only  $\text{poly}(n)$  time and with  $\frac{1}{\text{poly}(n)}$  success probability:



By contrast, they showed that a randomized algorithm needs  $2^{\Omega(n)}$  queries to an oracle encoding the graph to solve the same problem (improved by Fenner and Zhang [31] to  $\tilde{\Omega}(2^{n/2})$ ).

**Problem 10** *Suppose, however, that we actually want to find a path from left to right. Does even a quantum computer need  $2^{\Omega(n)}$  queries for that task?*

Certainly, if we try to measure the state of the quantum walk algorithm to reveal such a path, we'll destroy the quantum interference that causes that algorithm to succeed. But this, of course, doesn't show that no other quantum algorithm is possible. It seems to me that a lower bound—showing that a quantum algorithm can't efficiently find even a single left-right path, even though it can traverse exponentially many such paths in superposition—would be a striking algorithmic version of wave/particle duality.

## 9 Comparing Query Models

Given a function  $f$ , the usual model of quantum query complexity is that we get access to an oracle that maps basis states of the form  $|x, a\rangle$  to basis states of the form  $|x, a \oplus f(x)\rangle$ , where  $\oplus$  denotes

bitwise XOR and where I’m ignoring workspace registers. However, if  $f$  is injective, then another model is possible: namely, an oracle that simply maps basis states of the form  $|x\rangle$  to basis states of the form  $|f(x)\rangle$ , “erasing” the previous contents of the  $|x\rangle$  register. This second model has the great advantage of not leaving  $x$  around as garbage, but the disadvantage of not being inherently reversible.

In 2000, Elham Kashefi (personal communication) asked me the following question: are there sets of injective functions  $f$  for which a quantum computer can learn certain properties of  $f$  using few queries to an erasing oracle, but *not* using few queries to a standard oracle? I realized that a lower bound for the collision problem would naturally lead to an affirmative answer to this question. This provided a central motivation for my work on the collision problem [1], which, in an appendix, did give an affirmative answer to Kashefi’s question.

More recently, I became aware that the converse question is equally interesting:

**Problem 11** *Are there sets of injective functions  $f$  for which a quantum computer can learn certain properties of  $f$  using few queries to a standard oracle, but not using few queries to an erasing oracle?*

One natural candidate would be as follows:

$$f(x) = \langle h(x), \text{gar}(x) \rangle,$$

where  $h$  is a Simon function (that is, a function that’s either 1-to-1 or else satisfies the Simon promise, and for which the promise is to decide which), and  $\text{gar}(x)$  is a long string of random garbage depending on  $x$ . The inclusion of  $\text{gar}(x)$  makes  $f$  injective with overwhelming probability, but with a standard oracle is no bar to running Simon’s algorithm, since we can simply use a second oracle invocation to uncompute garbage:

$$|x\rangle \rightarrow |x, h(x), \text{gar}(x)\rangle \rightarrow |x, h(x), \text{gar}(x), h(x)\rangle \rightarrow |x, h(x)\rangle.$$

On the other hand, the garbage seems to make erasing queries no more useful than classical queries. A central reason I’m interested in this conjecture is that a proof of it seems likely to proceed by proving a much more general statement, about the presence of a sufficient amount of garbage, in an erasing oracle’s responses, being equivalent (under appropriate conditions) to decohering or measuring the responses.

## 10 The Linear Cross-Entropy Benchmark

In Fall 2019, a team at Google reported the achievement of quantum supremacy based on a sampling benchmark with superconducting qubits [27]. In Summer 2021, a team at USTC in China reported an independent replication [28]. Briefly, in these experiments, one generates a random quantum circuit  $C$  acting on  $n$  qubits (in the Google experiment,  $n = 53$ ; in the USTC experiment,  $n = 56$ ). One then uses a quantum computer to (hopefully) sample from  $\mathcal{D}_C$ , the probability distribution over  $n$ -bit strings induced by preparing the state  $C|0^n\rangle$  and then measuring all  $n$  qubits in the computational basis. Finally, having generated samples  $s_1, \dots, s_k \in \{0, 1\}^n$ , one uses a classical computer to calculate

$$\chi := \frac{2^n}{k} \sum_{i=1}^k |\langle 0^n | C | s_i \rangle|^2.$$

One can show that ideal sampling, with a noiseless quantum computer, would yield an expected value of  $\chi \approx 2$ , whereas classical random guessing would yield an expected value of  $\chi \approx 1$ . The test is considered a success if and only if  $\chi$  is sufficiently bounded above 1. Google’s experiment achieved a value of  $\chi \approx 1.002$ .

One can ask: what if we wanted to achieve  $\chi \gg 2$ ? Would that problem be intractable even for a quantum computer—analogueous to violating the Tsirelson inequality (i.e., the statement that even quantumly entangled players can win the CHSH game with probability at most  $\cos^2 \frac{\pi}{8}$ )? If we imagined a black box able to output samples with, say,  $\chi \approx 3$ , would that black box provide “beyond-quantum” computational abilities, and if so can we say anything about those abilities?

Here I ask a query complexity version of this question. Given a Boolean function  $f : \{0, 1\}^n \rightarrow \{-1, 1\}$ , an easy quantum algorithm samples a string  $s$  with probability equal to  $\widehat{f}(s)^2$ , where  $\widehat{f}(z)$  is the  $z^{\text{th}}$  Boolean Fourier coefficient of  $f$ . Define the quantity

$$\chi := \frac{2^n}{k} \sum_{i=1}^k \widehat{f}(s_i)^2.$$

Then one can show that repeating the Fourier sampling algorithm yields samples  $s_1, \dots, s_k$  that satisfy  $\chi \approx 3$ . We now ask:

**Problem 12** *What is the quantum query complexity of outputting samples  $s_1, \dots, s_k$  that satisfy, say,  $\chi \approx 4$ ?*

Very recently, Kretschmer [38] made significant progress on this problem, by showing that

- (1) given an  $n$ -qubit Haar-random *quantum* oracle,  $\widetilde{\Omega}(2^{n/4})$  queries are needed to violate the “quantum supremacy Tsirelson’s inequality” for that oracle (compared to an upper bound of  $O(2^{n/3})$ ), and
- (2) when  $k = 1$ , the obvious quantum algorithm for Fourier-sampling a Boolean function  $f$  is optimal among all 1-query quantum algorithms.

## 11 Acknowledgments

I thank Andrew Childs for the glued trees figure, and Travis Humble and Mingsheng Ying for commissioning this piece.

## References

- [1] S. Aaronson. Quantum lower bound for the collision problem. In *Proc. ACM STOC*, pages 635–642, 2002. quant-ph/0111102.
- [2] S. Aaronson. Quantum copy-protection and quantum money. In *Proc. Conference on Computational Complexity*, pages 229–242, 2009. arXiv:1110.5353.
- [3] S. Aaronson. BQP and the polynomial hierarchy. In *Proc. ACM STOC*, 2010. arXiv:0910.4698.



- [4] S. Aaronson. The complexity of quantum states and transformations: From quantum money to black holes, February 2016. Lecture Notes for the 28th McGill Invitational Workshop on Computational Complexity, Holetown, Barbados. With guest lectures by A. Bouland and L. Schaeffer. [www.scottaaronson.com/barbados-2016.pdf](http://www.scottaaronson.com/barbados-2016.pdf).
- [5] S. Aaronson, S. Beigi, A. Drucker, B. Fefferman, and P. Shor. The power of unentanglement. In *Proc. Conference on Computational Complexity*, pages 223–236, 2008. arXiv:0804.0802.
- [6] S. Aaronson, S. Ben-David, and R. Kothari. Separations in query complexity using cheat sheets. In *Proc. ACM STOC*, pages 863–876, 2016. arXiv:1511.01937.
- [7] S. Aaronson, S. Ben-David, R. Kothari, S. Rao, and A. Tal. Degree vs. approximate degree and quantum implications of Huang’s sensitivity theorem. In *Proc. ACM STOC*, pages 1330–1342, 2021. arXiv:2010.12629.
- [8] S. Aaronson and G. Kuperberg. Quantum versus classical proofs and advice. *Theory of Computing*, 3(7):129–157, 2007. Earlier version in CCC’2007. arXiv:quant-ph/0604056.
- [9] S. Aaronson and Y. Shi. Quantum lower bounds for the collision and the element distinctness problems. *J. of the ACM*, 51(4):595–605, 2004.
- [10] D. Aharonov, I. Arad, and T. Vidick. The quantum PCP conjecture. *SIGACT News*, 44(2):47–79, 2013. arXiv:1309.7495.
- [11] D. Aharonov, M. Ben-Or, E. Eban, and U. Mahadev. Interactive proofs for quantum computations. Earlier version in ICS’2010. arXiv:1704.04487, 2017.
- [12] D. Aharonov and T. Naveh. Quantum NP - a survey. [quant-ph/0210077](http://quant-ph/0210077), 2002.
- [13] W. Aiello, S. Goldwasser, and J. Håstad. On the power of interaction. *Combinatorica*, 10(1):3–25, 1990. Earlier version in FOCS’86.
- [14] A. Ambainis. Polynomial degree vs. quantum query complexity. *J. Comput. Sys. Sci.*, 72(2):220–238, 2006. Earlier version in FOCS’2003. [quant-ph/0305028](http://quant-ph/0305028).
- [15] A. Ambainis. Quantum walk algorithm for element distinctness. *SIAM J. Comput.*, 37(1):210–239, 2007. Earlier version in FOCS’2004. [quant-ph/0311001](http://quant-ph/0311001).
- [16] A. Ambainis, K. Balodis, A. Belovs, T. Lee, M. Santha, and J. Smotrovs. Separations in query complexity based on pointer functions. *J. of the ACM*, 64(5):1–24, 2017. Earlier version in STOC’2016. arXiv:1506.04719.
- [17] N. Bansal and M. Sinha. k-forrelation optimally separates quantum and classical query complexity. In *Proc. ACM STOC*, pages 1303–1316, 2021. arXiv:2008.07003.
- [18] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. *J. of the ACM*, 48(4):778–797, 2001. Earlier version in FOCS’1998, pp. 352–361. [quant-ph/9802049](http://quant-ph/9802049).
- [19] S. Ben-David. A super-Grover separation between randomized and quantum query complexities. arXiv:1506.08106, 2015.

- [20] C. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, 1997. quant-ph/9701001.
- [21] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM J. Comput.*, 26(5):1411–1473, 1997. Earlier version in STOC’1993.
- [22] G. Brassard, P. Høyer, and A. Tapp. Quantum algorithm for the collision problem. *ACM SIGACT News*, 28:14–19, 1997. quant-ph/9705002.
- [23] A. Broadbent, J. Fitzsimons, and E. Kashefi. Universal blind quantum computation. In *Proc. IEEE FOCS*, 2009. arXiv:0807.4154.
- [24] H. Buhrman and R. de Wolf. Complexity measures and decision tree complexity: a survey. *Theoretical Comput. Sci.*, 288:21–43, 2002.
- [25] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman. Exponential algorithmic speedup by a quantum walk. In *Proc. ACM STOC*, pages 59–68, 2003. quant-ph/0209131.
- [26] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proc. Roy. Soc. London*, A439:553–558, 1992.
- [27] F. Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574:505–510, 2019. arXiv:1910.11333.
- [28] Y. Wu et al. Strong quantum computational advantage using a superconducting quantum processor, 2021. arXiv:2106.14734.
- [29] M. Etinger, P. Høyer, and E. Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Inform. Proc. Lett.*, 91(1):43–48, 2004. quant-ph/0401083.
- [30] B. Fefferman and S. Kimmel. Quantum vs. classical proofs and subset verification. In *Mathematical Foundations of Computer Science*, pages 1–22, 2018. arXiv:1510.06750.
- [31] S. Fenner and Y. Zhang. A note on the classical lower bound for a quantum walk algorithm. quant-ph/0312230v1, 2003.
- [32] L. Fortnow and M. Sipser. Are there interactive protocols for co-NP languages? *Inform. Proc. Lett.*, 28:249–251, 1988.
- [33] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proc. ACM STOC*, pages 212–219, 1996. quant-ph/9605043.
- [34] D. Harlow and P. Hayden. Quantum computation vs. firewalls. *Journal of High Energy Physics*, (85), 2013. arXiv:1301.4504.
- [35] A. W. Harrow, A. Natarajan, and X. Wu. Limitations of semidefinite programs for separable states and entangled games. *Commun. Math. Phys.*, 366(2):423–468, 2019. arXiv:1612.09306.
- [36] H. Huang. Induced subgraphs of hypercubes and a proof of the Sensitivity Conjecture. *Annals of Mathematics*, 190(3):949–955, 2019. arXiv:1907.00847.

- [37] H. Klauck, R. Špalek, and R. de Wolf. Quantum and classical strong direct product theorems and optimal time-space tradeoffs. *SIAM J. Comput.*, 36(5):1472–1493, 2007. Earlier version in FOCS’2004. quant-ph/0402123.
- [38] W. Kretschmer. The quantum supremacy Tsirelson inequality. In *Proc. Innovations in Theoretical Computer Science (ITCS)*, pages 1–13, 2021. arXiv:2008.08721.
- [39] A. P. Lund, M. J. Bremner, and T. C. Ralph. Quantum sampling problems, BosonSampling, and quantum supremacy. *npj Quantum Information*, 3(15), 2017. arXiv:1702.03061.
- [40] A. Lutomirski. Component mixers and a hardness result for counterfeiting quantum money. arXiv:1107.0321, 2011.
- [41] U. Mahadev. Classical verification of quantum computations. In *Proc. IEEE FOCS*, pages 259–267, 2018. arXiv:1804.01082.
- [42] J. Preskill. Quantum computing in the NISQ era and beyond, 2018. arXiv:1801.00862.
- [43] R. Raz and A. Tal. Oracle separation of BQP and PH. In *Proc. ACM STOC*, pages 13–23, 2019. ECCC TR18-107.
- [44] B. Reichardt. Reflections for quantum query algorithms. In *Proc. ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 560–569, 2011. arXiv:1005.1601.
- [45] A. Shamir. IP=PSPACE. *J. of the ACM*, 39(4):869–877, 1992. Earlier version in FOCS’1990, pp. 11-15.
- [46] A. A. Sherstov, A. A. Storozhenko, and P. Wu. An optimal separation of randomized and query complexity. In *Proc. ACM STOC*, pages 1289–1302, 2021. arXiv:2008.10223.
- [47] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997. Earlier version in FOCS’1994. quant-ph/9508027.
- [48] D. Simon. On the power of quantum computation. In *Proc. IEEE FOCS*, pages 116–123, 1994.
- [49] A. Tal. Towards optimal separations between quantum and randomized query complexities. In *Proc. IEEE FOCS*, pages 228–239, 2020. arXiv:1912.12561.