

# Complexity-Theoretic Limitations on Blind Delegated Quantum Computation

**Scott Aaronson**

Department of Computer Science, University of Texas at Austin, USA  
aaronson@cs.utexas.edu

**Alexandru Cojocaru**

School of Informatics, University of Edinburgh, UK  
a.cojocaru@sms.ed.ac.uk

**Alexandru Gheorghiu**<sup>1</sup> 

Department of Computing and Mathematical Sciences, California Institute of Technology, USA  
School of Informatics, University of Edinburgh, UK  
andrugh@caltech.edu

**Elham Kashefi**

School of Informatics, University of Edinburgh, UK  
CNRS LIP6, Université Pierre et Marie Curie, Paris, France  
ekashefi@inf.ed.ac.uk

---

## Abstract

Blind delegation protocols allow a client to delegate a computation to a server so that the server learns nothing about the input to the computation apart from its size. For the specific case of *quantum computation* we know, from work over the past decade, that blind delegation protocols can achieve information-theoretic security (provided the client and the server exchange some amount of quantum information). In this paper we prove, provided certain complexity-theoretic conjectures are true, that the power of *information-theoretically secure* blind delegation protocols for quantum computation (ITS-BQC protocols) is in a number of ways constrained.

In the first part of our paper we provide some indication that ITS-BQC protocols for delegating polynomial-time quantum computations in which the client and the server interact only classically are unlikely to exist. We first show that having such a protocol in which the client and the server exchange  $O(n^d)$  bits of communication, implies that  $BQP \subseteq MA/O(n^d)$ . We conjecture that this containment is unlikely by proving that there exists an oracle relative to which  $BQP \not\subseteq MA/O(n^d)$ . We then show that if an ITS-BQC protocol exists in which the client and the server interact only classically and which allows the client to delegate quantum sampling problems to the server (such as *BOSONSAMPLING*) then there exist non-uniform circuits of size  $2^{n-\Omega(n/\log(n))}$ , making polynomially-sized queries to an  $NP^{NP}$  oracle, for computing the permanent of an  $n \times n$  matrix.

The second part of our paper concerns ITS-BQC protocols in which the client and the server engage in one round of quantum communication and then exchange polynomially many classical messages. First, we provide a complexity-theoretic upper bound on the types of functions that could be delegated in such a protocol by showing that they must be contained in  $QCMA/qpoly \cap coQCMA/qpoly$ . Then, we show that having such a protocol for delegating NP-hard functions implies  $coNP^{NP^{NP}} \subseteq NP^{NP^{PromiseQMA}}$ .

**2012 ACM Subject Classification** Theory of computation → Quantum computation theory; Theory of computation → Quantum complexity theory

**Keywords and phrases** Quantum cryptography, Complexity theory, Delegated quantum computation, Computing on encrypted data

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2019.6

---

<sup>1</sup> Corresponding author.



© Scott Aaronson, Alexandru Cojocaru, Alexandru Gheorghiu, and Elham Kashefi; licensed under Creative Commons License CC-BY

46th International Colloquium on Automata, Languages, and Programming (ICALP 2019).

Editors: Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi;

Article No. 6; pp. 6:1–6:13



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



**Category** Track A: Algorithms, Complexity and Games

**Related Version** A full version of the paper is available at <https://arxiv.org/abs/1704.08482>.

**Funding** *Scott Aaronson*: Vannevar Bush Fellowship from the US Department of Defense.

*Alexandru Cojocaru*: EPSRC grants EP/N003829/1, EP/M013243/1.

*Alexandru Gheorghiu*: MURI Grant FA9550-18-1-0161 and the IQIM, an NSF Physics Frontiers Center (NSF Grant PHY-1125565) with support of the Gordon and Betty Moore Foundation (GBMF-12500028).

*Elham Kashefi*: EPSRC grants EP/N003829/1, EP/M013243/1.

**Acknowledgements** We would like to thank the following people for useful discussions and comments: Petros Wallden, Matty J Hoban, Kousha Etesami, Marc Kaplan, Ronald de Wolf, Urmila Mahadev, Umesh Vazirani, Chris Heunen, Thomas Vidick, Ashley Montanaro, Tina Zhang and Pia Kullik. A.G. is in particular grateful to Petros Wallden and Matty J Hoban for their patience and for their help in clarifying several technical issues.

## 1 Introduction


An important area of research in modern cryptography is that of performing *computations on encrypted data*. The general idea is that a client wants to compute some function  $f$  on some input  $x$ , but lacks the computational power to do this in a reasonable amount of time. Luckily, the client has access to a computationally powerful server (cloud, cluster etc) which can compute  $f(x)$  quickly. However, because the computation might involve sensitive or classified information, or the server could be compromised remotely, we would like the input  $x$  to be hidden from the server at all times. The client can simply encrypt  $x$ , but this raises the question: how can the server compute  $f(x)$  if it doesn't know  $x$ ? The general problem of computing on encrypted data was first considered by Rivest, Adleman and Dertouzos [52]. Since then, instances of this problem have appeared in many areas of modern research including those of electronic voting, machine learning on encrypted data, program obfuscation and others [22, 32, 11, 28, 37, 41].

It was shown in 2009, when Gentry produced the first *fully homomorphic encryption scheme*, that performing classical computations on encrypted data is possible [29]. In homomorphic encryption the client has a pair of efficient algorithms  $(Enc, Dec)$ , which respectively perform encryption and decryption, and which satisfy the property  $Dec(f, x, Eval(f, Enc(x))) = f(x)$ , for any function  $f$  from some set  $\mathcal{C}$ . In other words, the server evaluates  $f$  on the encrypted input  $Enc(x)$  using  $Eval$  and returns this to the client which can then decrypt it to  $f(x)$ . Of course, the server should not be able to infer information about  $x$  from  $Enc(x)$ , a condition which is typically expressed through the criterion of *semantic security* [40]. If the set  $\mathcal{C}$  contains all polynomial-sized circuits then the scheme becomes a fully homomorphic encryption scheme, commonly abbreviated FHE. All known FHE schemes are secure under *cryptographic assumptions*.

Computing on encrypted data becomes particularly interesting when the server is a *quantum computer*. This is because efficient quantum algorithms have been found for various problems which are believed to be intractable for classical computers. In fact, it has been shown that if a classical computer and a quantum computer are both given black-box or oracle access to certain functions, then the quantum computer exponentially outperforms the classical computer [13, 55, 23, 5]. Classical clients would therefore be highly motivated to delegate problems to quantum computers. However, ensuring the privacy of their inputs is challenging. In particular, we'd have to solve the following problems:

- Devise an encryption scheme which is secure against quantum computers and does not leak information to the server about the client's input.
- Ensure that the encryption scheme allows the client to recover the output of the computation from the result provided by the quantum server.
- Ensure that the protocol is efficient for the client. Ideally, the number of rounds of interaction between the client and the server as well as the client's local computations, should scale at most polynomially with the size of the input.

In spite of these stringent requirements, protocols that achieve these properties already exist and are known collectively as *delegated blind quantum computing schemes* [26]. In such protocols, a probabilistic polynomial-time client is able to delegate polynomial-time quantum computations to a server in such a way that the client's input (apart from an upper bound on its size) is kept hidden from the server in an *information-theoretic* sense. All of the above schemes require the client and the server to share at least one round of quantum communication. *Universal Blind Quantum Computation* (UBQC), shown schematically in Figure 1, is an example of such a protocol [19].



■ **Figure 1** Universal Blind Quantum Computation [19]. In UBQC, a classical client augmented with the ability to prepare single-qubit states sends these qubits to the server along with instructions on how to entangle and measure them in order to perform a computation. The  $M(\delta_i)$  indicate measurement instructions and the  $b_i$  indicate the server's responses for these instructions (if he follows the protocol, these responses would represent the outcomes of the measurements that the client instructed him to perform).

The first blind delegation protocol was devised by Childs in [21], and since then these protocols have been improved and extended in various works [47, 31, 44, 27, 45, 46, 39, 38]. UBQC and related protocols require the client and the server to exchange only one quantum message, while the rest of the communication is classical [19, 9, 18]. This quantum message (which is sent by the client to the server) consists of a tensor product of single-qubit states. As such, the only quantum capability the client needs is the ability to prepare single-qubit states. In this paper, we explore two questions pertaining to blind delegation protocols:

- (1) Is there a scheme for blind quantum computing that is information-theoretically secure, and that requires only classical communication between client and server?
- (2) For schemes in which the client and the server are allowed one round of quantum communication, which functions can the client delegate to the server while maintaining information-theoretic security? In particular, could the client delegate the evaluation of NP-hard functions?

We provide some indication, based on complexity-theoretic conjectures, that the answer to the first question is no. In other words, provided these complexity-theoretic conjectures hold, a classical client running in polynomial time and communicating only classically with a

server cannot delegate arbitrary polynomial-time quantum computations to that server while keeping its input hidden in an information-theoretic sense. Importantly, our result does not contradict recent results on quantum fully homomorphic encryption with a classical client [43, 15], since those schemes are based on cryptographic assumptions: *we are interested only in information-theoretic security*.

In answer to the second question, we provide a complexity-theoretic upper bound on the types of functions that can be evaluated by UBQC-type protocols (i.e. protocols in which the client can send one quantum message to the server<sup>2</sup>). We show that, under plausible complexity-theoretic assumptions, this upper bound prevents the client from delegating NP-hard functions to the server. Thus, allowing for quantum communication between the client and the server expands the set of functions that the client can delegate to the server to include BQP, but not enough so as to include NP as well.

## 1.1 Main results

We phrase our results formally using the concept of a *generalised encryption scheme* (GES) introduced by Abadi, Feigenbaum and Killian [8]. Roughly speaking, a GES is a protocol between a probabilistic polynomial-time classical client and a computationally unbounded server for computing on encrypted data. The client sends the server a description of some function<sup>3</sup>  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ . Using some polynomial-time algorithm denoted  $E$ , the client encrypts its input  $x$ , and sends  $E(x)$  to the server. The server and the client then interact for a number of rounds which is polynomial in the length of  $x$ . Finally, using a polynomial-time decryption algorithm denoted  $D$ , the client decrypts the server's responses and obtains  $f(x)$  with probability  $1/2 + 1/\text{poly}(n)$ . Importantly throughout the protocol, the server learns no more than the length of  $x$ . Because the server is computationally unbounded, the scheme requires information-theoretic security. Abadi et al. gave a complexity theoretic upper bound on the types of functions that admit such a scheme. They showed that any function  $f$  that the client could delegate in a GES must be contained in the class  $\text{NP/poly} \cap \text{coNP/poly}$ .

The GES framework allows us to restate the questions we address in this paper as follows:

- (1) Can we design a GES for delegating BQP functions? Note that, by the Abadi et al. result, this is the same as asking whether  $\text{BQP} \subset \text{NP/poly} \cap \text{coNP/poly}$ . We will consider two variants on the GES framework: one which allows the client to delegate sampling problems to the server, and one in which the total communication between client and server is bounded by  $O(n^d)$ , for some constant  $d > 0$ . For the former, we show that having such a scheme for quantum sampling problems, like `BOSONSAMPLING`, implies that circuits exist which can compute the permanent of a matrix more efficiently than we believe is possible. For the latter, having a GES with bounded communication for polynomial-time quantum computation implies that  $\text{BQP} \subset \text{MA}/O(n^d)$ , and we argue that this containment is unlikely by providing an oracle separation between these classes.
- (2) If we change the GES framework to allow one round of quantum communication between the client and the server, what functions can the client delegate to the server? We answer this question by “quantising” the Abadi et al. result and showing that such

<sup>2</sup> In fact our result concerns protocols in which the client and the server start with *one round* of quantum communication, followed by polynomially-many rounds of classical communication. In other words, not only is there one quantum message from the client to the server, but the server is also allowed to respond with a quantum message.

<sup>3</sup> Unless otherwise specified, we restrict our attention to decision problems. This is why the function  $f$  has the codomain  $\{0, 1\}$ .

functions would be contained in  $\text{QCMA}/\text{qpoly} \cap \text{coQCMA}/\text{qpoly}$  (a quantum analogue of  $\text{NP}/\text{poly} \cap \text{coNP}/\text{poly}$ ). We also show that  $\text{QCMA}/\text{qpoly} \cap \text{coQCMA}/\text{qpoly}$  is unlikely to contain NP-hard functions.

The complete proofs for our results can be found in the full version of our paper [7].

### 1.1.1 Generalised encryption scheme for BQP decision problems

As we have mentioned, for the case of decision problems, Abadi et al. showed that the class of problems that a client can delegate to a server using the GES framework is contained in  $\text{NP}/\text{poly} \cap \text{coNP}/\text{poly}$ . They also observed that if NP-hard functions could be delegated by the client using a GES, then  $\text{NP} \subset \text{NP}/\text{poly} \cap \text{coNP}/\text{poly}$ , and, in particular,  $\text{NP} \subset \text{coNP}/\text{poly}$ . Yap showed that having such a containment leads to a collapse of the polynomial hierarchy at the third level [58]. In other words, it seems unlikely that NP-hard problems would admit a GES.

What about BQP-hard functions? The Abadi et al. result implies that having a GES for BQP-hard functions leads to  $\text{BQP} \subset \text{NP}/\text{poly} \cap \text{coNP}/\text{poly}$ . While we would like to argue, similarly, that such a containment leads to a collapse of the polynomial hierarchy, even  $\text{BQP} = \text{P}$  isn't known to lead to such a collapse. We instead consider a modified GES in which the total communication between the client and the server is upper bounded by a polynomial of *fixed* degree,  $d > 0$ , in the size of the input<sup>4</sup>. In that case, it can be shown that  $\text{BQP} \subset \text{MA}/\text{O}(n^d) \cap \text{coMA}/\text{O}(n^d)$ . We argue that this containment is unlikely based on the following result:

► **Theorem 1.** *For each  $d \in \mathbb{N}$ , there exists an oracle  $O_d$  such that  $\text{BQP}^{O_d}$  is not contained in  $(\text{MA}/\text{O}(n^d))^{O_d}$ .*

Essentially, the theorem shows that relative to an oracle  $O_d$ , there are problems that can be solved by a polynomial-time quantum algorithm, but which a classical client cannot delegate to a server in a GES with bounded communication. Since the oracle is parameterised by  $d$ , we are in fact defining a family of oracles. The specific problem on which the oracle  $O_d$  is based is a version of *Simon's problem* [55]. Simon's problem is the following: for an input of size  $n$ , and given oracle access to a function  $g : \{0, 1\}^n \rightarrow \{0, 1\}^n$  that is guaranteed to be either an injective function, or a 2-to-1 and periodic function<sup>5</sup>, the task is to decide which is the case. Simon provided a polynomial-time quantum algorithm for solving this problem, thus showing that it belongs to BQP (relative to the function oracle). For the case in which one should accept when the function is 2-to-1, the problem can be shown to be outside of MA (relative to the function oracle). As such, Simon's problem provides an oracle separation between BQP and MA.

In Simon's original construction, the oracle function is the same for all inputs of size  $n$ . Note that, this version of the problem can be solved with one bit of advice: for all inputs of size  $n$ , the advice bit simply specifies whether the function is 1-to-1 or 2-to-1 and periodic. Therefore such a setup would not be useful in our case. For this reason, in our proof of Theorem 1, the function that the oracle provides access to is input-dependent. The problem we define, relative to this oracle, is again to decide whether the function is 1-to-1 or the function is 2-to-1 and periodic. However, we can show that, by considering a sufficiently

<sup>4</sup> Note that we impose no such restriction on the running time of the client.

<sup>5</sup> In other words, there exists a period  $s \in \{0, 1\}^n$ ,  $s \neq 0^n$ , such that for all  $x, y \in \{0, 1\}^n$ ,  $x \neq y$ , it is the case that  $g(x) = g(y)$  iff.  $x = s \oplus y$ .

large domain for these functions – in other words, by letting  $g : \{0, 1\}^{n^D} \rightarrow \{0, 1\}^{n^D}$  for some  $D > d$  – the problem is not contained in  $(\text{MA}/\text{O}(n^d))^{O_d}$ , but is nevertheless contained in BQP. The proof uses a diagonalisation argument and can be found in the full version of our paper [7].

Unfortunately, the same oracle cannot be used to separate BQP from NP/poly. This is because  $D$  is a function of  $d$ ; to prove a separation with respect to NP/poly, where the length of the advice string can be any polynomial, we would have to find an oracle that works *for all* possible values of  $d$ . It would be interesting to see whether the oracle that Raz and Tal [51] recently used to prove a separation between BQP and PH could also be used in order to separate BQP from NP/poly, or even from PH/poly. We leave this as an open problem.

One can argue that oracle results do not constitute compelling evidence on the relationships between complexity classes. For example, it has been known for a while that there exist oracles  $O_1, O_2$  such that  $\text{P}^{O_1} \neq \text{NP}^{O_1}$  but  $\text{P}^{O_2} = \text{NP}^{O_2}$ , and that, while  $\text{IP} = \text{PSPACE}$ , there is an oracle such that  $\text{IP}^O \neq \text{PSPACE}^O$ . Nonetheless, oracles allow us to study the query complexity of problems in different models of computation. In fact, there are situations in practice where computer programs are restricted to making black-box calls to functions in order to determine their properties [36]. Apart from this, oracle results have also inspired a number of important developments in algorithms and complexity theory<sup>6</sup>. For more arguments concerning the usefulness of oracle results, see Section 1.3 of [2].

### 1.1.2 Generalised encryption scheme for BQP sampling problems

We consider what would happen if we have a generalised encryption scheme which allowed a client to delegate a sampling problem, such as BOSONSAMPLING, to the server. BOSONSAMPLING, defined by Aaronson and Arkhipov in [6], is essentially the problem of simulating the statistics of photons (bosons) passing through a linear optics network. One starts with a configuration of identical photons in known locations (referred to as *modes*). The photons then pass through the linear optics network, which consists of optical elements (beamsplitters and phase shifters). Finally, one performs a measurement to determine the new locations of the photons in the *output modes* of the system. The reason this is referred to as a sampling problem is because we have a probability distribution over the different configurations of photons in the output modes. In *exact* BOSONSAMPLING, which is the problem we consider, the task is to produce a sample from that probability distribution. Aaronson and Arkhipov showed that the probability of observing a particular configuration of photons is proportional to the squared permanent of a matrix that can be obtained efficiently from the description of the optical network. They also showed that no polynomial-time probabilistic algorithm can sample from this distribution, unless the polynomial hierarchy collapses at the third level [6]. As such, while a quantum computer could simulate the optical network and sample from the target distribution in polynomial time, it seems unlikely that classical computers could do the same.

Sampling problems, like BOSONSAMPLING, are of interest because of their potential use in demonstrating *quantum computational supremacy* [35]. This entails having a quantum device perform a computational task that no classical computer would be able to reproduce *efficiently*. Sampling problems are natural candidates for this task for two main reasons. Firstly, many of the quantum sampling problems that have been considered could in principle be performed on a small-scale quantum computer having up to (or on the order of) 100

<sup>6</sup> A notable example is the fact that Simon’s oracle separation between BPP and BQP led to Shor’s algorithm for factoring and computing the discrete logarithm [54]



qubits and not requiring fault-tolerance [50]. Secondly, it has been shown that having a polynomial-time classical algorithm that can sample from the distribution of the quantum sampling problem leads to a collapse of the polynomial hierarchy. Contrast this to tasks such as factoring for which the existence of an efficient classical algorithm is not known to lead to any “disastrous” complexity-theoretic consequences.

Given that sampling problems are primarily considered in the context of demonstrating quantum computational supremacy, one could ask why a client would like to delegate such a problem to the server via a GES. Firstly, all existing schemes for blind delegated quantum computation allow the client to delegate both decision and sampling problems [26]. It is therefore natural to ask whether such a scheme, involving only classical communication, can also exist. Secondly, there is currently no known way for a classical client to efficiently certify whether the server has sampled from the correct distribution (at least with an information-theoretic guarantee). In fact, in certain cases the client would require exponentially many samples from the server in order to perform this verification [34]. However, if a GES for quantum sampling problems existed, the client might be able to leverage it in order to perform the certification, in the same way that many delegated quantum computation protocols leverage blindness to achieve verifiability [30]. Finally, due to the equivalence between sampling and searching shown in [3], our result holds if we substitute sampling problems with search problems.

In a GES for exact BOSONSAMPLING, the client’s input would be a description of a linear optics network<sup>7</sup>. The client would like to delegate to the server the task of sampling from the BOSONSAMPLING distribution associated with this network, while keeping the description of the network hidden from the server. In other words, upon interacting with the server and decrypting its responses, the client should obtain a sample from the BOSONSAMPLING distribution. At the same time, the server learns at most an upper bound on the size of the network. We show the following:

► **Theorem 2.** *If exact BOSONSAMPLING admits a GES, then for any matrix  $X \in \{-1, 0, 1\}^{n \times n}$ , there exist circuits of size  $2^{n - \Omega(\frac{n}{\log n})}$ , making polynomially-sized queries to an  $\text{NP}^{\text{NP}}$  oracle, for computing the permanent of  $X$ .*

Computing the permanent of a matrix is a problem known to be #P-hard. By Toda’s theorem, this means that if computing the permanent were possible at any level of the polynomial hierarchy, the hierarchy would collapse at that level [56]. Moreover, the best known algorithm for computing the permanent, by Björklund, has a run-time of  $2^{n - \Omega(\sqrt{n/\log(n)})}$  [14]. Prior to that, the leading algorithm for computing the permanent was Ryser’s algorithm, developed over 50 years ago, which requires  $O(n2^n)$  arithmetic operations [53]. We conjecture that the circuits of Theorem 2 do not exist and, thus, that there can be no GES for BOSONSAMPLING.

### 1.1.3 Quantum generalised encryption scheme

While having a GES for delegating BQP computations seems unlikely, we know that giving the client some minimal quantum capabilities removes this limitation: schemes such as UBQC exist which allow for the information-theoretically secure blind delegation of quantum

<sup>7</sup> In principle, one could also specify the configuration of the photons in the input modes as part of the client’s input. Equivalently, however, one can always initialise the input modes to some fixed initial state, and produce whichever starting state is in fact desired by altering the linear optics network.

computations. In the spirit of the Abadi et al. result, it is natural to consider *quantum generalised encryption schemes* (or QGES), in which the client is no longer classical, and investigate the complexity-theoretic upper bounds on functions that admit such a protocol. For the QGES, we are still assuming information-theoretic security and that the encryption scheme leaks at most the size of the input. However, unlike the GES, the client is now assumed to be a quantum computer performing polynomial-time computations<sup>8</sup>. Additionally, the client and the server perform one round of quantum communication at the beginning of the protocol. The rest of the communication is classical.

We impose one other restriction on the QGES, known as *offline-ness*. Roughly speaking, an offline protocol is one in which the client does not need to commit to any particular input (of a given size), after having sent the quantum message to the server. The quantum message only depends on the size of the input. We note that offline-ness is a property which UBQC and all other currently known blind quantum computing protocols share. From a practical perspective, this presents the client with the option of sending the first quantum message to the server and deciding at a later time on which input the server should perform the computation. One could imagine, for instance, that the client and the server have access to a quantum channel for a limited amount of time. In practice, such a situation can occur if the communication between the parties is mediated by a satellite, as is the case with satellite-based quantum-key distribution [42]. In this case, the satellite is in the line of sight of the two parties for only a few minutes at a time. Our result is the following:

► **Theorem 3.** *The class of functions that a client can delegate to a server in an offline QGES is contained in  $\text{QCMA}/\text{qpoly} \cap \text{coQCMA}/\text{qpoly}$ .*

Note that the class  $\text{QCMA}/\text{qpoly} \cap \text{coQCMA}/\text{qpoly}$  can be seen as a quantum analogue of the class  $\text{NP}/\text{poly} \cap \text{coNP}/\text{poly}$  which we encounter in the GES case. We therefore view Theorem 3 as a “quantisation” of the Abadi et al. bound on the power of generalised encryption schemes.

Again, in the spirit of the Abadi et al. result, one can ask whether NP-complete functions are contained in  $\text{QCMA}/\text{qpoly} \cap \text{coQCMA}/\text{qpoly}$ . In other words: does giving quantum capabilities to the client increase the class of functions that it can securely delegate so that this class contains NP? We give an indication that the answer is no:

► **Theorem 4.**  $\text{NP} \subset \text{QCMA}/\text{qpoly} \cap \text{coQCMA}/\text{qpoly}$  *implies*  $\text{coNP}^{\text{NP}^{\text{NP}}} \subseteq \text{NP}^{\text{NP}^{\text{PromiseQMA}}}$ .

Note that if PromiseQMA in the above expression were replaced with NP, this would imply a collapse of the polynomial hierarchy at the third level. Our result is as close to a collapse of the polynomial hierarchy as one can reasonably hope to get, given a quantum hypothesis. Hence, while a QGES does allow the client to delegate BQP computations, it seems to be no more useful than the regular GES for delegating NP-hard functions.

One could ask why we would even be interested in delegating NP-hard problems to a quantum computer, given that we do not expect quantum computers to be able to solve such problems in polynomial time [1]. First of all, from a theoretical perspective, note that in the QGES formalism we are not limiting the server to polynomial-time quantum computations, but instead assuming that it has unbounded computational power. Therefore, the way to view this result is not as “how can a client blindly delegate the evaluation of NP-hard functions to a quantum computer?” but as “can quantum communication help in blindly delegating the evaluation of NP-hard functions to an unbounded server?”

From a practical perspective, while we do not believe that quantum computers can solve NP-complete problems in polynomial time, they could, in principle, solve such problems quadratically faster than classical computers, thanks to Grover’s algorithm [33]. Even

<sup>8</sup> It should be noted that our upper bound on the power of QGES schemes also holds if the client is restricted to BPP computations (as is the case in UBQC), since  $\text{BPP} \subseteq \text{BQP}$ .



though the speedup of Grover's algorithm is only quadratic, from (say)  $2^n$  to  $2^{n/2}$ , our result is only concerned with the length of the computation performed on the client side, and therefore applies to Grover's algorithm just as it would to a quantum algorithm achieving exponential speedup. In fact, as is mentioned in [4], there are NP-complete problems for which quantum computers provide a superpolynomial speedup, at least with respect to the best known classical algorithms. Our no-go theorem indicates that clients cannot exploit such speedups by delegating the computation to the server, even when allowing some quantum communication, if we also want to keep their inputs hidden in an information-theoretic sense.

Proofs of these results can be found in the full version of our paper [7].

## 1.2 Related work

As mentioned, the problem of computing on encrypted data was first considered by Rivest, Adleman and Dertouzos [52], which then led to the development of *homomorphic encryption* and eventually to fully homomorphic encryption with Gentry's scheme [29]. Since then there have been many other FHE protocols relying on more standard cryptographic assumptions and having more practical requirements [17, 16, 57].

While FHE is similar to the GES in many respects, there are also significant differences. For starters, FHE protocols have only one round of interaction between the client and the server, whereas a GES allows for polynomially many rounds. Additionally, the GES assumes the server is computationally unbounded and hence requires information-theoretic security. In contrast, FHE relies on computational security. More precisely FHE schemes have semantic security against polynomial-time (quantum) algorithms [29].

The problem of *quantum* computing on encrypted data was introduced by Childs [21] and Arrighi and Salvail [12]. Further development eventually led to UBQC [19, 9] and a scheme of Broadbent [18]. The latter was followed by the construction of the first schemes for quantum fully homomorphic encryption (QFHE) [20, 24]. For a review of blind quantum computing and related protocols see [26].

In the QFHE schemes of [20, 24], the server is a polynomial-time quantum computer and the client has some quantum capabilities of its own, although it is not able to perform universal quantum computations. Both the size of the exchanged messages and the number of operations of the client are polynomial in the size of the input. More recently, QFHE schemes have been proposed in which the client is completely classical [43, 15]. Similar to FHE, these protocols rely on computational assumptions for security [10] and involve one round of back and forth interaction between the client and the server. QFHE with information-theoretic security (and a computationally unbounded server) has been considered by Yu et al. in [59], where it is shown that it is impossible to have such a scheme for arbitrary unitary operations (or even arbitrary reversible classical computations). This result was later reproved by Newmann and Shi using quantum random-access codes [49]. In relation to our work, QFHE with information-theoretic security can be viewed as a one-round QGES in which the server responds with a quantum message. The complexity-theoretic upper bound we prove for QGES computable functions would then apply to QFHE as well (provided that in QFHE we only leak the size of the input to the server), since our proof allows a quantum message from the server just as well as a classical message.

The possibility of a classical client delegating a blind computation to a quantum server was considered by Morimae and Koshihara [48]. They showed that such a protocol in which the client leaks no information about its input to the server and there is only one round of interaction leads to  $\text{BQP} \subseteq \text{NP}$ , considered an unlikely containment. We consider the more general setting of a GES for BQP functions, where the number of rounds can be polynomial

in the size of the input and we allow the encryption to leak the size of the input. In fact, the question of whether a GES, as defined in Abadi et al. [8], can exist for quantum computations was raised before by Dunjko and Kashefi [25].

### 1.3 Future work

As we remarked in Section 1.1, in the case of decision problems, the existence of a GES with bounded communication, for polynomial-time quantum computations, leads to the inclusion  $\text{BQP} \subset \text{MA}/\text{O}(n^d)$ . We argue that this containment is unlikely based on the existence of an oracle separating the two complexity classes. A natural extension of this result would be to prove an oracle separation between BQP and NP/poly. This would provide more compelling evidence that a GES for quantum computations cannot exist.

In the case of sampling problems, we showed that a GES for `BOSONSAMPLING` implies the existence of circuits of size  $2^{n-\Omega(\frac{n}{\log n})}$ , making polynomially-sized queries to an  $\text{NP}^{\text{NP}}$  oracle, for computing matrix permanents. Can this result be strengthened so as to provide circuits for computing matrix permanents that would be ruled out by the *strong exponential-time hypothesis*? Alternatively, could one use other quantum sampling problems (such as random circuit sampling or IQP problems [35]) to show that having a GES for such a problem leads to a collapse of the polynomial hierarchy?

We also defined the QGES, which extends the GES by allowing the client to send one quantum message to the server, and gave an upper bound for the set of functions that can be delegated using an offline version of such a scheme. The immediate question one could ask is: what upper bound can we give for an online QGES? A related question is: what upper bound can we give for a QGES that allows all of the communication between the client and the server to be quantum? The difficulty in answering both of these questions is that the offline property of the QGES is what allowed us to relate the set of functions that can be delegated to advice classes. Without this property, it seems that a different approach would be needed to provide a complexity-theoretic upper bound.

Another direction that can be explored has to do with the size of the quantum communication between the client and the server. In a QGES in which the client's quantum message is logarithmic or poly-logarithmic in the size of the input (while the classical communication is still polynomial), is it still possible to delegate BQP functions to the server? Of course, this question only makes sense if we assume that the client is not able to perform BQP computations itself.

---

### References

- 1 Scott Aaronson. Limits on efficient computation in the physical world. *arXiv preprint*, 2004. [arXiv:quant-ph/0412143](https://arxiv.org/abs/quant-ph/0412143).
- 2 Scott Aaronson. BQP and the polynomial hierarchy. In *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, STOC '10, pages 141–150, New York, NY, USA, 2010. ACM. doi:10.1145/1806689.1806711.
- 3 Scott Aaronson. The Equivalence of Sampling and Searching. In *Proceedings of the 6th International Conference on Computer Science: Theory and Applications*, CSR'11, pages 1–14, Berlin, Heidelberg, 2011. Springer-Verlag. URL: <http://dl.acm.org/citation.cfm?id=2017990.2017991>.
- 4 Scott Aaronson.  $\text{P} \stackrel{?}{=} \text{NP}$ , 2017. URL: <https://www.scottaaronson.com/papers/pnp.pdf>.
- 5 Scott Aaronson and Andris Ambainis. Forrelation: A Problem That Optimally Separates Quantum from Classical Computing. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing*, STOC '15, pages 307–316, New York, NY, USA, 2015. ACM. doi:10.1145/2746539.2746547.

- 6 Scott Aaronson and Alex Arkhipov. The Computational Complexity of Linear Optics. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 333–342, New York, NY, USA, 2011. ACM. doi:10.1145/1993636.1993682.
- 7 Scott Aaronson, Alexandru Cojocaru, Alexandru Gheorghiu, and Elham Kashefi. Complexity-theoretic limitations on blind delegated quantum computation. *arXiv preprint*, 2017. arXiv:1704.08482.
- 8 M. Abadi, J. Feigenbaum, and J. Kilian. On Hiding Information from an Oracle. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, STOC '87, pages 195–203, New York, NY, USA, 1987. ACM. doi:10.1145/28395.28417.
- 9 Dorit Aharonov, Michael Ben-Or, and Elad Eban. Interactive Proofs For Quantum Computations. In *Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 5-7, 2010. Proceedings*, pages 453–469, 2010. URL: <http://conference.itcs.tsinghua.edu.cn/ICS2010/content/papers/35.html>.
- 10 Gorjan Alagic, Anne Broadbent, Bill Fefferman, Tommaso Gagliardoni, Christian Schaffner, and Michael St. Jules. *Computational Security of Quantum Encryption*, pages 47–71. Springer International Publishing, Cham, 2016. doi:10.1007/978-3-319-49175-2\_3.
- 11 Joël Alwen, Manuel Barbosa, Pooya Farshim, Rosario Gennaro, S. Dov Gordon, Stefano Tessaro, and David A. Wilson. *On the Relationship between Functional Encryption, Obfuscation, and Fully Homomorphic Encryption*, pages 65–84. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. doi:10.1007/978-3-642-45239-0\_5.
- 12 Pablo Arrighi and Louis Salvail. Blind quantum computation. *International Journal of Quantum Information*, 04(05):883–898, 2006. doi:10.1142/S0219749906002171.
- 13 Ethan Bernstein and Umesh Vazirani. Quantum Complexity Theory. *SIAM J. Comput.*, 26(5):1411–1473, October 1997. doi:10.1137/S0097539796300921.
- 14 Andreas Björklund. Below All Subsets for Some Permutational Counting Problems . In Rasmus Pagh, editor, *15th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2016)*, volume 53 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:11, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.SWAT.2016.17.
- 15 Zvika Brakerski. Quantum FHE (Almost) As Secure as Classical. Cryptology ePrint Archive, Report 2018/338, 2018. URL: <https://eprint.iacr.org/2018/338>.
- 16 Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) Fully Homomorphic Encryption Without Bootstrapping. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, pages 309–325, New York, NY, USA, 2012. ACM. doi:10.1145/2090236.2090262.
- 17 Zvika Brakerski and Vinod Vaikuntanathan. Efficient Fully Homomorphic Encryption from (Standard) LWE. In *Proceedings of the 2011 IEEE 52Nd Annual Symposium on Foundations of Computer Science*, FOCS '11, pages 97–106, Washington, DC, USA, 2011. IEEE Computer Society. doi:10.1109/FOCS.2011.12.
- 18 Anne Broadbent. Delegating private quantum computations. *Canadian Journal of Physics*, 93(9):941–946, 2015. doi:10.1139/cjp-2015-0030.
- 19 Anne Broadbent, Joseph Fitzsimons, and Elham Kashefi. Universal blind quantum computation. In *Proceedings of the 50th Annual Symposium on Foundations of Computer Science*, FOCS '09, pages 517–526. IEEE Computer Society, 2009. doi:10.1109/FOCS.2009.36.
- 20 Anne Broadbent and Stacey Jeffery. Quantum Homomorphic Encryption for Circuits of Low T-gate Complexity. In *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part II*, pages 609–629, 2015. doi:10.1007/978-3-662-48000-7\_30.
- 21 Andrew M. Childs. Secure Assisted Quantum Computation. *Quantum Info. Comput.*, 5(6):456–466, September 2005. URL: <http://dl.acm.org/citation.cfm?id=2011670.2011674>.

- 22 Ivan Damgård, Jens Groth, and Gorm Salomonsen. *The Theory and Implementation of an Electronic Voting System*, pages 77–99. Springer US, Boston, MA, 2003. doi:10.1007/978-1-4615-0239-5\_6.
- 23 J Niel De Beaudrap, Richard Cleve, John Watrous, et al. Sharp quantum versus classical query complexity separations. *Algorithmica*, 34(4):449–461, 2002.
- 24 Yfke Dulek, Christian Schaffner, and Florian Speelman. *Quantum Homomorphic Encryption for Polynomial-Sized Circuits*, pages 3–32. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016. doi:10.1007/978-3-662-53015-3\_1.
- 25 Vedran Dunjko and Elham Kashefi. Blind quantum computing with two almost identical states, 2016. arXiv:1604.01586.
- 26 Joseph F Fitzsimons. Private quantum computation: an introduction to blind quantum computing and related protocols. *npj Quantum Information*, 3(1):23, 2017.
- 27 Joseph F Fitzsimons and Elham Kashefi. Unconditionally verifiable blind quantum computation. *Physical Review A*, 96(1):012303, 2017.
- 28 Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate Indistinguishability Obfuscation and Functional Encryption for All Circuits. In *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, FOCS '13*, pages 40–49, Washington, DC, USA, 2013. IEEE Computer Society. doi:10.1109/FOCS.2013.13.
- 29 Craig Gentry. Fully Homomorphic Encryption Using Ideal Lattices. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing, STOC '09*, pages 169–178, New York, NY, USA, 2009. ACM. doi:10.1145/1536414.1536440.
- 30 Alexandru Gheorghiu, Theodoros Kapourniotis, and Elham Kashefi. Verification of quantum computation: An overview of existing approaches. *Theory of computing systems*, pages 1–94, 2018.
- 31 Vittorio Giovannetti, Lorenzo Maccone, Tomoyuki Morimae, and Terry G. Rudolph. Efficient Universal Blind Quantum Computation. *Phys. Rev. Lett.*, 111:230501, December 2013. doi:10.1103/PhysRevLett.111.230501.
- 32 Thore Graepel, Kristin Lauter, and Michael Naehrig. ML Confidential: Machine Learning on Encrypted Data. In *Proceedings of the 15th International Conference on Information Security and Cryptology, ICISC'12*, pages 1–21, Berlin, Heidelberg, 2013. Springer-Verlag. doi:10.1007/978-3-642-37682-5\_1.
- 33 Lov K. Grover. A Fast Quantum Mechanical Algorithm for Database Search. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing, STOC '96*, pages 212–219, New York, NY, USA, 1996. ACM. doi:10.1145/237814.237866.
- 34 Dominik Hangleiter, Martin Kliesch, Jens Eisert, and Christian Gogolin. Sample complexity of device-independently certified" quantum supremacy". *arXiv preprint*, 2018. arXiv:1812.01023.
- 35 Aram W Harrow and Ashley Montanaro. Quantum computational supremacy. *Nature*, 549(7671):203, 2017.
- 36 Thomas Jansen. On the Black-Box Complexity of Example Functions: The Real Jump Function. In *Proceedings of the 2015 ACM Conference on Foundations of Genetic Algorithms XIII, FOGA '15*, pages 16–24, New York, NY, USA, 2015. ACM. doi:10.1145/2725494.2725507.
- 37 Arjan Jeckmans, Andreas Peter, and Pieter Hartel. *Efficient Privacy-Enhanced Familiarity-Based Recommender System*, pages 400–417. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. doi:10.1007/978-3-642-40203-6\_23.
- 38 Elham Kashefi, Luka Music, and Petros Wallden. The Quantum Cut-and-Choose Technique and Quantum Two-Party Computation, 2017. arXiv:1703.03754.
- 39 Elham Kashefi and Petros Wallden. Garbled quantum computation. *Cryptography*, 1(1):6, 2017.
- 40 Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography*. CRC press, 2014.

- 41 Kristin E. Lauter. Practical Applications of Homomorphic Encryption. In *Proceedings of the 2012 ACM Workshop on Cloud Computing Security Workshop, CCSW '12*, pages 57–58, New York, NY, USA, 2012. ACM. doi:10.1145/2381913.2381924.
- 42 Sheng-Kai Liao, Wen-Qi Cai, Wei-Yue Liu, Liang Zhang, Yang Li, Ji-Gang Ren, Juan Yin, Qi Shen, Yuan Cao, Zheng-Ping Li, et al. Satellite-to-ground quantum key distribution. *Nature*, 549(7670):43, 2017.
- 43 Urmila Mahadev. Classical homomorphic encryption for quantum circuits. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 332–338. IEEE, 2018.
- 44 Atul Mantri, Tommaso F Demarie, and Joseph F Fitzsimons. Universality of quantum computation with cluster states and (X, Y)-plane measurements. *Scientific reports*, 7:42861, 2017.
- 45 Atul Mantri, Carlos A. Pérez-Delgado, and Joseph F. Fitzsimons. Optimal Blind Quantum Computation. *Phys. Rev. Lett.*, 111:230502, December 2013. doi:10.1103/PhysRevLett.111.230502.
- 46 Tomoyuki Morimae, Vedran Dunjko, and Elham Kashefi. Ground State Blind Quantum Computation on AKLT State. *Quantum Info. Comput.*, 15(3-4):200–234, March 2015. URL: <http://dl.acm.org/citation.cfm?id=2871393>. 2871395.
- 47 Tomoyuki Morimae and Keisuke Fujii. Blind quantum computation protocol in which Alice only makes measurements. *Phys. Rev. A*, 87:050301, May 2013. doi:10.1103/PhysRevA.87.050301.
- 48 Tomoyuki Morimae and Takeshi Koshihba. Impossibility of perfectly-secure delegated quantum computing for classical client, 2014. arXiv:1407.1636.
- 49 Michael Newman and Yaoyun Shi. Limitations on Transversal Computation through Quantum Homomorphic Encryption. *Quantum Information and Computation*, 18:0927–0948, 2018.
- 50 John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.
- 51 Ran Raz and Avishay Tal. Oracle Separation of BQP and PH. *eccc preprint TR18-107*, 2018. Eprint:<https://eccc.weizmann.ac.il/report/2018/107/>.
- 52 Ronald L Rivest, Len Adleman, and Michael L. Dertouzos. On data banks and privacy homomorphisms. *Foundations of Secure Computation*, 4(11):169–180, 1978.
- 53 Herbert John Ryser. *Combinatorial mathematics*, volume 14. JSTOR, 1963.
- 54 Peter W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Review*, 41(2):303–332, 1999. doi:10.1137/S0036144598347011.
- 55 Daniel R. Simon. On the Power of Quantum Computation. *SIAM Journal on Computing*, 26(5):1474–1483, 1997. doi:10.1137/S0097539796298637.
- 56 Seinosuke Toda. PP is as hard as the polynomial-time hierarchy. *SIAM Journal on Computing*, 20(5):865–877, 1991. doi:10.1137/0220053.
- 57 Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully Homomorphic Encryption over the Integers. In *Proceedings of the 29th Annual International Conference on Theory and Applications of Cryptographic Techniques, EUROCRYPT'10*, pages 24–43, Berlin, Heidelberg, 2010. Springer-Verlag. doi:10.1007/978-3-642-13190-5\_2.
- 58 Chee K. Yap. Some consequences of non-uniform conditions on uniform classes. *Theoretical Computer Science*, 26(3):287–300, 1983. doi:10.1016/0304-3975(83)90020-8.
- 59 Li Yu, Carlos A. Pérez-Delgado, and Joseph F. Fitzsimons. Limitations on information-theoretically-secure quantum homomorphic encryption. *Phys. Rev. A*, 90:050303, November 2014. doi:10.1103/PhysRevA.90.050303.