

Secure Quantum Computation with Classical Communication

James Bartusek

University of California, Berkeley

Abstract. The study of secure multi-party computation (MPC) has thus far been limited to the following two settings: every party is fully classical, or every party has quantum capabilities. This paper studies a notion of MPC that allows some classical and some quantum parties to securely compute a quantum functionality over their joint private inputs. In particular, we construct constant-round *composable* protocols for blind and verifiable classical delegation of quantum computation, and give applications to secure quantum computation with classical communication. Assuming QLWE (the quantum hardness of learning with errors), we obtain the following (maliciously-secure) protocols for computing any BQP (bounded-error quantum polynomial-time) functionality.

- A six-round protocol between one quantum server and multiple classical clients in the CRS (common random string) model.
- A three-round protocol between one quantum server and multiple classical clients in the PKI (public-key infrastructure) + QRO (quantum random oracle) model.
- A two-message protocol between quantum sender and classical receiver (a quantum non-interactive secure computation protocol), in the QRO model.

To enable composability of classical verification of quantum computation, we require the notion of *malicious blindness*, which stipulates that the prover does not learn anything about the verifier’s delegated computation, even if it is able to observe whether or not the verifier accepted the proof. To construct a protocol with malicious blindness, we use a classical verification protocol for sampBQP computation (Chung et al., Arxiv 2020), which in general has inverse polynomial soundness error, to prove honest evaluation of QFHE (quantum fully-homomorphic encryption) ciphertexts with negligible soundness error. Obtaining a constant-round protocol requires a strong parallel repetition theorem for classical verification of quantum computation, which we show following the “nearly orthogonal projector” proof strategy (Alagic et al., TCC 2020).

1 Introduction

Secure multi-party computation (MPC) is a fundamental cryptographic task that allows for multiple parties to securely evaluate a function on their joint private inputs. The study of MPC is foundational to the field of modern cryptography [48,32,12,19] and has since only increased in depth and scope.

Naturally, the vast majority of MPC literature considers the task of securely evaluating a classical functionality over classical inputs. However, the emergence of quantum computing technology raises several interesting and increasingly relevant questions for the field of secure computation. Indeed, secure computation of quantum functionalities over potentially quantum inputs also has a rich history of study, including many recent works [23,11,47,27,38,25,4,34,8,9,37].

One drawback of each of the above multi-party quantum computation (MPQC) protocols is that they require each party to operate a quantum computer, or at least be able to manipulate some quantum information. Indeed, personal quantum computers remain far from a reality, and it appears as if quantum computation will be concentrated in the hands of a few technologically-advanced entities for the foreseeable future. Thus, recent years have seen a major research effort towards the goal of *classical delegation of quantum computation*, which allows a classical client to enlist the resources of a quantum server, ideally without comprising the privacy of the client’s data or the integrity of the computation [40,41,14,30,3,20,21].

These works consider a single classical client with (potentially private) input, interacting with an input-less quantum server. Thus, they do not address the possibility of *multi-party* quantum computation with classical communication, and do not attempt to realize fully simulation-secure protocols, which is the gold standard notion of security for distributed computation. In this work, we address the following feasibility question for the first time, where “securely compute” refers to simulation security against arbitrarily malicious parties.

Can multiple parties, some of which do not have any quantum capabilities, securely compute a quantum functionality over their joint private inputs?

1.1 Results

We study the notion of MPQC with classical communication secure against a dishonest majority of arbitrarily malicious parties. We focus on MPQC for BQP (bounded-error quantum polynomial-time) computation.¹ We capture BQP by considering “pseudo-deterministic” quantum functionalities $D(\cdot)$ that on classical input x , produce a fixed classical output z except with negligible probability.

Composable blind CVQC. We begin by considering a simple two-party functionality between classical client and quantum server, defined by a pseudo-deterministic circuit $D(\cdot)$. It takes an input x from the client and a bit b from the server (indicating honest or dishonest behavior) and delivers the output $D(x)$

¹ This is as opposed to a more general class of quantum functionalities that may output an arbitrary distribution over classical strings (see discussion in Section 1.3). Note that this distinction generally does not arise in the classical setting, since one can make any randomized functionality deterministic by fixing the random coins. In the quantum setting, this strategy will not always work since randomness can come from measurement.

to the client if $b = 0$ and the output \perp if $b = 1$. We say that a protocol with classical communication that securely implements this ideal functionality is a *composable blind CVQC* protocol. Our first result is described in the following informal theorem.

Theorem 1. *(Informal) Assuming the quantum hardness of learning with errors (QLWE), there exists a four-round composable blind CVQC protocol. The protocol can be made two rounds in the quantum random oracle model.*

Next, we give applications of composable blind CVQC to secure quantum computation.

Multi-party results. In the multi-party setting, we construct two protocols, both of which only require a *single* party (called the server) to have quantum capabilities. This setting of one quantum server and several classical clients can be viewed as a quantum analogue of “cloud-assisted” MPC, introduced by [6]. In their setting, several clients wish to securely outsource the bulk of some computation to a single powerful server, and they require that the client computation is much smaller than the functionality to be computed. In our setting, we consider several classical clients that wish to outsource a quantum computation to a single quantum server, and we require that the client computation is entirely classical (though it may grow with the size of the functionality to be computed).

The features of our first protocol are described in the following theorem.

Theorem 2. *(Informal) Assuming the quantum hardness of learning with errors (QLWE), there exists a six-round protocol (in the common random string model²) between multiple classical clients and one quantum server for computing any pseudo-deterministic quantum functionality over the private inputs of the clients. The protocol tolerates any coalition (including client-server collusion) of malicious quantum polynomial-time adversaries.*

We next study cloud-assisted MPQC with the following interaction pattern.

- Round 1: each classical client P_i computes and broadcasts an encryption ct_i of their input x_i .
- Round 2: the quantum server computes and broadcasts an encryption $\tilde{\text{ct}}$ of the output.
- Round 3: the clients participate in a one-round decryption procedure that delivers output y_i to each client P_i .

The feasibility of this interaction pattern in the classical setting was established by [6] in the PKI (public-key infrastructure) model, where each client can publish a succinct and reusable public key before the protocol begins.³ Here, we

² A constant-round protocol in the plain model can also be obtained by using constant-round post-quantum MPC [2] to set up the CRS. However, this introduces more rounds and assumptions (in particular, a circular-security assumption).

³ It was later shown how to remove the PKI via multi-key fully homomorphic encryption [44].

show how to achieve three-round cloud-assisted MPQC in the QROM (quantum random oracle model),⁴ also assuming a PKI setup, as described in the following theorem.

Theorem 3. *(Informal) Assuming QLWE, there exists a three-round protocol (in the QRO + PKI model) between multiple classical clients and one quantum server for computing any psuedo-deterministic quantum functionality over the private inputs of the clients. The protocol tolerates any coalition (including client-server collusion) of malicious quantum polynomial-time adversaries.*

Two-party results. In the two-party setting, we show how to construct a *round-optimal* (two-message) protocol where one party receives output. That is, we consider a quantum sender \mathcal{S} with classical input $x_{\mathcal{S}}$ and a classical receiver \mathcal{R} with classical input $x_{\mathcal{R}}$, and we construct a maliciously-secure two-message protocol that delivers $D(x_{\mathcal{R}}, x_{\mathcal{S}})$ to the receiver. This can be seen as a (non-reusable) NISC (non-interactive secure computation) protocol for BQP. Both non-reusable and reusable NISC protocols have a long history of study in the classical setting [48,36,35,1,7,18,42], and we give the first construction that supports quantum functionalities while maintaining classical communication.

Theorem 4. *(Informal) Assuming QLWE, there exists a NISC for BQP with classical receiver in the quantum random oracle model.*

1.2 Technical overview

Background. Our starting point is two works of Mahadev [40,41] on classical delegation of quantum computation. Taken together, they show that a classical client can delegate a BQP computation to a quantum server while maintaining both privacy of the client’s input and integrity of the computation performed. Indeed, [41] shows how to obtain *soundness* via a construction of classical verification of quantum computation (CVQC), meaning that a (computationally bounded) cheating server won’t be able to convince the classical client of a false outcome. Furthermore, [40] shows how to obtain *privacy* of the client’s input via a construction of quantum fully-homomorphic encryption (QFHE) with classical keys. Executing CVQC under the hood of QFHE then provides both privacy and soundness, which was recently formalized by [21].

Our goal is to extend these results to the setting of fully-simulatable maliciously-secure computation, while also enabling *multiple* classical clients to outsource a quantum computation on their private inputs to a single quantum server. A natural idea would be to make use of post-quantum classical MPC to simulate the classical client in the above two-party client-server protocol. That is, n parties engage in classical MPC to set up a joint encryption $\text{QFHE.Enc}(x_1, \dots, x_n)$ of their inputs, and then they proceed to interact with the quantum server as a single entity.

⁴ We also require a common *random* string (CRS) setup, but this is subsumed by the random oracle model.

Unfortunately, the resulting protocol suffers from an input-dependent abort issue, rendering it insecure. Consider a malicious server that colludes with any one of the classical clients P_1 . Under QFHE, the server can decide whether to honestly complete the CVQC or force an abort *as a function of the clients' private inputs*. Then, P_1 's output will signal which was the case, allowing the server (and P_1) to learn any arbitrary predicate of the honest clients' inputs. In fact, the possibility of causing an input-dependent abort also prevents the original protocol between single client and single server from satisfying the standard notion of simulatable *two-party* computation. In other words, executing CVQC under QFHE does not result in a *composable* blind CVQC protocol.

CVQC with malicious blindness. In order to prevent such input-dependent abort attacks, what we need is a CVQC protocol where the prover cannot learn anything about the verifier's input, even if is able to learn whether or not the verifier aborted (rejected its proof). We will refer to this property as *malicious blindness*.

As explained above, executing CVQC under QFHE does not result in malicious blindness. But what if we switch the nesting of CVQC and QFHE, using CVQC to prove that a QFHE evaluation $\text{QFHE.Enc}(x) \rightarrow \text{QFHE.Enc}(y)$ is performed honestly? It is not immediately clear how to do this, since [41]'s CVQC protocol only works for BQP computation, and QFHE evaluation is a *sampBQP* computation. Indeed, the "encrypted CNOT" operation at the heart of [40]'s QFHE involves obviously sampling a classical FHE ciphertext by measuring a superposition over encryption random coins. Thus, performing QFHE evaluation of a pseudo-deterministic quantum functionality $y := D(x)$ will produce a *distribution* over ciphertexts $\text{QFHE.Enc}(y; r)$ with the same outcome y but with varying random coins r .

However, Chung et al. [21] recently showed how to extend the protocols of [28,41] to prove the correctness of *sampBQP* computations. The caveat is that the soundness error of their protocol is non-negligible, due to the following issue. Roughly, the prover prepares multiple copies of the history state of the computation, and the verifier chooses all but one of them to test and one of them to sample from. A malicious prover can always guess which state the verifier will sample from with inverse polynomial probability and cheat only in that copy of the history state, convincing the verifier to accept a completely invalid result. For general *sampBQP* problems, this issue appears somewhat inherent to their approach, since there is no meaningful way to combine multiple potentially invalid samples into a single valid sample.

We observe that for the special case of proving honest QFHE evaluation of some BQP computation, one *can* meaningfully combine multiple potentially invalid samples. Consider a verifier that requests multiple output ciphertexts, decrypts them all, and then outputs the most frequently occurring plaintext. If the prover can only cheat with some small (say 1/4) probability on each sample, then one should be able to drive the probability of accepting an invalid result down to negligible, with enough samples.

More abstractly, we consider any pseudo-deterministic circuit $D(\cdot)$ that can be written as $C(Q(\cdot))$, where Q is a quantum circuit and C is a classical cir-

cuit. On any classical input x , $C(Q(x))$ produces a well-defined output z (with overwhelming probability), while $Q(x)$ may produce any distribution over intermediate classical values y . The goal will be to obtain a protocol for delegating the computation of $C(Q(\cdot))$ where the prover’s computation and the verifier’s decision to accept or reject is independent of C .

We first use the one-round sampBQP verification protocol of [21] (with a quantum verifier that performs single-qubit computational and Hadamard basis measurements) to show how to verify such pseudo-deterministic computations with negligible soundness error (details in Section 4.2). The resulting protocol consists of a quantum proof and requires a quantum verifier to perform single-qubit measurements. The next step is then to incorporate Mahadev’s measurement protocol [41] in order to make the proof and the verifier fully classical.

The measurement protocol as described in [41] proceeds in a number of rounds, where in each round, the verifier issues a single bit challenge indicating either a “test” round or a “Hadamard” round. The test round is meant to check that the prover is behaving honestly (i.e. it is honestly “committing” to some particular quantum state), while the Hadamard round is meant to produce a sequence of classical measurement results that the verifier can use to produce its verdict. A single round consists of four classical messages between prover and verifier, and [41] shows that it satisfies the following property: Any prover that passes the test round with overwhelming probability will only be able to “cheat” in a Hadamard round with negligible probability, assuming QLWE.

In Section 4.3, we use this protocol to obtain a four message protocol for verifying circuits $D(\cdot) = C(Q(\cdot))$ that satisfies the following property. The verifier will either choose a test round, in which case they simply accept or reject, or they will choose a Hadamard round, in which case they either reject or obtain a purported sample $y \leftarrow Q(x)$. In the latter case, they compute and output $z := C(y)$. Any prover that passes the test round with overwhelming probability will only be able to force an incorrect output in a Hadamard round with negligible probability.

Parallel repetition. Now, we would like to obtain negligible soundness error, ideally while maintaining the four-message interaction. Recent works [3,20] have shown how to do this in the setting where the cheating prover is attempting to convince the verifier to accept some false BQP statement. They show that if the four-message protocol is run sufficiently many times n in parallel, then the probability that the verifier accepts on *all* repetitions is negligible. Phrased differently, they show that, conditioned on the verifier accepting each of the (roughly) $n/2$ test rounds, the prover will not be able to successfully “cheat” on *all* the Hadamard rounds, except with negligible probability.

However, this is not quite enough for our setting. Recall that in the i ’th Hadamard round, the verifier receives a purported sample $y_i \leftarrow Q(x)$ and computes $z_i := C(y_i)$. Crucially, we want the verifier to have already decided to accept by the time they invoke C to compute the outputs z_i . Thus, to combine these $\{z_i\}_i$ into a final output z , we will simply have the verifier output the most frequently occurring string z in the set (in particular, this final output compu-

tation cannot decide to accept/reject based on any properties of the set $\{z_i\}_i$. Now, a prover that successfully cheats in only *half* of the Hadamard rounds may be able to force an invalid output. Our goal is thus to show that this can only happen with negligible probability.

To do so, we take a closer look at the proof of the parallel repetition theorem from [3]. The following exposition will be simplified and not technically accurate, and is just meant to convey intuition. They consider any cheating prover state $|\psi\rangle$ right before the verifier's challenge $c \leftarrow \{0, 1\}^n$ is chosen (n is the number of repetitions, and $c_i = 0$ corresponds to a test round while $c_i = 1$ corresponds to a Hadamard round). Then, for each possible $c \in \{0, 1\}^n$, they consider a binary-valued projector Π_c that corresponds to running the prover's remaining strategy and then applying the verifier's verdict function on the resulting proof. Since c is chosen uniformly at random, it suffices to show that $\frac{1}{2^n} \langle \psi | \sum_{c \in \{0, 1\}^n} \Pi_c | \psi \rangle$ is negligible. To do so, they square the quantity $\langle \psi | \sum_c \Pi_c | \psi \rangle$ and then show that each cross term $\langle \psi | \Pi_c \Pi_{c'} | \psi \rangle$ for $c \neq c'$ is negligible.

This is possible in their setting because the verifier only accepts if *every* test and Hadamard round accepts. This means that if Π_c and $\Pi_{c'}$ are both accepting for $c \neq c'$, there must be some index i where the prover is being accepted on both a test round and a Hadamard round (any i such that $c_i \neq c'_i$). This can be used to contradict the key property of the single repetition protocol described above, that any prover accepted on a test round with overwhelming probability can only cheat on a Hadamard round with negligible probability.

In our setting, we say that the verifier accepts if *every* test round accepts and *at least half* of the Hadamard rounds accept. Thus, it is no longer true that any noticeably large cross term $\langle \psi | \Pi_c \Pi_{c'} | \psi \rangle$ will imply a contradiction to the single repetition protocol. Indeed, if c and c' are somewhat close in Hamming distance, a prover could be rejected in a Hadamard round on all indices i such that $c_i \neq c'_i$ without causing the overall verifier to reject.

However, we observe that (i) it suffices to bound an overwhelming fraction of the cross terms (rather than all), and (ii) for any c, c' that have sufficiently large Hamming distance, if $\langle \psi | \Pi_c \Pi_{c'} | \psi \rangle$ is large, then there must be some index where the prover is simultaneously passing both the test and Hadamard rounds. Thus, we will need an overwhelming fraction of cross terms to correspond to pairs of challenge strings c, c' with Hamming distance at least as large as the number of Hadamard rounds. To facilitate this, we alter the protocol, setting the number of Hadamard rounds to some fixed λ , the total number of rounds to $n = \lambda^{1+\epsilon}$, and having the verifier sample a challenge string $c \in \{0, 1\}^n$ with Hamming weight exactly λ . The full details and proof of this strengthened parallel repetition theorem can be found in Section 3.

Secure quantum computation. The above shows how to obtain negligible soundness for classical verification of quantum-classical circuits $C(Q(\cdot))$, where the verifier decides to accept or reject independently of C . In the remainder of the paper, we show how to use this primitive to construct composable blind CVQC, and then present applications to secure quantum computation with classical communication.

In Section 4.4 we show that four-message composable blind CVQC follows by letting Q correspond to QFHE evaluation, and C correspond to QFHE decryption. Crucially, the QFHE decryption key is not needed to determine whether the verifier accepts or rejects, which results in the malicious blindness needed to ensure composability. Then, in Section 5.1, we use post-quantum MPC for classical (reactive) functionalities to allow multiple parties to simulate a single verifier participating in the CVQC protocol with the server. This results in a constant-round MPQC protocol in the common random string model from QLWE.

Next, we consider the three-round interaction pattern described in Section 1.1. To implement this, we will need a CVQC protocol with (i) two total messages, and (ii) distributed setup, where multiple parties can encrypt their respective inputs without any interaction. We construct this primitive in the full version via the following observations. Property (i) can be obtained in the quantum random oracle model (QROM) by appealing to Fiat-Shamir in the QROM [24,39]. Property (ii) can be obtained via the use of quantum *multi-key* fully homomorphic encryption, which was recently constructed in [2]. In fact, we will also require the CVQC protocol to satisfy various *perfect correctness* properties to obtain security against verifiers that make a malicious choice of random coins, and we defer discussion of this to the body.

Now, in the full version, we show how to combine the two-message CVQC with distributed setup with a classical *multi-party reusable non-interactive secure computation* protocol (mrNISC) to obtain three-round MPQC in the public-key infrastructure (PKI) model. Maliciously-secure post-quantum mrNISC protocols have recently been constructed from QLWE [5,13]. The reason we need a PKI is the following. In the two-message CVQC protocol, the various inputs can be encrypted in a distributed fashion. However, the verifier also needs to set up some (input-independent) public parameters \mathbf{pp} along with secret parameters \mathbf{sp} kept private from the prover. This part cannot be fully distributed, so before the protocol begins, we will have each party publish a public key consisting of a mrNISC first-round message committing to a PRF key. This message is succinct (does not depend on the size of the functionality to be computed) and reusable across any varying subsets of parties (due to the reusability of the mrNISC protocol), and so it satisfies the requirements of the PKI model. These PRF keys can then be used to compute public parameters in the first round of the MPQC protocol, which are sent to the server along with the encryptions of each party's input.

Finally, in the full version, we show that the two-message CVQC with distributed setup primitive can also be used to construct a *two-message* protocol between two parties: a quantum sender \mathcal{S} and a classical receiver \mathcal{R} . Both parties have classical inputs, but can compute a (pseudo-deterministic) quantum functionality over these inputs. This follows by letting each of the sender and receiver independently encrypt their input, having the sender evaluate the CVQC prover, and then executing the CVQC verifier under a *classical* NISC (non-interactive

secure computation) protocol. This results in a quantum NISC protocol with classical communication.

1.3 Discussion and open problems

Quantum sampling circuits. The above describes how to achieve standard malicious security for secure multi-party computation of pseudo-deterministic circuits. That is, any adversary will only have negligible advantage in distinguishing the real and simulated worlds. A natural next question is whether this is achievable for the more general class of polynomial-time quantum *sampling* problems, i.e. functionalities that output an arbitrary distribution over classical strings.

It is straightforward to see that the sampBQP protocol of [21] can be combined with QFHE and post-quantum classical MPC to give MPQC for quantum sampling problems with classical communication, but with some inverse polynomial security. That is, the adversary will be able to distinguish the real and simulated worlds with at most some inverse polynomial probability (but where the communication complexity of the protocol grows with this polynomial). This follows by using their sampBQP protocol to prove the correct computation of $\text{QFHE.Enc}(y) \leftarrow \text{QFHE.Eval}(Q, \text{QFHE.Enc}(x))$, where Q is some quantum sampling circuit, and x is the joint private inputs of the clients.

Another potential approach to secure computation of sampling circuits would be to follow [30], which uses ideas from [15] and [41] to construct a blind CVQC protocol in the measurement-based quantum computing framework (avoiding the reduction to local Hamiltonian). This framework appears to naturally support quantum sampling circuits, though its soundness for sampling circuits has not been analyzed. However, the inverse soundness error of their protocol also grows with the communication complexity. Thus, we leave MPQC for sampling circuits with classical communication and standard negligible security as an open question. Indeed, the more basic question that remains open is whether it is possible to construct a CVQC protocol for sampling circuits where the verifier only accepts samples that are distributed negligibly close to the real distribution induced by the quantum circuit.

Obfuscation of quantum circuits. As also discussed in [10], one approach to obtaining (heuristic) obfuscation of quantum circuits involves the notion of blind CVQC. Given a two-message composable blind CVQC with delayed functionality (meaning that the circuit D can be chosen by the prover after the verifier’s message has been sampled) one could imagine obfuscating the (pseudo-deterministic) quantum circuit U as follows. Sample the verifier’s first message on input U and output this message along with a classical obfuscation of the verification circuit (with the verification secret key hard-coded). Then, the evaluator with input x can attempt to produce a proof π for the circuit $D_x(\cdot)$ that takes U as input and outputs $U(x)$. It can then query the obfuscated verifier on π to learn $U(x)$. Note that malicious blindness is crucial for making this approach work, as the evaluator can clearly see whether or not the verifier accepts or rejects its proof.

Unfortunately, this is not the only property that is required. It is also crucial that the blind CVQC protocol is *reusably* sound, meaning that it securely implements an ideal functionality that allows the prover to repeatedly query the verifier on different circuits of its choice, learning the verifier’s output each time. The composable blind CVQC constructed in this paper does not satisfy this reusable ideal functionality, and more discussion about the difficulties in obtaining reusable security can be found in [10]. Thus, the (one-time) composable blind CVQC protocol constructed in this paper can be seen a step towards heuristic obfuscation of quantum circuits, though the possibility of obtaining the crucial property of *reusability* remains open.

Quantum vs. classical simulation. Our definition of secure multi-party quantum computation (Definition 6) by default allows for a quantum simulator. However, in some settings it would be desirable to require a classical simulator in the case where the adversary is only corrupting classical parties, in order to argue that the malicious classical parties cannot obtain arbitrary quantum-computable information by interacting in the protocol. In fact, it is easy to see that our construction of constant-round MPQC from QLWE (Protocol 2) does satisfy this stricter requirement. On the other hand, our other results (three-round MPQC protocol and NISC protocol) require a quantum simulator, even for corrupted classical parties. It would be interesting to explore this question further, to see if such protocols can be constructed with a classical simulator.

1.4 Other related work

Composable security. The notion of *composable blind and verifiable quantum computation* has been studied previously both in the setting of a quantum verifier and a classical verifier. In particular, [26] showed that the blind and verifiable protocols of [29] and [43] with quantum verifier are composable. Next, [30] gave a construction of composable blind and verifiable delegation of quantum computation with a classical verifier, and noted that such a protocol should have implications to multi-party quantum computation (though they left this formalization to future work). However, their protocol only achieves inverse polynomial security and requires polynomially many rounds,⁵ while ours achieves standard negligible security and has constant rounds.

We also remark that the works of [26] and [30] achieve composable security by combining the properties of standard blindness and “independent” verifiability, which are very similar to the two properties of malicious blindness and standard soundness that we use to enable composable security.

Two-party quantum computation. There are a couple of recent works that study a relaxed variant of two-party secure computation, called *secure function evalua-*

⁵ This appears to be somewhat inherent to their approach, as they follow the measurement-based computing paradigm, which requires the prover and verifier to interact for each sequential gate being computed.

tion (SFE), in the quantum setting. This notion relaxes security from simulation-based to indistinguishability-based, and in particular does not require correctness against malicious senders, meaning that constructions of quantum SFE do not require (classical) verification of quantum computation.

First, the work of [22] constructs SFE for quantum functionalities with classical communication, and even achieves *one-sided* simulation security (against malicious receiver). In fact, they show that if the sender has a quantum input, then full simulation-security is *impossible* to achieve if only classical communication is allowed. We evade this barrier by constructing protocols where all parties have *classical* input, showing that full simulation-based security is indeed possible in this setting. Next, the work of [17] constructs SFE for quantum functionalities (where receiver may have a quantum input, and thus communication is quantum) from QLWE, and even show how to achieve rate-1 communication complexity.

Multi-party delegated quantum computation. Multi-party delegated quantum computation was first studied by [38], in the setting where multiple computationally weak but *quantum* clients would like to outsource some quantum computation to a powerful quantum server. This work only provides blindness (no correctness against a malicious server) and does not handle client-server collusions. Very recently, [37] showed how to achieve standard malicious security in this setting, though still with the requirement that the clients have quantum resources. The work of [33] studies the notion of multi-party delegated quantum computation in three rounds, though security is only semi-honest, and again the clients require quantum resources. In summary, our work is the first to construct multi-party delegated quantum computation with entirely classical clients.

2 Preliminaries

Let λ denote the security parameter. A function $f : \mathbb{N} \rightarrow [0, 1]$ is negligible if for every constant $c \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$, and we write $\text{negl}(\cdot)$ to denote such a function. Let $\mathcal{HW}_{n,m}$ denote the set of binary strings of length n with Hamming weight m . We will refer to pure quantum states with ket notation $|\psi\rangle$ and mixed quantum states with lowercase Greek letters such as ρ . Throughout, we will consider non-uniform quantum polynomial-time (QPT) adversaries, which are families of polynomial-size quantum circuits $\{\text{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$ along with some polynomial-size quantum advice $\{|\psi_\lambda\rangle\}_{\lambda \in \mathbb{N}}$, though we will often drop the indexing by λ when clear from context. We will also often refer to families of “psuedo-deterministic” quantum circuits, defined below. Again, we usually drop the indexing by λ when clear from context.

Definition 1 (Pseudo-Deterministic Quantum Circuit). *A family of psuedo-deterministic quantum circuits $\{D_\lambda\}_{\lambda \in \mathbb{N}}$ is defined as follows. The circuit D_λ takes as input a classical bit string $x \in \{0, 1\}^{n(\lambda)}$ and outputs a single classical string $z \leftarrow D(x)$. The circuit is pseudo-deterministic if there exists a negligible*

function ν such that for every sequence of classical inputs $\{x_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a sequence of outputs $\{z_\lambda\}_{\lambda \in \mathbb{N}}$ such that

$$\Pr[D_\lambda(x_\lambda) = z_\lambda] = 1 - \nu(\lambda).$$

The following notation will also be useful.

Definition 2 ($M_{XZ}(\rho, h)$). For a string $h \in \{0, 1\}^n$ and an n -qubit quantum state ρ , consider the following procedure. For each $i \in [n]$, measure the i 'th qubit of ρ in the standard basis if $h_i = 0$ and in the Hadamard basis if $h_i = 1$. Let the resulting random variable (over classical n -bit strings) be denoted by $M_{XZ}(\rho, h)$.

2.1 Delegation of quantum computation

We will consider protocols $\Pi = (\mathcal{P}, \mathcal{V})$ for delegating the computation of pseudo-deterministic quantum circuits D . In such a protocol, \mathcal{P} and \mathcal{V} interact on input the security parameter 1^λ , and \mathcal{V} has additional (possibly private) inputs D, x . After the interaction, \mathcal{V} outputs (v, z) , where $v \in \{\text{acc}, \text{rej}\}$ and $z \in \{0, 1\}^*$. We denote this by $(v, z) \leftarrow (\mathcal{P}, \mathcal{V}(D, x))(1^\lambda)$. In general, \mathcal{V} will satisfy some efficiency properties (e.g. it has limited or no quantum capabilities), making the protocol non-trivial.

Definition 3. A protocol $\Pi = (\mathcal{V}, \mathcal{P})$ for delegating the computation of a pseudo-deterministic quantum circuit D should satisfy the following properties.

- **Completeness:** For any circuit D , input x , and output z such that $\Pr[D(x) = z] = 1 - \text{negl}(\lambda)$, it holds that

$$\Pr[(\text{acc}, z) \leftarrow (\mathcal{P}, \mathcal{V}(D, x))(1^\lambda)] = 1 - \text{negl}(\lambda).$$

- **Soundness:** For any circuit D , input x , output z such that $\Pr[D(x) = z] = 1 - \text{negl}(\lambda)$, and cheating prover \mathcal{P}^* with advice $|\psi\rangle$, it holds that

$$\Pr[v = \text{acc} \wedge z' \neq z : (v, z') \leftarrow (\mathcal{P}^*(|\psi\rangle), \mathcal{V}(D, x))(1^\lambda)] = \text{negl}(\lambda).$$

We say that soundness is statistical if \mathcal{P}^* is unbounded, and that soundness is computational if \mathcal{P}^* is a QPT machine with polynomial-size advice.

2.2 Quantum fully-homomorphic encryption

We define quantum fully-homomorphic encryption (QFHE) with classical keys and classical encryption of classical messages. One could also define encryption for quantum states and decryption for quantum ciphertexts, but we will not need that in this work.

Definition 4 (Quantum Homomorphic Encryption). A quantum fully-homomorphic encryption scheme (QFHE.Gen, QFHE.Enc, QFHE.Eval, QFHE.Dec) consists of the following efficient algorithms.

- $\text{QFHE.Gen}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$: On input the security parameter, the PPT key generation algorithm returns public/secret key pair (pk, sk) .
- $\text{QFHE.Enc}(\text{pk}, x) \rightarrow \text{ct}$: On input the public key pk and a classical plaintext x , the PPT encryption algorithm returns a classical ciphertext ct .
- $\text{QFHE.Eval}(Q, \text{ct}) \rightarrow \tilde{\text{ct}}$: On input a quantum circuit Q , and a ciphertext ct , the QPT evaluation algorithm returns an evaluated ciphertext $\tilde{\text{ct}}$.
- $\text{QFHE.Dec}(\text{sk}, \text{ct}) \rightarrow x$: On input the secret key sk and a classical ciphertext ct , the decryption algorithm returns a message x .

The scheme should satisfy the standard notion of semantic security. We will require the following notion of correctness for evaluation of pseudo-deterministic quantum circuits. Note that this evaluation correctness holds over *all* key generation and encryption random coins.

Definition 5 (Evaluation Correctness). *A QFHE scheme $(\text{QFHE.Gen}, \text{QFHE.Enc}, \text{QFHE.Eval}, \text{QFHE.Dec})$ is correct if for every $\lambda \in \mathbb{N}$, every $(\text{pk}, \text{sk}) \in \text{QFHE.Gen}(1^\lambda)$, every input x , every $\text{ct} \in \text{QFHE.Enc}(\text{pk}, x)$, and every polynomial-size pseudo-deterministic quantum circuit Q and output y such that $\Pr[Q(x) = y] = 1 - \text{negl}(\lambda)$, it holds that*

$$\Pr[\text{QFHE.Dec}(\text{sk}, \text{QFHE.Eval}(Q, \text{ct})) = y] = 1 - \text{negl}(\lambda).$$

The works of Mahadev [40] and Brakerski [14] show that such a QFHE scheme can be constructed from QLWE (we will not consider unlevelled QFHE in this work, which requires circular-security assumptions).

Multi-key. We will also make use of a quantum *multi-key* fully-homomorphic encryption scheme $(\text{QMFHE.Gen}, \text{QMFHE.KeyGen}, \text{QMFHE.Enc}, \text{QMFHE.Eval}, \text{QMFHE.Dec})$, which was constructed in [2]. In such a scheme, the evaluation algorithm QMFHE.Eval now may take as input some n -input circuit Q along with n ciphertexts $(\text{ct}_1, \dots, \text{ct}_n)$, each encrypted under independently sampled public keys (assume that each ct_i contains a description of the public key pk_i it is encrypted under). Likewise, QMFHE.Dec can decrypt a ciphertext ct that is the result of evaluating ciphertexts encrypted under public keys $\text{pk}_1, \dots, \text{pk}_n$, given the corresponding secret keys $\text{sk}_1, \dots, \text{sk}_n$. We require the same evaluation correctness (Definition 5) to hold, except over n -input pseudo-deterministic functionalities Q . We have also added a QMFHE.Gen algorithm which samples a common random string crs , to be used by each KeyGen algorithm.

2.3 Multi-party quantum computation

Below we give a definition of maliciously-secure multi-party quantum computation for pseudo-deterministic quantum functionalities, following the standard real/ideal world paradigm for defining secure computation [31]. We assume that parties have access to a (classical) broadcast channel, and we aim for security with unanimous abort.

Consider an n -party quantum functionality specified by a family of pseudo-deterministic quantum circuits $\mathcal{Q} = \{Q_\lambda\}_{\lambda \in \mathbb{N}}$ where Q_λ has classical input of size $m_1(\lambda) + \dots + m_n(\lambda)$ and classical output of size $\ell_1(\lambda) + \dots + \ell_n(\lambda)$. We will consider a QPT adversary $\text{Adv} = \{\text{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$ that corrupts any subset $M \subset [n]$ of parties. Let $H := [n] \setminus M$.

Let Π be an n -party protocol for computing \mathcal{Q} . Consider any collection $(x_1, \dots, x_n, |\psi\rangle_{\text{Adv}, \mathcal{D}})$, where x_1, \dots, x_n are classical bitstrings and $|\psi\rangle_{\text{Adv}, \mathcal{D}}$ is a polynomial-size quantum state on two registers Adv and \mathcal{D} . Abusing notation, we let $|\psi\rangle_{\text{Adv}}$ be the part of the state on register Adv and likewise for $|\psi\rangle_{\mathcal{D}}$. Now define quantum random variable $\text{REAL}_{\Pi, \mathcal{Q}}(\text{Adv}_\lambda, \{x_i\}_{i \in [n]}, |\psi\rangle_{\text{Adv}})$ as follows. $\text{Adv}_\lambda(\{x_i\}_{i \in M}, |\psi\rangle_{\text{Adv}})$ interacts with honest party algorithms on inputs $\{x_i\}_{i \in H}$ participating in protocol Π , after which the honest parties output $\{y_i\}_{i \in H}$ and Adv outputs a final state $|\psi_{\text{out}}\rangle$ (an arbitrary function computed on an arbitrary subset of the registers that comprise its view). The random variable $\text{REAL}_{\Pi, \mathcal{Q}}(\text{Adv}_\lambda, \{x_i\}_{i \in [n]}, |\psi\rangle_{\text{Adv}})$ then consists of $\{y_i\}_{i \in H}$ along with $|\psi_{\text{out}}\rangle$.

For any Adv , we require the existence of a simulator $\text{Sim} = \{\text{Sim}_\lambda\}_{\lambda \in \mathbb{N}}$ that takes as input $(\{x_i\}_{i \in M}, |\psi\rangle_{\text{Adv}})$, has access to an ideal functionality $\mathcal{I}[\{x_i\}_{i \in H}](\cdot)$, and outputs a state $|\psi_{\text{out}}\rangle$. The ideal functionality accepts an input $\{x_i\}_{i \in M}$, applies Q_λ to (x_1, \dots, x_n) to recover (y_1, \dots, y_n) , and returns $\{y_i\}_{i \in M}$ to Sim_λ . Then, it waits for either an `abort` or `ok` message from Sim_λ . In the case of `ok` it includes $\{y_i\}_{i \in H}$ in its output and in the case of `abort` it includes $\{\perp\}_{i \in H}$. Now, we define the quantum random variable $\text{IDEAL}_{\Pi, \mathcal{Q}}(\text{Sim}_\lambda, \{x_i\}_{i \in [n]}, |\psi\rangle_{\text{Adv}})$ to consist of the output of $\mathcal{I}[\{x_i\}_{i \in H}](\cdot)$ and the final state $|\psi_{\text{out}}\rangle$ of $\text{Sim}_\lambda^{\mathcal{I}[\{x_i\}_{i \in H}](\cdot)}(\{x_i\}_{i \in M}, |\psi\rangle_{\text{Adv}})$.

Definition 6 (Secure Multi-Party Quantum Computation). *A protocol Π securely computes Q if for all QPT $\text{Adv} = \{\text{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$ corrupting subset of parties $M \subset [n]$, there exists a QPT $\text{Sim} = \{\text{Sim}_\lambda\}_{\lambda \in \mathbb{N}}$ such that for all $\{x_{1,\lambda}, \dots, x_{n,\lambda}, |\psi_\lambda\rangle_{\text{Adv}, \mathcal{D}}\}_{\lambda \in \mathbb{N}}$ and all QPT $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\nu(\cdot)$ such that*

$$\left| \Pr[\mathcal{D}_\lambda(|\psi_\lambda\rangle_{\mathcal{D}}, \text{REAL}_{\Pi, \mathcal{Q}}(\text{Adv}_\lambda, \{x_{i,\lambda}\}_{i \in [n]}, |\psi_\lambda\rangle_{\text{Adv}})) = 1] - \Pr[\mathcal{D}_\lambda(|\psi_\lambda\rangle_{\mathcal{D}}, \text{IDEAL}_{\Pi, \mathcal{Q}}(\text{Sim}_\lambda, \{x_{i,\lambda}\}_{i \in [n]}, |\psi_\lambda\rangle_{\text{Adv}})) = 1] \right| \leq \nu(\lambda).$$

2.4 Classical non-interactive secure computation

In the full version, we will make use of a particular type of classical two-party computation protocol, called non-interactive secure computation (NISC). A NISC protocol is a two-message protocol between sender \mathcal{S} with input $x_{\mathcal{S}}$ and receiver \mathcal{R} with input $x_{\mathcal{R}}$. The protocol consists of two total messages, and is defined by four algorithms ($\text{NISC}_{\text{Gen}}, \text{NISC}_1, \text{NISC}_2, \text{NISC}_{\text{out}}$). We require that the receiver's message can be computed independently of the functionality C to be computed. The syntax of these algorithms is as follows.

- $\text{crs} \leftarrow \text{NISC}_{\text{Gen}}(1^\lambda)$. The gen algorithm generates the crs.

- $(m_{\mathcal{R}}, \text{st}) \leftarrow \text{NISC}_1(\text{crs}, x_{\mathcal{R}})$. The first message algorithm takes as input crs and the receiver’s input $x_{\mathcal{R}}$, and outputs the receiver’s message $m_{\mathcal{R}}$ and private state st .
- $m_{\mathcal{S}} \leftarrow \text{NISC}_2(\text{crs}, C, m_{\mathcal{R}}, x_{\mathcal{S}})$. The second message algorithm takes as input crs , a circuit C , the receiver’s message $m_{\mathcal{R}}$, and the sender’s input $x_{\mathcal{S}}$, and outputs the sender’s message $m_{\mathcal{S}}$.
- $y \leftarrow \text{NISC}_{\text{out}}(\text{st}, C, m_{\mathcal{S}})$. The output algorithm takes as input the receiver’s private state st , the circuit C , and the sender’s message $m_{\mathcal{S}}$, and outputs the receiver’s output y .

A NISC protocol should satisfy standard simulation-based security against arbitrarily malicious adversaries. We note that post-quantum NISC protocols are known assuming simulation-secure two-message oblivious transfer [36]. Post-quantum two-message oblivious transfer (with a reusable common random string) can be obtained from QLWE by combining a semi-malicious oblivious transfer (such as [16]) with non-interactive zero-knowledge [45]. Alternatively, such an oblivious transfer can be obtained directly from QLWE with subexponential modulus-to-noise ratio [46].

3 Generalizing the Alagic et al. Parallel Repetition Theorem

Consider the following outline for a four-message protocol between a classical verifier \mathcal{V} and a quantum prover \mathcal{P} . For concreteness, one can think of the public key pk as encoding some computation that the verifier would like the prover to perform. However, for the purposes of this section, we will only need to consider a few generic properties of such a protocol.

- $\mathcal{V}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$: the verifier, on input the security parameter, generates a public/secret key pair (pk, sk) and sends pk to \mathcal{P} .
- $\mathcal{P}(\text{pk}) \rightarrow (y, |\text{st}\rangle)$: the prover, on input the public key, generates a classical string y and a quantum private state $|\text{st}\rangle$, and sends y to \mathcal{V} .
- \mathcal{V} samples a random challenge bit $c \leftarrow \{0, 1\}$ and sends c to \mathcal{P} .
- $\mathcal{P}(|\text{st}\rangle, c) \rightarrow \pi$: the prover, on input its quantum state and the verifier’s challenge bit, generates a classical proof π and sends π to \mathcal{V} .
- $\mathcal{V}(\text{pk}, \text{sk}, y, c, \pi) \rightarrow b$: the verifier, on input the transcript and its secret key, outputs a bit b where $b = 1$ indicates that it accepts and $b = 0$ indicates that it rejects.

If the verifier sampled $c = 0$, we refer to the execution as a “test round”, and if the verifier sampled $c = 1$, we refer to the execution as a “Hadamard round”. Now consider a generic prover strategy $(|\psi_{\text{init}}\rangle, U, V_0, V_1)$, where

- $|\psi_{\text{init}}\rangle$ is some initial state on three registers X, Y, Z ,
- U is a unitary on registers X, Y, Z, K , classically controlled on K ,

- and V_0, V_1 are unitaries on registers X, Y, Z, K , classically controlled on Y and K .

Given this generic strategy, the protocol proceeds as follows. The pair $(\mathbf{pk}, \mathbf{sk})$ is sampled by the verifier, and then U is applied to $|\psi_{\text{init}}\rangle |\mathbf{pk}\rangle$ to produce private state $|\psi_{\mathbf{pk}}\rangle$. Then, c is sampled and sent to the prover, who applies V_c to $|\psi_{\mathbf{pk}}\rangle$, and then measures Y, X to produce y, π .⁶ For any choice of $(\mathbf{pk}, \mathbf{sk}, c)$, let $\Pi_{\mathbf{pk}, \mathbf{sk}, c}$ be the binary-valued projector that, when applied to $V_c |\psi_{\mathbf{pk}}\rangle$, corresponds to measuring y, π and then applying the verifier’s verdict function on $(\mathbf{pk}, \mathbf{sk}, y, c, \pi)$.

Now suppose that the following two conditions hold. Whenever we take an expectation over $(\mathbf{pk}, \mathbf{sk})$, we mean over $(\mathbf{pk}, \mathbf{sk}) \leftarrow \mathcal{V}(1^\lambda)$.

1. $\Pi_{\mathbf{pk}, \mathbf{sk}, 0}$ does not depend on \mathbf{sk} (in the protocols we are interested in, it simply checks whether π is in some set of classical strings determined by \mathbf{pk} and y).
2. For any efficient prover strategy $(|\psi_{\text{init}}\rangle, U, V_0, V_1)$ ⁷ such that

$$\mathbb{E}_{\mathbf{pk}, \mathbf{sk}} [\langle \psi_{\mathbf{pk}} | V_0^\dagger \Pi_{\mathbf{pk}, \mathbf{sk}, 0} V_0 | \psi_{\mathbf{pk}} \rangle] = 1 - \text{negl}(\lambda),$$

it holds that $\mathbb{E}_{\mathbf{pk}, \mathbf{sk}} [\langle \psi_{\mathbf{pk}} | V_1^\dagger \Pi_{\mathbf{pk}, \mathbf{sk}, 1} V_1 | \psi_{\mathbf{pk}} \rangle] = \text{negl}(\lambda)$.

Consider now an (n, λ) -parallel repeated version of the protocol, with a challenge string d that is chosen uniformly at random from the set $\mathcal{HW}_{n, \lambda}$ of n -bit strings with Hamming weight exactly λ . That is, the protocol is repeated n times in parallel, except that a Hadamard round is performed in exactly λ of the n protocols. The following theorem shows that any efficient cheating prover in this multi-copy protocol cannot make the verifier accept all test rounds and more than half of the Hadamard rounds.

Theorem 5. *Let $n(\lambda) \geq \lambda^{1+\epsilon}$ for some constant $\epsilon > 0$. Then in the (n, λ) -parallel repeated protocol, the probability that the verifier accepts every test round (i such that $d_i = 0$) and $\geq \lambda/2$ of the Hadamard rounds (i such that $d_i = 1$) is $\text{negl}(\lambda)$.⁸*

The proof of this is given in the full version.

4 Composable Blind CVQC

In this section, we show how to construct (constant-round) composable blind classical verification of quantum computation. The main property we need to

⁶ Note that, since V_c must be classically controlled on Y , we can indeed consider Y to be measured after applying V_c rather than before.

⁷ By efficient, we mean that U, V_0, V_1 are unitaries that can be implemented in quantum polynomial-time, while $|\psi_{\text{init}}\rangle$ is any (potentially inefficiently preparable) state on a polynomial number of registers.

⁸ In fact, it is straightforward to adjust the proof to show that for *any* constant c , the verifier will only accept all test rounds and at least λ/c Hadamard rounds with $\text{negl}(\lambda)$ probability.

achieve composability (other than standard blindness and soundness) is that the prover cannot obtain information about the verifier’s input even if it gets to observe whether or not the verifier accepted its proof.

In fact, we will first construct a more general notion that we call *CVQC for quantum-classical circuits*, which allows for delegation of pseudo-deterministic quantum-classical circuits $C(Q(\cdot))$, where the verifier’s decision to accept or reject does not depend on the description of C . We will eventually make use of such a protocol in multiple ways, using the classical circuit C to compute various functionalities, such as QFHE decryption or the NISC receiver’s output algorithm.

In Section 4.1, we give a formal definition of this primitive. In Section 4.2, we use [21]’s sampBQP protocol to show how to delegate quantum-classical circuits with a quantum verifier that only performs single qubit measurements. In Section 4.3, we compile the protocol with [41]’s measurement protocol to make the verifier classical, and then apply the parallel repetition theorem from last section to obtain negligible soundness error, satisfying our definition in Section 4.1. In Section 4.4, we formally define the notion of composable blind CVQC, and show how to construct it from CVQC for quantum-classical circuits plus quantum fully-homomorphic encryption. If the full version, we also give an alternative CVQC protocol that satisfies various extra properties and will be useful for obtaining round-optimized secure computation protocols.

4.1 CVQC for quantum-classical circuits

Below we define CVQC for quantum-classical circuits. We crucially split the final verification into two parts, requiring that the classical part of the circuit is not needed to determine whether or not the verifier will accept the prover’s proof.

We require the protocol to satisfy the standard completeness and soundness properties, along with *semi-malicious correctness*, which will be useful when building secure computation from this primitive. This property guarantees that even if the verifier’s initial keys are computed honestly, but with a malicious choice of random coins, the output of the honest verification procedure applied to an honest proof must still be correct (if it doesn’t abort). Below we present the two-message syntax, but such a definition easily extends to more general interactive protocols.

Definition 7 (CVQC for Quantum-Classical Circuits). *A two-message CVQC protocol for a quantum-classical pseudo-deterministic circuit $D(\cdot) = C(Q(\cdot))$ has the following syntax.*

- $\mathcal{V}_{\text{QC}}^{\text{setup}}(1^\lambda, Q, x) \rightarrow (\text{pk}, \text{sk})$: the verifier takes the security parameter 1^λ , a quantum circuit Q , and a classical input x , and outputs a public key secret key pair (pk, sk) .
- $\mathcal{P}_{\text{QC}}(\text{pk}, Q, x) \rightarrow \pi$: the prover takes the public key pk , a quantum circuit Q , and classical input x , and outputs a classical proof π .
- $\mathcal{V}_{\text{QC}}(Q, C, x, \text{pk}, \text{sk}, \pi) \rightarrow (v, z)$: the final verification circuit is split up into the following two parts.

- $\mathcal{V}_{\text{QC}}^{\text{verify}}(Q, x, \text{pk}, \text{sk}, \pi) \rightarrow (v, \tilde{z})$: the first part of the final verification circuit takes the quantum circuit Q , the input x , public key secret key pair (pk, sk) , and proof π , and outputs (v, \tilde{z}) , where $v \in \{\text{acc}, \text{rej}\}$, and $\tilde{z} \in \{0, 1\}^*$. If $v = \text{acc}$, then the second part of the verification is invoked on \tilde{z} to produce z , and the final output is (acc, z) . Otherwise, the final output is (rej, \perp) .
- $\mathcal{V}_{\text{QC}}^{\text{out}}(C, \tilde{z}) \rightarrow z$: the second part of the verification takes as input a classical circuit C and a string \tilde{z} , and outputs a string z .

Crucially, both the prover and the first part of the final verification circuit do not depend on the classical circuit C . This protocol should satisfy the following standard completeness and soundness properties.

- **Completeness:** For any quantum-classical circuit $D(\cdot) = C(Q(\cdot))$, input x , and output z such that $\Pr[D(x) = z] = 1 - \text{negl}(\lambda)$, it holds that

$$\Pr \left[v = \text{acc} \wedge z' = z : \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \mathcal{V}_{\text{QC}}^{\text{setup}}(1^\lambda, Q, x) \\ \pi \leftarrow \mathcal{P}_{\text{QC}}(\text{pk}, Q, x) \\ (v, z') \leftarrow \mathcal{V}_{\text{QC}}(Q, C, x, \text{pk}, \text{sk}, \pi) \end{array} \right] = 1 - \text{negl}(\lambda).$$

- **Soundness:** For any quantum-classical circuit $D(\cdot) = C(Q(\cdot))$, input x , output z such that $\Pr[D(x) = z] = 1 - \text{negl}(\lambda)$, and cheating prover $\mathcal{P}_{\text{QC}}^*$ with advice $|\psi\rangle$, it holds that

$$\Pr \left[v = \text{acc} \wedge z' \neq z : \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \mathcal{V}_{\text{QC}}^{\text{setup}}(1^\lambda, Q, x) \\ \pi \leftarrow \mathcal{P}_{\text{QC}}^*(|\psi\rangle, \text{pk}) \\ (v, z') \leftarrow \mathcal{V}_{\text{QC}}(Q, C, x, \text{pk}, \text{sk}, \pi) \end{array} \right] = \text{negl}(\lambda).$$

In addition, we may want the protocol to satisfy the following property.

- **Semi-malicious correctness:** For any quantum-classical circuit $D(\cdot) = C(Q(\cdot))$, input x , output z such that $\Pr[D(x) = z] = 1 - \text{negl}(\lambda)$, and $(\text{pk}, \text{sk}) \in \mathcal{V}_{\text{QC}}^{\text{setup}}(1^\lambda, Q, x)$, it holds that

$$\Pr \left[v = \text{acc} \wedge z' \neq z : \begin{array}{l} \pi \leftarrow \mathcal{P}_{\text{QC}}(\text{pk}, Q, x) \\ (v, z') \leftarrow \mathcal{V}_{\text{QC}}(Q, C, x, \text{pk}, \text{sk}, \pi) \end{array} \right] = \text{negl}(\lambda).$$

4.2 Delegation of quantum-classical circuits with quantum verifier

Quantum sampling protocol. Recently, [21] constructed a one-message protocol $\Pi_{\text{Samp}} = (\mathcal{P}_{\text{Samp}}, \mathcal{V}_{\text{Samp}})$ for delegating the computation of a quantum circuit Q on classical input x . The soundness of their protocol stipulates that if the verifier accepts an output z , then z is ϵ -close to a sample from the classical distribution induced by running Q on x and then measuring the output.

- $\mathcal{P}_{\text{Samp}}(1^\lambda, 1^{1/\epsilon}, Q, x) \rightarrow |\psi\rangle$: the prover takes the security parameter 1^λ , the accuracy parameter $1^{1/\epsilon}$, a quantum circuit Q , and a classical input x , and samples a quantum proof $|\psi\rangle$.

- $\mathcal{V}_{\text{Samp}}(1^\lambda, 1^{1/\epsilon}, Q, x, |\psi\rangle) \rightarrow (v, z)$: the verifier samples a binary string h independently of $|\psi\rangle$ (where each position indicates a computational basis or Hadamard basis measurement), samples $e \leftarrow M_{XZ}(|\psi\rangle, h)$ (see Definition 2), and applies a classical circuit $\mathcal{V}_{\text{out}}(Q, x, h, e)$ to obtain output (v, z) where $v \in \{\text{acc}, \text{rej}\}$ and $z \in \{0, 1\}^*$.

[21] show that the protocol satisfies the following.

- **Completeness:** For any $\epsilon(\lambda) = 1/\text{poly}(\lambda)$ and quantum circuit Q with input x ,

$$\Pr[(\text{rej}, \perp) \leftarrow \mathcal{V}_{\text{Samp}}(1^\lambda, 1^{1/\epsilon}, Q, x, |\psi\rangle) : |\psi\rangle \leftarrow \mathcal{P}_{\text{Samp}}(1^\lambda, 1^{1/\epsilon}, Q, x)] = \text{negl}(\lambda).$$

- **Soundness:** For any circuit Q , input x , cheating prover $\mathcal{P}_{\text{Samp}}^*$, $\epsilon(\lambda) = 1/\text{poly}(\lambda)$, and sufficiently large $\lambda \in \mathbb{N}$,
 - let $(v, z) \leftarrow \mathcal{V}_{\text{Samp}}(1^\lambda, 1^{1/\epsilon}, Q, x, \mathcal{P}_{\text{Samp}}^*(1^\lambda, 1^{1/\epsilon}, Q, x))$,
 - and define $z_{\text{ideal}} = \perp$ if $v = \text{rej}$ and $z_{\text{ideal}} \leftarrow Q(x)$ if $v = \text{acc}$.
Then it holds that $\|(v, z) - (v, z_{\text{ideal}})\|_1 \leq \epsilon(\lambda)$.

Quantum-classical circuits. Now consider any pseudo-deterministic quantum circuit D that can be split into two parts Q, C , where Q is quantum and C is classical. That is, $D(x) = C(Q(x))$, where $Q(x)$ outputs a classical string \hat{z} , and then $C(\hat{z})$ outputs z . Since D is pseudo-deterministic, each input x results in some fixed z with overwhelming probability. However, x still may induce an arbitrary distribution over intermediate values \hat{z} .

We now show that parallel repetition of Π_{Samp} gives a one-message protocol for delegating computation of $D(x)$ with *negligible* soundness, and where the prover's algorithm depends *only* on Q . In particular, consider the following protocol $\Pi_{\text{QVer}} = (\mathcal{P}_{\text{QVer}}, \mathcal{V}_{\text{QVer}})$ for delegating the computation of $D(\cdot) = C(Q(\cdot))$.

- $\mathcal{P}_{\text{QVer}}(1^\lambda, Q, x) \rightarrow |\psi\rangle$: the prover obtains a description of Q and an input x , sets $\epsilon = 1/4$, and runs λ copies of $\mathcal{P}_{\text{Samp}}(1^\lambda, 1^{1/\epsilon}, Q, x)$ to produce a proof $|\psi\rangle := (|\psi_1\rangle, \dots, |\psi_\lambda\rangle)$.
- $\mathcal{V}_{\text{QVer}}(Q, C, x, |\psi\rangle) \rightarrow (v, z)$: we split this verifier into three parts $(\mathcal{V}_{\text{QC}}^{\text{meas}}, \mathcal{V}_{\text{QC}}^{\text{test}}, \mathcal{V}_{\text{QC}}^{\text{out}})$.
 - $\mathcal{V}_{\text{QC}}^{\text{meas}}(Q, x, |\psi\rangle)$ samples h_i according to the distribution defined by $\mathcal{V}_{\text{Samp}}$ and obtains $e_i \leftarrow M_{X,Z}(|\psi_i\rangle, h_i)$ for each $i \in [\lambda]$. Set $h := (h_1, \dots, h_\lambda)$ and $e := (e_1, \dots, e_\lambda)$.
 - $\mathcal{V}_{\text{QC}}^{\text{test}}(Q, x, h, e)$ applies $\mathcal{V}_{\text{out}}(Q, x, h_i, e_i)$ for each $i \in \lambda$ to obtain (v_i, \hat{z}_i) . If any $v_i = \text{rej}$, then output (rej, \perp) (and do not proceed to $\mathcal{V}_{\text{QC}}^{\text{out}}$), and otherwise set $\hat{z} := (\hat{z}_1, \dots, \hat{z}_\lambda)$ and continue.
 - $\mathcal{V}_{\text{QC}}^{\text{out}}(C, \hat{z})$ computes $z_i := C(\hat{z}_i)$ for each $i \in [\lambda]$, determines the most frequently occurring z in $\{z_i\}_{i \in [\lambda]}$, and outputs (acc, z) .

Lemma 1. Π_{QVer} satisfies completeness and statistical soundness as defined in Definition 3 for delegating the computation of a pseudo-deterministic circuit $D(\cdot) = C(Q(\cdot))$.

Proof. Fix an x and let z be such that $C(Q(x)) = z$ with probability $1 - \text{negl}(\lambda)$. Consider the projector P that corresponds to running the single-copy verifier $\mathcal{V}_{\text{Samp}}$ on some part $|\psi_i\rangle$ of the prover's proof and accepting if the output is (acc, z') for some $z' \neq z$. By the soundness of Π_{Samp} , this projector will accept with probability at most $1/4 + \text{negl}(\lambda)$ on *any* prover state $|\psi_i\rangle$.

Note that soundness of Π_{QVer} is violated only if the following procedure accepts. Partition the prover's proof $|\psi\rangle$ into λ registers $|\psi_1\rangle, \dots, |\psi_\lambda\rangle$, apply P to each, and accept if at least $\lambda/2$ of the projectors accept. Even though $|\psi_1\rangle, \dots, |\psi_\lambda\rangle$ may be entangled, it still holds that P accepts on each $|\psi_i\rangle$ individually with probability at most $1/4 + \text{negl}(\lambda)$, since the registers are disjoint. That is, conditioned on any sequence of previous results of measuring $|\psi_1\rangle, \dots, |\psi_i\rangle$, applying P to $|\psi_{i+1}\rangle$ will accept with probability at most $1/4 + \text{negl}(\lambda)$. Thus, the distribution on number of acceptances is stochastically dominated by the distribution arising from independent Bernoulli trials that each output 1 with probability $1/4 + \text{negl}(\lambda)$. Applying Chernoff to this distribution, we see that the probability that at least $\lambda/2$ projectors accept is $\text{negl}(\lambda)$.

To show completeness, we know from the completeness of Π_{Samp} and a union bound that $\mathcal{V}_{\text{QVer}}$ will accept all parts $|\psi_i\rangle$ of the honest prover's proof, except with negligible probability. Conditioned on this, soundness of Π_{Samp} implies that for each i , the verifier will obtain $z' \neq z$ with probability at most $1/4 + \text{negl}(\lambda)$, and so again by Chernoff, the verifier will output (acc, z) except with negligible probability.

4.3 Making the verifier classical

Measurement protocol. [41] constructed a four-message protocol $\Pi_{\text{meas}} = (\mathcal{P}_{\text{meas}}, \mathcal{V}_{\text{meas}})$ between a quantum prover and a classical verifier, with the following syntax.

- $\mathcal{V}_{\text{meas}}(1^\lambda, h) \rightarrow (\text{pk}, \text{sk})$: the verifier, on input a string of basis choices h , samples a public key pk and a secret key sk , and sends pk to the prover.
- $\mathcal{P}_{\text{meas}}(\text{pk}, |\psi\rangle) \rightarrow (y, |\text{st}\rangle)$: the prover generates a classical commitment y , which it sends to the verifier, and a quantum internal state $|\text{st}\rangle$.
- The verifier samples $c \leftarrow \{0, 1\}$ and sends c to the prover, where $c = 0$ indicates a *test* round, and $c = 1$ indicates a *Hadamard* round.
- $\mathcal{P}_{\text{meas}}(|\text{st}\rangle, c) \rightarrow \pi$: The prover generates a classical proof π and sends it to the verifier.
- $\mathcal{V}_{\text{meas}}(\text{pk}, \text{sk}, y, c, \pi) \rightarrow \text{out}$: If $c = 0$, the verifier computes some classical circuit $\mathcal{V}_{\text{meas}, T}(\text{pk}, y, \pi) \rightarrow \text{out}$, where $\text{out} \in \{\text{acc}, \text{rej}\}$, and if $c = 1$, the verifier computes some classical circuit $\mathcal{V}_{\text{meas}, H}(\text{sk}, y, \pi) \rightarrow \text{out}$, where $\text{out} \in \{0, 1\}^*$.

This protocol satisfies the following property.

Lemma 2 ([41]). *Let $(\mathcal{P}_{\text{meas}}^*, |\psi_{\text{init}}\rangle)$ be any polynomial-size cheating prover for Π_{meas} , and suppose that there exists an h such that the probability that the verifier accepts if their basis choice was h and $c = 0$ is $1 - \text{negl}(\lambda)$. Then there exists a state ρ such that for all h , the verifier's output if $c = 1$ is computationally indistinguishable from $M_{XZ}(\rho, h)$.*

As observed in [3], a similar lemma follows by combining the claims [41, Claim 5.7] and [41, Claim 7.3]. The lemma stated above is potentially stronger in that (i) it considers non-uniform cheating provers with advice $|\psi_{\text{init}}\rangle$, and (ii) in the premise, it is only required that $\mathcal{P}_{\text{meas}}^*$ is accepted in the test round with probability $1 - \text{negl}(\lambda)$ for a *single* basis choice h rather than all. However, it is easy to see that (i) follows from the (standard) assumption that LWE is hard against non-uniform QPT adversaries. Next, (ii) follows due to properties of the extended trapdoor claw-free function used in [41]’s protocol. Indeed, [41] shows that for any two basis choices h_0, h_1 , no QPT prover will be able to distinguish between pk sampled by $\mathcal{V}(1^\lambda, h_0)$ and pk sampled by $\mathcal{V}(1^\lambda, h_1)$, assuming QLWE. Then since the circuit $\mathcal{V}_{\text{meas}, T}$ does not depend on sk , it follows that the probability that $\mathcal{P}_{\text{meas}}^*$ is accepted in a test round is negligibly close for *all* values of h . Thus, if $\mathcal{P}_{\text{meas}}^*$ passes the test round with $1 - \text{negl}(\lambda)$ probability for any h , it will pass with $1 - \text{negl}(\lambda)$ for all h .

A single repetition. We now combine the measurement protocol with Π_{QVer} to obtain Π_{single} , which is a delegation protocol for $D(\cdot) = C(Q(\cdot))$ with a classical verifier and a non-trivial soundness property (though not negligibly sound).

- $\mathcal{V}_{\text{single}}(1^\lambda, Q, x) \rightarrow (\text{pk}, \text{sk})$: the verifier samples h according to the distribution defined by $\mathcal{V}_{\text{QVer}}$, samples $(\text{pk}, \text{sk}) \leftarrow \mathcal{V}_{\text{meas}}(1^\lambda, h)$, and sends (pk, Q, x) to the prover.
- $\mathcal{P}_{\text{single}}(\text{pk}, Q, x) \rightarrow (y, |\text{st}\rangle)$: the prover samples $|\psi\rangle \leftarrow \mathcal{P}_{\text{QVer}}(1^\lambda, Q, x)$, computes $(y, |\text{st}\rangle) \leftarrow \mathcal{P}_{\text{meas}}(\text{pk}, |\psi\rangle)$, and sends y to the verifier.
- The verifier samples $c \leftarrow \{0, 1\}$ and sends c to the prover.
- $\mathcal{P}_{\text{single}}(|\text{st}\rangle, c) \rightarrow \pi$: the prover samples $\pi \leftarrow \mathcal{P}_{\text{meas}}(|\text{st}\rangle, c)$ and sends π to the verifier.
- $\mathcal{V}_{\text{single}}(Q, C, x, \text{pk}, \text{sk}, y, c, \pi) \rightarrow (v, z)$: we split the verifier into two parts $\mathcal{V}_{\text{single}}^{\text{vrfy}}$ and $\mathcal{V}_{\text{single}}^{\text{out}}$.
 - $\mathcal{V}_{\text{single}}^{\text{vrfy}}(Q, x, \text{pk}, \text{sk}, y, c, \pi)$ does the following. If $c = 0$, compute $\text{out} \leftarrow \mathcal{V}_{\text{meas}, T}(\text{pk}, y, \pi)$, where $\text{out} \in \{\text{acc}, \text{rej}\}$, and output (out, \perp) . If $c = 1$, compute $e \leftarrow \mathcal{V}_{\text{meas}, H}(\text{sk}, y, \pi)$ and run $\mathcal{V}_{\text{QVer}}^{\text{test}}(Q, x, h, e)$ to obtain either rej or \hat{z} . In the first case, output (rej, \perp) , and in the second case, continue.
 - $\mathcal{V}_{\text{single}}^{\text{out}}(C, \hat{z})$ computes and outputs $(\text{acc}, z) := \mathcal{V}_{\text{QVer}}^{\text{out}}(C, \hat{z})$.

We show that this protocol satisfies the following.

Lemma 3. *Let $(\mathcal{P}_{\text{single}}^*, |\psi_{\text{init}}\rangle)$ be any polynomial-size cheating prover for Π_{single} , and let Q, C, x, z be such that $C(Q(x)) = z$ with probability $1 - \text{negl}(\lambda)$. Suppose that the verifier outputs (acc, \perp) when $c = 0$ with probability $1 - \text{negl}(\lambda)$. Then assuming QLWE, the probability that the verifier outputs (acc, z') for $z' \neq z$ when $c = 1$ is $\text{negl}(\lambda)$.*

Proof. Fix Q, C, x, z . Consider the predicate F that has Q, C, x, z hard-coded, and on input basis choice h and measurement results e , checks whether $\mathcal{V}_{\text{single}}$ outputs (acc, z') for $z' \neq z$. That is, F first runs $\mathcal{V}_{\text{QVer}}^{\text{test}}(Q, x, h, e)$. If this procedure

accepted and output \widehat{z} , F computes $(\text{acc}, z') := \mathcal{V}_{\text{QVer}}^{\text{out}}(C, \widehat{z})$ and outputs 1 if $z' \neq z$. Let $p(h)$ be the probability that basis choice h is sampled by $\mathcal{V}_{\text{single}}$. Then we want to show that

$$\sum_h p(h)F(h, e) = \text{negl}(\lambda),$$

where e is obtained by running the protocol with basis choice h and $c = 1$. Now since $\mathcal{P}_{\text{single}}^*$ is accepted with overwhelming probability in the test round, it satisfies the premise in Lemma 2. So by Lemma 2, the above expression is negligibly close to

$$\sum_h p(h)F(h, M_{X,Z}(\rho, h)),$$

since F is an efficient function outputting a single bit. Finally, this expression is negligible due to the soundness of the underlying information-theoretic protocol given by Lemma 1.

Parallel repetition. Now we repeat the above protocol in parallel to obtain a QCQC protocol Π_{QC} that satisfies Definition 7.

- $\mathcal{V}_{\text{QC}}^{\text{setup}}(1^\lambda, Q, x) \rightarrow (\text{pk}, \text{sk})$: set $n \geq \lambda^{1+\epsilon}$ and for each $i \in [n]$, sample $(\text{pk}_i, \text{sk}_i) \leftarrow \mathcal{V}_{\text{single}}(1^\lambda, Q, x)$. Set $\text{pk} := (\text{pk}_1, \dots, \text{pk}_n)$ and $\text{sk} := (\text{sk}_1, \dots, \text{sk}_n)$, and send (pk, Q, x) to the prover.
- $\mathcal{P}_{\text{QC}}(\text{pk}, Q, x) \rightarrow (y, |\text{st}\rangle)$: for each $i \in [n]$, the prover samples $(y_i, |\text{st}_i\rangle) \leftarrow \mathcal{P}_{\text{single}}(\text{pk}_i, Q, x)$, sets $y := (y_1, \dots, y_n)$ and $|\text{st}\rangle := (|\text{st}_1\rangle, \dots, |\text{st}_n\rangle)$, and sends y to the verifier.
- The verifier samples $d \leftarrow \mathcal{HW}_{n,\lambda}$ and sends d to the prover.
- $\mathcal{P}_{\text{QC}}(|\text{st}\rangle, d) \rightarrow \pi$: for each $i \in [n]$, the prover samples $\pi_i \leftarrow \mathcal{P}_{\text{single}}(|\text{st}_i\rangle, d_i)$ and sends $\pi := (\pi_1, \dots, \pi_n)$ to the verifier.
- $\mathcal{V}_{\text{QC}}(Q, C, x, \text{pk}, \text{sk}, y, d, \pi) \rightarrow \text{out}$: we split the verifier into two parts $\mathcal{V}_{\text{QC}}^{\text{vrfy}}$ and $\mathcal{V}_{\text{QC}}^{\text{out}}$.
 - $\mathcal{V}_{\text{QC}}^{\text{vrfy}}(Q, x, \text{pk}, \text{sk}, y, d, \pi)$ runs $\mathcal{V}_{\text{single}}^{\text{vrfy}}(Q, x, \text{pk}_i, \text{sk}_i, y_i, d_i, \pi_i)$ for each $i \in [n]$. If any outputs (rej, \perp) then output (rej, \perp) . Otherwise, obtain strings $\{\widehat{z}_i\}_{i:d_i=1}$.
 - $\mathcal{V}_{\text{QC}}^{\text{out}}(C, \{\widehat{z}_i\}_{i:d_i=1})$ computes the λ outputs $(\text{acc}, z_i) := \mathcal{V}_{\text{single}}^{\text{out}}(C, \widehat{z}_i)$. Then determine the most frequently occurring z in the resulting set and output (acc, z) .

Lemma 4. *Assuming QLWE, Π_{QC} satisfies Definition 7*

Proof. First, we show soundness. Fix C, Q, x, z such that $C(Q(x)) = z$ with probability $1 - \text{negl}(\lambda)$. Define a Hadamard-round verification projector for Π_{single} that checks whether $\mathcal{V}_{\text{single}}$ outputs (acc, z') for $z' \neq z$. Lemma 3 and Theorem 5 imply that, conditioned on all test rounds (i such that $d_i = 0$) accepting in Π_{QC} , this verifier only accepts at most $\lceil \lambda/2 \rceil - 1$ of the indices for which $d_i = 1$, except with negligible probability. Thus, conditioned on $\mathcal{V}_{\text{QC}}^{\text{vrfy}}$ accepting, a strict

majority of the Hadamard rounds will result in the outcome z , except with negligible probability.

It remains to argue semi-malicious correctness. We observe that the strings $\{\widehat{z}_i\}_{i:d_i=1}$ are obtained via computational basis measurements (which follows from the description of [21]’s sampBQP protocol). Since we are considering an honest prover, it suffices to argue that computational basis measurements are always performed correctly under $(\mathbf{pk}, \mathbf{sk})$. The parts of $(\mathbf{pk}, \mathbf{sk})$ that are used for computational basis measurements are injective trapdoor keys, and the measurements will be correct if the keys are indeed injective. The sampling procedure for injective keys given in [41, Section 9.2] can be made to produce an injective key with probability 1 over its random coins, by outputting a fixed injective key in the case that the initial sampling failed to produce an injective key.

4.4 Four-message CVQC

In this section, we use CVQC for quantum-classical circuits in combination with quantum fully-homomorphic encryption to construct a four-message blind CVQC protocol that enjoys composability. That is, the protocol that we construct satisfies an ideal functionality that either delivers the correct output $D(x)$ or a \perp symbol to the verifier, depending on a bit b input by the prover. Security is argued in the simulation sense according to Definition 6, which in particular implies that the prover’s bit b cannot depend on the verifier’s input x . We first give our definition of composable blind CVQC, which is essentially the same as the notion of *blind and verifiable delegated quantum computation* studied by [26], but adapted to the classical verifier setting.

Definition 8 (Composable Blind CVQC). *Consider the two-party functionality \mathcal{F}_D between classical verifier and quantum prover defined by a pseudo-deterministic circuit $D(\cdot)$. It takes an input string x from the verifier and an input bit b from the prover (which is always 0 if the prover is honest). If $b = 0$, it delivers $D(x)$ to the verifier (and no output to the prover), and if $b = 1$, it delivers \perp to the verifier (and no output to the prover). Protocol Π is a composable blind CVQC protocol if for any pseudo-deterministic D , it satisfies Definition 6 for functionality \mathcal{F}_D .*

To construct such a protocol, we use the following ingredients.

- A four-message CVQC protocol for quantum-classical circuits Π_{QC} (Definition 7).
- Quantum fully-homomorphic encryption with classical keys (QFHE.Gen, QFHE.Enc, QFHE.Eval, QFHE.Dec) (Definition 4).

The construction is given in Protocol 1.

Theorem 6. *Assuming QLWE, Protocol 1 is a composable blind CVQC protocol satisfying definition Definition 8.*

Protocol 1: Composable Blind CVQC

- $\mathcal{V}_{\text{blind}}^{\text{setup}}(1^\lambda, D, x)$: the verifier samples

$$(\text{pk}_{\text{QFHE}}, \text{sk}_{\text{QFHE}}) \leftarrow \text{QFHE.Gen}(1^\lambda), \text{ct} \leftarrow \text{QFHE.Enc}(\text{pk}_{\text{QFHE}}, x),$$
 then it samples $(\text{pk}_{\text{eval}}, \text{sk}_{\text{eval}}) \leftarrow \mathcal{V}_{\text{QC}}^{\text{setup}}(1^\lambda, \text{QFHE.Eval}(D, \cdot), \text{ct})$, sets

$$\text{pk} := (\text{pk}_{\text{eval}}, \text{QFHE.Eval}(D, \cdot), \text{ct}), \text{sk} := (\text{sk}_{\text{eval}}, \text{sk}_{\text{QFHE}}),$$
 and sends pk to the prover.
- $\mathcal{P}_{\text{blind}}(\text{pk})$: the prover samples $(y, |\text{st}\rangle) \leftarrow \mathcal{P}_{\text{QC}}(\text{pk}_{\text{eval}}, \text{QFHE.Eval}(D, \cdot), \text{ct})$, and sends y to the verifier.
- The verifier samples d according to Π_{QC} .
- $\mathcal{P}_{\text{blind}}(|\text{st}\rangle, d)$: the prover samples $\pi \leftarrow \mathcal{P}_{\text{QC}}(|\text{st}\rangle, d)$, and sends π to the verifier.
- $\mathcal{V}_{\text{blind}}(\text{pk}, \text{sk}, y, d, \pi)$: the verifier first runs

$$\mathcal{V}_{\text{QC}}^{\text{verify}}(\text{QFHE.Eval}(D, \cdot), \text{ct}, \text{pk}_{\text{eval}}, \text{sk}_{\text{eval}}, y, d, \pi)$$
 to obtain either rej or \tilde{z} . In the first case, it outputs \perp . In the second case, it computes

$$z := \mathcal{V}_{\text{QC}}^{\text{out}}(\text{QFHE.Dec}(\text{sk}_{\text{QFHE}}, \cdot), \tilde{z}),$$
 and outputs z .

Fig. 1. A four-message composable blind CVQC from QLWE.

Proof. First, we argue that in case neither party is corrupted, the verifier’s output is correct with overwhelming probability. This follows from the completeness of Π_{QC} , since for any pseudo-deterministic circuit D , input x , $(\text{pk}_{\text{QFHE}}, \text{sk}_{\text{QFHE}}) \in \text{QFHE.Gen}(1^\lambda)$, and $\text{ct} \in \text{QFHE.Enc}(\text{pk}_{\text{QFHE}}, x)$, it holds that $\text{QFHE.Dec}(\text{sk}_{\text{QFHE}}, \text{QFHE.Eval}(\text{pk}_{\text{QFHE}}, D, \cdot))$ applied to ct outputs $D(x)$ with overwhelming probability, so the computation performed by Π_{QC} is pseudo-deterministic. Next, note that simulation in case the verifier is corrupted is trivial, since the honest prover has no input or output.

It remains to argue security when the prover is corrupted, for which we define the following simulator. The simulator samples a QFHE encryption of 0, and then interacts with the prover as the honest verifier delegating the circuit $\text{QFHE.Eval}(D, \cdot)$. At the end of the protocol, the simulator first runs $\mathcal{V}_{\text{QC}}^{\text{verify}}$. If this results in rej , send $b = 1$ to the ideal functionality, and otherwise send $b = 0$ to the ideal functionality. Now consider the following sequence of hybrids.

- \mathcal{H}_0 : This is the real interaction between honest verifier and malicious prover, resulting in a verifier output in $\{0, 1\}^* \cup \{\perp\}$ and a final prover state.
- \mathcal{H}_1 : In this hybrid, we change how the final verifier output is computed. In particular, in the case that $\mathcal{V}_{\text{QC}}^{\text{verify}}$ does not reject, compute the output $z = D(x)$. This is computationally indistinguishable from \mathcal{H}_0 due to the

- soundness of Π_{QC} , which guarantees that, conditioned on the verifier accepting, the output will be equal to $D(x)$ except with negligible probability.
- \mathcal{H}_2 : In this hybrid, we change how the verifier’s first message is computed. In particular, set ct to be a QFHE encryption of 0 rather than x . This is computationally indistinguishable from \mathcal{H}_1 due to the semantic security of QFHE, since the QFHE secret key is no longer needed to compute the output of the experiment.
 - \mathcal{H}_3 : In this hybrid, if $\mathcal{V}_{\text{QC}}^{\text{verify}}$ rejects, then send $b = 1$ to the ideal functionality, and otherwise send $b = 0$. This is identical to \mathcal{H}_2 , since we have just moved the computation of $z = D(x)$ to the ideal functionality. This is the simulator, completing the proof.

5 Secure Quantum Computation

In this section, we show applications of the CVQC protocols constructed in Section 4 to secure quantum computation with classical communication (where only a single party requires quantum capabilities).

In Section 5.1, we give a generic compiler from composable blind CVQC to multi-party quantum computation with classical computation, assuming the existence of post-quantum oblivious transfer. In fact, assuming two-message post-quantum oblivious transfer, the compiler only adds two rounds of interaction to the composable blind CVQC protocol.

In the full version, we show how to optimize the round complexity of multi-party quantum computation, achieving a three-round protocol. This protocol requires a (succinct and reusable) PKI setup and security follows from QLWE in the quantum random oracle model. In addition, we show how to construct a two-message two-party protocol between quantum sender and classical receiver (a quantum NISC protocol). Security of this protocol follows from QLWE in the quantum random oracle model.

5.1 A generic construction of multi-party quantum computation

In this section, we show that, assuming post-quantum two-message oblivious transfer, k -message composable blind CVQC implies $k + 2$ -round multi-party quantum computation between any n classical clients and a single quantum server. The protocol (in fact, all protocols in this section) is described for functionalities with a single public output, but this is easy to generalize to multiple private outputs using secret-key encryption.

Ingredients.

- A post-quantum round-optimal MPC protocol for classical reactive functionalities in the CRS model, to be treated as an oracle called MPC. Such a protocol is known from post-quantum two-message oblivious transfer, which is known from QLWE.

- A k -message composable blind CVQC protocol (Definition 8). We assume without loss of generality that k is even and the verifier sends the first message. The transcript of messages between verifier and prover are denoted $\text{msg}_1^{(V)}, \text{msg}_2^{(P)}, \dots, \text{msg}_{k-1}^{(V)}, \text{msg}_k^{(P)}$.

Protocol 2: Multi-Party Quantum Computation with Classical Communication

Public Information: An n -party pseudo-deterministic quantum functionality $D(\cdot, \dots, \cdot)$, security parameter λ , n classical client parties P_1, \dots, P_n , and a designated server S with quantum capabilities.

- **Round 1:** Each party P_i sends their input x_i to MPC.
- **Round 2:** MPC computes $\text{msg}_1^{(V)}$ and sends it to S .
- **Round 3:** S computes and broadcasts $\text{msg}_2^{(P)}$.
- ⋮
- **Round k :** MPC computes $\text{msg}_{k-1}^{(V)}$ and sends it to S .
- **Round $k+1$:** S computes and broadcasts $\text{msg}_k^{(P)}$.
- **Round $k+2$:** MPC computes the output $z \in \{0, 1\}^* \cup \{\perp\}$ of the CVQC protocol and delivers z to each P_i .

Fig. 2. An MPQC protocol with classical communication.

Theorem 7. *Protocol 2 satisfies Definition 6.*

Proof. First consider any adversary $\{\text{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$ that corrupts a set of parties M such that $S \in M$, and let $H := [n] \setminus M$. The simulator is defined below. We allow Sim to maintain the MPC oracle, intercepting the adversary’s inputs and computing the outputs.

$\text{Sim}(\{x_i\}_{i \in M}, |\psi\rangle_{\text{Adv}})$:

- Obtain $\{x_i\}_{i \in M}$ from Adv ’s initial query to MPC.
- Invoke the malicious prover simulator for the composable blind CVQC protocol, which computes messages on behalf of MPC, and interacts with Adv (controlling the server) until after Round $k+1$. The simulator outputs a bit b .
- If $b = 0$, query the ideal functionality with $\{x_i\}_{i \in M}$ to obtain z and deliver z to Adv . Otherwise, if $b = 1$, send **abort** to the ideal functionality and deliver \perp to Adv .

Now consider the following hybrids.

- \mathcal{H}_0 : This is the real distribution $\text{REAL}_{II,Q}(\text{Adv}_\lambda, \{x_i\}_{i \in [n]}, |\psi\rangle_{\text{Adv}})$, where MPC is implemented honestly as the CVQC verifier with input x_1, \dots, x_n .
- \mathcal{H}_1 : Invoke the malicious prover simulator for the composable blind CVQC protocol to simulate the interaction between Adv and MPC through Round $k + 1$. Also, change how the final message from MPC is computed. That is, if the CVQC simulator output $b = 0$, compute $z = D(x_1, \dots, x_n)$ (where $\{x_i\}_{i \in M}$ were obtained from Adv’s first query to MPC, and $\{x_i\}_{i \in H}$ are the honest party inputs) and deliver z to Adv, and otherwise, deliver \perp to Adv. This is indistinguishable from \mathcal{H}_0 due to the security of the CVQC protocol.
- \mathcal{H}_2 : Obtain the output z by querying the ideal functionality with $\{x_i\}_{i \in M}$, which is perfectly indistinguishable from \mathcal{H}_1 . This is the simulator.

Now consider any adversary $\{\text{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$ that corrupts a set of parties $M \subset \{P_1, \dots, P_n\}$. Security in this case is almost immediate. Indeed, the simulator will obtain $\{x_i\}_{i \in M}$ from Adv’s query to MPC in Round 1, query the ideal functionality to obtain z and then deliver z to Adv in Round $k + 2$. This is indistinguishable from the real distribution due to the correctness of the CVQC protocol, that is, the requirement that the correct output is generated when neither of the parties is corrupted.

Combining with our four-message composable blind CVQC protocol from Section 4.4, we obtain the following corollary.

Corollary 1. *Assuming QLWE, there exists a six-round multi-party quantum computation protocol in the CRS model between n classical clients and one quantum server for computing any n -party pseudo-deterministic functionality $D(\cdot, \dots, \cdot)$.*

Acknowledgements

Thank you to Dakshita Khurana and Giulio Malavolta for helpful discussions, and to anonymous reviewers for comments and suggestions.

References

1. Afshar, A., Mohassel, P., Pinkas, B., Riva, B.: Non-interactive secure computation based on cut-and-choose. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 387–404. Springer, Heidelberg (May 2014). https://doi.org/10.1007/978-3-642-55220-5_22
2. Agarwal, A., Bartusek, J., Goyal, V., Khurana, D., Malavolta, G.: Post-quantum multi-party computation. Eurocrypt (2021)
3. Alagic, G., Childs, A.M., Grilo, A.B., Hung, S.H.: Non-interactive classical verification of quantum computation. pp. 153–180. LNCS, Springer, Heidelberg (Mar 2020). https://doi.org/10.1007/978-3-030-64381-2_6
4. Alon, B., Chung, H., Chung, K.M., Huang, M.Y., Lee, Y., Shen, Y.C.: Round efficient secure multiparty quantum computation with identifiable abort. Cryptology ePrint Archive, Report 2020/1464 (2020), <https://eprint.iacr.org/2020/1464>

5. Ananth, P., Jain, A., Jin, Z., Malavolta, G.: Unbounded multi-party computation from learning with errors. Eurocrypt (2021)
6. Asharov, G., Jain, A., López-Alt, A., Tromer, E., Vaikuntanathan, V., Wichs, D.: Multiparty computation with low communication, computation and interaction via threshold FHE. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 483–501. Springer, Heidelberg (Apr 2012). https://doi.org/10.1007/978-3-642-29011-4_29
7. Badrinarayanan, S., Garg, S., Ishai, Y., Sahai, A., Wadia, A.: Two-message witness indistinguishability and secure computation in the plain model from new assumptions. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part III. LNCS, vol. 10626, pp. 275–303. Springer, Heidelberg (Dec 2017). https://doi.org/10.1007/978-3-319-70700-6_10
8. Bartusek, J., Coladangelo, A., Khurana, D., Ma, F.: On the round complexity of secure quantum computation. CRYPTO (2021)
9. Bartusek, J., Coladangelo, A., Khurana, D., Ma, F.: One-way functions imply secure computation in a quantum world. CRYPTO (2021)
10. Bartusek, J., Malavolta, G.: Candidate obfuscation of null quantum circuits and witness encryption for qma. Cryptology ePrint Archive, Report 2021/421 (2021), <https://eprint.iacr.org/2021/421>
11. Ben-Or, M., Crépeau, C., Gottesman, D., Hassidim, A., Smith, A.: Secure multi-party quantum computation with (only) a strict honest majority. In: 47th FOCS. pp. 249–260. IEEE Computer Society Press (Oct 2006). <https://doi.org/10.1109/FOCS.2006.68>
12. Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In: 20th ACM STOC. pp. 1–10. ACM Press (May 1988). <https://doi.org/10.1145/62212.62213>
13. Benhamouda, F., Jain, A., Komargodski, I., Lin, H.: Multi-party reusable non-interactive secure computation from lwe. Eurocrypt (2021)
14. Brakerski, Z.: Quantum FHE (almost) as secure as classical. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part III. LNCS, vol. 10993, pp. 67–95. Springer, Heidelberg (Aug 2018). https://doi.org/10.1007/978-3-319-96878-0_3
15. Brakerski, Z., Christiano, P., Mahadev, U., Vazirani, U.V., Vidick, T.: A cryptographic test of quantumness and certifiable randomness from a single quantum device. In: Thorup, M. (ed.) 59th FOCS. pp. 320–331. IEEE Computer Society Press (Oct 2018). <https://doi.org/10.1109/FOCS.2018.00038>
16. Brakerski, Z., Döttling, N.: Two-message statistically sender-private OT from LWE. In: Beimel, A., Dziembowski, S. (eds.) TCC 2018, Part II. LNCS, vol. 11240, pp. 370–390. Springer, Heidelberg (Nov 2018). https://doi.org/10.1007/978-3-030-03810-6_14
17. Chardouvelis, O., Döttling, N., Malavolta, G.: Rate-1 secure function evaluation for bqp. Cryptology ePrint Archive, Report 2020/1454 (2020), <https://eprint.iacr.org/2020/1454>
18. Chase, M., Dodis, Y., Ishai, Y., Kraschewski, D., Liu, T., Ostrovsky, R., Vaikuntanathan, V.: Reusable non-interactive secure computation. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part III. LNCS, vol. 11694, pp. 462–488. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-26954-8_15

19. Chaum, D., Crépeau, C., Damgård, I.: Multiparty unconditionally secure protocols (abstract) (informal contribution). In: Pomerance, C. (ed.) CRYPTO'87. LNCS, vol. 293, p. 462. Springer, Heidelberg (Aug 1988). https://doi.org/10.1007/3-540-48184-2_43
20. Chia, N.H., Chung, K.M., Yamakawa, T.: Classical verification of quantum computations with efficient verifier. pp. 181–206. LNCS, Springer, Heidelberg (Mar 2020). https://doi.org/10.1007/978-3-030-64381-2_7
21. Chung, K.M., Lee, Y., Lin, H.H., Wu, X.: Constant-round blind classical verification of quantum sampling (2020)
22. Ciampi, M., Cojocaru, A., Kashefi, E., Mantri, A.: Secure quantum two-party computation: Impossibility and constructions. Cryptology ePrint Archive, Report 2020/1286 (2020), <https://eprint.iacr.org/2020/1286>
23. Crépeau, C., Gottesman, D., Smith, A.: Secure multi-party quantum computation. In: 34th ACM STOC. pp. 643–652. ACM Press (May 2002). <https://doi.org/10.1145/509907.510000>
24. Don, J., Fehr, S., Majenz, C., Schaffner, C.: Security of the Fiat-Shamir transformation in the quantum random-oracle model. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 356–383. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-26951-7_13
25. Dulek, Y., Grilo, A.B., Jeffery, S., Majenz, C., Schaffner, C.: Secure multi-party quantum computation with a dishonest majority. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part III. LNCS, vol. 12107, pp. 729–758. Springer, Heidelberg (May 2020). https://doi.org/10.1007/978-3-030-45727-3_25
26. Dunjko, V., Fitzsimons, J., Portmann, C., Renner, R.: Composable security of delegated quantum computation. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014, Part II. LNCS, vol. 8874, pp. 406–425. Springer, Heidelberg (Dec 2014). https://doi.org/10.1007/978-3-662-45608-8_22
27. Dupuis, F., Nielsen, J.B., Salvail, L.: Actively secure two-party evaluation of any quantum operation. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 794–811. Springer, Heidelberg (Aug 2012). https://doi.org/10.1007/978-3-642-32009-5_46
28. Fitzsimons, J.F., Hajdusek, M., Morimae, T.: Post hoc verification of quantum computation. Phys. Rev. Lett. **120**, 040501 (Jan 2018). <https://doi.org/10.1103/PhysRevLett.120.040501>, <https://link.aps.org/doi/10.1103/PhysRevLett.120.040501>
29. Fitzsimons, J.F., Kashefi, E.: Unconditionally verifiable blind quantum computation. Phys. Rev. A **96**, 012303 (Jul 2017). <https://doi.org/10.1103/PhysRevA.96.012303>, <https://link.aps.org/doi/10.1103/PhysRevA.96.012303>
30. Gheorghiu, A., Vidick, T.: Computationally-secure and composable remote state preparation. In: Zuckerman, D. (ed.) 60th FOCS. pp. 1024–1033. IEEE Computer Society Press (Nov 2019). <https://doi.org/10.1109/FOCS.2019.00066>
31. Goldreich, O.: The Foundations of Cryptography - Volume 2, Basic Applications. Cambridge University Press (2004)
32. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game or A completeness theorem for protocols with honest majority. In: Aho, A. (ed.) 19th ACM STOC. pp. 218–229. ACM Press (May 1987). <https://doi.org/10.1145/28395.28420>
33. Goyal, R.: Quantum multi-key homomorphic encryption for polynomial-sized circuits. Cryptology ePrint Archive, Report 2018/443 (2018), <https://eprint.iacr.org/2018/443>

34. Grilo, A.B., Lin, H., Song, F., Vaikuntanathan, V.: Oblivious transfer is in minicrypt. Eurocrypt (2021)
35. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Prabhakaran, M., Sahai, A.: Efficient non-interactive secure computation. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 406–425. Springer, Heidelberg (May 2011). https://doi.org/10.1007/978-3-642-20465-4_23
36. Ishai, Y., Prabhakaran, M., Sahai, A.: Founding cryptography on oblivious transfer - efficiently. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 572–591. Springer, Heidelberg (Aug 2008). https://doi.org/10.1007/978-3-540-85174-5_32
37. Kapourniotis, T., Kashefi, E., Music, L., Ollivier, H.: Delegating multi-party quantum computations vs. dishonest majority in two quantum rounds (2021)
38. Kashefi, E., Pappa, A.: Multiparty delegated quantum computing. Cryptography **1**, 12 (07 2017). <https://doi.org/10.3390/cryptography1020012>
39. Liu, Q., Zhandry, M.: Revisiting post-quantum Fiat-Shamir. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 326–355. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-26951-7_12
40. Mahadev, U.: Classical homomorphic encryption for quantum circuits. In: Thorup, M. (ed.) 59th FOCS. pp. 332–338. IEEE Computer Society Press (Oct 2018). <https://doi.org/10.1109/FOCS.2018.00039>
41. Mahadev, U.: Classical verification of quantum computations. In: Thorup, M. (ed.) 59th FOCS. pp. 259–267. IEEE Computer Society Press (Oct 2018). <https://doi.org/10.1109/FOCS.2018.00033>
42. Morgan, A., Pass, R., Polychriadiou, A.: Succinct non-interactive secure computation. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part II. LNCS, vol. 12106, pp. 216–245. Springer, Heidelberg (May 2020). https://doi.org/10.1007/978-3-030-45724-2_8
43. Morimae, T.: Verification for measurement-only blind quantum computing. Phys. Rev. A **89**, 060302 (Jun 2014). <https://doi.org/10.1103/PhysRevA.89.060302>, <https://link.aps.org/doi/10.1103/PhysRevA.89.060302>
44. Mukherjee, P., Wichs, D.: Two round multiparty computation via multi-key FHE. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 735–763. Springer, Heidelberg (May 2016). https://doi.org/10.1007/978-3-662-49896-5_26
45. Peikert, C., Shiehian, S.: Noninteractive zero knowledge for NP from (plain) learning with errors. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part I. LNCS, vol. 11692, pp. 89–114. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-26948-7_4
46. Quach, W.: Uc-secure ot from lwe, revisited. In: Galdi, C., Kolesnikov, V. (eds.) Security and Cryptography for Networks. pp. 192–211. Springer International Publishing, Cham (2020)
47. Unruh, D.: Universally composable quantum multi-party computation. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 486–505. Springer, Heidelberg (May / Jun 2010). https://doi.org/10.1007/978-3-642-13190-5_25
48. Yao, A.C.C.: How to generate and exchange secrets. In: FOCS (1986)