Circuit Privacy for Quantum Fully Homomorphic Encryption

Giulio Malavolta*

Abstract

We study the problem of enforcing circuit privacy for the homomorphic evaluation of quantum circuits. We present a generic transformation from semi-honest to malicious circuit privacy for quantum fully homomorphic encryption (QFHE). Our compiler assumes minimal structural properties of the underlying QFHE scheme, which are satisfied by recent candidate constructions [Mahadev, FOCS 2018]. Thus we obtain a maliciously circuit private QFHE scheme, assuming the quantum hardness of (a circular variant of) the learning with errors (LWE) problem. This immediately implies the existence of a two-round secure function evaluation protocol for quantum circuits with input-indistinguishable security for the client (holding a quantum state $|\psi\rangle$) and unbounded simulation security for the server (evaluating a quantum circuit C).

*Max Planck Institute for Security and Privacy

Contents

1 Introduction

Fully homomorphic encryption (FHE) [Gen09] allows one to encrypt some message m and homomorphically evaluate a circuit C to produce

$$\mathsf{Enc}(\mathsf{pk}, m) \xrightarrow{\mathsf{Eval}(C,\cdot)} \mathsf{Enc}(\mathsf{pk}, C(m))$$

without the need to know the secret key. FHE realizes the vision of computing over encrypted data and, as a natural application, yields the following secure function evaluation (SFE) protocol: A client sends an encrypted message $\mathsf{Enc}(\mathsf{pk}, m)$ to a powerful server, who homomorphically evaluate some circuit C without learning the client's message. The server can then send the resulting $\mathsf{Enc}(\mathsf{pk}, C(m))$ to the client who decrypts it to recover the output of the computation. The communication complexity of this protocol is essentially optimal, as it depends (polynomially) only on the size of the message $|m|$, the size of the output $|C(m)|$ and the security parameter λ.

In terms of security, it is follows from the semantic security of the FHE scheme that the message m is hidden to the eyes of the server. Conversely, it is less clear whether the client is oblivious to which circuit C has been computed by the server. The compactness of the FHE guarantees that some information is lost, but there might be still some residual leakage from C left in the evaluated ciphertext. Ensuring that the client learns nothing beyond the result of the computation requires one to construct an FHE scheme that satisfies *circuit privacy*. This notion has been originally studied in the semi-honest settings [Gen09, vGHV10], i.e. assuming that the client faithfully executes the key generation and the encryption algorithm of the FHE scheme. More recently, the work of Ostrovsky et al. [OPP14] shows how an FHE scheme can be lifted to guarantee circuit privacy also in the malicious settings (where the client can behave arbitrarily). Loosely speaking, *malicious* circuit privacy guarantees that any key-ciphertext pair (pk, c), induces some (but not necessarily efficiently computable) encrypted input m^* such that the client learns nothing beyond $C(m^*)$. As an immediate application, maliciously circuit-private FHE implies the first low-communication two-round SFE protocol in the plain model with meaningful security guarantees on both sides.

The promise of quantum computing has motivated the study of quantum fully homomorphic encryption (QFHE) [BJ15], which extends the classical notion of FHE to the homomorphic computation of quantum circuits. In its most general form, QFHE allows one to evaluate the transformation

$$\mathsf{Enc}(\mathsf{pk}, |\psi\rangle) \xrightarrow{\mathsf{Eval}(C,\cdot)} \mathsf{Enc}(\mathsf{pk}, C(|\psi\rangle))$$

where $|\psi\rangle$ is some arbitrary quantum state and C is a quantum circuit. Analogously to the classical settings, QFHE enables natural communication-efficient SFE protocols for computing quantum circuits over encrypted quantum states. However, in contrast to the classical settings, the notion of malicious circuit privacy for QFHE has so far remained elusive. Motivated by the unsatisfactory state of affairs, we ask the following question:

Can we construct maliciously circuit-private QFHE?

1.1 Our Results

In this work we give a positive answer to the above question and we show how to construct maliciously circuit-private QFHE assuming the quantum hardness of the learning with errors (LWE) problem. As standard in the (Q)FHE settings, to enable the homomorphic computation

of *unbounded* quantum circuits, we additionally assume that the scheme is circularly secure. More precisely, we prove the following theorem.

Theorem 1.1 (Informal)**.** *Assuming the quantum hardness of the (circular) LWE problem, there exists a maliciously circuit-private QFHE scheme for all quantum circuits.*

Our scheme is presented in the form of a generic transformation: We show how to combine a semi-honest circuit private QFHE (e.g. the schemes from [Mah18a, Bra18]) with a maliciously circuit-private *classical* FHE to obtain a maliciously circuit private QFHE. In the quantum settings, the notion of malicious circuit privacy can be (roughly) interpreted as follows: For all key-ciphertext pairs $(\mathsf{pk}, |\phi\rangle)$ there exists some well-defined (but not necessarily efficiently computable) quantum state $|\psi^*\rangle$ such that an evaluated ciphertext carries no information beside $C(|\psi^*\rangle)$. As a direct application of this notion, we obtain the following SFE protocol for quantum computation.

- **1st Round:** The client samples a QFHE key pair $(\mathsf{sk}, \mathsf{pk})$ and sends to the server $\mathsf{Enc}(\mathsf{pk}, |\psi\rangle)$.
- **2nd Round:** The server computes homomorphically $\mathsf{Enc}(\mathsf{pk}, C(|\psi\rangle))$ and returns the resulting ciphertext.
- **Output:** The client decrypts the ciphertext and recovers $C(|\psi\rangle)$.

The state of the client $|\psi\rangle$ is hidden (in an indistinguishability sense) by the semantic security of the QFHE scheme, whereas the malicious circuit privacy guarantees that the circuit C is not leaked to the client, beyond what is already revealed by $C(|\psi\rangle)$.

It is worth mentioning that our QFHE scheme does not assume a trusted setup or a common reference string (CRS), i.e. it is in the plain model. We stress however that, as opposed to the classical case, in the quantum settings even having a CRS does not immediately trivialize the problem: This is because, at present, we do not know of any (publicly verifiable) quantum counterpart of non-interactive zero-knowledge (NIZK) proofs for certifying the validity of a QFHE ciphertext. Finally, as a bonus, we also discuss how to extend our techniques to multi-hop and multi-key homomorphic evaluation of quantum circuits.

1.2 Technical Overview

In the following we give a cursory overview of the main ideas behind our result. For further details, we refer the reader to the technical sections.

Circuit Privacy for Classical FHE. As our approach is intimately related with the transformation of Ostrovsky et al. [OPP14], it is useful to briefly recall the main idea of their work. On a high level, their (simplified) approach to construct maliciously circuit-private FHE relies on the conditional disclosure of secret (CDS) paradigm. A CDS protocol allows a receiver to compute a commitment $\mathsf{Com}(w)$ encoding a certain witness w for a statement x of an NP language $\mathcal{L}$. Given such a commitment, the sender can transfer a message m, conditioned of the fact that $x \in \mathcal{L}$, i.e. the receiver will learn the message m (in a statistical sense) only if the committed w is a valid witness for x.

Equipped with a CDS protocol, the authors show how to lift a *semi-honest* circuit-private FHE scheme into a maliciously secure one: In addition to the public key and the ciphertext (pk, c),

the encrypter also includes a commitment to the random coins used to compute pk and c. This information is handed over to the evaluator, who homomorphically computes $\tilde{c} = \mathsf{Eval}(\mathsf{pk}, C, c)$. Note that at this point we have no guarantees about the circuit-privacy of the evaluated ciphertext $\tilde{c}$, since the encrypter might decide to commit to some garbage, instead of the correct random coins. For this reason, the evaluator does not hand over $\tilde{c}$ directly to the encrypter, instead it transfers $\tilde{c}$ using the CDS protocol, conditioned on the fact that the pair (pk, c) is well formed. This way, if (pk, c) is valid, then C is hidden by the semi-honest circuit privacy of the FHE, whereas if (pk, c) is malformed, no information at all is leaked by the (statistical) security of the CDS protocol.

A two-round CDS protocol can be constructed from any two-round oblivious transfer [BD18]. Note that, while the CDS protocol is non-compact, this does not affect the compactness of the resulting FHE scheme, since the condition checked by the CDS is anyway independent of the size of C. Thus, one interpretation of the [OPP14] approach is that it allows to combine a *non-compact* maliciously circuit-private FHE (the CDS protocol) with a compact *semi-honestly* circuit-private FHE, to obtain the best of both worlds. Unfortunately, the same strategy does not seem to apply to the quantum settings, because of the lack of a clear quantum counterpart of the CDS protocol. In contrast with the classical case, sacrificing compactness does not seem to ease the task of achieving malicious circuit privacy for quantum computation.

Background of QFHE. Our approach is inspired by recent advancements in classically verifiable quantum computation [Mah18b, Mah18a]. Our main idea is to constrain the (quantum) encrypter with a *classical leash* that prevents it from generating malformed keys and ciphertexts. In order to understand our transformation in more details, it is instructive to recall how QFHE schemes are constructed. At a very high level, QFHE schemes follow a paradigm introduced by Broadbent and Jeffery [BJ15], which exploits the properties of the quantum one-time pad (QOTP). A QOTP allows one to unconditionally hide a qubit $(\alpha_0 |0\rangle + \alpha_1 |1\rangle)$ by applying the Pauli transformation $X^x Z^z$, which corresponds to the following unitary:

$$(\alpha_0 |0\rangle + \alpha_1 |1\rangle) \rightarrow (\alpha_0 |x\rangle + (-1)^z \alpha_1 |x_i \oplus 1\rangle)$$

where x and z are two uniformly sampled classical bits. Computational security is achieved by also including a classical FHE encryption of the bits (x, z), which allows the owner of the secret key to invert the Pauli operator and recover the encrypted qubit. To homomorphically evaluate quantum gates, one can apply the gate to the encrypted quantum state, and update the classical encryption of the one-time pad appropriately. The original work of Broadbent and Jeffery [BJ15] supported a somewhat limited class of quantum circuits that could be homomorphically evaluated, however this limitation was later removed by Mahadev [Mah18a]. We highlight two properties that are going to be crucial for our approach:

(1) The scheme has completely classical keys and classical encryptions of classical messages.

(2) The scheme (and variant thereof [Bra18]) can be shown to be semi-honestly circuit-private.

Interestingly, the latter requirement is also necessary in order to guarantee the correct evaluation of quantum gates. The connection between homomorphic evaluation of quantum circuits and (semi-honest) circuit privacy is explored in more details in [Bra18].

A Classical Leash. We have now the tools needed to construct a maliciously circuit-private QFHE. Our main observation is that property (1) guarantees that the validity of QFHE ciphertext can be *classically* checked. This suggests the following template for bootstrapping a semi-honest to malicious circuit privacy in QFHE, using an additional maliciously circuit-private *classical* FHE: The evaluator will compute homomrphically the circuit of interest to obtain some state $(\tilde{c}, |\tilde{\phi}\rangle)$. Then it will transmit this information back to the encrypter, only if the initial keys and the ciphertexts are well-formed. The latter will turn out to be classically-checkable condition and therefore implementable via a quantum CDS for classical relations, which we will show how to construct.

More concretely, a ciphertext encrypting a quantum state $|\psi\rangle$ consists of:

- A QOTP of the state $|\phi\rangle = \mathsf{QOTP}((x,z), |\psi\rangle)$, where (x,z) is the corresponding one-time key.
- A QFHE encryption of the one-time key $c = \mathsf{QEnc}(\mathsf{pk}, (x,z))$, which is a classical string since (x,z) are classical bits.
- A classical FHE encryption $\widehat{c}$ of the random coins used to compute pk and c.

Note that an honestly computed $(|\phi\rangle, c)$ is a valid QFHE ciphertext and thus the evaluator can homomorphically evaluate a quantum circuit C to obtain an evaluated ciphertext $(\tilde{c}, |\tilde{\phi}\rangle)$. Recall however that QFHE only guarantees circuit privacy if both the keys and the ciphertexts are in the support of the corresponding algorithms (i.e. the encrypter is semi-honest). Following the template of [OPP14], we would then like to transmit $(\tilde{c}, |\tilde{\phi}\rangle)$ back to the encrypter only if the above condition is satisfied. Towards achieving this goal, observe that for verifying the validity of the QFHE ciphertext it suffices to check whether (pk, c) is well-formed, since the Pauli transformation is reversible. This means that all we need to do is to implement the CDS of a quantum state under a *classical* condition.

Quantum CDS for Classical Relations. What is left to be discussed is how to implement the above channel, i.e. a CDS for a quantum state under a classically-checkable condition. We achieve this by encrypting the evaluated state $(\tilde{c}, |\tilde{\phi}\rangle)$ under a QOTP (where encryption is done qubit-by-qubit) with a classical one-time key otk. Then we evaluate homomorphically the circuit Γ_{otk} over $\widehat{c}$ (the encryption of the random coins used to sample pk and c), where Γ_{otk} is defined as follows: On input some random coins, it checks whether pk and c are well-formed (by recomputing them) and if this is the case it returns otk, otherwise it returns 0. Here we crucially exploit the fact that the QOTP has a classical one-time key, which allows us to evaluate the above circuit under a classical FHE. The evaluator finally returns

$$\mathsf{QOTP}(\mathsf{otk}, (\tilde{c}, |\tilde{\phi}\rangle)) \text{ and } \mathsf{Eval}(\Gamma_{\mathsf{otk}}, \widehat{c}).$$

To see why the QFHE scheme satisfies (malicious) circuit privacy, we consider two cases: If (pk, c) are well-formed, then the encrypter can recover otk, but the semi-honest circuit privacy of QFHE guarantees that nothing is learned about C. On the other hand, if (pk, c) are malformed, then the malicious circuit privacy of the classical FHE scheme guarantees that otk is statistically hidden, and therefore the QOTP unconditionally hides the evaluated ciphertext $(\tilde{c}, |\tilde{\phi}\rangle)$. It follows that no information at all is leaked to the encrypter.

Multi-Key and Multi-Hop Evaluation. As described above, our techniques suffer from two major limitations:

- The evaluated ciphertexts are syntactically different from fresh encryptions (i.e. the scheme supports single-hop homomorphic evaluation).
- The homomorphic computation is restricted to ciphertexts encrypted under the same keys.

Fortunately, none of the above limitations is really inherent and our template can be naturally modified to support multi-hop and multi-key evaluation. We refer the curious reader to the technical sections for more details.

2 Preliminaries

We denote by λ the security parameter. A function $f : \mathbb{N} \to [0, 1]$ is negligible if for every constant $c \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$. We recall some standard notation for classical Turing machines and Boolean circuits:

- We say that a Turing machine (or algorithm) is PPT if it is probabilistic and runs in polynomial time in λ.
- We sometimes think about PPT Turing machines as polynomial-size uniform families of circuits. A polynomial-size circuit family C is a sequence of circuits $C = \{C_\lambda\}_{\lambda\in\mathbb{N}}$, such that each circuit C_λ is of polynomial size $\lambda^{O(1)}$ and has $\lambda^{O(1)}$ input and output bits. We say that the family is uniform if there exists a polynomial-time deterministic Turing machine M that on input 1^λ outputs C_λ.
- For a PPT Turing machine (algorithm) M, we denote by $M(x; r)$ the output of M on input x and random coins r. For such an algorithm, and any input x, we write $m \in M(x)$ to denote that m is in the support of $M(x; \cdot)$. Finally we write $y \leftarrow\$ M(x)$ to denote the computation of M on input x with some uniformly sampled random coins.

2.1 Quantum Adversaries

We recall some notation for quantum computation and we define the notions of computational and statistical indistinguishability for quantum adversaries. Various parts of what follows are taken almost in verbatim from [BS20].

- We say that a Turing machine (or algorithm) is QPT if it is quantum and runs in polynomial time.
- We sometimes think about QPT Turing machines as polynomial-size uniform families of quantum circuits (as these are equivalent models). We call a polynomial-size quantum circuit family $C = \{C_\lambda\}_{\lambda\in\mathbb{N}}$ uniform if there exists a polynomial-time deterministic Turing machine M that on input 1^λ outputs C_λ.
- Classical communication channels in the quantum setting are identical to classical communication channels in the classical setting, except that when a set of qubits is sent through a classical communication channel, then the qubits decohere and are automatically measured in the standard basis.

- A quantum interactive algorithm (in the two-party setting) has input divided into two registers and output divided into two registers. For the input qubits, one register is for an input message from the other party, and a second register is for a potential inner state the machine holds. For the output, one register is for the message to be sent to the other party, and another register is for a potential inner state for the machine to keep for itself.

Throughout this work, we model efficient adversaries as quantum circuits with non-uniform quantum advices. This is denoted by $\mathcal{A}^* = \{\mathcal{A}^*_\lambda, \rho_\lambda\}_{\lambda\in\mathbb{N}}$, where $\{\mathcal{A}^*_\lambda\}_{\lambda\in\mathbb{N}}$ is a polynomial-size non-uniform sequence of quantum circuits, and $\{\rho_\lambda\}_{\lambda\in\mathbb{N}}$ is some polynomial-size sequence of mixed quantum states. We now define the formal notion of computational indistinguishability in the quantum settings.

Definition 2.1 (Computational Indistinguishability)**.** *Two ensembles of quantum random variables $\mathcal{X} = \{X_\lambda\}_{\lambda\in\mathbb{N}}$ and $\mathcal{Y} = \{Y_\lambda\}_{\lambda\in\mathbb{N}}$ are said to be computationally indistinguishable (denoted by $\mathcal{X} \approx_c \mathcal{Y}$) if there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$ and all non-uniform QPT distinguishers with quantum advice $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda\in\mathbb{N}}$, it holds that*

$$|\Pr[\mathcal{A}(X;\rho) = 1] - \Pr[\mathcal{A}(Y;\rho) = 1]| \leq \mu(\lambda)$$

where $X \leftarrow\!\$\, X_\lambda$ and $Y \leftarrow\!\$\, Y_\lambda$.

The trace distance between two quantum distributions (X_λ, Y_λ), denoted by $\mathsf{TD}(X_\lambda, Y_\lambda)$, is a generalization of statistical distance to the quantum setting and represents the maximal distinguishing advantage between two quantum distributions by an unbounded quantum algorithm. We define below the notion of statistical indistinguishability.

Definition 2.2 (Statistical Indistinguishability)**.** *Two ensembles of quantum random variables $\mathcal{X} = \{X_\lambda\}_{\lambda\in\mathbb{N}}$ and $\mathcal{Y} = \{Y_\lambda\}_{\lambda\in\mathbb{N}}$ are said to be statistically indistinguishable (denoted by $\mathcal{X} \approx_s \mathcal{Y}$) if there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$, it holds that*

$$\mathsf{TD}(X_\lambda, Y_\lambda) \leq \mu(\lambda).$$

2.2 Learning with Errors

We recall the definition of the learning with errors (LWE) problem [Reg05].

Definition 2.3 (Learning with Errors)**.** *The LWE problem is parametrized by a modulus $q = q(\lambda)$, polynomials $n = n(\lambda)$ and $m = m(\lambda)$, and an error distribution χ. The LWE problem is hard if it holds that*

$$(\mathbf{A}, \mathbf{A}\cdot\mathbf{s} + \mathbf{e}) \approx_c (\mathbf{A}, \mathbf{u})$$

where $\mathbf{A} \leftarrow\!\$\, \mathbb{Z}_q^{m\times n}$, $\mathbf{s} \leftarrow\!\$\, \mathbb{Z}_q^n$, $\mathbf{u} \leftarrow\!\$\, \mathbb{Z}_q^m$, and $\mathbf{e} \leftarrow\!\$\, \chi^m$.

As shown in [Reg05, PRS17], for any sufficiently large modulus q the LWE problem where χ is a discrete Gaussian distribution with parameter $\sigma = \xi q \geq 2\sqrt{n}$ (i.e. the distribution over $\mathbb{Z}$ where the probability of x is proportional to $e^{-\pi(|x|/\sigma)^2}$), is at least as hard as approximating the shortest independent vector problem (SIVP) to within a factor of $\gamma = \tilde{O}(n/\xi)$ in *worst case* dimension n lattices.

2.3 Quantum One-Time Pad

We recall the quantum one-time pad (QOTP) construction [AMTDW00] for quantum states. We explicitly consider the scheme that allows one to encrypt an n-qubit quantum state with unconditional security.

Definition 2.4 (Quantum One-Time Pad)**.** *A quantum one-time pad* $(\mathsf{QOTP.Gen}, \mathsf{QOTP.Enc}, \mathsf{QOTP.Dec})$ *consists of the following efficient algorithms.*

- $\mathsf{QOTP.Gen}(1^n)$*: For all* $i = 1 \dots n$ *sample two classical bits* $(x_i, z_i) \leftarrow\!\!\$\ \{0,1\}^2$*. Return the one-time key* $\mathsf{otk} = (x_1, z_1, \dots, x_n, z_n)$*.*
- $\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle)$*: On input a one-time key* otk *and an* n*-qubit state* $|\psi\rangle$*, apply the Pauli transformation* $X^{x_i} Z^{z_i}$ *to the* i*-th qubit, for all* $i = 1 \dots n$*. Return the resulting state* $|\phi\rangle$*.*
- $\mathsf{QOTP.Dec}(\mathsf{otk}, |\phi\rangle)$*: On input a one-time key* otk *and an* n*-qubit state* $|\phi\rangle$*, apply the reverse Pauli transformation* $Z^{z_i} X^{x_i}$ *qubit-by-qubit to recover the original state.*

More explicitly, the (single qubit) Pauli transformation $X^{x_i} Z^{z_i}$ is the following unitary:

$$(\alpha_0 |0\rangle + \alpha_1 |1\rangle) \to (\alpha_0 |x_i\rangle + (-1)^{z_i} \alpha_1 |x_i \oplus 1\rangle).$$

As shown in [AMTDW00], the above scheme can be used to transform *any* n-qubit quantum state into a totally mixed state (no matter if some of its initial qubits are in an entangled state).

3 Homomorphic Encryption

In the following we define the main object of interest of our work, namely homomorphic encryption that allows one to evaluate classical and/or quantum circuits over encrypted data.

3.1 Classical Homomorphic Encryption

We recall the notion of classical homomorphic encryption [Gen09].

Definition 3.1 (Homomorphic Encryption)**.** *A homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ *consists of the following efficient algorithms.*

- $\mathsf{FHE.Gen}(1^\lambda)$*: On input the security parameter, the key generation algorithm returns secret/public key pair* $(\mathsf{sk}, \mathsf{pk})$*.*
- $\mathsf{FHE.Enc}(\mathsf{pk}, m)$*: On input the public key* pk *and a message* m*, the encryption algorithm returns a ciphertext* c*.*
- $\mathsf{FHE.Eval}(\mathsf{pk}, C, c)$*: On input the public key* pk*, a (classical) circuit* C*, and a ciphertext* c*, the evaluation algorithm returns an evaluated ciphertext* $\tilde{c}$*.*
- $\mathsf{FHE.Dec}(\mathsf{sk}, c)$*: On input the secret key* sk *and a ciphertext* c*, the decryption algorithm returns a message* m*.*

We say that a scheme is fully homomorphic (FHE) if the evaluation algorithm supports all polynomial-size classical circuits (without posing an a-priori bound on the size of $|C|$). If the size of C needs to be fixed at the time of key generation, then we say that the scheme is levelled homomorphic. It is well-known that levelled FHE schemes can be based on the hardness of the (plain) LWE problem [BV11, BV14]. We recall the notion of single-hop evaluation correctness in the following and we refer the reader to [GHV10] for a more general definition of multi-hop evaluation correctness.

Definition 3.2 (Single-Hop Evaluation Correctness)**.** *A homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ *is correct if for all* $\lambda \in \mathbb{N}$*, all* $(\mathsf{sk}, \mathsf{pk}) \in \mathsf{FHE.Gen}(1^\lambda)$*, all messages* m*, and all polynomial-size circuits* C*, it holds that*

$$\Pr\left[\mathsf{FHE.Dec}(\mathsf{sk}, \mathsf{FHE.Eval}(\mathsf{pk}, C, \mathsf{FHE.Enc}(\mathsf{pk}, m))) = C(m)\right] = 1$$

We recall the notion of semantic security for public-key encryption.

Definition 3.3 (Semantic Security)**.** *A homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ *is semantically secure if for all* $\lambda \in \mathbb{N}$ *and all pairs of messages* (m_0, m_1)*, it holds that*

$$\mathsf{FHE.Enc}(\mathsf{pk}, m_0) \approx_c \mathsf{FHE.Enc}(\mathsf{pk}, m_1)$$

where $(\mathsf{sk}, \mathsf{pk}) \leftarrow\!\!\$\, \mathsf{FHE.Gen}(1^\lambda)$.

Finally we define the notion of (malicious) statistical circuit privacy for FHE [OPP14].

Definition 3.4 (Statistical Circuit Privacy)**.** *A homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ *is (malicious) statistically circuit private if there exists a pair of unbounded algorithms* $\mathsf{FHE.Ext}$ *and* $\mathsf{FHE.Sim}$ *such that for all* $\lambda \in \mathbb{N}$*, all public keys* pk^**, all ciphertexts* c^**, and all circuits* C*, it holds that*

$$\mathsf{FHE.Eval}(\mathsf{pk}^*, C, c^*) \approx_s \mathsf{FHE.Sim}(1^\lambda, \mathsf{pk}^*, c^*, C(x^*))$$

where $x^* = \mathsf{FHE.Ext}(1^\lambda, \mathsf{pk}^*, c^*)$.

It is shown in [OPP14] that any FHE scheme can be converted into one with malicious circuit privacy generically, by additionally assuming a two-round statistically sender-private oblivious transfer. The latter can in turn be instantiated from LWE [BD18, DGI+19, BDGM19]. Taken together, these results give us the following implication.

Lemma 3.5 ([OPP14, BD18])**.** *Assuming the hardness of the (circular) LWE problem, there exists an FHE scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ *with (malicious) statistical circuit privacy.*

3.2 Quantum Homomorphic Encryption

We extend the notion of classical FHE to the evaluation of quantum circuits [BJ15]. In this work we consider only quantum FHE (QFHE) schemes with completely classical key generation algorithms. We extend the syntax of classical FHE below.

Definition 3.6 (Quantum Homomorphic Encryption)**.** *A quantum homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.QEnc}, \mathsf{FHE.QEval}, \mathsf{FHE.QDec})$ *consists of the following efficient algorithms.*

- $\mathsf{FHE.Gen}(1^\lambda)$*: Same as in Definition 3.1.*
- $\mathsf{FHE.QEnc}(\mathsf{pk}, |\psi\rangle)$*: On input the public key* pk *and a quantum state* $|\psi\rangle$*, the encryption algorithm returns a quantum ciphertext* $|\phi\rangle$*.*
- $\mathsf{FHE.QEval}(\mathsf{pk}, C, |\phi\rangle)$*: On input the public key* pk*, a quantum circuit* C*, and a quantum ciphertext* $|\phi\rangle$*, the evaluation algorithm returns an evaluated quantum ciphertext* $|\tilde{\phi}\rangle$*.*
- $\mathsf{FHE.QDec}(\mathsf{sk}, |\phi\rangle)$*: On input the secret key* sk *and a quantum ciphertext* $|\phi\rangle$*, the decryption algorithm returns a quantum state* $|\psi\rangle$*.*

Analogously to the classical case, we say that the scheme is fully homomorphic if the evaluation algorithm supports all polynomial-size quantum circuits. Next we define the notion of single-hop evaluation correctness for QFHE.

Definition 3.7 (Single-Hop Evaluation Correctness)**.** *A quantum homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.QEnc}, \mathsf{FHE.QEval}, \mathsf{FHE.QDec})$ *is correct if for all* $\lambda \in \mathbb{N}$*, all* $(\mathsf{sk}, \mathsf{pk}) \in \mathsf{FHE.Gen}(1^\lambda)$*, all quantum states* $|\psi\rangle$*, and all polynomial-size quantum circuits* C*, it holds that*

$$\mathsf{FHE.QDec}(\mathsf{sk}, \mathsf{FHE.QEval}(\mathsf{pk}, C, \mathsf{FHE.QEnc}(\mathsf{pk}, |\psi\rangle))) \approx_s C(|\psi\rangle).$$

The notion of semantic security is defined analogously to the classical case, and we refer the reader to [BJ15] for a formal definition. We define the main notion of interest of this work, namely, malicious statistical circuit privacy for QFHE.

Definition 3.8 (Statistical Circuit Privacy)**.** *A quantum homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.QEnc}, \mathsf{FHE.QEval}, \mathsf{FHE.QDec})$ *is (malicious) statistically circuit private if there exists a pair of unbounded algorithms* $\mathsf{FHE.Ext}$ *and* $\mathsf{FHE.Sim}$ *such that for all* $\lambda \in \mathbb{N}$*, all public keys* pk^**, all quantum ciphertexts* $|\phi^*\rangle$*, and all quantum circuits* C*, it holds that*

$$\mathsf{FHE.QEval}(\mathsf{pk}^*, C, |\phi^*\rangle) \approx_s \mathsf{FHE.Sim}(1^\lambda, \mathsf{pk}^*, \alpha, C(|\psi^*\rangle))$$

where $(|\psi^*\rangle, \alpha) = \mathsf{FHE.Ext}(1^\lambda, \mathsf{pk}^*, |\phi^*\rangle)$*.*

In addition, we say that a scheme satisfies the weaker *semi-honest* circuit privacy if the above indistinguishability is required to hold only for well-formed (i.e. in the support of the respective algorithms) public keys pk^* and ciphertexts $|\phi^*\rangle$. For completeness, we present the definition below.

Definition 3.9 (Semi-Honest Statistical Circuit Privacy)**.** *A quantum homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.QEnc}, \mathsf{FHE.QEval}, \mathsf{FHE.QDec})$ *is (semi-honest) statistically circuit private if there exists an unbounded algorithm* $\mathsf{FHE.Sim}$ *such that for all* $\lambda \in \mathbb{N}$*, all public keys* $\mathsf{pk} \in \mathsf{FHE.Gen}(1^\lambda)$*, all quantum states* $|\psi\rangle$*, all quantum ciphertexts* $|\phi\rangle \in \mathsf{FHE.QEnc}(\mathsf{pk}, |\psi\rangle)$*, and all quantum circuits* C*, it holds that*

$$\mathsf{FHE.QEval}(\mathsf{pk}, C, |\phi\rangle) \approx_s \mathsf{FHE.Sim}(1^\lambda, \mathsf{pk}, C(|\psi\rangle)).$$

The works of Mahadev [Mah18a] and Brakerski [Bra18] show that QFHE with classical keys can be constructed from the quantum hardness of the LWE problem. For the evaluation of unbounded circuits, an additional circularity assumption is required due to an application of the bootstrapping theorem [Gen09]. Both schemes follow the *hybrid encryption* approach where each

ciphertext consists of (i) a QOTP of a given quantum state and (ii) a (classical) FHE encryption of the corresponding one-time key.

It follows almost immediately from the analysis of these works that the schemes satisfy semi-honest circuit privacy: It is shown (see e.g. Theorem 4.1 in [Bra18]) that the classical part of each ciphertext after the evaluation procedure is statistically close to a fresh FHE ciphertext. Somewhat curiously, this is actually needed to ensure correctness for the evaluation of a quantum circuit. However, nothing prevents the one-time key of the QOTP to carry some information about the evaluated circuit. We can easily fix this by re-randomizing the QOTP encryption of a given quantum state by computing

$$X^v Z^w \mathsf{QOTP.Enc}((x,z),|\psi\rangle) = X^v Z^w X^x Z^z |\psi\rangle = X^{v\oplus x} Z^{w\oplus z} |\psi\rangle$$

where $(v,w) \leftarrow\!\!\$\, \{0,1\}^2$ and the equality above holds up to a global phase. Correctness is then preserved by homomorphically propagating the changes to the one-time key encrypted under the classical FHE. This allows us to state the following fact.

Lemma 3.10 ([Mah18a, Bra18])**.** *Assuming the quantum hardness of the (circular) LWE problem, there exists a QFHE scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.QEnc}, \mathsf{FHE.QEval}, \mathsf{FHE.QDec})$ *with semi-honest statistical circuit privacy.*

4 Malicious Circuit Privacy for Quantum Computation

In the following we describe the main result of this work, namely the constrution of a (malicious) statistically circuit private QFHE scheme.

4.1 Our Bootstrapping Theorem

We describe our scheme in the form of a generic transformation, starting from the following building blocks:

- A maliciously circuit private classical FHE scheme $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$.
- A QFHE scheme $(\mathsf{QFHE.Gen}, \mathsf{QFHE.QEnc}, \mathsf{QFHE.QEval}, \mathsf{QFHE.QDec})$ that satisfies the following properties:
  (1) Has a classical key generation $\mathsf{QFHE.Gen}$ algorithm and classical keys $(\mathsf{qsk}, \mathsf{qpk})$.
  (2) The encryption algorithm $\mathsf{QFHE.QEnc}$ for a classical message is entirely classical.
  (3) Is semi-honest statistically circuit-private.

Our transformation is presented formally in Figure 1. If the above schemes are levelled homomorphic, then so is the resulting QFHE scheme is also levelled homomorphic. In contrast, if the underlying building blocks are fully homomorphic, then the resulting QFHE can evaluate (unbounded) polynomial-size quantum circuits.

Maliciously Circuit Private QFHE

- **Key Generation:** On input the security parameter 1^λ, the (classical) key generation algorithm samples two key pairs $(\mathsf{sk}, \mathsf{pk}) \leftarrow\!\$\, \mathsf{FHE.Gen}(1^\lambda)$ and $(\mathsf{qsk}, \mathsf{qpk}) = \mathsf{QFHE.Gen}(1^\lambda; r)$, where $r \leftarrow\!\$\, \{0,1\}^\lambda$. Then it computes an encryption $c_r \leftarrow\!\$\, \mathsf{FHE.Enc}(\mathsf{pk}, r)$ of the (classical) random coins used in the QFHE key generation. The secret key of the scheme is set to $(\mathsf{sk}, \mathsf{qsk})$ and he public key consists of $(\mathsf{pk}, \mathsf{qpk}, c_r)$.

- **Encryption:** On input the public key $(\mathsf{pk}, \mathsf{qpk}, c_{\mathsf{Gen}})$ and an n-qubit state $|\psi\rangle$, the encryption algorithm samples a QOPT key $\mathsf{otk} \leftarrow\!\$\, \mathsf{QOTP.Gen}(1^n)$ and some classical random coins $s \leftarrow\!\$\, \{0,1\}^\lambda$. Set the ciphertext as

$$(|\phi\rangle, c, c_s) = (\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QFHE.QEnc}(\mathsf{qpk}, \mathsf{otk}; s), \mathsf{FHE.Enc}(\mathsf{pk}, (\mathsf{otk}, s))).$$

- **Evaluation:** On input the public key $(\mathsf{pk}, \mathsf{qpk}, c_{\mathsf{Gen}})$, a quantum circuit C, and a ciphertext $(|\phi\rangle, c, c_s)$, define the quantum circuit $\Gamma_{|\phi\rangle}$ as

$$\Gamma_{|\phi\rangle}(\mathsf{otk}) : \text{Return } C(\mathsf{QOTP.Dec}(\mathsf{otk}, |\phi\rangle)).$$

  Then evaluate homomorphically $|\tilde{\phi}\rangle = \mathsf{QFHE.QEval}(\mathsf{qpk}, \Gamma_{|\phi\rangle}, c)$, which results in some $\tilde{n}$ qubit state $|\tilde{\phi}\rangle$. Sample a fresh quantum one-time key $\tilde{\mathsf{otk}} \leftarrow\!\$\, \mathsf{QOTP.Gen}(1^{\tilde{n}})$ and let $|\xi\rangle = \mathsf{QOTP.Enc}(\tilde{\mathsf{otk}}, |\tilde{\phi}\rangle)$. Let $\Theta_{(\mathsf{qpk}, c, \tilde{\mathsf{otk}})}$ be a (classical) circuit defined as

$$\Theta_{(\mathsf{qpk}, c, \tilde{\mathsf{otk}})}(r, \mathsf{otk}, s): \begin{array}{l} \text{If } (\cdot, \mathsf{qpk}) = \mathsf{QFHE.Gen}(1^\lambda; r) \text{ and } c = \mathsf{QFHE.Enc}(\mathsf{qpk}, \mathsf{otk}; s) \\ \quad \text{then return } \tilde{\mathsf{otk}}. \\ \text{Else return } 0. \end{array}$$

  Return the evaluated ciphertext $(|\xi\rangle, \mathsf{FHE.Eval}(\mathsf{pk}, \Theta_{(\mathsf{qpk}, c, \tilde{\mathsf{otk}})}, (c_r, c_s)))$.

- **Decryption:** On input a secret key $(\mathsf{sk}, \mathsf{qsk})$ and (without loss of generality) an evaluated ciphertext $(|\xi\rangle, \tilde{c})$, the decryption algorithm returns

$$\mathsf{QFHE.Dec}(\mathsf{qsk}, \mathsf{QOTP.Dec}(\mathsf{FHE.Dec}(\mathsf{sk}, \tilde{c}), |\xi\rangle)).$$

Figure 1: Description of a (malicious) statistically circuit private QFHE scheme.

Analysis. To see why the scheme satisfies (single-hop) evaluation correctness, recall that

$$\begin{aligned} \tilde{c} &= \mathsf{FHE.Eval}(\mathsf{pk}, \Theta_{(\mathsf{qpk}, c, \tilde{\mathsf{otk}})}, (c_r, c_s)) \\ &= \mathsf{FHE.Enc}(\mathsf{pk}, \tilde{\mathsf{otk}}) \end{aligned}$$

since qpk is in the support of $\mathsf{QFHE.Gen}$ and c is computed as $\mathsf{QFHE.Enc}(\mathsf{qpk}, \mathsf{otk}; s)$. Note that, by property (1) and (2), the key generation and the encryption of classical messages of the QFHE scheme are completely classical. Therefore, the circuit $\Theta_{(\mathsf{qpk}, c, \tilde{\mathsf{otk}})}$ is a well-defined classical circuit and the above equality follows from the evaluation correctness of the FHE scheme. Thus it follows

that

$$
\begin{aligned}
&\mathsf{QFHE.Dec}(\mathsf{qsk}, \mathsf{QOTP.Dec}(\mathsf{FHE.Dec}(\mathsf{sk}, \tilde{c}), |\xi\rangle)) \\
&= \mathsf{QFHE.Dec}(\mathsf{qsk}, \mathsf{QOTP.Dec}(\tilde{\mathsf{otk}}, |\xi\rangle)) \\
&= \mathsf{QFHE.Dec}(\mathsf{qsk}, \mathsf{QOTP.Dec}(\tilde{\mathsf{otk}}, \mathsf{QOTP.Enc}(\tilde{\mathsf{otk}}, |\tilde{\phi}\rangle))) \\
&= \mathsf{QFHE.Dec}(\mathsf{qsk}, |\tilde{\phi}\rangle) \\
&= \mathsf{QFHE.Dec}(\mathsf{qsk}, \mathsf{QFHE.QEval}(\mathsf{pk}, \Gamma_{|\phi\rangle}, c)) \\
&= \mathsf{QFHE.Dec}(\mathsf{qsk}, \mathsf{QFHE.Enc}(\mathsf{qpk}, C(\mathsf{QOTP.Dec}(\mathsf{otk}, |\phi\rangle)))) \\
&= \mathsf{QFHE.Dec}(\mathsf{qsk}, \mathsf{QFHE.Enc}(\mathsf{qpk}, C(\mathsf{QOTP.Dec}(\mathsf{otk}, \mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle))))) \\
&= \mathsf{QFHE.Dec}(\mathsf{qsk}, \mathsf{QFHE.Enc}(\mathsf{qpk}, C(|\psi\rangle))) \\
&= C(|\psi\rangle)
\end{aligned}
$$

by the (single-hop) evaluation correctness of the QFHE scheme. Next we show that the scheme satisfies semantic security.

Lemma 4.1 (Semantic Security)**.** *Assuming that the FHE and the QFHE schemes are semantically secure, the scheme in Figure 1 satisfies semantic security.*

Proof. Let $(|\phi\rangle, c, c_s)$ be an honestly computed ciphertext

$$(\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QFHE.QEnc}(\mathsf{qpk}, \mathsf{otk}; s), \mathsf{FHE.Enc}(\mathsf{pk}, (\mathsf{otk}, s))).$$

We define a series of hybrid distributions and we argue that they are computationally indistinguishable from the original ciphertext. We begin by substituting the FHE ciphertext with an encryption of 0 (padded to the appropriate length), thus obtaining

$$(\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QFHE.QEnc}(\mathsf{qpk}, \mathsf{otk}; s), \mathsf{FHE.Enc}(\mathsf{pk}, 0)).$$

This distribution is computationally indistinguishable from the above one by an invocation of the semantic security of the FHE scheme. Next, we switch the second ciphertext to a uniformly sampled encryption of 0 (again padded to the appropriate length). This gives us the following distribution

$$(\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QFHE.QEnc}(\mathsf{qpk}, 0), \mathsf{FHE.Enc}(\mathsf{pk}, 0)).$$

Indistinguishability follows from the semantic security of QFHE for classical messages. At this point, the state $|\psi\rangle$ is information theoretically hidden by the one-time key otk and thus it is identical to a completely mixed state from the eyes of the adversary. This concludes our proof. □

Finally, we show that the scheme satisfies statistical circuit privacy in the malicious settings.

Lemma 4.2 (Circuit Privacy)**.** *Assuming that FHE is malicious statistically circuit private and that QFHE is semi-honest statistically circuit private, the scheme in Figure 1 satisfies malicious statistical circuit privacy.*

Proof. First we define the algorithms for the extractor Ext and the simulator Sim and then we argue that the output of the simulator is statistically indistinguishable from the output of the honest evaluation algorithm. In the following we preset the extraction algorithm.

- Ext: On input the public key $(\mathsf{pk}, \mathsf{qpk}, c_r)$ and a ciphertext $(|\phi\rangle, c, c_s)$, the extractor runs the extractor of the FHE scheme on $(\mathsf{otk}^*, s^*) = \mathsf{FHE.Ext}(1^\lambda, \mathsf{pk}, c_s)$ and on $r^* = \mathsf{FHE.Ext}(1^\lambda, \mathsf{pk}, c_r)$. Then it checks whether

  (a) $(\cdot, \mathsf{qpk}) = \mathsf{QFHE.Gen}(1^\lambda; r^*)$ and

  (b) $c = \mathsf{QFHE.Enc}(\mathsf{qpk}, \mathsf{otk}^*; s^*)$

  and returns $|\psi^*\rangle = \mathsf{QOTP.Dec}(\mathsf{otk}^*, |\phi\rangle)$ and $\alpha = 1$ if both equalities are satisfied. Otherwise it returns a totally mixed state $|\psi^*\rangle = 1/2^n \cdot \mathcal{I}_n$ and the auxiliary bit $\alpha = 0$.

Next we describe the simulator.

- Sim: On input the public key $(\mathsf{pk}, \mathsf{qpk}, c_r)$, an auxiliary information bit α, and a quantum state $|\theta\rangle$, the simulator proceeds as follows. First it computes $\mathsf{QFHE.Sim}(1^\lambda, \mathsf{qpk}, |\theta\rangle)$ and sets $|\xi\rangle$ to be a QOTP encryption of the resulting state with some uniformly sampled one-time key $\tilde{\mathsf{otk}}$. If $\alpha = 0$, then it sets $\tilde{c} \leftarrow\$ \mathsf{FHE.Sim}(1^\lambda, \mathsf{pk}, (c_s, c_r), 0)$, otherwise if $\alpha = 1$ it sets $\tilde{c} \leftarrow\$ \mathsf{FHE.Sim}(1^\lambda, \mathsf{pk}, (c_s, c_r), \tilde{\mathsf{otk}})$. The simulator returns $(|\xi\rangle, \tilde{c})$.

Let $(|\xi_0\rangle, \tilde{c}_0)$ be the simulated ciphertext as computed above. We define $(|\xi_1\rangle, \tilde{c}_1)$ identically except that if $\alpha = 0$ we compute $|\xi_1\rangle$ as

$$|\xi_1\rangle = \mathsf{QOTP.Enc}(\tilde{\mathsf{otk}}, \mathsf{QFHE.QEval}(\mathsf{qpk}, \Gamma_{|\phi\rangle}, c)).$$

Recall that if $\alpha = 0$, then $\tilde{c}_1$ is defined to be a simulated encryption of 0 and thus the quantum state $|\xi_1\rangle$ is totally mixed from the point of view of the adversary, by the unconditional security of the QOTP. Thus we have that

$$(|\xi_0\rangle, \tilde{c}_0) \equiv (|\xi_1\rangle, \tilde{c}_1).$$

Next we define $(|\xi_2\rangle, \tilde{c}_2)$ analogously, except that if $\alpha = 1$, then we compute the state $|\xi_2\rangle$ as

$$|\xi_2\rangle = \mathsf{QOTP.Enc}(\tilde{\mathsf{otk}}, \mathsf{QFHE.QEval}(\mathsf{qpk}, \Gamma_{|\phi\rangle}, c)).$$

Note that if $\alpha = 1$, then it holds that conditions (a) and (b) are satisfied, which in particular means that the public key of the QFHE scheme is in the support of the key generation algorithm and that the ciphertext c is in the support of the QFHE.QEnc algorithm (invoked on input some classical string otk^*). Thus, by the semi-honest circuit privacy of the QFHE scheme, we have that

$$(|\xi_1\rangle, \tilde{c}_1) \approx_s (|\xi_2\rangle, \tilde{c}_2).$$

Finally, we define $(|\xi_3\rangle, \tilde{c}_3)$ as before except that we compute $\tilde{c}_3$ as

$$\tilde{c}_3 = \mathsf{FHE.Eval}(\mathsf{pk}, \Theta_{(\mathsf{qpk}, c, \tilde{\mathsf{otk}})}, (c_r, c_s)).$$

Recall that the function $\Theta_{(\mathsf{qpk}, c, \tilde{\mathsf{otk}})}$ takes as input two random coins r and s and a one-time-key otk and returns $\tilde{\mathsf{otk}}$ if conditions (a) and (b) are satisfied and returns 0 otherwise. This is exactly the circuit computed by the simulator on input the extracted messages. Thus, by the malicious circuit privacy of the FHE scheme, it holds that

$$(|\xi_2\rangle, \tilde{c}_2) \approx_s (|\xi_3\rangle, \tilde{c}_3).$$

Observe that the state $(|\xi_3\rangle, \tilde{c}_3)$ is computed exactly as in the evaluation algorithm, whereas the state $(|\xi_0\rangle, \tilde{c}_0)$ is the output of the simulator. Combining the above implications we have that

$$(|\xi_0\rangle, \tilde{c}_0) \equiv (|\xi_1\rangle, \tilde{c}_1) \approx_s (|\xi_2\rangle, \tilde{c}_2) \approx_s (|\xi_3\rangle, \tilde{c}_3)$$

which concludes our proof. □

Combining Lemma 4.1 and Lemma 4.2 we obtain the following main theorem.

Theorem 4.3 (Malicious Circuit Privacy)**.** *Assuming the quantum hardness of the (circular) LWE problem, there exists a QFHE scheme with malicious statistical circuit privacy.*

4.2 Multi-Key and Multi-Hop Evaluation

We briefly sketch how our template can be applied to QFHE schemes with multi-key and multi-hop evaluation.

Multi-Key QFHE. Recall that a multi-key QFHE allows one to evaluate quantum circuits over ciphertexts encrypted under independent keys. We can naturally combine a maliciously circuit private classical multi-key FHE [CO17] together with a semi-honest circuit private quantum multi-key FHE [ABG+20] to construct a maliciously circuit private QFHE with multi-key evaluation. The only difference is that we want to make sure that *all* of the N public keys and ciphertexts (for some polynomial number of parties N) are well-formed. This can be easily accomplished by computing a QOTP of the evaluated (multi-key) ciphertext $|\tilde{\phi}\rangle$ using a fresh one-time key $\tilde{\mathsf{otk}}$, then computing an N-out-of-N secret sharing

$$\bigoplus_{i=1}^{N} \rho_i = \tilde{\mathsf{otk}}$$

and finally encrypting the i-th share ρ_i under the i-th (classical) public key conditioned on the fact that the key and ciphertext are well formed. This conditional encryption can be done (as in the single-key settings) by homomorphically recomputing the desired relation and returning ρ_i only if the check succeeds. The malicious circuit privacy of the classical (multi-key) FHE scheme guarantees that no information about ρ_i is revealed if the condition is not satisfied.

Multi-Hop Evaluation. We then discuss how to lift our scheme to support multi-hop evaluation of ciphertexts (for simplicity we only consider the single-key settings). Recall that a fresh (i.e. non-evaluated) ciphertext consists of the following components:

- A QOTP $\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle)$ of the state $|\psi\rangle$.
- A QFHE encryption of the one-time key $\mathsf{QFHE.QEnc}(\mathsf{qpk}, \mathsf{otk}; s)$.
- A classical FHE encryption $\mathsf{FHE.Enc}(\mathsf{pk}, (\mathsf{otk}, s))$ of the plaintext and random coins of the above QFHE ciphertext.

On the other hand, an evaluated ciphertext (in the honest case) consists of:

- A QOTP $\mathsf{QOTP.Enc}(\tilde{\mathsf{otk}}, |\tilde{\phi}\rangle)$ of the state $|\tilde{\phi}\rangle$, which is the output of the $\mathsf{QFHE.QEval}$ algorithm.

- A classical FHE encryption $\mathsf{FHE.Enc}(\mathsf{pk}, \tilde{\mathsf{otk}})$ of the one-time key $\tilde{\mathsf{otk}}$.

To transform an evaluated ciphertext into a non-evaluated one, we additionally assume that (i) the classical component of the QFHE scheme is randomness recoverable (which is the case for example for [Mah18a]) and that (ii) FHE and QFHE are secure in the presence of a two-key cycle, i.e. we additionally publish $\mathsf{QFHE.QEnc}(\mathsf{qpk}, \mathsf{sk})$. Using the latter, we can homomorphically recover $\tilde{\mathsf{otk}}$ under the hood of QFHE, peel off the QOTP layer, and decrypt $|\tilde{\phi}\rangle$. By the evaluation correctness of the QFHE, the resulting QFHE ciphertext consists of a QOTP of the new state $|\widehat{\psi}\rangle$ and a QFHE encryption $\widehat{c}$ of the corresponding one-time key $\widehat{\mathsf{otk}}$. Thus, all we are missing is a classical FHE encryption of $\widehat{\mathsf{otk}}$ and the random coins of its encryption $\widehat{c}$. Using the fact that QFHE is randomness recoverable, we can homomorphically compute the desired ciphertext using $\mathsf{FHE.Enc}(\mathsf{pk}, r)$, where r are the random coins used in the key generation algorithm of the QFHE scheme. It is not hard to show that the resulting ciphertext is in the correct form, and in particular is amenable to execute the evaluation algorithm as described in Figure 1.

Acknowledgments

The author wishes to thank Nico Döttling for insightful discussions on quantum fully homomorphic encryption.

References

[ABG+20] Amit Agarwal, James Bartusek, Vipul Goyal, Dakshita Khurana, and Giulio Malavolta. Post-quantum multi-party computation. Cryptology ePrint Archive, Report 2020/1395, 2020.

[AMTDW00] Andris Ambainis, Michele Mosca, Alain Tapp, and Ronald De Wolf. Private quantum channels. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 547–553. IEEE, 2000.

[BD18] Zvika Brakerski and Nico Döttling. Two-message statistically sender-private OT from LWE. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 370–390. Springer, Heidelberg, November 2018.

[BDGM19] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Leveraging linear decryption: Rate-1 fully-homomorphic encryption and time-lock puzzles. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 407–437. Springer, Heidelberg, December 2019.

[BJ15] Anne Broadbent and Stacey Jeffery. Quantum homomorphic encryption for circuits of low T-gate complexity. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 609–629. Springer, Heidelberg, August 2015.

[Bra18] Zvika Brakerski. Quantum FHE (almost) as secure as classical. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 67–95. Springer, Heidelberg, August 2018.

[BS20] Nir Bitansky and Omri Shmueli. Post-quantum zero knowledge in constant rounds. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 269–279, 2020.

[BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 97–106. IEEE Computer Society Press, October 2011.

[BV14] Zvika Brakerski and Vinod Vaikuntanathan. Lattice-based FHE as secure as PKE. In Moni Naor, editor, *ITCS 2014*, pages 1–12. ACM, January 2014.

[CO17] Wutichai Chongchitmate and Rafail Ostrovsky. Circuit-private multi-key FHE. In Serge Fehr, editor, *PKC 2017, Part II*, volume 10175 of *LNCS*, pages 241–270. Springer, Heidelberg, March 2017.

[DGI+19] Nico Döttling, Sanjam Garg, Yuval Ishai, Giulio Malavolta, Tamer Mour, and Rafail Ostrovsky. Trapdoor hash functions and their applications. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 3–32. Springer, Heidelberg, August 2019.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.

[GHV10] Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. i-Hop homomorphic encryption and rerandomizable Yao circuits. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 155–172. Springer, Heidelberg, August 2010.

[Mah18a] Urmila Mahadev. Classical homomorphic encryption for quantum circuits. In Mikkel Thorup, editor, *59th FOCS*, pages 332–338. IEEE Computer Society Press, October 2018.

[Mah18b] Urmila Mahadev. Classical verification of quantum computations. In Mikkel Thorup, editor, *59th FOCS*, pages 259–267. IEEE Computer Society Press, October 2018.

[OPP14] Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private FHE. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 536–553. Springer, Heidelberg, August 2014.

[PRS17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *49th ACM STOC*, pages 461–473. ACM Press, June 2017.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

[vGHV10] Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 24–43. Springer, Heidelberg, May / June 2010.