Rate-1 Quantum Fully Homomorphic Encryption

Orestis Chardouvelis*[1], Nico Döttling†[2], and Giulio Malavolta[3]

[1]National Technical University of Athens
[2]CISPA Helmholtz Center for Information Security
[3]Max Planck Institute for Security and Privacy

Abstract

Secure function evaluation (SFE) allows Alice to publish an encrypted version of her input m such that Bob (holding a circuit C) can send a single message that reveals $C(m)$ to Alice, and nothing more. Security is required to hold against malicious parties, that may behave arbitrarily. In this work we study the notion of SFE in the quantum setting, where Alice outputs an encrypted quantum state $|\psi\rangle$ and learns $C(|\psi\rangle)$ after receiving Bob's message.

We show that, assuming the quantum hardness of the learning with errors problem (LWE), there exists an SFE protocol for quantum computation with communication complexity

$$(|\,|\psi\rangle\,| + |C(|\psi\rangle)|) \cdot (1 + o(1))$$

which is nearly optimal. This result is obtained by two main technical steps, which might be of independent interest. Specifically, we show (i) a construction of a rate-1 quantum fully-homomorphic encryption and (ii) a generic transformation to achieve malicious circuit privacy in the quantum setting.

1 Introduction

Secure function evaluation (SFE) [Yao86] allows Alice to encrypt some message m such that later Bob (holding a circuit C) can compute

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

which allows Alice to recover $C(m)$ and nothing beyond that. Since standard simulation security is impossible in two rounds [KO04] (without assuming a trusted setup), the canonical notion of security for SFE [NP99, NP01, AIR01] requires the scheme to satisfy the following properties.

- Semantic Security: Alice's input m must be hidden in an indistinguishability sense.
- Circuit Privacy: The message of Bob must be statistically independent of C, conditioned on the output $C(m)$, for any choice of Alice's first message.

[*] Work done while the author was an intern at the Max Planck Institute for Security and Privacy.

[†] This work is partially funded by the Helmholtz Association within the project "Trustworthy Federated Data Analytics" (TFDA) (funding number ZT-I-OO1 4).

Among other applications, SFE realizes the vision of computation over encrypted data, where a computationally constrained client uploads some data to a powerful server that can perform expensive computation, while preserving data privacy. In this setting, it is important to ensure that the communication overhead introduced by the SFE protocol does not nullify the efficiency gains of outsourcing the computation to a server. Minimizing the communication complexity of this class of protocols has lead to the development of fully-homomorphic encryption (FHE) [Gen09, BV11], one of the cornerstones of modern cryptography. More recently, it was shown [BDGM19] that there exist SFE protocols where the communication complexity approaches that of the *insecure protocol* (where Alice sends her input m in plain), assuming the hardness of the learning with errors (LWE) problem.

Quantum SFE. In contrast to the classical setting, much less is known about SFE for quantum circuits. In its most general form, quantum SFE allows anyone to evaluate the transformation

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

where $|\psi\rangle$ is some arbitrary quantum state and C is some unitary matrix (ignoring ancillas). Despite the fact that this problem has received far less attention, we believe that this question is even more pressing than the classical case, due to the large gap between quantum capabilities of regular users and servers sitting on the cloud. Even in a future where regular users will be equipped with quantum-capable computers, it is likely that intensive quantum computations will be exclusive to large computer clusters.

For the semi-honest case, solutions exist based on quantum fully homomorphic encryption (QFHE) [BJ15], even assuming a completely classical client (Alice) [Mah18a]. However, to the best of our knowledge, the question of maliciously-secure SFE with compact (i.e. independent of the size of the circuit) communication complexity has not been considered in the literature. Motivated by the unsatisfactory state of affairs, we ask the following question:

Can we construct quantum SFE with minimal communication complexity?

1.1 Our Results

In this work we initiate the study of the communication complexity of SFE for quantum circuits (quantum SFE) in the malicious setting. Our main result is a protocol to compute any quantum circuit with communication complexity

$$(|\,|\psi\rangle\,| + |C(|\psi\rangle)|) \cdot (1 + o(1))$$

to compute some quantum circuit C over some state $|\psi\rangle$. This approaches the communication complexity of the insecure protocol, where Alice sends the state $|\psi\rangle$ in plain, and it is (asymptotically) optimal. Our protocol assumes the quantum hardness of the LWE problem (with polynomial modulo-to-noise ratio) in addition to a circular security assumption to apply the bootstrapping theorem [Gen09]. Our main result stems from a combination of two main technical steps that we outline below.

Rate-1 QFHE. As we discussed before, for the semi-honest setting, QFHE schemes [Mah18a, Bra18] constitute a valid solution to the SFE problem. However, all known QFHE schemes blow up the ciphertext by a polynomial factor $\mathsf{poly}(\lambda)$ for evaluated ciphertexts, i.e. they have low (inverse polynomial) rate. This means that the communication complexity of the resulting SFE would be at least $|C(|\psi\rangle)| \cdot \mathsf{poly}(\lambda)$. Our first step is to reduce this gap by constructing a QFHE scheme with nearly optimal ciphertext expansion.

Lemma 1.1 (Informal)**.** *Assuming the quantum hardness of the LWE problem, there exists a (leveled) QFHE scheme with rate-1.*

Malicious Circuit Privacy. We then lift the protocol based on QFHE to the malicious setting. Here the challenge is to ensure that the ciphertext computed by Bob does not contain any residual information about C, besides the output $C(|\psi\rangle)$. In other words, we want a QFHE scheme that satisfies circuit privacy [OPP14] for *any choice* of Alice's first message. Our second step is to give a generic transformation from any QFHE scheme (satisfying some natural structural properties) to a QFHE scheme with malicious circuit privacy. This allows us to state the following lemma.

Lemma 1.2 (Informal)**.** *Assuming the quantum hardness of the LWE problem, there exists a maliciously circuit private (leveled) QFHE scheme.*

In the quantum setting, 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 besides $C(|\psi^*\rangle)$. We remark that our transformation is *in the plain model*, i.e. it does not assume any form of trusted setup or common reference string.

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 Related Work

The problem of secure (i.e. blind) computation of quantum circuits [BFK09, DNS10, DNS12] has a strong tradition in the quantum cryptography literature. To the best of our knowledge, the only two-round protocol was given in the recent work of Bartusek et al. [BCKM20]. In contrast to our work, their protocol assumes a trusted setup and the resulting communication complexity is proportional to the size of the circuit (i.e. it is not compact). On the flip side, they achieve the strong notion of simulation security and they assume any post-quantum two-round oblivious transfer, whereas we crucially rely on the LWE assumption.

We also mention a line of work on *verifiability* of quantum computation (see [Mah18b] and references therein) where it is required that a malicious Bob must prove to Alice that he evaluated the "correct" circuit C.[1] This notion is orthogonal to our settings and can be seen as the complement of malicious circuit privacy, where the roles of the corrupted parties are reversed.

2 Technical Overview

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

[1] Clearly, this notion only makes sense when the circuit C is public and the resources needed by Alice to check Bob's proof are less than those required to evaluate C.

2.1 Malicious Circuit Privacy

We begin by outlining our transformation to add malicious circuit privacy. Our approach is generic and works with almost any existing QFHE scheme, and in particular will also be compatible with the rate-1 QFHE scheme (described later in this overview).

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 on 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 setting, 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:

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

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 operators 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 classical component of the ciphertext satisfies (semi-honest) circuit privacy.

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].

Semi-Honest Circuit Privacy. Our first observation is that the schemes from [Mah18a, Bra18] can be lifted almost generically to satisfy semi-honest circuit privacy. As we discussed before, an evaluated (single qubit) ciphertext consists of the pair

$$\mathsf{QOTP}((x,z), |\psi\rangle), \mathsf{FHE.Enc}(\mathsf{pk}, (x,z))$$

where the second component (i.e. the classical fully homomorphic part of the ciphertext) is already statistically close to a uniformly sampled encryption of (x,z), by property (2). However, nothing prevents the one-time key (x,z) from carrying some information about the circuit being computed. Fortunately, we can re-randomize the QOTP key 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. To obtain a consistent QFHE ciphertext, we need to propagate this key switch to the classical component, which can be simply done by evaluating the function

$$f_{(v,w)} : (x,z) \to (v \oplus x, w \oplus z)$$

homomorphically over the ciphertext $\mathsf{FHE.Enc}(\mathsf{pk}, (x,z))$. We again rely on the classical (semi-honest) circuit privacy of the classical FHE scheme to establish that the resulting QFHE ciphertext is statistically close to a fresh encryption of $|\psi\rangle$.

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 homomorphically 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 a 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 classical FHE encryption of the one-time key $c = \mathsf{FHE.Enc}(\mathsf{pk}, (x,z))$ -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. The classical part of the ciphertext can trivially be interpreted as quantum, where the bit z in the Pauli key (x, z) has no effect. 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) is well-formed, then the encrypter can recover otk, but the semi-honest circuit privacy of the FHE guarantees that nothing is learned about C. On the other hand, if (pk, c) is 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.

2.2 Rate-1 Quantum Fully-Homomorphic Encryption

We now turn to the description of the other ingredient for our final protocol, namely a rate-1 QFHE scheme.

What Makes This a Non-Trivial Problem? Before describing our solution, it is instructive to understand why existing schemes fail to achieve good ciphertext expansions and have low (inverse polynomial) rate. In the schemes from [Mah18a, Bra18], a ciphertext encrypting an ℓ-qubit state $|\psi\rangle$ is of the form

$$\mathsf{QOTP}((x_1, z_1, \dots, x_\ell, z_\ell), |\psi\rangle), \mathsf{QEnc}(\mathsf{pk}, (x_1, z_1, \dots, x_\ell, z_\ell))$$

where the QOTP is applied qubit-by-qubit and the classical string $\mathsf{otk} = (x_1, z_1, \dots, x_\ell, z_\ell)$ is encrypted bit-by-bit. It is not hard to see that this scheme has inverse polynomial rate, due to the blow-up introduced by the (classical) FHE encryption.

One obvious solution to improve the rate would be to adopt the *hybrid encryption* approach and sample the QOTP key using a cryptographic PRG with polynomial stretch. That is, we could improve the rate of the ciphertexts by computing

$$\mathsf{QOTP}(\mathsf{PRG}(\mathsf{seed}), |\psi\rangle), \mathsf{QEnc}(\mathsf{pk}, \mathsf{seed})$$

for some uniformly sampled $\mathsf{seed} \leftarrow\!\!\$ \{0,1\}^\lambda$. Note that we can still homomorphically compute a function in the resulting scheme, since one can always convert the ciphertexts back to their original form by evaluating the PRG homomorphically.

While this generic approach suffices for fresh ciphertexts, the troubles start once we begin to evaluate functions homomorphically: Depending on the gate that we apply to the quantum state, the one-time key otk changes accordingly to otk'. For the case of the encrypted CNOT operation, the modification is even non-deterministic. While [Mah18a] shows a way to update the classical component consistently, this method conflicts with our hybrid encryption strategy. This is because the modified otk' will most likely lie outside the support of the PRG and thus a string seed' such that $\mathsf{PRG}(\mathsf{seed}') = \mathsf{otk}'$ might simply not exist. Thus we are stuck with a classical encryption $\mathsf{QEnc}(\mathsf{pk}, \mathsf{otk})$, which brings us back to our original problem. Even assuming an ideal case where the classical FHE scheme has optimal rate, we still have a constant (> 2) ciphertext blow-up. Since two classical bits are necessary to encrypt a qubit [AMTDW00], we seem to have encountered a roadblock.

Spooky Interactions. On a high-level, our solution will leverage the structure of a special classical FHE scheme to refresh our QFHE ciphertext to the hybrid (i.e. rate-1) state. More in details, we observe that certain recent FHE schemes [BDGM19] pack k classical bits in ciphertexts of the form $c = (\mathbf{c}_0, c_1, \dots, c_k) \in \mathbb{Z}_q^{n+1} \times \{0,1\}^k$, for some modulus q and $n = \mathsf{poly}(\lambda)$. The interesting property for us is that the last k-bits of the ciphertexts are *non-locally* correlated with the secret key sk. Specifically, the decryption recovers the plaintext by computing

$$\mathsf{Dec}(\mathsf{sk}, c) = F(\mathsf{sk}, \mathbf{c}_0) \oplus (c_1, \dots, c_k)$$

for some function F, whose exact description is irrelevant for us. This property, that we refer to as *spooky decryption*,[2] will be the key to our solution.

The Solution. Equipped with the tool described above, we can convert evaluated QFHE ciphertexts of the form $(\mathsf{QOTP}(\mathsf{otk}', |\psi'\rangle), \mathsf{QEnc}(\mathsf{pk}, \mathsf{otk}'))$ back to a rate-1 form using the following procedure:

- Convert $\mathsf{QEnc}(\mathsf{pk}, \mathsf{otk}')$ into an FHE ciphertext with spooky decryption via bootstrapping (i.e. evaluating the decryption circuit of QEnc homomorphically).

[2] The name is inspired by a similar phenomenon happening in multi-key FHE schemes [DHRW16].

- Parse the resulting ciphertext as

$$c = (\mathbf{c}_0, c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z}) \in \mathbb{Z}_q^{n+1} \times \{0,1\}^{2\ell}.$$

- Return $\mathbf{c}_0$ and $\bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot \mathsf{QOTP}(\mathsf{otk}', |\psi'\rangle)$.

Since $|\mathbf{c}_0| = \mathsf{poly}(\lambda)$, the size of the compressed ciphertext is ℓ qubits plus $\mathsf{poly}(\lambda)$ bits of classical information. This rate is optimal (up to polynomial additive terms), given that any public-key encryption scheme must have ciphertexts of size at least λ bits, so an additive term in the security parameter is unavoidable. This is the exact situation here, except that we have a larger additive term, which is however asymptotically insignificant.

To see why this procedure gives us a decryptable ciphertext, re-arrange the equation above to obtain

$$F(\mathsf{sk}, \mathbf{c}_0) = (x_1', z_1', \ldots, x_\ell', z_\ell') \oplus (c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z})$$

which is the correct one-time key of the quantum state

$$\begin{aligned} &\bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot \mathsf{QOTP}\left(\mathsf{otk}', |\psi'\rangle\right) \\ &= \bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot \bigotimes_{i\in[l]} \left(X^{x_i'} Z^{z_i'}\right) \cdot |\psi'\rangle \\ &= \bigotimes_{i\in[l]} \left(X^{c_{i,x}\oplus x_i'} Z^{c_{i,z}\oplus z_i'}\right) \cdot |\psi'\rangle . \end{aligned}$$

A Non-Generic Approach. The savvy reader might have noticed that the above solution introduces an additional secret key in the scheme. In the transformation from leveled to fully homomorphic this results in a different circularity assumption: Instead of the plain circular security of the QFHE scheme, we now need to assume that semantic security is retained in the presence of a two-key cycle. While formally the two assumptions are incomparable, this motivates us to investigate on whether we can achieve full homomorphism and rate-1 under the plain one-key circularity. We show that this is fact the case, by constructing a packed version of the dual-GSW FHE scheme [Mah18a] and we prove that it is quantum capable (i.e. it supports the homomorphic evaluation of quantum circuits). Next, using the shrinking algorithm from [BDGM19], we end up with a rate-1 quantum capable scheme with the same *spooky decryption* introduced above. Thus, following a similar technique, we again obtain a rate-1 quantum fully homomorphic encryption scheme.

Packed Dual-GSW scheme. The construction of the packed dual-GSW scheme is essentially the dual of the scheme from Hiromasa et al. [HAO15]. Recall that, in the (non-packed) dual-GSW scheme, the ciphertext of a plaintext μ is of the form

$$\mathbf{C} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \mu\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times(m+1)\log q}$$

where $\mathbf{A}' \in \mathbb{Z}_q^{(m+1)\times n}$, $\mathbf{S} \in \mathbb{Z}_q^{n\times(m+1)\log q}$ and $\mathsf{sk}\cdot\mathbf{A}' = 0$, with sk being the secret key of the scheme. The plaintext information is encoded in the last row of the ciphertext. In a packed scheme, we want to encrypt ℓ-bit messages, so we interpret the plaintext as a diagonal matrix $\mathbf{M} \in \{0,1\}^{\ell \times \ell}$ containing ℓ bits, and we define the ciphertext to be

$$\mathbf{C} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \mathbf{Y} \cdot \mathbf{G} \in \mathbb{Z}_q^{(m+\ell)\times(m+\ell)\log q}$$

where $\mathbf{Y} \in \{0,1\}^{(m+\ell)\times(m+\ell)}$ is an encoding of the message, $\mathbf{A}' \in \mathbb{Z}_q^{(m+\ell)\times n}$, and $\mathbf{S} \in \mathbb{Z}_q^{n\times(m+\ell)\log q}$.

In order to maintain the scheme's homomorphic properties and be able to compute a NAND gate without altering the structure of the ciphertext, we select a message encoding that preserves plaintext-point-wise addition and multiplication, as well as the relation $\mathbf{Y} \cdot \mathbf{A}' = 0$ to cancel out the mixed term of the multiplication. To achieve this, the secret key is defined as $\left[\begin{array}{c|c} \mathbf{E}_{sk} & \mathbf{I}_l \end{array}\right]$, for a matrix $\mathbf{E}_{sk} \in \{0,1\}^{\ell \times m}$ and $\mathbf{Y}$ is defined as $\left[\begin{array}{c} \mathbf{0} \\ \hline \mathbf{M} \cdot \mathsf{sk} \end{array}\right]$. Note that, in order to produce said form of $\mathbf{Y}$, the key-generation algorithm needs to provide encryptions of $\mathbf{P}_i$ for $i \in \{0, \ldots, \ell\}$, where $\mathbf{P}_i$ is a diagonal matrix with 1 in slot (i,i) and zero everywhere else. Then, the encryption algorithm sums all the encryptions corresponding to the input message and re-randomizes the result.

To see why the scheme is quantum capable, observe that by summing up columns $(m+i)\log q$ for $i \in \{1, \ldots, \ell\}$ in our ciphertext, we end up with

$$\mathbf{c}^* = \mathbf{A}'\mathbf{s}^* + \mathbf{e}^* + \left[\begin{array}{c|c} \mathbf{0} & \frac{q}{2}\mu_1 \cdots \frac{q}{2}\mu_\ell \end{array}\right]^T \in \mathbb{Z}_q^{m+\ell}$$

where $(\mu_1, \ldots, \mu_\ell)$ are the entries in $\mathbf{M}$. Next, by isolating the first m rows of the result, alongside the $(m+i)$-th row, we obtain a dual-Regev ciphertext encrypting μ_i. This is the same scheme that Mahadev [Mah18a] converts dual-GSW to (by isolating the last column), and shows that it is quantum capable. Thus, we can apply the encrypted CNOT operation from [Mah18a] using each of the ℓ ciphertexts in parallel and then bootstrap back into the packed scheme to continue the homomorphic computations. We refer the reader to Section 7 for further details.

2.3 Putting Things Together

Applying the malicious circuit privacy transformation to the newly obtained rate-1 QHFE scheme, we obtain our main result as a straightforward implication. For completeness, we outline the protocol below.

- **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 semantic security of the QFHE scheme ensures that the ℓ-qubit state $|\psi\rangle$ is computationally indistinguishable from an encryption of the state $|0\rangle^{\otimes \ell}$. Malicious circuit privacy guarantees that no information about the circuit C is leaked to the client, beyond what is already revealed by $C(|\psi\rangle)$.

The ciphertext size of the rate-1 QFHE scheme is that of the underlying message, plus an additive term $\mathsf{poly}(\lambda)$. The size of the public key pk roughly corresponds to the size of a ciphertext, although this can be amortized by splitting the output of the computation in large enough blocks [BDGM19]. Thus we obtain a total communication complexity of

$$(|\,|\psi\rangle\,| + |C(|\psi\rangle)|) \cdot (1 + o(1))$$

which is nearly optimal.

3 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.

3.1 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 setting.

Definition 3.1 (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 3.2 (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*

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

3.2 Learning with Errors

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

Definition 3.3 (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.

3.3 Pauli Operators

The Pauli Operators X, Y, Z are 2×2 matrices that are unitary and Hermitian. More specifically:

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \; Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \; Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

3.4 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 3.4 (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).

4 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.

4.1 Classical Homomorphic Encryption

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

Definition 4.1 (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 4.2 (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 4.3 (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 4.4 (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 4.5 ([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.*

4.2 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 4.6 (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 4.1.*
- $\mathsf{FHE.QEnc}(\mathsf{pk}, |\psi\rangle)$: *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 4.7 (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 4.8 (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 $(|\psi^*\rangle, \alpha) = \mathsf{FHE.Ext}(1^\lambda, \mathsf{pk}^*, |\phi^*\rangle)$*.*

5 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.

5.1 Semi-Honest Circuit Privacy

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 $|\phi^*\rangle$. We present the definition for QFHE below as a more general case for classical FHE.

Definition 5.1 (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. This is captured by the following Lemma.

Lemma 5.2 ([Mah18a, Bra18])**.** *Assuming the quantum hardness of the circular LWE problem, there exists a QFHE scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.QEnc}, \mathsf{FHE.QEval}, \mathsf{FHE.QDec})$ *where (evaluated) ciphetexts are of the form*

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

where $\mathsf{FHE.Enc}$ *is the encryption algorithm of a classical semi-honest circuit-private FHE scheme.*

In the following we show a generic transformation that transform such schemes into semi-honest circuit private QFHE schemes [DSS16]. More formally, we have the following statement and include a proof for completeness.

Lemma 5.3 (Semi-Honest Circuit Privacy)**.** *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.*

Proof. The proof proceeds by describing an augmented evaluation algorithm, that internally runs the original evaluation algorithm from Lemma 5.2 to obtain

$$\begin{aligned}(|\phi\rangle, c) &= (\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{FHE.Enc}(\mathsf{pk}, \mathsf{otk}))\\ &= (\mathsf{QOTP.Enc}((x, z), |\psi\rangle), \mathsf{FHE.Enc}(\mathsf{pk}, (x, z)).\end{aligned}$$

Then it samples $(v, w) \leftarrow\$ \{0,1\}^2$ and outputs

$$(X^v Z^w |\phi\rangle, \mathsf{FHE.Eval}(\mathsf{pk}, f_{(v,w)}, c))$$

where $f_{(v,w)} : (x, z) \to (v \oplus x, w \oplus z)$. First observe that, by the semi-honest circuit privacy of the classical FHE it holds that

$$\begin{aligned}\mathsf{FHE.Eval}(\mathsf{pk}, f_{(v,w)}, c) &= \mathsf{FHE.Eval}(\mathsf{pk}, f_{(v,w)}, \mathsf{FHE.Enc}(\mathsf{pk}, (x, z)))\\ &\approx_s \mathsf{FHE.Enc}(\mathsf{pk}, (v \oplus x, w \oplus z))\end{aligned}$$

is statistically close to a fresh encryption of $(v \oplus x, w \oplus z)$. We now rewrite the first term as

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

where the equality above holds up to a global phase. Finally, observe that $(v \oplus x, w \oplus z)$ is a uniformly sampled one-time key. Thus, the output of the evaluation algorithm is statistically close to a fresh encryption of $|\psi\rangle$. The algorithm naturally generalizes to encryption of multiple qubits. □

5.2 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.

Abusing the notation, instead of a classical FHE scheme (as presented in the technical overview), we use the above quantum FHE scheme for the classical part of the ciphertext. 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.

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

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

since qpk is in the support of $\mathsf{QFHE.Gen}$ and c is computed as $\mathsf{QFHE.QEnc}(\mathsf{qpk}, \mathsf{otk}; s)$. 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 $\Theta_{(\mathsf{qpk}, c, \widetilde{\mathsf{otk}})}$ is a well-defined classical circuit and the above equality follows from the evaluation correctness of the FHE scheme. Thus it follows that

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

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

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 the 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 $s \leftarrow\$\, \{0,1\}^\lambda$. It sets the ciphertext as
$$(|\phi\rangle, c, c_s) = (\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QFHE.Enc}(\mathsf{qpk}, \mathsf{otk}; s), \mathsf{FHE.Enc}(\mathsf{pk}, (\mathsf{otk}, s))).$$

- **Evaluation:** On input the public key $(\mathsf{pk}, \mathsf{qpk}, c_{\mathsf{Gen}})$, a quantum circuit C, and a ciphertext $(|\phi\rangle, c, c_s)$, the evaluation algorithm defines the quantum circuit $\Gamma_{|\phi\rangle}$ as
$$\Gamma_{|\phi\rangle}(\mathsf{otk}) : \text{Return } C(\mathsf{QOTP.Dec}(\mathsf{otk}, |\phi\rangle)).$$
Then it evaluates homomorphically $|\tilde{\phi}\rangle = \mathsf{QFHE.QEval}(\mathsf{qpk}, \Gamma_{|\phi\rangle}, c)$, which results in some $\tilde{n}$ qubit state $|\tilde{\phi}\rangle$. It samples 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
$$\Theta_{(\mathsf{qpk}, c, \tilde{\mathsf{otk}})}(r, \mathsf{otk}, s) : \begin{cases} \text{If } (\cdot, \mathsf{qpk}) = \mathsf{FHE.Gen}(1^\lambda; r) \text{ and } c = \mathsf{FHE.Enc}(\mathsf{qpk}, \mathsf{otk}; s) \\ \quad \text{then return } \tilde{\mathsf{otk}}. \\ \text{Else return } 0. \end{cases}$$
It returns 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.

Lemma 5.4 (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 setting.

Lemma 5.5 (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) $c = \mathsf{QFHE.QEnc}(\mathsf{qpk}, \mathsf{otk}^*; s^*)$

  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 5.4 and Lemma 5.5 we obtain the following main theorem.

Theorem 5.6 (Malicious Circuit Privacy)**.** *Assuming the quantum hardness of the LWE problem, there exists a leveled QFHE scheme with malicious statistical circuit privacy. Additionally, assuming that the scheme is circularly secure, there exists a QFHE scheme with malicious statistical circuit privacy.*

5.3 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][3] 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

[3] The scheme from [ABG$^+$20] follows the same template as for the single-key case and can therefore be shown to achieve semi-honest circuit privacy analogously.

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 i-th key and i-th ciphertext are well formed. This conditional encryption can be done (as in the single-key setting) 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 setting). 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.

6 Rate-1 Quantum Fully Homomorphic Encryption

In the following we construct a QFHE scheme with rate approaching 1, as the security parameter (and consequently the message space) grows.

6.1 Definition

We begin by formally defining the notion of rate for a quantum homomorphic encryption scheme.

Definition 6.1 (Rate). *We say that a quantum homomorphic encryption scheme* $(\mathsf{QFHE.Gen}, \mathsf{QFHE.QEnc}, \mathsf{QFHE.QEval}, \mathsf{QFHE.QDec})$ *has rate* $\rho = \rho(\lambda)$*, if for all* pk *in the support of* $\mathsf{QFHE.Gen}(1^\lambda)$*, all supported quantum circuits* C *with sufficiently large output size, all polynomials* $\ell = \ell(\lambda)$*, all* ℓ*-qubit quantum states* $|\psi\rangle$*, and all states* $|\phi\rangle$ *where* $|\phi\rangle \in \mathsf{QFHE.QEnc}(\mathsf{pk}, |\psi\rangle)$*, it holds that*

$$\frac{|C(|\psi\rangle)|}{|\mathsf{QFHE.QEval}\,(pk, C, |\phi\rangle)|} \geq \rho$$

where $|\cdot|$ *is the size in qubits for quantum information and bits for classical information. We also say that a scheme has rate 1, if it holds that*

$$\lim_{\lambda \to \infty} \rho(\lambda) = 1$$

The notation $|\cdot|$ generally corresponds to the size of the input. In the classical setting, this translates to the number of bits that the information consists of. Similarly, in the quantum setting, we can extend the definition and measure the size in the basic unit of quantum information, a qubit. For constructing rate-1 QFHE schemes, it is convenient to define an additional ciphertext compression algorithm, together with a corresponding compressed decryption algorithm. The following are definitions from [BDGM19], extended to the quantum setting.

Definition 6.2 (Compression). *Let* $\mathsf{QFHE} = (\mathsf{QFHE.Gen}, \mathsf{QFHE.QEnc}, \mathsf{QFHE.QEval}, \mathsf{QFHE.QDec})$ *be a QFHE scheme and let* $\ell = \ell(\lambda)$ *be a polynomial. We say that* QFHE *supports* ℓ*-qubits ciphertext compression if there exist two algorithms* $\mathsf{QFHE.Compress}$ *and* $\mathsf{QFHE.CompressDec}$ *with the following syntax:*

- $\mathsf{QFHE.Compress}(\mathsf{pk}, |\phi\rangle)$*: Takes as input a public key* pk *and an encrypted* ℓ*-qubit state* $|\phi\rangle$ *and outputs a compressed ciphertext* $|\phi^*\rangle$*.*

- $\mathsf{QFHE.CompressDec}(\mathsf{sk}, |\phi^*\rangle)$*: Takes as input a secret key* sk *and a compressed ciphertext* $|\phi^*\rangle$ *and outputs an* ℓ*-qubit state* $|\psi\rangle$*.*

We require the following notion of correctness to hold for compressed ciphertexts.

Definition 6.3 (Compressed Correctness). *A quantum homomorphic encryption scheme* $(\mathsf{QFHE.Gen}, \mathsf{QFHE.QEnc}, \mathsf{QFHE.QEval}, \mathsf{QFHE.QDec}, \mathsf{QFHE.Compress}, \mathsf{QFHE.CompressDec})$ *satisfies compressed correctness if for all* $\lambda \in \mathbb{N}$*, all* $\ell = \ell(\lambda)$*, all* $(\mathsf{sk}, \mathsf{pk})$ *in the support of* $\mathsf{FHE.Gen}(1^\lambda)$*, all* ℓ*-qubit quantum states* $|\psi\rangle$*, all* $|\phi\rangle$ *such that* $|\psi\rangle = \mathsf{QDec}(\mathsf{sk}, |\phi\rangle)$*, it holds that*

$$\mathsf{CompressDec}\,(\mathsf{sk}, \mathsf{Compress}(\mathsf{pk}, |\phi\rangle)) = |\psi\rangle\,.$$

The definition of rate is unchanged, except that we consider the size of compressed ciphertexts. For the case of classical FHE, it was recently shown by Brakerski et al. [BDGM19] that a leveled scheme with rate-1 exists under the standard LWE assumption (with polynomial modulo-to-noise ratio), which can be converted to fully homomorphic by an additional circularity assumption. The scheme satisfies an additional structural property that we call *spooky decryption* and we formally define below.

Lemma 6.4 ([BDGM19]). *Assuming the hardness of the circular LWE problem, there exists a rate-1 FHE scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ *and a function* F *such that for all ciphertexts* $c = (\mathbf{c}_0, c_1, \ldots, c_k) \in \mathbb{Z}_q^{n+1} \times \{0,1\}^k$ *it holds that*

$$\mathsf{FHE.Dec}(\mathsf{sk}, c) = F(\mathsf{sk}, \mathbf{c}_0) \oplus (c_1, \ldots, c_k).$$

6.2 Our Construction

Our scheme is again described as a generic transformation, assuming the existence of the following primitives:

- A rate-1 classical FHE scheme $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ with spooky decryption (see Lemma 6.4).
- A quantum fully homomorphic encryption scheme $(\mathsf{QFHE.Gen}, \mathsf{QFHE.QEnc}, \mathsf{QFHE.QEval}, \mathsf{QFHE.QDec})$ with classical keys and hybrid ciphertexts of the form $(\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QFHE.Enc}(\mathsf{qpk}, \mathsf{otk}))$ (see Lemma 5.2)

Our transformation is presented formally in Figure 2. As before, the scheme is fully homomorphic if both ingredients are also fully homomorphic and it is otherwise leveled homomorphic.

Analysis. We proceed by analyzing the security and the correctness of our scheme.

Lemma 6.5 (Security)**.** *Assuming that QFHE and FHE are semantically secure, the scheme in Figure 2 is semantically secure.*

Proof. Let $\mathcal{A}$ be a QPT adversary against the semantic security of the rate-1 QFHE scheme. Let $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$ be a public key in support of the key generation algorithm where $\mathsf{ck} = \mathsf{FHE.Enc}(\mathsf{pk}, \mathsf{qsk})$ and $(|\phi\rangle, c)$ be an honestly computed ciphertext, where

$$\mathsf{ck} = \mathsf{FHE.Enc}(\mathsf{pk}, \mathsf{qsk}) \text{ and } (|\phi\rangle, c) = (\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QFHE.Enc}(\mathsf{qpk}, \mathsf{otk}))\,.$$

We define a series of hybrid distributions and argue that they are indistinguishable from the original ciphertext. First, we substitute the computation of the compression key with an encryption of 0 (padded to the appropriate length), obtaining

$$\mathsf{FHE.Enc}(\mathsf{pk}, 0)$$

The resulting distribution is computationally indistinguishable due to the semantic security of FHE. Next, we substitute the classical part of the ciphertext with an encryption of 0 (padded to the appropriate length), obtaining the ciphertext

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

Computational indistinguishability follows from the semantic security of QFHE. Then, we replace the quantum one-time-padded state with a totally mixed ℓ-qubit state $|u\rangle$ and get

$$(|u\rangle, \mathsf{QFHE.Enc}(\mathsf{qpk}, 0))\,.$$

This distribution is indistinguishable from the above due to the information-theoretic security of the QOTP. $\mathcal{A}$'s advantage in this experiment is 0, given that the ciphertext consists of a maximally mixed state and an encryption of 0, whereas the public key no longer includes any information about the secret key. Since this last distribution is computationally indistinguishable from the original ciphertext, it follows that $\mathcal{A}$'s advantage in the original experiment is negligible. □

Next we show that the scheme satisfies single-hop evaluation correctness. We remark that, making an additional 2-key circularity assumption, we can extend the scheme to multi-hop (for any number of hops) homomorphic via the techniques outlined in Section 5.3.

Rate-1 QFHE

- **Key Generation:** On input the security parameter 1^λ, the key generation algorithm samples two key pairs

$$(\mathsf{pk}, \mathsf{sk}) \leftarrow\!\$\, \mathsf{FHE.Gen}(1^\lambda) \text{ and } (\mathsf{qpk}, \mathsf{qsk}) \leftarrow\!\$\, \mathsf{QFHE.Gen}(1^\lambda).$$

  Then it samples a compression key $\mathsf{ck} \leftarrow\!\$\, \mathsf{FHE.Enc}(\mathsf{pk}, \mathsf{qsk})$. The secret key of the scheme is set to $(\mathsf{sk}, \mathsf{qsk})$ and the public key consists of $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$.

- **Encryption:** On input the public key $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$ and a quantum state $|\psi\rangle$, the algorithm computes and outputs $(|\phi\rangle, c) \leftarrow\!\$\, \mathsf{QFHE.QEnc}(\mathsf{qpk}, |\psi\rangle)$.

- **Evaluation:** On input the public key $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$, a quantum circuit C, and a ciphertext $\mathsf{ct} = (|\phi\rangle, c)$, the algorithm computes and outputs the evaluated ciphertext $(|\xi\rangle, \tilde{c}) = \mathsf{QFHE.QEval}(\mathsf{qpk}, C, \mathsf{ct})$.

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

- **Compression:** On input the public key $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$ and (without loss of generality) an evaluated ciphertext $(|\xi\rangle, \tilde{c})$, the compression algorithm key-switches from QFHE to FHE, by homomorphically decrypting the classical part of the ciphertext, computing

$$(\mathbf{c}_0, c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z}) = \mathsf{FHE.Eval}\left(\mathsf{pk}, \mathsf{QFHE.Dec}(\cdot, \tilde{c}), \mathsf{ck}\right)$$

  Then, it computes an ℓ-qubit state

$$|\phi\rangle = \bigotimes_{i \in [l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot |\xi\rangle$$

  and outputs $(|\phi\rangle, \mathbf{c}_0)$.

- **Compressed Decryption:** On input the secret key $(\mathsf{sk}, \mathsf{qsk})$ and a compressed ciphertext $(|\phi\rangle, \mathbf{c}_0)$, where $|\phi\rangle$ is an ℓ-qubit state, the algorithm proceeds as follows. It computes $F(\mathsf{sk}, \mathbf{c}_0) = ((f_{1,x}, f_{1,z}), \ldots, (f_{\ell,x}, f_{\ell,z}))$ and outputs the ℓ-qubit state

$$|\psi\rangle = \bigotimes_{i \in [l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) \cdot |\phi\rangle .$$

Figure 2: Description of a rate-1 QFHE scheme.

Lemma 6.6 (Correctness)**.** *Assuming that the schemes* FHE *and* QFHE *are correct, the scheme in Figure 2 satisfies compressed correctness.*

Proof. Fix a public key $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$ and a secret key $(\mathsf{sk}, \mathsf{qsk})$ and an input ciphetext $(|\xi\rangle, \tilde{c})$ where

$$|\xi\rangle = \mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle)$$

for some quantum state $|\psi\rangle$, where $\mathsf{otk} = (x_1, z_1, \ldots, x_\ell, z_\ell)$ and $\tilde{c}$ is a classical encryption of otk. Recall that the compression algorithm defines

$$(\mathbf{c}_0, c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z}) = \mathsf{FHE.Eval}\left(\mathsf{pk}, \mathsf{QFHE.Dec}(\cdot, \tilde{c}), \mathsf{ck}\right)$$

which is also a classical encryption of otk, and

$$\begin{aligned}
|\phi\rangle &= \bigotimes_{i\in[l]} \left(X^{c_{i,x}} Z^{c_{i,z}}\right) \cdot |\xi\rangle \\
&= \bigotimes_{i\in[l]} \left(X^{c_{i,x}} Z^{c_{i,z}}\right) \cdot \mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle) \\
&= \bigotimes_{i\in[l]} \left(X^{c_{i,x}\oplus x_i} Z^{c_{i,z}\oplus z_i}\right) \cdot |\psi\rangle \\
&= \bigotimes_{i\in[l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) |\psi\rangle
\end{aligned}$$

by the spooky decryption property of the rate-1 FHE scheme. The compressed decryption algorithm then returns

$$\begin{aligned}
&\bigotimes_{i\in[l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) \cdot |\phi\rangle \\
&= \bigotimes_{i\in[l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) \cdot \bigotimes_{i\in[l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) \cdot |\psi\rangle \\
&= |\psi\rangle
\end{aligned}$$

which is the correct state. □

Parameters. We calculate the rate of the above scheme. Assuming that the plaintext $|\psi\rangle$ is an ℓ-qubit state, the compressed ciphertext consists of an ℓ-qubit state $|\phi\rangle$ and the classical information $\mathbf{c}_0 \in \mathbb{Z}_q^{n+1}$. Thus we obtain a rate of

$$\rho(\lambda) = \frac{\ell}{(n+1)\log(q) + \ell} = 1 - \frac{(n+1)\log(q)}{(n+1)\log(q) + \ell}.$$

Recall that q is some polynomial in λ and thus we can bound $\log(q) \leq \log(\lambda)^2$. Setting $\ell = \Omega(\lambda(n+1)\log(\lambda)^2)$ we obtain a rate of $\rho(\lambda) = 1 - O(1/\lambda)$.

Combining Lemma 6.5 and Lemma 6.6 we obtain the following result.

Theorem 6.7 (Rate-1 QFHE). *Assuming the quantum hardness of the LWE problem, there exists a leveled QFHE scheme with rate-1. Additionally assuming that the scheme is circularly secure, there exists a QFHE scheme with rate-1.*

7 Rate-1 QFHE via Packed Dual-GSW

In Section 6 we constructed a rate-1 QFHE scheme by key switching between the classical part of a quantum FHE scheme and a rate-1 FHE scheme. In this section, we present a different, and somewhat more direct, approach to construct rate-1 QFHE. Our generic approach required us to augment the encryption scheme with a two-key cycle in order to obtain full (as opposed to leveled) homomorphism. This non-generic approach has the advantage of requiring only a one-key cycle. While these two assumptions are formally incomparable, one-key circularity is arguably more studied and it is the same assumption that was used in [Mah18a].

7.1 Definitions

For our construction we will make use of the notion of a lattice trapdoor, along with the following theorem.

Theorem 7.1 ([MP12]). *There is an efficient algorithm* $\mathsf{GenTrap}(1^n, 1^m, q)$ *that, given* $n, m \geq 1$ *and* $q \geq 2$ *such that* $m = \Omega(n\ log(q))$*, returns a matrix* $\mathbf{A} \in \mathbb{Z}_q^{m\times n}$ *and a trapdoor* τ_A *such that the distribution of* $\mathbf{A}$ *is negligibly (in n) close to the uniform distribution. Moreover, there is an efficient algorithm* Invert *that on input* $\mathbf{A}$*,* τ_A *and* $\mathbf{As} + \mathbf{e}$*, where* $\mathbf{s}$ *is arbitrary in* $\mathbb{Z}_q^n$ *and* $\|e\| \leq q/\left(O(n\ log(q))\right)$*, returns* $\mathbf{s}$ *and* $\mathbf{e}$ *with overwhelming probability.*

We recall the definition of the ciphertext shrinking algorithm.

Definition 7.2 ([BDGM19]). *Let* $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ *be an encryption scheme where the public key specifies a message space* $\mathbb{Z}_q^\ell$*, the secret key* sk *is a matrix in* $\mathbb{Z}_q^{\ell\times m}$ *and ciphertexts are of the form* $(\mathbf{c}_a, \mathbf{c}_b)$*. Assume publicly known functions* $J,\ H,\ F$ *and let noisy decryption compute* $\mathsf{Dec}(\mathsf{sk}, (\mathbf{c}_a, \mathbf{c}_b)) = J(\mathbf{c}_b) - \mathsf{sk} \cdot H(\mathbf{c}_a)$*, where* $J(\mathbf{c}_b) \in \mathbb{Z}_q^\ell,\ H(\mathbf{c}_a) \in \mathbb{Z}_q^m$*. Then the following algorithms exist:*

- Shrink : *On input the public key and a ciphertext* $(\mathbf{c}_a, \mathbf{c}_b)$ *that encrypts* $\mathbf{m} = (m_1, \ldots, m_\ell)$*, it computes and outputs the shrunken ciphertext* $(\mathbf{c}_0, c_1, \ldots, c_\ell)$ $\in \mathbb{Z}_q^{m+1} \times \{0,1\}^\ell$*.*
- $\mathsf{ShrinkDec}$: *On input the secret key and a shrunken ciphertext* $(\mathbf{c}_0, c_1, \ldots, c_\ell)$*, it computes and outputs* $F\left(\mathsf{sk}, \mathbf{c}_a\right) \oplus (c_1, \ldots, c_\ell) = (m_1, \ldots, m_\ell)$*.*

The classical FHE used to encrypt classical information in a quantum scheme in Lemma 5.2 is referred to as a quantum capable scheme. Below we present the definition and requirements of a quantum capable scheme.

Definition 7.3 ([Mah18a]). *Let FHE be a classical leveled fully-homomorphic encryption scheme. FHE is quantum capable if there exists an encryption scheme AltHE such that:*

1. *There exists an algorithm FHE.Convert that takes as input an encryption* c *under FHE and outputs an encryption* $\hat{c}$ *under AltHE, where both ciphertexts encrypt the same value.*
2. *AltHE allows the operation* $\oplus_H$*, which is the XOR operation that can be performed homomorphically. This operation should also be easily invertible using only the public key of AltHE.*
3. *There exists a distribution* D *(which may depend on parameters of FHE) that satisfies the following conditions:*
   (a) *For all ciphertexts* c *that can arise during homomorphic evaluation,* $\{AltHE.Enc(pk, x; r) \mid (x, r) \leftarrow D\} \approx_s \{AltHE.Enc(pk, x; r) \oplus_H c \mid (x, r) \leftarrow D\}$*, where* x *is the plaintext and* r *is the chosen randomness.*

(b) *There exists a bounded-error quantum polynomial time procedure to, given AltHE's public key, construct the superposition*

$$\sum_{x \in \{0,1\}, r} \sqrt{D(x,r)}\, |x, r\rangle$$

(c) *Given a ciphertext $c = AltHE.Enc(pk, x; r) \mid (x, r) \leftarrow D$, the secret key and some trapdoor information, it must be possible to compute (x, r)*

7.2 Packed Dual GSW

The first step to construct a rate-1 QFHE scheme is to present a packed version of the dual-GSW FHE scheme.

Let λ be the security parameter of the scheme. For simplicity, we assume that q is a power of 2. Let m and n be polynomially bounded functions of λ, where $m = \Omega(n \log(q))$, ℓ be the number of bits encrypted in a packed ciphertext and $N = (m + \ell) \log(q)$. Also, let B be a positive integer that will constitute the bound of the distribution χ that the error will be sampled from (see [Mah18a] for more details). We require that $q \geq \omega(\mathsf{poly}(\lambda) \cdot B)$.

For our construction we will make use of two operations from [GSW13] as shown in [Mah18a]: The linear operator $\mathbf{G}$ and the inverse operator G^{-1}. $\mathbf{G}$ is the matrix $(1, 2, \ldots, 2^{\log(q)}) \otimes \mathbf{I}_{m+\ell}$, which converts a binary representation back to its original representation. The operator $\mathbf{G}$ is well defined even for non-binary vectors. The non linear operator G^{-1} is the inverse of $\mathbf{G}$ and converts a vector (or each column of a matrix) to its binary representation. It is important to note that $\mathbf{G}G^{-1}$ results in the identity operation.

Our construction is inspired by the work of Hiromasa et al. [HAO15] and it is shown in Figure 3.

Ciphertext Conversion. To show the quantum capability of our scheme and to obtain a rate-1 scheme, we need some additional algorithms. Let dGSW be the packed dual GSW scheme we introduced in Figure 3. Consider the following scheme QCdGSW, which is identical to dGSW, except that it has an additional conversion algorithm that converts the ciphertext to an alternate scheme, as well as a different (noisy) decryption algorithm, as described below.

- **Conversion:** On input a ciphertext $\mathbf{C}$, the algorithm sums up columns $(m + i) \log(q)$ for $i \in \{1, \ldots, \ell\}$ and outputs the one column ciphertext

$$\mathbf{c}^* = \mathbf{A}'\mathbf{s}^* + \mathbf{e}^* + \begin{bmatrix} \mathbf{0} \\ \hline \frac{q}{2}\mu_1 \\ \vdots \\ \frac{q}{2}\mu_\ell \end{bmatrix} \in \mathbb{Z}_q^{m+\ell}$$

It continues by extracting individual ciphertexts

$$\mathbf{c}_i^* = \begin{bmatrix} \mathbf{A} \\ \hline -\mathbf{e}_{sk_i}\mathbf{A} \end{bmatrix} \mathbf{s}^* + \mathbf{e}_i^* + \begin{bmatrix} \mathbf{0} \\ \hline \frac{q}{2}\mu_i \end{bmatrix} \in \mathbb{Z}_q^{m+1}$$

for $i \in \{1, \ldots, \ell\}$, by keeping the first m rows and the $(m + i)$-th row of $\mathbf{C}^*$, where $\mathbf{e}_{sk_i}$ is the i-th row of $\mathbf{E}_{sk}$.

Packed Dual GSW

- **Key Generation:** On input the security parameter 1^λ, the key generation algorithm chooses $\mathbf{E}_{sk} \in \{0,1\}^{\ell \times m}$. Using the procedure $\mathsf{GenTrap}(1^n, 1^m, q)$ it samples a random trapdoor matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$. Let $\mathsf{sk} = \left[\begin{array}{c|c} \mathbf{E}_{sk} & \mathbf{I}_l \end{array}\right] \in \{0,1\}^{\ell \times (m+\ell)}$ and $\mathbf{A}' = \left[\begin{array}{c} \mathbf{A} \\ \hline -\mathbf{E}_{sk}\mathbf{A} \end{array}\right] \in \mathbb{Z}_q^{(m+\ell)\times n}$. Let $\mathbf{P}_i \in \{0,1\}^{\ell \times \ell}$ ($i \in \{1,\ldots,\ell\}$) be the matrix with 1 in the (i,i)-th position and zero everywhere else. For $i \in \{1,\ldots,\ell+1\}$ it samples $\mathbf{S}_i \leftarrow\!\$\, \mathbb{Z}_q^{n \times N}$, $\mathbf{E}_i \leftarrow\!\$\, \chi^{(m+\ell)\times N}$ and then calculates
$$\mathbf{X}_i = \mathbf{A}'\mathbf{S}_i + \mathbf{E}_i + \left[\begin{array}{c} \mathbf{0} \\ \hline \mathbf{P}_i\mathsf{sk} \end{array}\right] \cdot \mathbf{G} \in \mathbb{Z}_q^{(m+\ell)\times N}$$
for $i \in \{1,\ldots,\ell\}$, and
$$\mathbf{C}_I = \mathbf{A}'\mathbf{S}_{\ell+1} + \mathbf{E}_{\ell+1} + \left[\begin{array}{c} \mathbf{0} \\ \hline \mathbf{I}_\ell\mathsf{sk} \end{array}\right] \cdot \mathbf{G} \in \mathbb{Z}_q^{(m+\ell)\times N}.$$
It outputs the secret key of the scheme set to sk, the public key set to $(\mathbf{A}', \{\mathbf{X}_i\}_{i\in\{1....,\ell\}}, \mathbf{C}_I)$ and the trapdoor τ_A.

- **Encryption:** On input the public key $(\mathbf{A}', \{\mathbf{X}_i\}_{i\in\ell}, \mathbf{C}_I)$ and messages $(\mu_1,\ldots,\mu_\ell)$, the algorithm samples $\mathbf{S} \leftarrow\!\$\, \mathbb{Z}_q^{n\times N}$, $\mathbf{E} \leftarrow\!\$\, \chi^{(m+\ell)\times N}$ and it computes and outputs
$$\mathbf{C} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \sum_{i=1}^{\ell} \mu_i \cdot \mathbf{X}_i$$

- **Evaluation:** On input the public key $(\mathbf{A}', \{\mathbf{X}_i\}_{i\in\{0,\ldots,\ell\}}, \mathbf{C}_I)$, a circuit C, and a ciphertext $\mathbf{C}$, the algorithm computes and outputs the evaluated ciphertext $\mathbf{C}'$. In order to apply the NAND gate on input $\mathbf{C}_1, \mathbf{C}_2$, it computes $\mathbf{C}_I - \mathbf{C}_1\mathbf{G}^{-1}(\mathbf{C}_2)$.

- **Decryption:** On input a secret key sk and (without loss of generality) an evaluated ciphertext $\mathbf{C}$, the algorithm computes
$$\mathbf{M} = \begin{bmatrix} \mu'_1 & & \\ & \ddots & \\ & & \mu'_\ell \end{bmatrix} = \mathsf{sk} \cdot \mathbf{C}.$$
For each entry μ'_i it returns 0 if the result is closer to 0 than $q/2$ and 1 otherwise.

Figure 3: Description of the packed dual GSW scheme.

- **Noisy Decryption:** On input a secret key sk and a converted ciphertext $\mathbf{c}^*$, the algorithm creates two ciphertexts $\mathbf{c}^*_a$ and $\mathbf{c}^*_b$ by keeping the first m rows and the last ℓ rows of $\mathbf{c}^*$ respectively. Then it computes and outputs

$$\mathbf{c}_b^* + \mathbf{E}_{sk}\mathbf{c}_a^*.$$

Observe that the structure of the (noisy) decryption algorithm allows us to apply the shrinking algorithm from Definition 7.2.

7.3 Analysis

We proceed by proving that the scheme satisfies some properties of interest.

Homomorphic Evaluation. First we show that throughout homomorphic evaluations the ciphertext preserves the form

$$\mathbf{A}'\mathbf{S} + \mathbf{E} + \mathbf{Y}\cdot\mathbf{G} \in \mathbb{Z}_q^{(m+\ell)\times N}$$

where $\mathbf{Y} = \left[\begin{array}{c}\mathbf{0}\\ \hline \mathbf{M}\cdot\mathsf{sk}\end{array}\right] = \left[\begin{array}{c|c}\mathbf{0} & \mathbf{0}\\ \hline \mathbf{M}\cdot\mathbf{E}_{sk} & \mathbf{M}\end{array}\right] \in \{0,1\}^{(m+\ell)\times(m+\ell)}$ and $\mathbf{M} = \begin{bmatrix}\mu_1 & & \\ & \ddots & \\ & & \mu_\ell\end{bmatrix} \in \{0,1\}^{m\times m}$.

A freshly encrypted ciphertext corresponds to this structure since

$$\begin{aligned}
\mathbf{C} &= \mathbf{A}'\mathbf{S} + \mathbf{E} + \sum_{i=1}^{\ell}\mu_i\cdot\mathbf{X}_i\\
&= \mathbf{A}'\left(\mathbf{S} + \sum_{i=1}^{\ell}\mu_i\mathbf{S}_i\right) + \left(\mathbf{E} + \sum_{i=1}^{\ell}\mu_i\mathbf{E}_i\right) + \left[\begin{array}{c}\mathbf{0}\\ \hline \sum_{i=1}^{\ell}\mu_i\mathbf{P}_i\cdot\mathsf{sk}\end{array}\right]\cdot\mathbf{G}\\
&= \mathbf{A}'\left(\mathbf{S} + \sum_{i=1}^{\ell}\mu_i\mathbf{S}_i\right) + \left(\mathbf{E} + \sum_{i=1}^{\ell}\mu_i\mathbf{E}_i\right) + \left[\begin{array}{c}\mathbf{0}\\ \hline \mathbf{M}\cdot\mathsf{sk}\end{array}\right]\cdot\mathbf{G}
\end{aligned}$$

Now, we show that the same structure is maintained during a NAND operation.

$$\begin{aligned}
\mathbf{C}_I - \mathbf{C}_1\mathbf{G}^{-1}(\mathbf{C}_2) &= \mathbf{C}_I - \left(\mathbf{A}'\mathbf{S}_1 + \mathbf{E}_1 + \mathbf{Y}_1\cdot\mathbf{G}\right)\mathbf{G}^{-1}(\mathbf{C}_2)\\
&= \mathbf{C}_I - \mathbf{A}'\mathbf{S}_1\mathbf{G}^{-1}(\mathbf{C}_2) - \mathbf{E}_1\mathbf{G}^{-1}(\mathbf{C}_2) - \mathbf{Y}_1\mathbf{A}'\mathbf{S}_2 - \mathbf{Y}_1\mathbf{E}_2 - \mathbf{Y}_1\mathbf{Y}_2\mathbf{G}\\
&= \mathbf{C}_I - \mathbf{A}'\mathbf{S}_1\mathbf{G}^{-1}(\mathbf{C}_2) - \mathbf{E}_1\mathbf{G}^{-1}(\mathbf{C}_2) - \mathbf{Y}_1\mathbf{E}_2 - \mathbf{Y}_1\mathbf{Y}_2\mathbf{G}\\
&= \mathbf{A}'\tilde{\mathbf{S}} + \tilde{\mathbf{E}} + \tilde{\mathbf{Y}}\mathbf{G}
\end{aligned}$$

where $\tilde{\mathbf{S}} = \left(\mathbf{S}_I - \mathbf{S}_1\mathbf{G}^{-1}(\mathbf{C}_2)\right)$, $\tilde{\mathbf{E}} = \left(\mathbf{E}_I - \mathbf{E}_1\mathbf{G}^{-1}(\mathbf{C}_2) - \mathbf{Y}_1\mathbf{E}_2\right)$ and

$$\begin{aligned}
\tilde{\mathbf{Y}} &= \left(\left[\begin{array}{c}\mathbf{0}\\ \hline \mathbf{I}_\ell\mathsf{sk}\end{array}\right] - \mathbf{Y}_1\mathbf{Y}_2\right)\mathbf{G}\\
&= \left(\left[\begin{array}{c}\mathbf{0}\\ \hline \mathbf{I}_\ell\cdot\mathsf{sk}\end{array}\right] - \left[\begin{array}{c}\mathbf{0}\\ \hline \mathbf{M}_1\cdot\mathsf{sk}\end{array}\right]\left[\begin{array}{c}\mathbf{0}\\ \hline \mathbf{M}_2\cdot\mathsf{sk}\end{array}\right]\right)\mathbf{G}\\
&= \left[\begin{array}{c}\mathbf{0}\\ \hline (\mathbf{I}_\ell - \mathbf{M}_1\mathbf{M}_2)\cdot\mathsf{sk}\end{array}\right]\mathbf{G}
\end{aligned}$$

The second equality strands true because $\mathbf{G}G^{-1}(\mathbf{C}_2) = \mathbf{C}_2$ whereas the third equality because $\mathbf{Y}_1\mathbf{A}' = \left[\begin{array}{c|c} \mathbf{0} & \mathbf{0} \\ \hline \mathbf{M}_1 \cdot \mathbf{E}_{sk} & \mathbf{M}_1 \end{array}\right] \cdot \left[\begin{array}{c} \mathbf{A} \\ \hline -\mathbf{E}_{sk}\mathbf{A} \end{array}\right] = \mathbf{0}$. The form of $\tilde{\mathbf{Y}}$ is correct since

$$\mathbf{Y}_1\mathbf{Y}_2 = \left[\begin{array}{c|c} \mathbf{0} & \mathbf{0} \\ \hline \mathbf{M}_1 \cdot \mathbf{E}_{sk} & \mathbf{M}_1 \end{array}\right] \cdot \left[\begin{array}{c|c} \mathbf{0} & \mathbf{0} \\ \hline \mathbf{M}_2 \cdot \mathbf{E}_{sk} & \mathbf{M}_2 \end{array}\right] = \left[\begin{array}{c|c} \mathbf{0} & \mathbf{0} \\ \hline \mathbf{M}_1\mathbf{M}_2 \cdot \mathbf{E}_{sk} & \mathbf{M}_1\mathbf{M}_2 \end{array}\right]$$

Note that in [Mah18a], the error parameter entries are sampled from a truncated discrete Gaussian distribution with bound B. Hence, in a freshly encrypted ciphertext the error entries are bounded by $(\ell+1) \cdot B$, in $\mathbf{E}_I$ by B, and $\mathbf{Y}_1$ has at most $(m+1)$ non-zero entries in each row. As a result the entries in $\tilde{\mathbf{E}}$ are bounded by

$$B((\ell+1)N + (\ell+1)(m+1)+1) = B \cdot ((\ell+1)(N+m+1)+1) \leq B \cdot (\ell+1)(N+m+2).$$

If the scheme is used to compute an L-depth circuit, the error entries of $\tilde{\mathbf{E}}$ are bounded by $B' = B \cdot (\ell+1)(N+m+2)^L$ across all ciphertexts during computations.

Quantum Capability. Here we proceed to prove the quantum capability of our scheme.

Theorem 7.4. *The scheme* QCdGSW *described before is quantum capable.*

Proof. The individual ciphertexts $\mathbf{c}_i^*$ produced by the conversion algorithm are exactly the dual-Regev encryptions of $\mu_1, \ldots, \mu_\ell$, each with secret key $\left[\begin{array}{c|c} \mathbf{e}_{sk_i} & 1 \end{array}\right] \in \mathbb{Z}_q^{m+1}$. This constitutes the AltHE scheme in Definition 7.3, which is the same as the one used in [Mah18a]. As a result, the first two requirements in Definition 7.3 are satisfied immediately. Requirement 3c is also immediately satisfied, as the matrix A with trapdoor τ_A is intact in the converted ciphertexts.

For requirements 3a and 3b, we need to appropriately bound the error parameter. Specifically, in [Mah18a], the requirements hold if the parameter of the truncated discrete Gaussian distribution that the error was sampled from is super-polynomially larger than the error entries of $\mathbf{E}^*$. From the assumption parameters we know that modulus q is super-polynomially larger than the original error entries. When proving the invariability of the ciphertext form, we proved that the error entries of an evaluated ciphertext are bounded by $B' = B \cdot (\ell+1)(N+m+2)^L$. This means that in a converted ciphertext, they are bounded by $\ell \cdot B'$. As a result, it is possible to define a large enough error bound so as to satisfy these requirements. □

Security and Correctness. We proceed by analyzing the security and the correctness of our scheme. The security is proven against the the standard LWE assumption in the presence of encryptions of the secret key (circular LWE).

Lemma 7.5 (Security). *Assuming circular LWE, the* QCdGSW *scheme introduced above is semantically secure.*

Proof. Let $\mathcal{A}$ be a QPT adversary against the semantic security of QCdGSW. Let $(\mathbf{A}', \{\mathbf{X}_i\}_{i\in\{1\ldots,\ell\}}, \mathbf{C}_I)$ be a public key in support of the key generation algorithm and $\mathbf{C} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \sum_{i=1}^{\ell} \mu_i \cdot \mathbf{X}_i \in \mathbb{Z}_q^{(m+\ell)\times N}$ be an honestly computed ciphertext where

$$\mathbf{X}_i = \mathbf{A}'\mathbf{S}_i + \mathbf{E}_i + \left[\begin{array}{c} \mathbf{0} \\ \hline \mathbf{P}_i\mathsf{sk} \end{array}\right] \cdot \mathbf{G},\ \mathbf{C}_I = \mathbf{A}'\mathbf{S}_I + \mathbf{E}_I + \left[\begin{array}{c} \mathbf{0} \\ \hline \mathbf{I}_\ell\mathsf{sk} \end{array}\right] \cdot \mathbf{G} \text{ and } \mathbf{A}' = \left[\begin{array}{c} \mathbf{A} \\ \hline -\mathbf{E}_{sk}\mathbf{A} \end{array}\right].$$

We define a series of hybrid distributions and argue that they are indistinguishable from the original ciphertext. First we substitute the matrix $\mathbf{A}$ with a uniformly random matrix $\mathbf{U}$, getting

$$\mathbf{A}' = \left[\begin{array}{c} \mathbf{U} \\ \hline -\mathbf{E}_{sk}\mathbf{U} \end{array}\right]$$

This distribution is statistically indistinguishable from the original due to Theorem 7.1. Then, we substitute $\mathbf{A}'$ with a uniformly random matrix $\mathbf{U}'$. The resulting distribution is statistical indistinguishable from the previous one due to the the leftover hash lemma [HILL99].

We proceed with the next ℓ hybrid distributions. For $i \in \{1, \dots, \ell\}$, Hybrid_{i+2} is the distribution where we substitute each

$$\mathbf{X}_i = \mathbf{U}'\mathbf{S}_i + \mathbf{E}_i + \left[\begin{array}{c} \mathbf{0} \\ \hline \mathbf{P}_i\mathsf{sk} \end{array}\right] \cdot \mathbf{G}$$

with a uniform matrix $\mathbf{V}_i$. Each distribution is computationally indistinguishable from the previous one on account of the hardness of the circular LWE. Similarly, for the next hybrid distribution we replace $\mathbf{C}_I$ with a uniformly random matrix $\mathbf{V}_I$. At last, we substitute the ciphertext

$$\mathbf{C} = \mathbf{U}'\mathbf{S} + \mathbf{E} + \sum_{i=1}^{\ell} \mu_i \cdot \mathbf{V}_i$$

with a uniformly random matrix $\mathbf{W}$. The resulting distribution is again computationally indistinguishable from the previous by an invocation of the hardness of the circular LWE.

$\mathcal{A}$'s advantage in the last experiment is 0, given that the ciphertext consists of a uniformly random matrix and the public key no longer includes any information of the secret key. Since this last distribution is computationally indistinguishable from the original ciphertext, it follows that $\mathcal{A}$'s advantage in the original experiment is negligible. □

Next we proceed to show the single-hop evaluation correctness of the scheme. We can extend the scheme to multi-hop (for any number of hops) homomorphic via bootstrapping.

Lemma 7.6 (Correctness). *Assuming that the algorithms* Shrink *and* ShrinkDec *in Definition 7.2 are correct, then the* QCdGSW *scheme introduced above satisfies decryption correctness.*

Proof. Fix a public key $\mathsf{pk} = \mathbf{A}'$ and a secret key $\mathsf{sk} = \left[\begin{array}{c|c} \mathbf{E}_{sk} & \mathbf{I}_l \end{array}\right]$. It was proven that a QCdGSW ciphertext has form

$$\mathbf{A}'\mathbf{S} + \mathbf{E} + \mathbf{Y} \cdot \mathbf{G} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \left[\begin{array}{c|c} \mathbf{0} & \mathbf{0} \\ \hline \mathbf{M} \cdot \mathbf{E}_{sk} & \mathbf{M} \end{array}\right] \cdot \mathbf{G}$$

It is easily shown that column $(m+i)\log(q)$ of $\mathbf{Y} \cdot \mathbf{G}$ is a column with zeros everywhere else and $\frac{q}{2}\mu_i$ in the $(m+i)$-th row, for $i \in \{1, \dots, \ell\}$. Hence, adding those columns results into the converted ciphertext

$$\mathbf{c}^* = \mathbf{A}'\mathbf{s}^* + \mathbf{e}^* + \left[\begin{array}{c} \mathbf{0} \\ \hline \frac{q}{2}\mu_1 \\ \vdots \\ \frac{q}{2}\mu_\ell \end{array}\right] \in \mathbb{Z}_q^{m+\ell}$$

It is important to notice that the matrix $\mathbf{A}'$ remains intact after this transformation. Recall that the decryption algorithm creates $\mathbf{c}_a^*$ and $\mathbf{c}_b^*$ by keeping the first m rows and the last ℓ rows of $\mathbf{c}^*$ respectively. Then it outputs

$$\begin{aligned}\mathbf{c}_b^* + \mathbf{E}_{sk}\mathbf{c}_a^* &= -\mathbf{E}_{sk}\mathbf{A}\cdot\mathbf{s}^* + \mathbf{e}_b^* + \begin{bmatrix} \frac{q}{2}\mu_1 \\ \vdots \\ \frac{q}{2}\mu_\ell \end{bmatrix} + \mathbf{E}_{sk}\cdot\mathbf{A}\mathbf{s}^* + \mathbf{E}_{sk}\cdot\mathbf{e}_a^* \\ &= (\mathbf{E}_{sk}\mathbf{e}_a^* + \mathbf{e}_b^*) + \begin{bmatrix} \frac{q}{2}\mu_1 \\ \vdots \\ \frac{q}{2}\mu_\ell \end{bmatrix}\end{aligned}$$

which consist of the correct (noisy) plaintexts.

The rounding (in order to get the plaintexts) operates correctly as long as the entries in $(\mathbf{E}_{sk}\mathbf{e}_a^* + \mathbf{e}_b^*)$ are bounded by $\frac{q}{4}$. The entries in $\mathbf{e}^*$ are bounded by $\ell B'$ and $\mathbf{E}_{sk}$ contributes at most by a factor m. As a result, the entries in $(\mathbf{E}_{sk}\mathbf{e}_a^* + \mathbf{e}_b^*)$ are bounded by $\ell(m+1)B' = \ell(m+1)(\ell+1)B\cdot(N+m+2)^L$. Given that q is super-polynomially larger than B, we can meet the necessary requirement. □

7.3.1 Rate-1 Quantum FHE

The rate-1 scheme we construct is identical to a QFHE with classical keys and hybrid ciphertexts of the form

$$(\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QCdGSW.Enc}(\mathsf{pk}, \mathsf{otk}))$$

(see Lemma 5.2), with the addition of a Compression and a Compressed Decryption algorithm, as described below, in order to support ℓ-qubits ciphertext compression. We use the same techniques as the rate-1 QFHE scheme described Section 6.

- **Compression:** On input the public key and (without loss of generality) an evaluated ciphertext $(|\xi\rangle, \mathbf{c}^*)$, where $|\xi\rangle$ is a quantum one time padded ℓ-qubit state and $\mathbf{c}^* \in \mathbb{Z}_q^{m+2\ell}$ is a converted packed dual GSW ciphertext, the compression algorithm uses the Shrink algorithm described in Definition 7.2 and computes

  $$(\mathbf{c}_0, c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z})) \ \in \mathbb{Z}_q^{m+1} \times \{0,1\}^{2\ell}$$

  Then, it computes an ℓ-qubit state

  $$|\phi\rangle = \bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot |\xi\rangle$$

  and outputs $(|\phi\rangle, \mathbf{c}_0)$.

- **Compressed Decryption:** On input the secret key sk and a compressed ciphertext $(|\phi\rangle, \mathbf{c}_0)$, where $|\phi\rangle$ is an ℓ-qubit state, the algorithm computes $F(\mathsf{sk}, \mathbf{c}_0) = ((f_{1,x}, f_{1,z}), \ldots, (f_{\ell,x}, f_{\ell,z}))$ and outputs the ℓ-qubit state

  $$|\psi\rangle = \bigotimes_{i\in[l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) \cdot |\phi\rangle\,.$$

Parameters. Assume the quantum fully homomorphic encryption scheme $(\mathsf{QFHE.Gen}, \mathsf{QFHE.QEnc}, \mathsf{QFHE.QEval}, \mathsf{QFHE.QDec})$ with classical keys and hybrid ciphertexts of the form $(\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QCdGSW.Enc}(\mathsf{pk}, \mathsf{otk}))$. We can see that the Compression algorithm is identical with one described in Section 6, and similarly we obtain a rate of $\rho(\lambda) = 1 - O(1/\lambda)$.

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.

[AIR01] William Aiello, Yuval Ishai, and Omer Reingold. Priced oblivious transfer: How to sell digital goods. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 119–135. Springer, Heidelberg, May 2001.

[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.

[BCKM20] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. On the round complexity of two-party quantum computation. Cryptology ePrint Archive, Report 2020/1471, 2020. `https://eprint.iacr.org/2020/1471`.

[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.

[BFK09] Anne Broadbent, Joseph Fitzsimons, and Elham Kashefi. Universal blind quantum computation. In *50th FOCS*, pages 517–526. IEEE Computer Society Press, October 2009.

[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.

[DHRW16] Yevgeniy Dodis, Shai Halevi, Ron D. Rothblum, and Daniel Wichs. Spooky encryption and its applications. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 93–122. Springer, Heidelberg, August 2016.

[DNS10] Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Secure two-party quantum evaluation of unitaries against specious adversaries. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 685–706. Springer, Heidelberg, August 2010.

[DNS12] Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Actively secure two-party evaluation of any quantum operation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 794–811. Springer, Heidelberg, August 2012.

[DSS16] Yfke Dulek, Christian Schaffner, and Florian Speelman. Quantum homomorphic encryption for polynomial-sized circuits. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 3–32. Springer, Heidelberg, August 2016.

[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.

[GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Heidelberg, August 2013.

[HAO15] Ryo Hiromasa, Masayuki Abe, and Tatsuaki Okamoto. Packing messages and optimizing bootstrapping in gsw-fhe. In Jonathan Katz, editor, *Public-Key Cryptography – PKC 2015*, pages 699–715, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg.

[HILL99] Johan HÅstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, March 1999.

[KO04] Jonathan Katz and Rafail Ostrovsky. Round-optimal secure two-party computation. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 335–354. Springer, Heidelberg, August 2004.

[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.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, April 2012.

[NP99] Moni Naor and Benny Pinkas. Oblivious transfer and polynomial evaluation. In *31st ACM STOC*, pages 245–254. ACM Press, May 1999.

[NP01] Moni Naor and Benny Pinkas. Efficient oblivious transfer protocols. In S. Rao Kosaraju, editor, *12th SODA*, pages 448–457. ACM-SIAM, January 2001.

[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.

[Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986.