

Quantum homomorphic encryption for circuits of low T-gate complexity

Anne Broadbent* and Stacey Jeffery†

Abstract

Fully homomorphic encryption is an encryption method with the property that any computation on the plaintext can be performed by a party having access to the ciphertext only. Here, we formally define and give schemes for *quantum* homomorphic encryption, which is the encryption of *quantum* information such that *quantum* computations can be performed given the ciphertext only. Our schemes allow for arbitrary Clifford group gates, but become inefficient for circuits with large complexity, measured in terms of the non-Clifford portion of the circuit (we use the “ $\pi/8$ ” non-Clifford group gate, also known as the T-gate).

More specifically, two schemes are proposed: the first scheme has a decryption procedure whose complexity scales with the square of the *number* of T-gates (compared with a trivial scheme in which the complexity scales with the total number of gates); the second scheme uses a quantum evaluation key of length given by a polynomial of degree exponential in the circuit’s T-gate depth, yielding a homomorphic scheme for quantum circuits with constant T-depth. Both schemes build on a classical fully homomorphic encryption scheme.

A further contribution of ours is to formally define the security of encryption schemes for quantum messages: we define *quantum indistinguishability under chosen plaintext attacks* in both the public- and private-key settings. In this context, we show the equivalence of several definitions.

Our schemes are the first of their kind that are secure under modern cryptographic definitions, and can be seen as a quantum analogue of classical results establishing homomorphic encryption for circuits with a limited number of *multiplication* gates. Historically, such results appeared as precursors to the breakthrough result establishing classical fully homomorphic encryption.

*Department of Mathematics and Statistics, University of Ottawa, Ottawa, Ontario, Canada; abroadbe@uottawa.ca.

†Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, California, USA; sjeffery@caltech.edu.

Contents

1	Introduction	3
1.1	Summary of Contributions and Techniques	3
2	Preliminaries and Notation	5
2.1	Notation	5
2.2	Quantum Circuits	6
2.3	Quantum One-time Pad	7
3	Definitions	7
3.1	Classical and Quantum Homomorphic Encryption	7
3.2	Security of Quantum Homomorphic Encryption	8
3.3	Correctness and Compactness of Quantum Homomorphic Encryption	10
3.4	Indivisible Schemes	11
3.5	Quantum Homomorphic Encryption in the Symmetric-Key Setting	12
4	Main Contributions	14
5	Scheme CL: Homomorphic Encryption for Clifford Circuits	15
5.1	Analysis of CL	17
6	Scheme EPR: T-gate Computation Using Entanglement	19
6.1	Analysis of EPR	22
7	Scheme AUX: T-gate Computation Using Auxiliary States	22
7.1	Analysis of AUX	26
8	Conclusions and Open Problems	30
	Acknowledgements	30
	References	30
A	Classical Fully Homomorphic Encryption	33
B	Equivalence of Definitions for q-IND-CPA	34
C	Key Update Rules for Stabilizer Elements	37
D	Correctness of the T-gate Gadget	38
E	Proof of Lemma 7.6	39

1 Introduction

An encryption scheme is *homomorphic* over some set of circuits \mathcal{S} if any circuit in \mathcal{S} can be evaluated on an encrypted input. That is, given an encryption of the message m , it is possible to produce a ciphertext that decrypts to the output of the circuit C on input m , for any $C \in \mathcal{S}$. In *fully homomorphic encryption (FHE)*, \mathcal{S} is the set of all classical circuits. FHE was introduced in 1978 by Rivest, Adleman and Dertouzos [RAD78], but the existence of such a scheme was an open problem for over 30 years. Some early public-key encryption schemes were homomorphic over the set of circuits consisting of only additions [GM84, Pai99] or over the set of circuits consisting of only multiplications [ElG85]. Several steps were made towards FHE, with schemes that were homomorphic over increasingly large circuit classes, such as circuits containing additions and a single multiplication [BGN05], or of logarithmic depth [SYY99], until finally in 2009, Gentry established a breakthrough result by giving the first fully homomorphic encryption scheme [Gen09b]. Follow-up work showed that FHE could be simplified [DGHV10], and based on standard assumptions, such as *learning with errors* [BV11]. The advent of FHE has unleashed a series of far-reaching consequences, such as delegating computations in a cloud architecture, and functional encryption [GKP⁺13]. For a survey on fully homomorphic encryption, see [Vai11].

Quantum cryptography is the study of cryptography in light of quantum information. One branch of quantum cryptography revisits *classical* primitives in the light of quantum information, establishing either no-go results (*e.g.* [LC97, May97]), or qualitative improvements achieved with quantum information (*e.g.* [BB84]). Another branch of quantum cryptography seeks to establish *quantum* cryptographic functionality, for instance in multiparty quantum computation [BOCG⁺06] or quantum message authentication [BCG⁺02]. The study of quantum cryptography is notorious for its subtleties and challenges, ranging from dealing with “purification attacks” [LC97, May97] to dealing with situations that are unique to the quantum world (such as “quantum rewinding” [Wat06, Unr12]).

A number of works have studied the cryptographic implications of the secure delegation of quantum computation, including: Childs [Chi05]; Broadbent, Fitzsimons and Kashefi [BFK09]; Aharonov, Ben-Or and Eban [ABOE10]; Vedran, Fitzsimons, Portmann and Renner [VFPR14]; Broadbent, Gutoski and Stebila [BGS13]; Fisher *et al.* [FBS⁺14]; and Broadbent [Bro15]. None of these works, however directly address the question of quantum homomorphic encryption, since they are interactive schemes, and the work of the client is proportional to the size of the circuit being evaluated (and thus, they do not satisfy the *compactness* requirement of fully homomorphic encryption, even if we allow interaction). Non-interactive approaches are given by Arrighi and Salvail [AS06], Rohde, Fitzsimons and Gilchrist [RFG12] and Tan, Kettlewell, Ouyang, Chen and Fitzsimons [TKO⁺14]. However, none of these approaches are applicable to universal circuit families. Furthermore, in the case of [AS06], security is given only in terms of cheat sensitivity, while both [RFG12] and [TKO⁺14] only bound the leakage of their encoding schemes.

Recent work by Yu, Pérez-Delgado and Fitzsimons [YPDF14] examines the question of perfect security and correctness for quantum fully homomorphic encryption (QFHE), concluding that the trivial scheme is optimal in this context. In light of this result, it is natural to consider computational assumptions in achieving QFHE. Indeed, the question of computationally secure QFHE remains an open problem; our contribution makes progress in this direction by presenting the first schemes that are homomorphic for a large class of quantum circuits.

1.1 Summary of Contributions and Techniques

We introduce schemes for *quantum homomorphic encryption (QHE)*, the quantum version of homomorphic encryption; we are thus interested in establishing functionality for the evaluation of *quantum* circuits on encrypted *quantum* data. In terms of definitions, we contribute by giving the

first definition of quantum homomorphic encryption (QHE) in the computational setting, in the case of both public-key and symmetric-key cryptosystems. As a consequence, we give the first formal definition (and scheme) for the public-key encryption of quantum information, where security is given in terms of *quantum indistinguishability under chosen plaintext attacks*—for which we show the equivalence of a number of definitions, including security for multiple messages. Prior work considered the computational setting for quantum encryption of classical plaintexts only [OTU00, Kos07, XY12].

In terms of QHE schemes, we start by using straightforward techniques to construct a scheme that is homomorphic for Clifford circuits (or, more generally, stabilizer circuits). This can be seen as an analogue to a classical scheme that is homomorphic for linear circuits (circuits performing only additions). While Clifford circuits are not universal for quantum computation, this already yields a range of applications for quantum information processing, including encoding and decoding into stabilizer codes. Our quantum public-key encryption scheme is a hybrid of a classical public-key fully homomorphic encryption scheme and the quantum one-time pad [AMTW00]. Intuitively, the scheme works by encrypting the quantum register with a quantum one-time pad, and then encrypting the one-time pad encryption keys with a classical public-key FHE scheme. Since Clifford circuits conjugate Pauli operators to Pauli operators, any Clifford circuit can be directly applied to the encrypted quantum register; the homomorphic property of the classical encryption scheme is used to update the encryption key. Of course, we specify that the classical FHE scheme should be secure against quantum adversaries. By using, *e.g.*, the scheme from [BV11], we get security based on the *learning with errors* (LWE) assumption [Reg05, Reg09]; this has been equated with worst-case hardness of “short vector problems” on arbitrary lattices [MR09], which is widely believed to be a quantum-safe (or “post-quantum”) assumption.

For universal quantum computations, we must evaluate a non-Clifford gate, for which we choose the “T” gate (also known as “R” or “ $\pi/8$ ”). Applying the above principle we run into trouble, since $\text{TX}^a\text{Z}^b = \text{X}^a\text{Z}^{a\oplus b}\text{P}^a\text{T}$. That is, conditioned on the quantum one-time pad encryption key $a, b \in \{0, 1\}$, the output picks up an undesirable non-Pauli error. Our main contribution is to present two schemes, EPR and AUX, that deal with this situation in two different ways:

EPR: The main idea of EPR (named after the famous Einstein-Podolski-Rosen trio [EPR35]) is to use entangled quantum registers to enable corrections *within the circuit* at the time of decryption. This scheme is efficient for any quantum circuit, however, it fails to meet a requirement for fully homomorphic encryption called *compactness*, which requires that the complexity of the decryption procedure be independent of the evaluated circuit. More specifically, the complexity of the decryption procedure for EPR scales with the square of the number of T-gates. This gives an advantage over the trivial scheme whenever the number of T-gates in the evaluated circuit is less than the squareroot of the number of gates. (The *trivial* scheme consists of appending to the ciphertext a description of the circuit to be evaluated, and specifying that it should be applied as part of the decryption procedure.)

AUX: Compared to EPR, the scheme AUX takes a more proactive approach to performing the correction required for a T-gate: to do this, it uses a number of auxiliary qubits that are given as part of the evaluation key. Intuitively, these auxiliary qubits encode the required corrections. In order to ensure universality, a large number of possible corrections must be available — the length of the evaluation key is thus given by a polynomial of degree exponential in the circuit’s T-gate *depth*, yielding a homomorphic scheme that is efficient for quantum circuits with constant T-depth.

The two main schemes EPR and AUX are incomparable; for some circuits, EPR is more desirable, while for others, it is preferable to use AUX. The scheme EPR becomes less *compact* (and therefore less interesting, since it approaches the trivial scheme), as the *number* of T-gates increases, while the scheme AUX becomes inefficient (*extremely* rapidly) as the *depth* of T-gates increases.

Our results can be viewed as a quantum analogue of precursory results to classical fully homomorphic encryption, which established the homomorphic property of encryption schemes that tolerate a limited amount of operations. One difference is that, while these schemes started with the modest goal of just a *single* multiplication (the addition operation being “easy”), we have already allowed for at the very least a *constant* number, and, depending on the circuit, up to a polynomial number of “hard” operations, namely of T-gates.

Our schemes use the existence of classical FHE, although at the expense of a slightly more complicated exposition, a classical scheme that is homomorphic only for linear circuits would actually suffice. We see the relationship between our schemes and classical FHE as a strength of our result, via the following interpretation: classical FHE is sufficient to enable QHE for a large family of circuits, and perhaps by taking greater advantage of the *fully* homomorphic property of the classical scheme in some as yet unknown way, our ideas might be extended to larger classes of quantum circuits. With this in mind, and for ease of exposition, we use a classical fully homomorphic encryption scheme for all of our quantum homomorphic encryption schemes.

An additional contribution of ours is conceptual: in the context of quantum circuits, it had been known for some time now that the non-Clifford part of a quantum computation is the “difficult” one (this phenomena appears, *e.g.* in the context of quantum simulations [Got98], fault-tolerant quantum computation [BK05] and quantum secure function evaluation [DNS10, DNS12, BOCG⁺06]). This has motivated a series of theoretical work seeking to optimize quantum circuits in terms of their T-gate complexity [Sel13, KMM13]. In particular, Amy, Maslov, Mosca, and Roetteler [AMMR13] recently proposed T-depth as a cost function, the idea being to count the number of T-layers in a quantum circuit and optimize over this parameter. Our contribution adds to this understanding, showing that, in the context of quantum homomorphic encryption, the main challenge is to evaluate non-Clifford gates, the bottleneck being, more precisely, the *depth* of the T-gate part of the circuit.

Organization. Some preliminaries and notation are given in Sec. 2. We give formal definitions of quantum homomorphic encryption and related concepts, including security definitions, in Sec. 3; this allows us to formally state our results in Sec. 4. Sec. 5 contains a basic quantum homomorphic encryption scheme, CL, for Clifford circuits that is used as a basis for EPR, the entanglement-based quantum homomorphic encryption scheme (Sec. 6), and for AUX, the auxiliary-qubit based quantum homomorphic encryption scheme (Sec. 7).

2 Preliminaries and Notation

2.1 Notation

A negligible function, $\eta(\cdot)$, is a function such that for every polynomial $p(\cdot)$, there exists an N such that for all integers $n > N$ it holds that $\eta(n) < \frac{1}{p(n)}$. As a convention, if a is a classical plaintext, we denote its encryption by \tilde{a} . Throughout this work we use κ to indicate the security parameter.

For a detailed and rigorous introduction to quantum information theory, we refer the reader to [Wat13]. In the remainder of this section, we give a brief overview of some of the necessary concepts, as well as our specific notation.

A *quantum register* is a quantum system, which we view as a physical object that stores quantum information. The contents of a quantum register are mathematically modelled as the set of trace-1,

positive semidefinite operators, called *density operators*, on \mathcal{X} , where \mathcal{X} is a complex Euclidean space. We denote the set of density operators on any space \mathcal{X} by $D(\mathcal{X})$.

Quantum registers are denoted with calligraphic typeset, such as \mathcal{X} , \mathcal{Y} . Two (or more) quantum systems, \mathcal{X} and \mathcal{Y} , form a composite system by the tensor product of the subsystems, $\mathcal{X} \otimes \mathcal{Y}$. If $\rho \in D(\mathcal{X} \otimes \mathcal{Y})$ is a state on the joint system, we write $\rho^{\mathcal{X}}$ to denote $\text{Tr}_{\mathcal{Y}}(\rho)$. If \mathcal{X} and \mathcal{Y} have the same dimension, we denote this by $\mathcal{X} \equiv \mathcal{Y}$.

The *trace distance* between two states, ρ and σ , is defined $\Delta(\rho, \sigma) := \text{Tr} \left(\sqrt{(\rho - \sigma)^\dagger (\rho - \sigma)} \right)$.

A density matrix that is diagonal in the computational basis corresponds to a classical random variable. For a random variable X on some set Σ_X , we define $\rho(X) := \sum_{x \in \Sigma_X} \text{Pr}[X = x] |x\rangle\langle x|$, the density matrix corresponding to X . A *classical-quantum* state is a state of the form $\rho^{\mathcal{M}\mathcal{A}} = \sum_x \text{Pr}[X = x] |x\rangle\langle x|^{\mathcal{M}} \otimes \rho_x^{\mathcal{A}}$.

One special quantum state on any system \mathcal{X} is the *completely mixed state*, $\frac{1}{\dim \mathcal{X}} \mathbb{I}_{\mathcal{X}}$, which we will sometimes denote by $\$$ (where \mathcal{X} should be implicit from the context). When \mathcal{X} is interpreted as \mathbb{C}^S for some finite set S , then $\$$ corresponds to the uniform distribution on S .

A *quantum channel* $\Phi : D(\mathcal{A}) \rightarrow D(\mathcal{B})$ refers to any physically-realizable mapping on quantum registers. The identity channel on register \mathcal{R} is denoted $\mathbb{I}_{\mathcal{R}}$. Let Φ be a quantum channel acting on register \mathcal{A} , and $\rho^{\mathcal{A}\mathcal{E}}$ a quantum system held in the joint registers $\mathcal{A} \otimes \mathcal{E}$. Then to simplify notation, when it is clear from the context, we write $\Phi(\rho^{\mathcal{A}\mathcal{E}})$ to mean $(\Phi \otimes \mathbb{I})(\rho^{\mathcal{A}\mathcal{E}})$.

We mention a special type of channel, a *conditional quantum channel*, which, on input the classical-quantum state $\sum_x \text{Pr}[x] |x\rangle\langle x|^{\mathcal{M}} \otimes \rho_x^{\mathcal{A}}$, outputs the quantum state:

$$\text{Tr}_{\mathcal{M}} \left(\sum_x \text{Pr}[x] |x\rangle\langle x|^{\mathcal{M}} \otimes \Phi_x(\rho_x^{\mathcal{A}}) \right)$$

for quantum channels $\Phi_x : D(\mathcal{A}) \rightarrow D(\mathcal{B})$.

Unless otherwise specified, a quantum *measurement* refers to a measurement in the computational basis. A quantum *algorithm* is a polynomial-time uniform family of quantum circuits, implementing a family of quantum channels.

2.2 Quantum Circuits

We work with the set of quantum gates consisting of single-qubit preparation in the $|0\rangle$ state, single-qubit measurements, as well as the gates the following unitary gates:

$$\begin{aligned} \mathbf{X} &= \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, & \mathbf{Z} &= \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, & \mathbf{P} &= \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}, & \mathbf{T} &= \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}, \\ \mathbf{H} &= \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, & \text{and} & & \mathbf{CNOT} &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \end{aligned}$$

The set $\{\mathbf{X}, \mathbf{Z}, \mathbf{P}, \mathbf{CNOT}, \mathbf{H}\}$ applied to arbitrary wires (redundantly) generates the Clifford group, and adding any non-Clifford gate, such as \mathbf{T} , gives a generating set for all quantum circuits. We note the following relations between these gates:

$$\mathbf{XZ} = -\mathbf{ZX}, \quad \mathbf{T}^2 = \mathbf{P}, \quad \mathbf{P}^2 = \mathbf{Z}, \quad \mathbf{HXH} = \mathbf{Z}, \quad \mathbf{TP} = \mathbf{PT}, \quad \mathbf{PZ} = \mathbf{ZP}.$$

A classical circuit is *layered* if it consists of alternating layers of either all ‘+’ gates or all ‘ \times ’ gates. The *multiplicative depth* of a layered circuit is the number of ‘ \times ’ layers. As we see in this

work, a natural quantum analogue of ‘+’ gates are Clifford group gates, while the analogue of the ‘×’ gate is the T-gate.¹ Thus, a *layered quantum circuit* consists of alternating layers of either all Clifford group gates, or T-gates. Then the T-depth of a layered quantum circuit is the number of such T layers [AMMR13].

2.3 Quantum One-time Pad

For a single-qubit system ρ in register \mathcal{R} , and $a, b \in \{0, 1\}$, we denote by $\text{QEnc}_{a,b} : \mathcal{R} \rightarrow \mathcal{R}$ the quantum one-time pad encryption and by $\text{QDec}_{a,b} : \mathcal{R} \rightarrow \mathcal{R}$ the quantum one-time pad decryption [AMTW00], namely:

$$\text{QEnc}_{a,b} : \rho \mapsto X^a Z^b \rho Z^b X^a \quad \text{and} \quad \text{QDec}_{a,b} : \rho \mapsto X^a Z^b \rho Z^b X^a. \quad (1)$$

It is easy to see that $\text{QDec}_{a,b} \circ \text{QEnc}_{a,b} = \mathbb{I}_{\mathcal{R}}$. By specifying that (a, b) be chosen uniformly at random, we get that the encryption maps any input to the completely mixed state (from the point of view of the adversary), since for all ρ ,

$$\frac{1}{4} \sum_{a,b} X^a Z^b \rho Z^b X^a = \frac{\mathbb{I}_2}{2}. \quad (2)$$

3 Definitions

In this section, we formally define QHE schemes and their properties. In Sec. 3.1, we first review classical FHE, and then define QHE in the public-key setting. Sec. 3.2 carefully defines the security of QHE, by considering two definitions for security under chosen plaintext attacks, and showing that they are equivalent. Sec. 3.3 defines correctness and compactness for QHE, culminating in a complete definition of quantum fully homomorphic encryption (Def. 3.8). Sec. 3.4 deals with an important subtlety that arises in the quantum case: due to the no-cloning theorem, when a large system is encrypted with some auxiliary quantum information needed for decryption, that auxiliary information cannot be copied and given to every subsystem, but rather, the system must now be decrypted as a whole, rather than subsystem-by-subsystem. We also define compactness and quasi-compactness in this context. Finally, for technical reasons, one of our schemes (AUX) must be used in the symmetric-key setting, which we define in Sec. 3.5. We do not address the issue of *circuit privacy* [GHV10], leaving this question for future work.

3.1 Classical and Quantum Homomorphic Encryption

Our schemes rely on a classical fully homomorphic encryption scheme; for completeness, we include a definition in App. A. Since our adversaries are modelled as being *quantum* polynomial-time, we need a further security guarantee on the classical scheme, namely that it is secure against *quantum* adversaries (see Def. 3.1). Fortunately, much of classical fully homomorphic encryption uses lattice-based cryptography, which exploits one of the few conjectured “quantum-safe” assumptions [MR09]. Among all known solutions, the scheme of [BV11] appears to be the best for our purposes, as it bases its security on the *learning with errors* (LWE) assumption [Reg05, Reg09], which has been equated with worst-case hardness of “short vector problems” on arbitrary lattices.

¹The analogy is due to the “easiness” of performing Clifford group computations on encrypted data, versus the “hardness” of performing non-Clifford group computations. Another way of seeing this is that the (reversible) quantum analogue of multiplication is the Toffoli gate: $|x, y, z\rangle \mapsto |x, y, x \cdot y \oplus z\rangle$. The Toffoli is a non-Clifford group gate that can be expressed in terms of T-gates [Sel13].

Definition 3.1 (q-IND-CPA). A classical homomorphic encryption scheme HE is q-IND-CPA secure if for any quantum polynomial-time adversary \mathcal{A} , there exists a negligible function η such that for $(pk, evk, sk) \leftarrow \text{HE.Keygen}(1^\kappa)$:

$$|\Pr[\mathcal{A}(pk, evk, \text{HE.Enc}_{pk}(0)) = 1] - \Pr[\mathcal{A}(pk, evk, \text{HE.Enc}_{pk}(1)) = 1]| \leq \eta(\kappa).$$

We note that a number of recent works examine the security of classical schemes against quantum superposition attacks [Zha12, BDF⁺11, BZ13]. In this context, our definition of q-IND-CPA above models security for classical plaintexts only (with an arbitrary learning phase, since the public key is given). Furthermore, we note that a classical homomorphic encryption scheme that is q-IND-CPA, is also IND-CPA. The converse, however, may not be true (in particular, if the IND-CPA property depends on a computational assumption that is broken by quantum computers). Note, however, that any proof that a scheme is IND-CPA can potentially be turned into a proof for q-IND-CPA if all statements still hold when “probabilistic polynomial-time adversary” is replaced by “quantum polynomial-time adversary” (see [Son14]).

We now give our new definitions for quantum homomorphic encryption. In our definitions, both pk , the public encryption key, and sk , the secret decryption key, are classical, whereas the evaluation key is allowed to be a quantum state. This choice is simply based on what is needed by our schemes.

Definition 3.2 (QHE). A quantum homomorphic encryption scheme is a 4-tuple of quantum algorithms (QHE.KeyGen, QHE.Enc, QHE.Eval, QHE.Dec):

Key Generation. QHE.KeyGen : $1^\kappa \rightarrow (pk, sk, \rho_{evk})$. This algorithm takes a unary representation of the security parameter as input and outputs a classical public encryption key pk , a classical secret decryption key sk and a quantum evaluation key $\rho_{evk} \in D(\mathcal{R}_{evk})$.

Encryption. QHE.Enc $_{pk}$: $D(\mathcal{M}) \rightarrow D(\mathcal{C})$. For every possible value of pk , the quantum channel Enc $_{pk}$ maps a state in the message space \mathcal{M} to a state (the cipherspace) in the cipherspace \mathcal{C} .

Homomorphic Evaluation. QHE.Eval $^{\mathcal{C}}$: $D(\mathcal{R}_{evk} \otimes \mathcal{C}^{\otimes n}) \rightarrow D(\mathcal{C}'^{\otimes m})$. For every quantum circuit \mathcal{C} , with induced channel $\Phi_{\mathcal{C}} : D(\mathcal{M}^{\otimes n}) \rightarrow D(\mathcal{M}^{\otimes m})$, we define a channel Eval $^{\mathcal{C}}$ that maps an n -fold cipherspace to an m -fold cipherspace, consuming the evaluation key in the process.²

Decryption. QHE.Dec $_{sk}$: $D(\mathcal{C}') \rightarrow D(\mathcal{M})$. For every possible value of sk , Dec $_{sk}$ is a quantum channel that maps the state in $D(\mathcal{C}')$ to a quantum state in $D(\mathcal{M})$.

3.2 Security of Quantum Homomorphic Encryption

We now define a notion of security for QHE analogous to the classical notion of indistinguishability under chosen plaintext attack. As in the classical case, there are several possible definitions, ranging from a relatively simple experiment (Def. 3.3) to multiple messages (Def. 3.4). As evidence of the robustness of these definitions, we show that they are equivalent; this strengthens our results since security in the most general case follows from security for the simplest definition. The proof of equivalence is similar to the classical case (see, e.g. [KL08]), and is included in App. B for completeness. We note that, by taking the evaluation key to be empty, our definitions and theorems are trivially applicable to the scenario of quantum public-key encryption (*i.e.* without a homomorphic property).

²Since we have not specified any requirement on the behaviour of this channel, we can define Eval $^{\mathcal{C}}$ to have some trivial behaviour on some, or even all quantum circuits \mathcal{C} . However, for the scheme to have the \mathcal{S} -homomorphic property (Def. 3.6), this cannot be the case for any circuit in \mathcal{S} .

CPA security. The CPA indistinguishability experiment is given below and illustrated in Fig. 1. The experiment interacts with an adversary \mathcal{A} , which is a pair of polynomial-time quantum algorithms $(\mathcal{A}_1, \mathcal{A}_2)$ (which we also refer to as adversaries). The first algorithm \mathcal{A}_1 implements a quantum channel $D(\mathcal{R}_{evk}) \rightarrow D(\mathcal{M} \otimes \mathcal{E})$ conditioned on pk , where \mathcal{E} is an arbitrary environment. The second algorithm \mathcal{A}_2 maps $D(\mathcal{C} \otimes \mathcal{E})$ to a bit.




Figure 1: The quantum CPA indistinguishability experiment.

The quantum CPA indistinguishability experiment $\text{PubK}_{\mathcal{A}, \text{QHE}}^{\text{cpa}}(\kappa)$

1. $\text{KeyGen}(1^\kappa)$ is run to obtain keys (pk, sk, ρ_{evk}) .
2. Adversary \mathcal{A}_1 is given (pk, ρ_{evk}) and outputs a quantum state on $\mathcal{M} \otimes \mathcal{E}$.
3. For $r \in \{0, 1\}$, let $\Xi_{\text{QHE}}^{\text{cpa}, r} : D(\mathcal{M}) \rightarrow D(\mathcal{C})$ be: $\Xi_{\text{QHE}}^{\text{cpa}, 0}(\rho) = \text{QHE.Enc}_{pk}(|\mathbf{0}\rangle\langle\mathbf{0}|)$ and $\Xi_{\text{QHE}}^{\text{cpa}, 1}(\rho) = \text{QHE.Enc}_{pk}(\rho)$. A random bit $r \in \{0, 1\}$ is chosen and $\Xi_{\text{QHE}}^{\text{cpa}, r}$ is applied to the state in \mathcal{M} (the output being a state in \mathcal{C}).
4. Adversary \mathcal{A}_2 obtains the system in $\mathcal{C} \otimes \mathcal{E}$ and outputs a bit r' .
5. The output of the experiment is defined to be 1 if $r' = r$ and 0 otherwise. In case $r = r'$, we say that \mathcal{A} wins the experiment.

Definition 3.3 (Quantum Indistinguishability under Chosen Plaintext Attack (q-IND-CPA)). A quantum homomorphic encryption scheme QHE is q-IND-CPA secure if for any quantum polynomial-time adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ there exists a negligible function η such that:

$$\Pr[\text{PubK}_{\mathcal{A}, \text{QHE}}^{\text{cpa}}(\kappa) = 1] \leq \frac{1}{2} + \eta(\kappa).$$

CPA-mult security. The CPA-mult indistinguishability experiment is similar to the CPA scenario above, but in this case the adversary chooses two t -tuples of messages, for any $t \geq 1$, and the challenger returns encryptions corresponding to one of the t -tuples. The adversary's task is then to guess which of the two t -tuples of messages has been encrypted. The experiment is given below; the illustration follows closely the one in Fig. 12 of App. B (but with single messages replaced by t -fold messages).

The quantum CPA-mult indistinguishability experiment $\text{PubK}_{\mathcal{A}, \text{QHE}}^{\text{cpa-mult}}(\kappa)$

1. $\text{KeyGen}(1^\kappa)$ is run to obtain keys (pk, sk, ρ_{evk}) .
2. For $r \in \{0, 1\}$, and $t \in O(\text{poly}(\kappa))$, let $\mathcal{M}_r = \mathcal{M}_r^1 \otimes \cdots \otimes \mathcal{M}_r^t$, where $\mathcal{M}_0^i \equiv \mathcal{M}_1^i \equiv \mathcal{M}$ (for all i). Adversary \mathcal{A}_1 is given (pk, ρ_{evk}) and outputs a quantum state ρ in $\mathcal{M}_0 \otimes \mathcal{M}_1 \otimes \mathcal{E}$.
3. For $r \in \{0, 1\}$, let $\Xi_{\text{QHE}}^{\text{cpa-mult}, r} : D(\mathcal{M}_0 \otimes \mathcal{M}_1) \rightarrow D(\mathcal{C}^1 \otimes \cdots \otimes \mathcal{C}^t)$ be given by $\Xi_{\text{QHE}}^{\text{cpa-mult}, 0}(\rho) = \text{Tr}_{\mathcal{M}_1}(\text{Enc}_{pk}^{\otimes t} \otimes \mathbb{I}_{\mathcal{M}_1})(\rho)$ and $\Xi_{\text{QHE}}^{\text{cpa-mult}, 1}(\rho) = \text{Tr}_{\mathcal{M}_0}(\mathbb{I}_{\mathcal{M}_0} \otimes \text{Enc}_{pk}^{\otimes t})(\rho)$. A random bit $r \in \{0, 1\}$ is chosen and $(\Xi_{\text{QHE}}^{\text{cpa-mult}, r} \otimes \mathbb{I}_{\mathcal{E}})$ is applied to ρ (the output being a state in $\mathcal{C}^{\otimes t} \otimes \mathcal{E}$).

4. Adversary \mathcal{A}_2 obtains the system in $\mathcal{C}^{\otimes t} \otimes \mathcal{E}$ and outputs a bit r' .
5. The output of the experiment is defined to be 1 if $r' = r$ and 0 otherwise. In case $r = r'$, we say that \mathcal{A} wins the experiment.

Definition 3.4 (Quantum Indistinguishability under Multiple Chosen Plaintext Attack). *A quantum homomorphic scheme QHE is q-IND-CPA-mult secure if for all quantum polynomial-time adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ there exists a negligible function η such that:*

$$\Pr[\text{PubK}_{\mathcal{A}, \text{QHE}}^{\text{cpa-mult}}(\kappa) = 1] \leq \frac{1}{2} + \eta(\kappa).$$

Theorem 3.5 (Equivalence of q-IND-CPA and q-IND-CPA-mult). *Let QHE be a quantum homomorphic encryption scheme. Then QHE is q-IND-CPA if and only if QHE is q-IND-CPA-mult.*

The proof of Thm. 3.5 is given in App. B.

3.3 Correctness and Compactness of Quantum Homomorphic Encryption

Next, we give, in Def. 3.6, a notion that encapsulates correctness of both encryption and evaluation, with respect to a class \mathcal{S} of quantum circuits (when \mathcal{S} is a strict subset of all computations, the literature sometimes refers to this as a “somewhat homomorphic” scheme). In the classical context, it is common to restrict attention to circuits that output a single bit, since any deterministic string can be computed bit-by-bit. We cannot do this quantumly, as a quantum state cannot be described, or generated, qubit-by-qubit. We therefore consider correctness as a global property of the output. Furthermore, as quantum data can be entangled, we require that a correct scheme preserve this entanglement and thus explicitly include an auxiliary space in the definition below.

Definition 3.6 (\mathcal{S} -homomorphic). *Let $\mathcal{S} = \{\mathcal{S}_\kappa\}_{\kappa \in \mathbb{N}}$ be a class of quantum circuits. A quantum encryption scheme QHE is \mathcal{S} -homomorphic (or homomorphic for the class \mathcal{S}) if for any sequence of circuits $\{\mathcal{C}_\kappa \in \mathcal{S}_\kappa\}_\kappa$ with induced channels $\Phi_{\mathcal{C}_\kappa} : \mathcal{M}^{\otimes n(\kappa)} \rightarrow \mathcal{M}^{\otimes m(\kappa)}$, and input $\rho \in D(\mathcal{M}^{\otimes n(\kappa)} \otimes \mathcal{E})$, there exists a negligible function η such that for $(pk, sk, \rho_{evk}) \leftarrow \text{QHE.Keygen}(1^\kappa)$:*

$$\Delta \left(\text{QHE.Dec}_{sk}^{\otimes m(\kappa)} \left(\text{QHE.Eval}^{\mathcal{C}_\kappa} \left(\rho_{evk}, \text{QHE.Enc}_{pk}^{\otimes n(\kappa)}(\rho) \right) \right), \Phi_{\mathcal{C}_\kappa}(\rho) \right) = \eta(\kappa). \quad (3)$$

We point out two properties of the above definition. First, we do not require that ciphertexts be decryptable themselves, only that they become decryptable after homomorphic evaluation, however, as long as QHE is homomorphic for the class of identity circuits, we can effectively decrypt a ciphertext by first homomorphically evaluating the identity. Second, we do not require that the output of QHE.Eval be able to undergo additional homomorphic evaluations; indeed, in the case that the evaluation key ρ_{evk} is quantum, it will in general be “consumed” by the QHE.Eval process, rendering any future applications of QHE.Eval impossible.

Analogously to the classical case, we define compactness (also parametrized by a class of circuits \mathcal{S}), which requires that the complexity of QHE.Dec be independent of the evaluated circuit, ruling out trivial quantum fully homomorphic encryption schemes where applying the circuit is delayed until after decryption (see the text following Def. 3.10 for an informal description of the trivial scheme, TRIV).

Definition 3.7 (\mathcal{S} -compactness). *Let $\mathcal{S} = \{\mathcal{S}_\kappa\}_{\kappa \in \mathbb{N}}$ be a class of quantum circuits. A quantum encryption scheme QHE is \mathcal{S} -compact if there exists a polynomial p such that for any sequence of circuits $\{\mathcal{C}_\kappa \in \mathcal{S}_\kappa\}_\kappa$, the circuit complexity of applying QHE.Dec to the output of QHE.Eval $^{\mathcal{C}_\kappa}$ is at most $p(\kappa)$. (That is, the circuit complexity of decryption does not depend on the circuit complexity of \mathcal{C}_κ).*

If QHE is \mathcal{S} -compact for \mathcal{S} the class of all quantum circuits over some universal gate set, then we simply say that QHE is compact.

Although this work leaves open the central problem of quantum fully homomorphic encryption, we have established all the machinery relevant for a formal definition, which we include below.

Definition 3.8 (Quantum Fully Homomorphic Encryption). *A scheme is a quantum fully homomorphic encryption scheme if it is both compact and homomorphic for the class of all quantum circuits over some universal gate set.*

3.4 Indivisible Schemes

In general, a quantum system is not equal to the sum of its parts. Because of this, for one of our schemes (as given in Sec. 6), it is convenient (if not necessary, by the no-cloning theorem [WZ82]) to define the output of QHE.Eval as containing, in addition to a series of cipherstates corresponding to each qubit, some auxiliary quantum register, possibly entangled with each cipherstate. Then the decryption operation, QHE.Dec must operate on the entire quantum system, rather than qubit-by-qubit. This is in contrast to a classical scheme, in which we could make a copy of the auxiliary register for each encrypted bit, enabling the decryption of individual bits, without decrypting the entire system.

Definition 3.9. *An indivisible quantum homomorphic encryption scheme is a quantum homomorphic encryption scheme with QHE.Eval and QHE.Dec re-defined as:*

Homomorphic Evaluation. $\text{QHE.Eval}^C : D(\mathcal{R}_{\text{evk}} \otimes \mathcal{C}^{\otimes n}) \rightarrow D(\mathcal{R}_{\text{aux}} \otimes \mathcal{C}'^{\otimes m})$. *Compared to QHE.Eval in a standard QHE, this algorithm outputs an additional auxiliary quantum register \mathcal{R}_{aux} . This extra information is used in the decryption phase. Since the state of \mathcal{R}_{aux} may be entangled with the state of each \mathcal{C}' , the system in $\mathcal{R}_{\text{aux}} \otimes \mathcal{C}'^{\otimes m}$ can no longer be considered subsystem-by-subsystem.*

Decryption. $\text{QHE.Dec}_{sk} : D(\mathcal{R}_{\text{aux}} \otimes \mathcal{C}'^{\otimes m}) \rightarrow D(\mathcal{M}^{\otimes m})$. *For every possible value of sk , Dec_{sk} is a quantum channel that maps an auxiliary register, together with an m -fold cipherstate, to an m -fold message in $D(\mathcal{M}^{\otimes m})$.*

We need to define compactness for an indivisible scheme (recall that here, there is no notion of separating the individual output systems).

Definition 3.10 (\mathcal{S} -compactness for an indivisible scheme). *Fix a class of quantum circuits, $\mathcal{S} = \{\mathcal{S}_\kappa\}_{\kappa \in \mathbb{N}}$. An indivisible quantum homomorphic encryption scheme QHE is \mathcal{S} -compact if there exists a polynomial p such that for any sequence of circuits $\{\mathcal{C}_\kappa \in \mathcal{S}_\kappa\}_\kappa$ with induced channels $\Phi_{\mathcal{C}_\kappa} : \mathcal{M}^{\otimes n(\kappa)} \rightarrow \mathcal{M}^{\otimes m(\kappa)}$, the circuit complexity of applying $\text{QHE.Dec}^{\otimes m(\kappa)}$ to the output of $\text{QHE.Eval}^{\mathcal{C}_\kappa}$ is at most $p(\kappa, m(\kappa))$. (That is, the circuit complexity of decryption does not depend on the circuit complexity of \mathcal{C}_κ).*

The trivial quantum fully homomorphic encryption scheme, TRIV, is easily phrased as an indivisible scheme. Informally, TRIV is the following:

1. The algorithms TRIV.KeyGen and TRIV.Enc are taken from any quantum public-key encryption scheme.
2. The algorithm TRIV.Eval simply sets \mathcal{R}_{aux} to be the target circuit, \mathcal{C} , and otherwise outputs the cipherstates corresponding to the encrypted inputs.

3. The algorithm TRIV.Dec first decrypts the cipherstates, then applies C and outputs the result.

Clearly, TRIV is homomorphic, but it is not compact, since TRIV.Dec must evaluate the quantum circuit C , and so its complexity scales with $G(C)$, the number of gates in C .

Although a decryption procedure with any dependence on G , or any other property of C , is not compact, it is still interesting to consider schemes whose decryption procedure has complexity that scales sublinearly in G (such schemes are called *quasi-compact* schemes [Gen09a]). We give a formal definition that quantifies this notion for indivisible quantum homomorphic encryption schemes.

Definition 3.11 (quasi-compactness). *Let $\mathcal{S} = \{\mathcal{S}_\kappa\}_\kappa$ be the set of all quantum circuits over some fixed universal gate set. Let $f : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ be some function on the circuits in \mathcal{S} . An indivisible quantum homomorphic encryption scheme QHE is f -quasi-compact if there exists a polynomial p such that for any sequence of circuits $\{C_\kappa \in \mathcal{S}_\kappa\}_\kappa$ with induced channels $\Phi_{C_\kappa} : \mathcal{M}^{\otimes n(\kappa)} \rightarrow \mathcal{M}^{\otimes m(\kappa)}$, the circuit complexity of decrypting the output of $\text{QHE.Eval}^{C_\kappa}$ is at most $f(C_\kappa)p(\kappa, m(\kappa))$.*

This definition allows us to consider schemes whose decryption complexity scales with some property of the evaluated circuit. We consider such a scaling non-trivial when it is smaller than $G(C)$, the number of gates in C .

3.5 Quantum Homomorphic Encryption in the Symmetric-Key Setting

We have defined quantum homomorphic encryption as a *public-key* encryption scheme. For technical reasons, our final scheme, AUX is given in the symmetric-key setting, so in this section we define functionality and security for symmetric-key quantum homomorphic encryption. In the case of classical fully homomorphic encryption, symmetric-key encryption is known to be *equivalent* to public-key encryption [Rot11]. In the quantum case, this is not known. This section also contains the definition of a *bounded* QHE scheme, which we again require for technical reasons in our symmetric-key scheme, AUX.

Definition 3.12. *A symmetric-key quantum homomorphic encryption scheme is a quantum homomorphic encryption scheme with QHE.KeyGen and QHE.Enc re-defined as:*

Key Generation. $\text{QHE.KeyGen} : 1^\kappa \rightarrow (sk, \rho_{evk})$. *This algorithm takes a unary representation of the security parameter as input and outputs a secret encryption/decryption key sk and a quantum evaluation key $\rho_{evk} \in D(\mathcal{R}_{evk})$.*

Encryption. $\text{QHE.Enc}_{sk} : D(\mathcal{M}) \rightarrow D(\mathcal{C})$. *For every possible value of sk , the quantum channel Dec_{sk} maps a state in the message space \mathcal{M} to a state (the cipherstate) in the cipherspace \mathcal{C} .*

Next, we define a quantum homomorphic encryption scheme that is *bounded* by n , which forces the number of ciphertexts encrypted by sk to be at most n . Furthermore, the scheme maintains a counter, d , of the number of previous encryptions, which can be thought of as allowing the scheme to avoid key reuse.

Definition 3.13. *A bounded symmetric-key quantum homomorphic encryption scheme is a symmetric-key quantum homomorphic encryption scheme with QHE.KeyGen, QHE.Enc, and QHE.Dec re-defined as:*

Key Generation. $\text{QHE.KeyGen} : (1^\kappa, 1^n) \rightarrow (sk, \rho_{evk})$.

Encryption. $\text{QHE.Enc}_{sk,d} : D(\mathcal{M}) \rightarrow D(\mathcal{C})$. *Every time $\text{QHE.Enc}_{sk,d}$ is called, the register containing d is incremented: $d \leftarrow d + 1$. If $d > n$, $\text{QHE.Enc}_{sk,d}$ outputs \perp , indicating an error.*

Decryption. $\text{QHE.Dec}_{sk,d} : D(\mathcal{C}') \rightarrow D(\mathcal{M})$.

Security of Symmetric Key Schemes. In order to define *indistinguishability under chosen plaintext attacks* in the symmetric-key setting, we must equip the adversary with an encryption oracle $\text{Enc}_{sk}(\cdot)$. An *adversary with access to an encryption oracle*, \mathcal{A} is a tuple of quantum channels $(\mathcal{A}^{(1)}, \dots, \mathcal{A}^{(q+1)})$, such that $\mathcal{A}^{(1)} : D(\mathcal{X}) \rightarrow D(\mathcal{M} \otimes \mathcal{E})$ for some space \mathcal{X} , for $i = 2, \dots, q$, $\mathcal{A}^{(i)} : D(\mathcal{C} \otimes \mathcal{E}) \rightarrow D(\mathcal{M} \otimes \mathcal{E})$, and $\mathcal{A}^{(q+1)} : D(\mathcal{C} \otimes \mathcal{E}) \rightarrow D(\mathcal{Y})$ for some space \mathcal{Y} . The interaction of the adversary and the encryption oracle is shown in Fig. 2, and for the case of a bounded encryption scheme, in which the oracle also updates a counter, in Fig. 3.




Figure 2: An adversary \mathcal{A} that makes at most q encryption oracle calls is a list of quantum channels $\mathcal{A}^{(1)}, \dots, \mathcal{A}^{(q+1)}$ such that for $j = 1, \dots, q$, $\mathcal{A}^{(j)}$ sends a message to an encryption oracle, and $\mathcal{A}^{(j+1)}$ receives the output. The full interaction is shown on the right, but we use the figure on the left as a short-hand for this interaction.




Figure 3: An adversary \mathcal{A} that makes at most q encryption oracle calls to a bounded encryption oracle is a list of quantum channels $\mathcal{A}^{(1)}, \dots, \mathcal{A}^{(q+1)}$ with the interaction shown on the right. We use the figure on the left as a short-hand for this interaction.

Just as in the public-key setting, we can define a quantum CPA indistinguishability experiment for the symmetric-key setting, $\text{SymK}_{\mathcal{A}, \text{QHE}}^{\text{cpa}}(\kappa)$. An adversary for $\text{SymK}_{\mathcal{A}, \text{QHE}}^{\text{cpa}}(\kappa)$ is a pair of adversaries with access to an encryption oracle $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2) = (\mathcal{A}_1^{(1)}, \dots, \mathcal{A}_1^{(q+1)}, \mathcal{A}_2^{(1)}, \dots, \mathcal{A}_2^{(q'+1)})$ (q is the number of oracle calls before the challenger is called, and q' is the number of oracle calls after the challenger is called). The experiment $\text{SymK}_{\mathcal{A}, \text{QHE}}^{\text{cpa}}(\kappa)$ is defined below, and shown in Fig. 4.

The quantum symmetric-key CPA indistinguishability experiment $\text{SymK}_{\mathcal{A}, \text{QHE}}^{\text{cpa}}(\kappa)$

1. $\text{KeyGen}(1^\kappa)$ is run to obtain keys (sk, ρ_{evk}) .
2. \mathcal{A}_1 is given ρ_{evk} , and may make a polynomial number of calls to an encryption oracle QHE.Enc_{sk} before outputting a quantum state in message space \mathcal{M} and environment register \mathcal{E} .
3. A random bit $r \in \{0, 1\}$ is chosen and $\Xi_{\text{QHE}}^{\text{cpa}, r}$ is applied to the state in \mathcal{M} (the output being a state in \mathcal{C}).
4. Adversary \mathcal{A}_2 obtains the system $\mathcal{C} \otimes \mathcal{E}$ and may make a polynomial number of calls to an encryption oracle QHE.Enc_{sk} before outputting a bit r' .

5. The output of the experiment is defined to be 1 if $r = r'$ and 0 otherwise. In case $r = r'$, we say that \mathcal{A} wins the experiment.




Figure 4: The quantum CPA experiment for symmetric-key systems (left) and bounded symmetric-key systems (right).

Definition 3.14 (Quantum Indistinguishability under Chosen Plaintext Attack (q-IND-CPA) for Symmetric Key Schemes). *A symmetric-key quantum homomorphic encryption scheme QHE is q -IND-CPA secure if for all quantum polynomial-time adversaries with oracle access, $\mathcal{A} = (\mathcal{A}_1^{(1)}, \dots, \mathcal{A}_1^{(q+1)}, \mathcal{A}_2^{(1)}, \dots, \mathcal{A}_2^{(q'+1)})$, there exists a negligible function η such that:*

$$\Pr[\text{SymK}_{\mathcal{A}, \text{QHE}}^{\text{cpa}}(\kappa) = 1] \leq \frac{1}{2} + \eta(\kappa).$$

Similar to the case of public-key encryption (Sec. 3.2), it is straightforward to give the seemingly stronger variant of q-IND-CPA, q -IND-CPA-mult, which is defined identically to the public-key case (Def. 3.4) but with an adversary having access to an encryption oracle. However, just as in the public-key case, it turns out that these definitions are equivalent.

Theorem 3.15 (Equivalence of q-IND-CPA and q-IND-CPA-mult in symmetric-key schemes). *Let QHE be a symmetric-key quantum homomorphic scheme. Then QHE is q -IND-CPA if and only if QHE is q -IND-CPA-mult.*

The proof of Thm. 3.15 is virtually identical to that of Thm. 3.5, given in App. B.

4 Main Contributions

We now formally state our main results (formal schemes and proofs are given in Sec. 5–7). Our first theorem, Thm. 4.1, establishes quantum homomorphic encryption for Clifford circuits.

Theorem 4.1. (Clifford scheme, CL) *Let \mathcal{S} be the class of Clifford circuits. Then assuming the existence of a classical fully homomorphic encryption scheme that is q -IND-CPA secure, there exists a quantum homomorphic encryption scheme that is q -IND-CPA, compact and \mathcal{S} -homomorphic.*

Next, we consider two variants of the scheme given by Thm. 4.1. Each variant deals with non-Clifford group gates — in our case T-gates — in a different way. The first scheme, described in Thm. 4.2 and formally defined in Sec. 6, uses entanglement to implement T-gates, resulting in a quantum homomorphic encryption scheme in which the complexity of decryption scales with the number of T-gates in the homomorphically evaluated circuit.

Theorem 4.2. (*entanglement-based scheme, EPR*) Let \mathcal{S} be the set of all quantum circuits over the universal gate set $\{X, Z, P, H, \text{CNOT}, T\}$ (as well as single-qubit preparation and measurement). Then assuming the existence of a classical fully homomorphic encryption scheme that is q -IND-CPA secure, there exists an indivisible quantum homomorphic encryption scheme that is q -IND-CPA, \mathcal{S} -homomorphic and R^2 -quasi-compact, where $R(C)$ is the number of T -gates in a circuit C .

The compactness of the scheme EPR is nontrivial for all circuits in which $R^2 \ll G$, where G is the number of gates.

Our second scheme, formally defined in Sec. 7, is based on the use of auxiliary qubits to implement T -gates, resulting in a quantum homomorphic encryption scheme that is homomorphic for circuits with constant T -depth, as described in the following theorem:

Theorem 4.3. (*auxiliary-qubit scheme, AUX*) Fix a constant L . Let \mathcal{S} be the set of quantum circuits over the universal gate set $\{X, Z, P, H, \text{CNOT}, T\}$ (as well as single-qubit preparation and measurement) with T -depth at most L . Then assuming the existence of a classical fully homomorphic encryption scheme that is q -IND-CPA secure, there exists a bounded symmetric-key quantum homomorphic encryption scheme that is q -IND-CPA, \mathcal{S} -homomorphic and compact.

The QHE scheme in Thm. 4.3 can be seen as somewhat analogous to an important building block in classical fully homomorphic encryption: a *levelled* fully homomorphic scheme, which is a scheme that takes a parameter L , which is an a-priori bound on the *depth* of the circuit that can be evaluated. However, we note that in contrast to a levelled fully homomorphic scheme, in which operations are polynomial in L , the complexity of our scheme is a polynomial of degree exponential in L , so we really require L to be constant.

As previously noted, Thm. 4.2 and 4.3 are complementary: the scheme EPR becomes less compact as the *number* of T -gates increases, while the scheme AUX becomes inefficient as the *depth* of T -gates increases.

5 Scheme CL: Homomorphic Encryption for Clifford Circuits

In this section, we present CL, a compact quantum homomorphic encryption scheme for stabilizer circuits, which consist of Clifford circuits combined with measurements and single-qubit preparation. This is a building block for the schemes that follow in Sec. 6 and 7. The main theorem we prove is Thm. 4.1, which follows directly from Thm. 5.1, 5.2 and 5.3.

By definition, Clifford circuits conjugate Pauli operators to Pauli operators [Got98]. In other words, for any Clifford C , and any Pauli, Q , there exists a Pauli Q' such that $CQ = Q'C$. Furthermore, applying a random Pauli operator is a perfectly secure symmetric-key quantum encryption scheme: the quantum one-time pad (see Sec. 2.3). Combining these observations, we see that it is possible to perform any Clifford circuit on quantum data that is encrypted using the quantum one-time pad. We can apply the desired Clifford, C , to the encrypted state $Q|\psi\rangle$ to get $Q'(C|\psi\rangle)$. Now decrypting the state requires applying the Pauli Q' . If Q can be described by the encryption key $(a_1, \dots, a_n, b_1, \dots, b_n)$ — that is, $Q = X^{a_1}Z^{b_1} \otimes \dots \otimes X^{a_n}Z^{b_n}$ — then Q' can be described by some key $(a'_1, \dots, a'_n, b'_1, \dots, b'_n)$ depending on C and $(a_1, \dots, a_n, b_1, \dots, b_n)$. We describe this dependence by a function $f^C : \mathbb{F}_2^{2n} \rightarrow \mathbb{F}_2^{2n}$, which we call a *key update rule*. We need only consider key update rules for each gate in our gate set, which consists of single-qubit measurement, single-qubit preparation, and the one- and two-qubit gates in $\{X, Z, P, \text{CNOT}, H\}$. For a single-qubit gate C , since the only keys that are affected are those corresponding to the wire to which C is applied, an update rule can be more succinctly described by a pair of functions $f_a^C, f_b^C : \mathbb{F}_2 \rightarrow \mathbb{F}_2$ such that when C is applied to the i^{th} wire, $a'_i = f_a^C(a_i, b_i)$ and $b'_i = f_b^C(a_i, b_i)$:

$$\mathcal{X}^{a_i} \mathcal{Z}^{b_i} |\psi\rangle \xrightarrow{\boxed{\text{C}}} \mathcal{X}^{a'_i} \mathcal{Z}^{b'_i} \text{C} |\psi\rangle \quad a_i \leftarrow a'_i = f_a^{\text{C}}(a_i, b_i), \quad b_i \leftarrow b'_i = f_b^{\text{C}}(a_i, b_i)$$

For the two-qubit CNOT-gate, the update rule is described by a 4-tuple of functions, since CNOT acts on two wires. We give the key update rules for all gates in App. C. By applying these rules after each gate, we can update the key so that the output is correctly decrypted. Such a technique was already used, *e.g.* in [Chi05, FBS⁺14, Bro15].

This solution, however, requires that the key updates be executed by the party holding the encryption keys: an “easy” classical computation, but nevertheless a computation that is polynomial in the *size* of the circuit. In the context of quantum homomorphic encryption, the challenge is therefore to allow the execution of *arbitrary* Clifford circuits, while maintaining the compactness condition. Here, we present a quantum public-key encryption scheme which is a hybrid of the quantum one-time pad and of a classical fully homomorphic encryption scheme. This encryption scheme is used to perform key updates on encrypted quantum one-time pad keys, enabling the computation of arbitrary Clifford group circuits on the encrypted quantum states, while maintaining the compactness condition. More precisely, to homomorphically evaluate a Clifford circuit consisting of a sequence of gates c_1, \dots, c_G , we apply the gates to the quantum one-time pad encrypted message, and homomorphically evaluate the function $f^{c_1} \circ \dots \circ f^{c_G}$ on the encrypted one-time pad keys $a_1, \dots, a_n, b_1, \dots, b_n$, where \circ denotes function composition. To accomplish this, we keep track of functions for each bit of the quantum one-time pad encryption key, $\{f_{a,i}, f_{b,i}\}_{i=1}^n$. Since each of the key update rules presented in App. C is linear, each $f_{a,i}$ and $f_{b,i}$ is a linear polynomial in $\mathbb{F}_2[a_1, \dots, a_n, b_1, \dots, b_n]$ (from the perspective of the evaluation procedure, $a_1, \dots, a_n, b_1, \dots, b_n$ are unknowns), so we refer to them as *key-polynomials*. Before we begin to evaluate the circuit, the key polynomials are the monomials $f_{a,i} = a_i$ and $f_{b,i} = b_i$. As we evaluate each gate c_j , we update the key-polynomials corresponding to the affected wires by composing them with the key update rules. To compute the new encrypted one-time pad keys once the circuit is complete, we homomorphically evaluate each key-polynomial on the old encrypted one-time pad keys. It is interesting to note that since the key update rules (App. C) for stabilizer circuit elements are all linear, for the scheme CL, the underlying classical fully homomorphic scheme only needs to be additively homomorphic.

We define our scheme CL as a QHE scheme. Here and throughout, we assume HE to be a classical fully homomorphic encryption scheme that is q-IND-CPA secure (see Def. 3.1 and App. A). As noted, such a scheme (based on the LWE assumption) could be derived from [BV11]. All of our schemes operate on qubit circuits, and encrypt qubit-by-qubit. Thus we fix $\mathcal{M} = \mathbb{C}^{\{0,1\}}$. Ciphertexts consist of quantum states in $\mathbb{C}^{\{0,1\}}$, combined with classical strings. Specifically, if \mathcal{C} is the output space of HE.Enc, and \mathcal{C}' is the output space of HE.Eval, then we define $\mathcal{C} = \mathbb{C}^{\mathcal{C} \times \mathcal{C}} \otimes \mathcal{X}$, where $\mathcal{X} \equiv \mathbb{C}^{\{0,1\}}$, and $\mathcal{C}' = \mathbb{C}^{\mathcal{C}' \times \mathcal{C}'} \otimes \mathcal{X}$.

Key Generation. CL.KeyGen(1^κ). For key generation, execute $(pk, sk, evk) \leftarrow \text{HE.Keygen}(1^\kappa)$. Output the obtained secret key, sk , and public key, pk . The evaluation key ρ_{evk} takes the value of the classical state $\rho(evk)$.

Encryption. CL.Enc $_{pk} : D(\mathcal{M}) \rightarrow D(\mathcal{C})$. Encryption is defined as the quantum channel that outputs the classical-quantum state:

$$\text{CL.Enc}_{pk}(\rho^{\mathcal{M}}) = \sum_{a,b \in \{0,1\}} \frac{1}{4} \rho(\text{HE.Enc}_{pk}(a), \text{HE.Enc}_{pk}(b)) \otimes \text{QEnc}_{a,b}(\rho^{\mathcal{M}}).$$

Homomorphic Evaluation. CL.Eval $^{\text{C}} : D(\mathcal{R}_{evk} \otimes \mathcal{C}^{\otimes n}) \rightarrow D(\mathcal{C}'^{\otimes m})$.

Suppose $C = c_1, \dots, c_G$ is a Clifford circuit. For every $j = 1, \dots, G$ such that c_j initializes a fresh qubit, we initialize a new qubit $\text{CL.Enc}_{pk}(|0\rangle\langle 0|)$ and append it to the system. Let $\rho \in D(\mathcal{X}_1 \otimes \dots \otimes \mathcal{X}_m)$, be the composite system consisting of the input quantum system and the initialized qubits.

1. For all $i \in [n]$, set $f_{a,i}, f_{b,i} \in \mathbb{F}_2[a_1, \dots, a_n, b_1, \dots, b_n]$ as $f_{a,i} \leftarrow a_i, f_{b,i} \leftarrow b_i$.
2. For $j = 1, \dots, G$ such that c_j is a gate or a measurement:
 - (a) Apply the gate c_j to the state: $\rho \leftarrow c_j \rho c_j^{-1}$.
 - (b) Compose the key update rules with the key-polynomials of the affected wires: if c_j is a single qubit gate or measurement acting on the i^{th} wire, update as $(f_{a,i}, f_{b,i}) \leftarrow (f_{a,i} \circ f_a^{c_j}, f_{b,i} \circ f_b^{c_j})$. Otherwise, if c_j is a CNOT-gate acting on wires i and i' , update $(f_{a,i}, f_{a,i'}, f_{b,i}, f_{b,i'})$ analogously.
3. Update the classical encryptions by computing

$$c_i = (\text{HE.Eval}_{evk}^{f_{a,i}}(\tilde{a}_i), \text{HE.Eval}_{evk}^{f_{b,i}}(\tilde{b}_i)).$$

4. Output (c_1, \dots, c_m, ρ) (with registers permuted to fit the prescribed form).

Decryption. $\text{CL.Dec}_{sk} : D(\mathcal{C}') \rightarrow D(\mathcal{M})$. For $\tilde{a}, \tilde{b} \in \mathcal{C}'$, the output space of HE.Eval , decryption is given by the conditional quantum channel:

$$\text{CL.Dec}_{sk} : |\tilde{a}\rangle\langle \tilde{a}| \otimes |\tilde{b}\rangle\langle \tilde{b}| \otimes \rho^{\mathcal{X}} \mapsto \text{QDec}_{\text{HE.Dec}_{sk}(\tilde{a}), \text{HE.Dec}_{sk}(\tilde{b})}(\rho^{\mathcal{X}}),$$

which can be implemented by first decoding the classical registers to obtain $a = \text{HE.Enc}_{sk}(\tilde{a})$ and $b = \text{HE.Enc}_{sk}(\tilde{b})$, applying $\text{QDec}_{a,b}$, and then tracing out $\mathbb{C}^{\mathcal{C}' \times \mathcal{C}'}$.

We have chosen to present CL.Enc_{pk} and CL.Dec_{sk} as quantum channels, since they are easily seen to be polynomial-time implementable. Note, however, that for more complicated quantum channels such as CL.Eval we will generally prefer their description in terms of a high-level algorithmic description.

5.1 Analysis of CL

We now analyse the various properties of CL.

Theorem 5.1. *Let \mathcal{S} be the class of Clifford circuits. Then CL is \mathcal{S} -homomorphic.*

Proof. This follows from the circuits in App. C, as well as the homomorphic property of HE. In particular, since the decrypted values of the ciphertexts are correct (except with exponentially small probability), then Equation (3) is satisfied. \square

Theorem 5.2. *CL is compact.*

Proof. Let p be a polynomial such that the complexity of applying HE.Dec to the output of HE.Eval is at most $p(\kappa)$ — such a polynomial exists by the compactness of HE. Then decrypting a single qubit of the output of CL.Eval has complexity at most $2p(\kappa) + 2$, since we must decrypt two keys a and b and then apply X^a and Z^b , so CL is also compact. \square

Theorem 5.3. *Assuming a classical fully homomorphic encryption scheme HE that is q -IND-CPA secure, the quantum homomorphic scheme CL is q -IND-CPA secure.*

Proof. The main part of this proof will be to show that the classical ciphertexts $\text{HE.Enc}_{pk}(a)$ and $\text{HE.Enc}_{pk}(b)$ give at most a negligible advantage. We will then see that without these classical ciphertexts, the quantum CPA Indistinguishability experiment is independent of r from the perspective of the adversary.

Let CL' be the quantum homomorphic encryption scheme with $\text{CL}'.\text{KeyGen} = \text{CL}.\text{KeyGen}$, $\text{CL}'.\text{Eval} = \text{CL}.\text{Eval}$, $\text{CL}'.\text{Dec} = \text{CL}.\text{Dec}$, and

$$\begin{aligned} \text{CL}'.\text{Enc}_{pk}(\rho) &= \sum_{a,b \in \{0,1\}} \frac{1}{4} \rho(\text{HE.Enc}_{pk}(0), \text{HE.Enc}_{pk}(0)) \otimes (X^a Z^b \rho Z^b X^a) \\ &= \rho(\text{HE.Enc}_{pk}(0), \text{HE.Enc}_{pk}(0)) \otimes \frac{1}{2} \mathbb{I}_2. \end{aligned}$$

Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an adversary for $\text{PubK}_{\mathcal{A}, \text{CL}}^{\text{cpa}}(\kappa)$. We will define an adversary $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2)$ for $\text{PubK}_{\mathcal{A}', \text{HE}}^{\text{cpa-mult}}(\kappa)$. Essentially, \mathcal{A}' will simulate $\text{PubK}_{\mathcal{A}, \text{CL}}^{\text{cpa-mult}}(\kappa)$, except that when it simulates $\Xi_{\text{CL}}^{\text{cpa},r}$, it will use $\Xi_{\text{HE}}^{\text{cpa-mult},s}$ in place of HE.Enc , so that it will actually be running either $\Xi_{\text{CL}}^{\text{cpa},r}$ (if $s = 1$) or $\Xi_{\text{CL}'}^{\text{cpa},r}$ (if $s = 0$) (see Fig. 5).




Figure 5: The new adversary \mathcal{A}' for $\text{PubK}_{\mathcal{A}', \text{HE}}^{\text{cpa-mult}}$, where Ψ is the channel that replaces the system with $|0\rangle\langle 0|$. Here $\$$ denotes the channel that outputs a completely mixed state, or equivalently, a uniform random variable. If $s = 1$, the middle dashed box is $\Xi_{\text{CL}}^{\text{cpa},r}$, and if $s = 0$, $\Xi_{\text{CL}'}^{\text{cpa},r}$.

$\mathcal{A}'_1(pk, evk)$: Run $\mathcal{A}_1(pk, evk)$ to get a state $\rho^{\mathcal{M}, \mathcal{E}}$. Choose a uniform random bit r . If $r = 0$, discard the \mathcal{M} subsystem and replace it with the state $|0\rangle\langle 0|$. Choose uniform random bits a and b , and apply $\text{QEnc}_{a,b}$, the quantum one-time pad, to \mathcal{M} , relabelling the resulting system by \mathcal{X} . Input (a, b) and $(0, 0)$ to $\Xi_{\text{HE}}^{\text{cpa-mult},s}$.

\mathcal{A}'_2 : Run \mathcal{A}_2 to get a bit r' . Output 1 if $r = r'$ and 0 otherwise.

We now compute the probability that \mathcal{A}' correctly guesses s , which we know must be at most $\frac{1}{2} + \eta(\kappa)$ for some negligible function, since HE is q-IND-CPA. If $s = 1$, then \mathcal{A}' is simulating $\text{PubK}_{\mathcal{A}, \text{CL}}^{\text{cpa}}$, so the probability that $r' = r$ (and thus that $s' = 1 = s$) is $\Pr[\text{PubK}_{\mathcal{A}, \text{CL}}^{\text{cpa}}(\kappa) = 1]$.

On the other hand, if $s = 0$, \mathcal{A}'_2 gets encryptions of 0 rather than $\text{HE.Enc}(a), \text{HE.Enc}(b)$, so \mathcal{A}' is simulating $\text{PubK}_{\mathcal{A}, \text{CL}'}^{\text{cpa}}$, so the probability that $r \neq r'$, and thus $s' = 0 = s$, is $\Pr[\text{PubK}_{\mathcal{A}, \text{CL}'}^{\text{cpa}}(\kappa) = 0]$.

Then since the total probability that $s = s'$ is at most $\frac{1}{2} + \eta(\kappa)$, we have:

$$\begin{aligned} \frac{1}{2} \Pr[\text{PubK}_{\mathcal{A}, \text{CL}}^{\text{cpa}}(\kappa) = 1] + \frac{1}{2} \Pr[\text{PubK}_{\mathcal{A}, \text{CL}'}^{\text{cpa}}(\kappa) = 0] &\leq \frac{1}{2} + \eta(\kappa) \\ \Pr[\text{PubK}_{\mathcal{A}, \text{CL}}^{\text{cpa}}(\kappa) = 1] + 1 - \Pr[\text{PubK}_{\mathcal{A}, \text{CL}'}^{\text{cpa}}(\kappa) = 1] &\leq 1 + 2\eta(\kappa) \\ \Pr[\text{PubK}_{\mathcal{A}, \text{CL}}^{\text{cpa}}(\kappa) = 1] - \Pr[\text{PubK}_{\mathcal{A}, \text{CL}'}^{\text{cpa}}(\kappa) = 1] &\leq 2\eta(\kappa). \end{aligned} \tag{4}$$

We complete the proof by noting that when $s = 0$, since $c = (\text{HE.Enc}_{pk}(0), \text{HE.Enc}_{pk}(0))$, it is independent of a, b (see Fig. 6).




Figure 6: When $s = 0$, \mathcal{A}' is simulating $\text{PubK}_{\mathcal{A}, \text{CL}'}^{\text{cpa}}$. In this case, c is independent of a, b , and so the only dependence on a, b is the quantum-one-time-pad encrypted message in \mathcal{X} .

Then from the perspective of \mathcal{A}_2 , since a, b is uniform random, the system \mathcal{X} just contains the completely mixed state $\$$ (see Fig. 7).




Figure 7: The circuit from Figure 6 is equivalent to the above circuit, in which the system in \mathcal{M} is replaced with the completely mixed state. Then from the perspective of \mathcal{A} , the experiment is independent of r .

Since the experiment $\text{PubK}_{\mathcal{A}, \text{CL}'}^{\text{cpa}}$ is independent of r from the perspective of \mathcal{A} , it follows that $\Pr[\text{PubK}_{\text{advA}, \text{CL}'}^{\text{cpa}}(\kappa) = 1] = \frac{1}{2}$. Combining this with Equation (4), we get

$$\Pr[\text{PubK}_{\mathcal{A}, \text{CL}'}^{\text{cpa}}(\kappa) = 1] \leq \frac{1}{2} + 2\eta(\kappa),$$

which completes the proof, since 2η is still a negligible function. \square

6 Scheme EPR: T-gate Computation Using Entanglement

In order to achieve universality for quantum circuits, we need to add a non-Clifford group gate, such as the T-gate. As noted in Sec. 1.1, if we apply the same technique as in Sec. 5 (*i.e.* to apply the T-gate on the encrypted quantum data) we run into a problem, since:

$$\text{T}X^a Z^b = X^a Z^{a \oplus b} \text{P} \text{T}. \quad (5)$$

That is, conditioned on a , the output picks up an undesirable P error, which cannot be corrected by applying Pauli corrections. In [Chi05], Childs arrives at the same conclusion, and makes the observation that, in the case where $a = 1$, the evaluation algorithm could be made to *correct* this erroneous P-gate by executing a correction (which consists of ZP). As long as the evaluation




Figure 8: Functionality of the T-gate gadget.




Figure 9: Evaluation protocol for the t^{th} T-gate, applied to the i^{th} wire. The key-polynomials $f_{a,i}$ and $f_{b,i}$ are in $\mathbb{F}_2[V]$. After the protocol, V gains a new variable corresponding to the unknown measurement result k_t . The dashed box shows part of the decryption procedure, which happens at some point in the future, after the complete evaluation is finished.

algorithm does not find out if this correction is being executed or not, security holds. The solution in [Chi05] involves quantum interaction; this was recently improved to a single auxiliary qubit, coupled with classical interaction [FBS⁺14, Bro15]. In this section, we base the evaluation of the T-gate on a modification of this technique, as presented in Fig. 8. The modification is that we allow the auxiliary qubit to be prepared in a state dependent on the X-encryption key, whereas [FBS⁺14, Bro15] explicitly avoids this since it requires the auxiliary qubits to be prepared independently of the computation. Correctness of Fig. 8 is proven in App. D.

As a proof technique (for establishing security), [FBS⁺14, Bro15] considers an equivalent, entanglement-based protocol. Here, we use the idea of exploiting entanglement in order to *delay* the correction required for the evaluation of the T-gate on encrypted data. The protocol is illustrated in Fig. 9.

Fig. 9 shows that, using the entangled state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, the conditional P correction can be delayed. The cost of this is that the value of the measurement result, k_t , on auxiliary register \mathcal{R}_t , is undetermined until later, when it is measured as part of the decryption algorithm. Thus we view the key updates as a symbolic computation: each time a T-gate is applied, an extra variable, k_t , is introduced.

For the first T-gate evaluation ($t = 1$), the evaluation procedure does not have the knowledge to evaluate $f_1 = f_{a,i}$, where i is the wire upon which the gate is performed, in order to perform the correction. It is possible (using the classical scheme HE), to compute a classical ciphertext \tilde{f}_1 that decrypts to $f_1(a_1, b_1, \dots, a_n, b_n)$. Thus, for this T-gate, the output part of the auxiliary system contains both f_1 and the register \mathcal{R}_1 . As part of the decryption operation, compute $f_1 \leftarrow \text{HE.Dec}(\tilde{f}_1)$, and apply P^{f_1} on \mathcal{R}_1 before measuring in the Hadamard basis and obtaining k_1 . From the point of view of the evaluation procedure, k_1 is unknown and so it becomes an *unknown* part of the encryption key (in contrast with the previous keys, which are also “unknown”, but to a lesser degree, since we have access to the classical encrypted values of these keys). The algorithm Eval

continues in this fashion for values of t up to R ; each time, the set of unknown variables increasing by one. Note that, according to Fig. 9, as well as the linearity of the key update rules, for all t , $f_t \in \mathbb{F}_2[a_1, \dots, a_n, b_1, \dots, b_n, k_1, \dots, k_{t-1}]$ is linear (since c is a known constant), so we can write $f_t = f_t^k + f_t^{ab}$ for $f_t^k \in \mathbb{F}_2[k_1, \dots, k_{t-1}]$ and $f_t^{ab} \in \mathbb{F}_2[a_1, \dots, a_n, b_1, \dots, b_n]$.

The cost of this construction is that each T-gate adds to the complexity of the decryption procedure, since, in particular, for each T-gate, we must perform a possible P-correction and a measurement on an auxiliary qubit. In addition, we cannot evaluate the key-polynomials, nor the f_t , until the variables k_t have been measured, so this evaluation must take place in the decryption phase, increasing the dependence on R , the number of T-gates, to $O(R^2)$. We make this dependence precise in Thm. 6.3.

We now formally define the indivisible quantum homomorphic encryption scheme, EPR. As in CL, we have message space $\mathcal{M} = \mathbb{C}^{\{0,1\}}$ and cipherspace $\mathcal{C} = \mathbb{C}^{C \times C} \otimes \mathcal{X}$, where C is the output space of HE.Enc and $\mathcal{X} \equiv \mathbb{C}^{\{0,1\}}$. Since EPR is indivisible, the output space of EPR.Eval^C has the form $\mathcal{R}_{aux} \otimes \mathcal{C}'^{\otimes m}$. We require an indivisible scheme, because decryption of any one of the output qubits requires access to the auxiliary system. In our case, we have $\mathcal{R}_{aux} = \mathcal{R}_1 \otimes \dots \otimes \mathcal{R}_R \otimes (\mathbb{C}^{\{0,1\}})^{\otimes R+1} \otimes \mathcal{R}' \otimes (\mathbb{C}^{C'})^{\otimes R}$, where R is the number of T-gates in C , C' is the output space of HE.Eval, and $\mathcal{R}_t \equiv \mathbb{C}^{\{0,1\}}$ for each t . The classical parts of the auxiliary space allow us to output R linear polynomials in $\mathbb{F}_2[k_1, \dots, k_R]$ corresponding to $\{f_t^k\}_{t=1}^R$, each of which can be represented with $R+1$ bits; as well as R HE.Eval outputs, corresponding to encryptions of $\{f_t^{ab}(a_1, \dots, a_n, b_1, \dots, b_n)\}_{t=1}^R$. Similarly, we have $\mathcal{C}' = (\mathbb{C}^{\{0,1\}})^{\otimes 2} \otimes \mathbb{C}^{C' \times C'} \otimes \mathcal{X}$.

Key Generation. EPR.KeyGen(1^κ). The key generation procedure is the same as CL.KeyGen(1^κ).

Encryption. EPR.Enc_{pk} : $D(\mathcal{M}) \rightarrow D(\mathcal{C})$. The encryption procedure is the same as CL.Enc_{pk}.

Evaluation. EPR.Eval_{evk}. As in CL, apply gates in $\{X, Z, P, H, \text{CNOT}\}$ directly on the encrypted quantum registers. For the T-gate, use the gadget defined in Fig. 9. This gadget differs from previous gadgets in that it uses an auxiliary Bell state, $|\Phi^+\rangle$. After the system of the i^{th} wire, \mathcal{X}_i , is measured, relabel half of the Bell state as \mathcal{X}_i , and the other half as \mathcal{R}_t , which is returned as part of \mathcal{R}_{aux} . The full evaluation procedure is as follows.

1. Set $V \leftarrow \{a_i, b_i\}_{i \in [n]}$, and $\forall i \in [n]$, set $f_{a,i}, f_{b,i} \in \mathbb{F}_2[V]$ as $f_{a,i} \leftarrow a_i, f_{b,i} \leftarrow b_i$.
2. Let $\mathbf{g}_1, \dots, \mathbf{g}_G$ be a topological ordering of the gates in C . For $j = 1, \dots, G$, evaluate \mathbf{g}_j using the appropriate gadget.
3. Let S be the set of output wires. Let \mathcal{L} be the set of labels $\mathcal{L} = \{(a, i), (b, i) : i \in S\} \cup \{1, \dots, R\}$. For each $\alpha \in \mathcal{L}$, we want to homomorphically evaluate f_α to obtain the actual (encrypted) key, but we can only actually evaluate the part of f_α that is in the variables $\{a_i, b_i\}_i$ — the $\{k_t\}_t$ are still unknown. Recall that we can write $f_\alpha = f_\alpha^k + f_\alpha^{ab}$ for $f_\alpha^k \in \mathbb{F}_2[k_1, \dots, k_R]$ and $f_\alpha^{ab} \in \mathbb{F}_2[a_1, \dots, a_n, b_1, \dots, b_n]$. Compute $\widetilde{f}_\alpha^{ab} \leftarrow \text{HE.Eval}_{evk}^{f_\alpha^{ab}}(\tilde{a}_1, \dots, \tilde{a}_n, \tilde{b}_1, \dots, \tilde{b}_n)$.
4. Output:
 - The $m = |S|$ qubit registers $\{\mathcal{X}_i : i \in S\}$ corresponding to the encrypted output of the circuit;
 - The R qubit registers $\mathcal{R}_1, \dots, \mathcal{R}_R$ corresponding to auxiliary states created by T-gadgets;
 - The polynomials $\{f_\alpha^k\}_{\alpha \in \mathcal{L}} \subset \mathbb{F}_2[k_1, \dots, k_R]$ and the homomorphically evaluated polynomials $\{\widetilde{f}_\alpha^{ab}\}_{\alpha \in \mathcal{L}}$.

Decryption. EPR.Dec_{sk}. In order to decrypt, measure the \mathcal{R}_t in order from 1 to R , computing $f_t(k_1, \dots, k_{t-1})$ as required. Formally:

1. For $t = 1, \dots, R$:
 - (a) Decrypt $f_t^{ab} \leftarrow \text{HE.Dec}_{sk}(\widetilde{f}_t^{ab})$.
 - (b) Compute $a \leftarrow f_t^k(k_1, \dots, k_{t-1}) \oplus f_t^{ab}$ and apply HP^a to \mathcal{R}_t .
 - (c) Measure \mathcal{R}_t to get k_t .
2. Let S be the set of indices of the output qubit registers. For $i \in S$:
 - (a) Decrypt $f_{a,i}^{ab} \leftarrow \text{HE.Dec}_{sk}(\widetilde{f}_{a,i}^{ab})$ and $f_{b,i}^{ab} \leftarrow \text{HE.Dec}_{sk}(\widetilde{f}_{b,i}^{ab})$.
 - (b) Compute $a_i \leftarrow f_{a,i}^k(k_1, \dots, k_t) \oplus f_{a,i}^{ab}$ and $b_i \leftarrow f_{b,i}^k(k_1, \dots, k_t) \oplus f_{b,i}^{ab}$.
3. To each register \mathcal{X}_i , apply the map QDec_{a_i, b_i} . Output registers $\mathcal{X}_1, \dots, \mathcal{X}_m$.

6.1 Analysis of EPR

We now analyse the various properties of EPR. Since the scheme EPR uses the same KeyGen and Enc procedures as CL, the following theorem follows from Thm. 5.3.

Theorem 6.1. *If HE is q -IND-CPA secure, then EPR is q -IND-CPA secure.*

The next theorem shows the homomorphic property for all circuits (recall that this property is independent of compactness).

Theorem 6.2. *Let \mathcal{S} be the class of all quantum circuits. Then EPR is \mathcal{S} -homomorphic.*

The proof follows from the circuits in App. C, Fig. 9, as well as the homomorphic property of HE.

Since the complexity of the decryption procedure depends on R , the number of T-gates in the circuit, it is clear that the scheme EPR is not compact. However, by analysing the circuit's dependence on R , we can see that for a very large class of quantum circuits, EPR is non-trivially quasi-compact. The following theorem is immediate from the decryption procedure.

Theorem 6.3. *Let p be a polynomial such that HE.Dec has complexity $O(p(\kappa))$. Then the decryption procedure EPR.Dec has complexity $O(R^2 + Rp(\kappa) + mp(\kappa) + mR)$.*

Thus, the dependence of the complexity of EPR.Dec on the evaluated circuit C is R^2 :

Corollary 6.4. *Let $R(C)$ denote the number of T-gates in a circuit C . Then EPR is R^2 -quasi-compact.*

This beats the compactness of the trivial scheme for all circuits C such that the number of T-gates is less than the squareroot of the number of gates; that is $R \ll \sqrt{G}$.

7 Scheme AUX: T-gate Computation Using Auxiliary States

In the previous QHE scheme, we solved the problem of performing the P correction (Eq. (5)) by *delaying* the correction via entanglement. In this section, we present a quantum homomorphic encryption scheme, AUX, that takes a more proactive approach to dealing with the P correction. At a high level, AUX can be understood as the following: as part of the evaluation key, AUX.Keygen outputs a number of auxiliary states. These states “encode” parts of the original encryption key, and are used to correct for the errors induced by the straightforward application of the T-gate on the cipherstates. In more details, the auxiliary states encode hidden versions of P corrections, such





Figure 10: A T-gadget for the scheme AUX consists of the above circuit and key-update rules. We use $\text{var}(k)$ to denote the set of variables in the polynomial k , which depends on the construction of the auxiliary state $|+_{f_{a,i},k}\rangle$, described below.

as $|+_{a,k}\rangle := Z^k P^a |+\rangle$ (where k is a random bit and a is an encryption key) that are useful for the evaluation of the T-gate (see Fig. 10). In general (after having applied prior gates), the exact auxiliary state will not be available; instead, the Eval procedure combines a number of auxiliary states in order to create a single copy of a state that is useful for performing the correction. This combination operation, however, is expensive as it introduces new unknowns (in terms of new variables as well as “cross-terms”), that need to be corrected in any future T-gate. Thus the size of the evaluation key grows rapidly, as a polynomial whose degree is exponential in the T-depth. We can thus tolerate only a *constant* T-gate depth for this scheme to be efficient.

We further specify that AUX is a symmetric-key encryption scheme. This is because AUX.KeyGen generates auxiliary qubits that depend on the quantum one-time pad encryption keys. Also, KeyGen takes an extra parameter 1^n , where n is an upper bound on the total number of qubits that can be encrypted (AUX acts much like a classical one-time pad scheme that picks a fixed-length encryption key ahead of time). After this bound on the number of encryptions has been attained, no further qubits can be encrypted. We will suppose without loss of generality that a circuit being homomorphically evaluated is on n wires. Furthermore, the number and type of auxiliary qubits will depend on the T-depth of the circuit to be evaluated, L . The scheme will not be able to homomorphically evaluate circuits with T-depth greater than L . We will see that the number of required auxiliary states grows super-exponentially in L , so we will require that L be a constant. Fix a constant L . We will now define a scheme $\text{AUX} = \text{AUX}_L$ that is homomorphic for all circuits with T-depth at most L .

Auxiliary Qubit Construction. In general, providing the necessary auxiliary states for each T-gate would require advance knowledge of the key $f_{a,i}$ at the time a T-gate is applied to the i^{th} wire. Since this depends on both the circuit being applied and on the prior measurement results, we appear to be at an impasse. The key observation that allows us to continue with this approach is that, given auxiliary states $|+_{f_1,k_1}\rangle$ and $|+_{f_2,k_2}\rangle$, we can combine them to get $|+_{f_1 \oplus f_2, k}\rangle$, for some k , using the following circuit:



By iterating this procedure, given auxiliary states $|+_{f_1,k_1}\rangle, \dots, |+_{f_r,k_r}\rangle$, we can construct the auxiliary state $|+_{f_1 \oplus \dots \oplus f_r, k}\rangle$, where $k = \bigoplus_{i=1}^m k_i \oplus \bigoplus_{i=2}^r c_i f_i \oplus \bigoplus_{i=1}^r \bigoplus_{j=1}^{i-1} f_i f_j$ for known values c_i . Thus, if we give many initial auxiliary states of the form $\{|+_{a_i, k_{a,i}}\rangle, |+_{b_i, k_{b,i}}\rangle\}_i$ (with different keys for different copies), we can construct $|+_{f,k}\rangle$ for f a linear function of $\{a_i, b_i\}_{i \in [n]}$. However, using an auxiliary state $|+_{f_{a,i},k}\rangle$ to facilitate a T-gate on the i^{th} wire introduces the unknown k into $f_{b,i}$. In particular, suppose $f_{a,i} = \bigoplus_{j=1}^r t_j$ for some monomial terms $t_j \in \mathbb{F}_2[V]$. Then we will need to construct it from auxiliary states $|+_{t_1,k_1}\rangle, \dots, |+_{t_r,k_r}\rangle$, to get $|+_{f_{a,i},k}\rangle$ for $k = \bigoplus_{i=1}^m k_i \oplus \bigoplus_{i=2}^r c_i t_i \oplus \bigoplus_{i=1}^r \bigoplus_{j=1}^{i-1} t_i t_j$. Thus, after the T-gadget, the new keys $f'_{a,i}, f'_{b,i}$ are in unknowns $V \cup \{k_1, \dots, k_r\}$. Furthermore, because of the cross terms $t_i t_j$, the degree of the key-

polynomials increases, so we can no longer assume they are linear. Since we can't produce $|+_{f_1 f_2, k}\rangle$ from $|+_{f_1, k_1}\rangle$ and $|+_{f_2, k_2}\rangle$, we need to provide additional auxiliary states for every possible term. We discuss this more formally below and in Sec. 7.1.

Spaces. As in CL and EPR, we work with qubits: $\mathcal{M} \equiv \mathbb{C}^{\{0,1\}}$. In contrast to our previous schemes, the classical encryptions of quantum one-time pad keys is part of the evaluation key (for convenience only), so we have $\mathcal{C} \equiv \mathbb{C}^{\{0,1\}}$. However, after evaluation, the classical encryption of the new one-time pad keys is needed for decryption, so as in CL, we have $\mathcal{C}' \equiv \mathbb{C}^{C' \times C'} \otimes \mathcal{X}$, where C' is the output space of HE.Eval, and $\mathcal{X} \equiv \mathbb{C}^{\{0,1\}}$.

Key Generation. $\text{AUX.Keygen}(1^\kappa, 1^n)$. The evaluation key contains auxiliary states that allow each of L layers of T-gates to be implemented. Thus, for each layer, since every wire must have the possibility to implement a T-gate, for each wire, we need to be able to construct an auxiliary state $|+_{f_{a,i}, k}\rangle$ for some k . Since we can add auxiliary states, we can construct this auxiliary state if we have an auxiliary state for each term in $f_{a,i}$. Since $f_{a,i}$ depends on the circuit, which we do not know in advance, we need to provide an auxiliary state for every term that could possibly be in $f_{a,i}$ at the ℓ^{th} layer of T-gates, for $\ell = 1, \dots, L$.

We now define sets of monomials T_1, \dots, T_L such that the keys in the ℓ^{th} layer consist of sums of terms from T_ℓ (as proven in Lemma 7.1). Let $V_1 := \{a_i, b_i\}_{i \in [n]}$, and define $T_1 \subset \mathbb{F}_2[V_1]$ by

$$T_1 := \{a_1, \dots, a_n, b_1, \dots, b_n\}.$$

The monomials in T_1 represent the possible terms in the key-polynomials before the first layer of T-gates. Each of the up to n T-gates in the first layer requires a copy of each of $\{|+_{t, k_t^{(1)}}\rangle\}_{t \in T_1}$, with independent random keys for each, for a total of $n|T_1|$ auxiliary states. More generally, for the ℓ^{th} layer of T-gates, we let T_ℓ be the set of possible terms in the key-polynomials before applying the ℓ^{th} layer of T-gates. We can see from the T-gadget, as well as the construction for adding auxiliary states that the keys from the previous layer's auxiliary states, $\{k_{1,i}^{(\ell-1)}, \dots, k_{|T_{\ell-1}|,i}^{(\ell-1)}\}_{i=1}^n$, may now be variables in the key-polynomials, and that products of terms from the previous layer may now be terms in the key-polynomials of the current layer. (This is caused by auxiliary state addition. See Lemma 7.1 for details). Thus, for $\ell > 1$, we can define $T_\ell \subset \mathbb{F}_2[V_\ell]$, where

$$V_\ell := V_{\ell-1} \cup \left\{ k_{1,i}^{(\ell-1)}, \dots, k_{|T_{\ell-1}|,i}^{(\ell-1)} \right\}_{i=1}^n,$$

by

$$T_\ell := T_{\ell-1} \cup \{tt' : t, t' \in T_{\ell-1}, t \neq t'\} \cup \left\{ k_{1,i}^{(\ell-1)}, \dots, k_{|T_{\ell-1}|,i}^{(\ell-1)} \right\}_{i=1}^n.$$

We then provide each of the n wires with an auxiliary state for each term in T_ℓ , for $\ell = 1, \dots, L$. We now make this more precise.

To each T_ℓ , we associate a family of strings $\{s^{(\ell)}(x)\}_{x \in \{0,1\}^{V_\ell}}$ in $\{0,1\}^{T_\ell}$, defined so that for every $f \in T_\ell$, the f -entry of $s^{(\ell)}(x)$ is $s_f^{(\ell)}(x) = f(x)$. That is, $s^{(\ell)}(x)$ represents evaluating every monomial in T_ℓ at x . For instance, we have, for any strings $a, b \in \{0,1\}^n$, $s^{(1)}(a, b) = (a_1, \dots, a_n, b_1, \dots, b_n)$.

For any strings $s, k \in \{0,1\}^n$, define

$$\sigma(s, k) := \bigotimes_{i=1}^n |+_{s_i, k_i}\rangle \langle +_{s_i, k_i}|.$$

For any string s , let s^{*n} denote the concatenation of n copies of s . For any $a, b \in \{0,1\}^n$ and $k = (k^{(1)}, \dots, k^{(L)}) \in \{0,1\}^{n|T_1|} \times \dots \times \{0,1\}^{n|T_L|}$, define

$$\sigma_{aux}^{a,b,k} := \sigma(s^{(1)}(a, b)^{*n}, k^{(1)}) \otimes \dots \otimes \sigma(s^{(L)}(a, b, k^{(1)}, \dots, k^{(L-1)})^n, k^{(L)}).$$

We can now define the procedure $\text{AUX.KeyGen}(1^\kappa, 1^n)$:

1. Execute $(pk, sk, evk) \leftarrow \text{HE.KeyGen}(1^{\kappa+n})$.
2. Choose uniform random $a, b \in \{0, 1\}^n$ and $k = (k^{(1)}, \dots, k^{(L)}) \in \{0, 1\}^{n|T_1|} \times \dots \times \{0, 1\}^{n|T_L|}$.
3. Output secret key (sk, a, b, k) .
4. Output evaluation key: $pk, evk, \tilde{a}_1 = \text{HE.Enc}_{pk}(a_1), \dots, \tilde{a}_n = \text{HE.Enc}_{pk}(a_n)$,
 $\tilde{b}_1 = \text{HE.Enc}_{pk}(b_1), \dots, \tilde{b}_n = \text{HE.Enc}_{pk}(b_n)$, $\left(\tilde{k}_i^{(\ell)} = \text{HE.Enc}_{pk} \left(k_{j,i}^{(\ell)} \right) \right)_{\substack{\ell \in [L] \\ i \in [n] \\ j \in [|T_\ell|]}}$, and $\sigma_{aux}^{a,b,k}$.

Encryption. $\text{AUX.Enc}_{(sk,a,b,k),d} : D(\mathcal{M}) \rightarrow D(\mathcal{C})$. The encryption procedure takes an extra parameter d that keeps track of the number of qubits already encrypted (we assume d is initially 1 and not modified outside of AUX.Enc). If $d \leq n$, for a single-qubit register \mathcal{M} , it applies the quantum one-time pad channel $\text{QEnc}_{a_d,b_d} : D(\mathcal{M}) \rightarrow D(\mathcal{C})$. The output is the cipherstate in register \mathcal{C} ; the parameter d is updated as $d \leftarrow d + 1$. If $d > n$, then output \perp to indicate an error.

Decryption. $\text{AUX.Dec}_{(sk,a,b,k),d} : D(\mathcal{C}') \rightarrow D(\mathcal{M})$. The decryption is defined the same as CL.Dec_{sk} .

Homomorphic Evaluation. $\text{AUX.Eval}^{\mathcal{C}} : D(\mathcal{R}_{evk} \otimes \mathcal{C}^{\otimes n}) \rightarrow D(\mathcal{C}'^{\otimes m})$. For Clifford group gates, we apply the gadgets as in CL.Eval . For T-gates, we apply the gadget in Fig. 10. The full evaluation procedure is as follows:

1. Set $V \leftarrow \{a_i, b_i\}_{i \in [n]}$, and $\forall i \in [n]$, set $f_{a,i}, f_{b,i} \in \mathbb{F}_2[V]$ as $f_{a,i} \leftarrow a_i, f_{b,i} \leftarrow b_i$.
2. Let $\mathbf{g}_1, \dots, \mathbf{g}_G$ be a topological ordering of the gates in \mathcal{C} . For $i = 1, \dots, G$, evaluate \mathbf{g}_i using the appropriate gadget.
3. Let S be the set of output wire labels. For each $i \in S$:
 - (a) Homomorphically evaluate $f_{a,i}$ and $f_{b,i}$ to obtain updated (encrypted) keys: $\tilde{a}_i \leftarrow \text{HE.Eval}_{evk}^{f_{a,i}}(\tilde{v} : v \in V)$ and $\tilde{b}_i \leftarrow \text{HE.Eval}_{evk}^{f_{b,i}}(\tilde{v} : v \in V)$.
4. Output in \mathcal{C}'_i the classical-quantum system given by:
 - The encrypted keys $\{\tilde{a}_i, \tilde{b}_i\}_{i \in S}$.
 - The output register corresponding to the encrypted output qubit i of the circuit.

The correctness of this scheme depends on two facts, which we prove in Sec. 7.1. First, for every unknown $v \in V$, we have an encrypted copy of \tilde{v} , encrypted using HE.Enc . We need these to compute the final keys $\{\tilde{a}_i, \tilde{b}_i\}$ using $f_{a,i}, f_{b,i} \in \mathbb{F}_2[V]$. Finally, for each level ℓ , for each wire label i , we need an auxiliary state $|+_{t,k}\rangle$ for every term that may appear in the key $f_{a,i}$ going into the ℓ^{th} level. This allows us to construct the auxiliary qubit required to execute each T-gadget.

We remark that if we only had a classical encryption scheme that was homomorphic over linear circuits, and not fully homomorphic, then we could get the same functionality from a slightly modified version of this scheme, in which we include with every auxiliary qubit $|+_{s,k}\rangle\langle +_{s,k}|, \text{HE.Enc}_{pk}(s)$ — at the moment we only include some of these, but not those auxiliary states arising from *products* of terms, since we can compute products homomorphically. Since we have classical fully homomorphic encryption, we use this to slightly simplify the scheme, however the observation that the fully homomorphic property is not fully taken advantage of strengthens the idea that Clifford circuits are analogous to classical linear circuits in the context of QHE.

7.1 Analysis of AUX

We now analyse the various properties of AUX. Consider a layered quantum circuit C with L layers of T-gates. To simplify the analysis, we assume that the ordering of gates g_1, \dots, g_G has the property that if g_i is a T-gate in level ℓ , and g_j is a T-gate in level $\ell + 1$, then $i < j$; that is, we completely evaluate level ℓ before we begin to evaluate level $\ell + 1$.

Lemma 7.1. *Let $f_{a,i}$ be a key-polynomial going into the ℓ^{th} layer of T-gates. Then $f_{a,i}$ is a sum of terms in T_ℓ .*

Proof. We prove this statement by induction on ℓ . Before any gates have been applied, the key-polynomial are $f_{a,i} = a_i$ and $f_{b,i} = b_i$ for $i = 1, \dots, n$. We can easily see from the update rules that applying Clifford gates results in keys of the form f or $f + f'$, where f and f' were previous keys. Thus, after a Clifford circuit has been applied, all key-polynomial are sums of terms from $\{a_1, \dots, a_n, b_1, \dots, b_n\} = T_1$.

Let $f_{a,1}, \dots, f_{a,n}, f_{b,1}, \dots, f_{b,n}$ be the key-polynomials going into the ℓ^{th} layer, and suppose they are sums of terms in T_ℓ . Let $f'_{a,1}, \dots, f'_{a,n}, f'_{b,1}, \dots, f'_{b,n}$ be the key-polynomials right after the ℓ^{th} layer of T-gates has been applied. If no T is applied on the i^{th} wire, then $f'_{a,i} = f_{a,i}$ and $f'_{b,i} = f_{b,i}$, so $f'_{a,i}, f'_{b,i}$ are both sums of terms in $T_\ell \subset T_{\ell+1}$. Suppose on the other hand that we apply a T-gate to the i^{th} wire at level ℓ . From the T-gadget (Fig. 10), we see that after applying a T to the i^{th} wire, we have new keys $f'_{a,i} = f_{a,i} \oplus c$ for a known constant c , so $f'_{a,i}$ is a sum of terms in $T_\ell \subset T_{\ell+1}$; and $f'_{b,i} = (1 \oplus c)f_{a,i} \oplus f_{b,i} \oplus k$, where k is the auxiliary state key of the auxiliary state used to implement the gadget. If $f_{a,i} = t_1 \oplus \dots \oplus t_r$, for $t_1, \dots, t_r \in T_\ell$, then we construct $|+_{f_{a,i},k}\rangle$ from auxiliary states $|+_{t_1,k_1}\rangle, \dots, |+_{t_r,k_r}\rangle$ for some $k_1, \dots, k_r \in \{k_{q,i}^{(\ell)}\}_{q=1}^{|T_\ell|} \subset T_{\ell+1}$, so we have $k = \bigoplus_{j=1}^r k_j \oplus \bigoplus_{j=2}^r c_j t_j \oplus \bigoplus_{j=1}^r \bigoplus_{j'=1}^{j-1} t_j t_{j'}$ for known c_2, \dots, c_r , which is the sum of terms in $T_{\ell+1}$, since $t_1, \dots, t_r \in T_\ell$. Thus, $f'_{b,i}$ is the sum of terms in $T_{\ell+1}$.

Thus, after applying the ℓ^{th} layer of T-gates, all key-polynomials are sums of terms from $T_{\ell+1}$. To complete the proof, we simply observe again that Clifford circuits act additively on the keys, and so do not introduce new terms, so just before the $(\ell + 1)^{\text{th}}$ layer of T-gates, the key-polynomials are still sums of terms in $T_{\ell+1}$. \square

The bottleneck in this scheme is the number of auxiliary states required:

Lemma 7.2. *The number of auxiliary qubits output by $\text{AUX.KeyGen}(1^\kappa, 1^n)$ grows as $O(n^{2^{\ell-1}+1})$ in n .*

Proof. The number of qubits encoded in $\sigma_{aux}^{a,b,k}$ is

$$|k^{(1)}| + |k^{(2)}| + \dots + |k^{(L)}| = n|T_1| + n|T_2| + \dots + n|T_L| = n \sum_{\ell=1}^L |T_\ell|.$$

From the definition of T_ℓ , we see that:

$$|T_1| = 2n, \quad \text{and for } \ell > 1, \quad |T_\ell| = |T_{\ell-1}| + \binom{|T_{\ell-1}|}{2} + n|T_{\ell-1}|.$$

So certainly for all $\ell > 1$, $|T_\ell| \leq c|T_{\ell-1}|^2$ for some constant c , and thus $|T_\ell| \leq c^{\ell-1}(2n)^{2^{\ell-1}} \in O(n^{2^{\ell-1}})$. Thus $n \sum_{\ell=1}^L |T_\ell| \in O(n^{2^{\ell-1}+1})$. \square

We thus have the following theorem:

Theorem 7.3. *Let \mathcal{S}_n be the class of all quantum circuits on n wires with T-depth at most L , and let $\mathcal{S} = \{\mathcal{S}_n\}_{n \in \mathbb{N}}$. Then AUX is \mathcal{S} -homomorphic and compact.*

We now consider the security of the scheme.

Theorem 7.4. *If HE is q -IND-CPA secure, then AUX is q -IND-CPA secure.*

We will prove Thm. 7.4 in several parts. To begin, we will show that an adversary that interacts with AUX.KeyGen can't do much better than an adversary that interacts instead with an altered version of AUX.KeyGen, KeyGen', in which every classical encryption has been replaced with HE.Enc_{pk}(0) (Lemma 7.5). Then we will be able to complete the proof by showing that an adversary interacting with KeyGen' instead of AUX.KeyGen can't win the q -IND-CPA experiment for AUX with probability better than $\frac{1}{2}$.

Lemma 7.5. *Define a QHE scheme AUX' such that AUX'.KeyGen($1^\kappa, 1^n$) = KeyGen'($1^\kappa, 1^n$), where KeyGen' behaves identically to AUX.KeyGen, except it replaces every classical encryption HE.Enc_{pk}(x) with HE.Enc_{pk}(0). Let AUX'.Enc = AUX.Enc, AUX'.Dec = AUX.Dec and AUX'.Eval = AUX.Eval. Then for any quantum polynomial-time adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ with encryption oracle access, there exists a negligible function η such that:*

$$\Pr[\text{SymK}_{\mathcal{A}, \text{AUX}}^{\text{cpa}}(\kappa) = 1] - \Pr[\text{SymK}_{\mathcal{A}, \text{AUX}'}^{\text{cpa}}(\kappa) = 1] \leq \eta(\kappa).$$

Thus, we can restrict our attention to adversaries that make no use of the classical encryptions, since they add at most a negligible advantage.

Proof. We will define an adversary $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2)$ for the quantum CPA-mult indistinguishability experiment for HE, PubK _{\mathcal{A}', HE} ^{cpa-mult}(κ). Essentially, \mathcal{A}' will run AUX.KeyGen, except it will use the challenger $\Xi_{\text{HE}}^{\text{cpa-mult}}$ in place of HE.Enc_{pk}, so that it is either running AUX.KeyGen or KeyGen'. It will then simulate the SymK experiment, and if \mathcal{A}' wins, it will guess that it ran the original version of AUX.KeyGen, and otherwise it will guess that it ran KeyGen'.

$\mathcal{A}'_1(pk, evk)$: \mathcal{A}'_1 chooses uniform random bit strings $a, b \in \{0, 1\}^n$ and $k \in \{0, 1\}^N$, where $N = n|T_1| + \dots + n|T_L|$, and gives $m_0 = \mathbf{0} = 0^{2n+N}$ and $m_1 = (a, b, k)$ to the challenger $\Xi_{\text{HE}}^{\text{cpa-mult}}$, which outputs either $c_1 = \text{HE.Enc}_{pk}(a, b, k)$, or $c_0 = \text{HE.Enc}_{pk}(\mathbf{0})$.

$\mathcal{A}'_2(c)$: \mathcal{A}'_2 computes $\sigma_{aux}^{a,b,k}$ and gives $\sigma_{aux}^{a,b,k}, pk, evk, c$ to \mathcal{A}'_1 . \mathcal{A}'_1 may make several oracle calls, which \mathcal{A}'_2 can simulate, because it has a, b and so can run AUX.Enc. When \mathcal{A}'_1 outputs a message to the challenger, \mathcal{A}'_2 samples a random bit r , and runs $\Xi_{\text{AUX}}^{\text{cpa}, r}$, which it can simulate, since it has a, b , and so can run AUX.Enc. \mathcal{A}'_2 then gives the challenge to \mathcal{A}'_1 , and if \mathcal{A}'_1 outputs r , \mathcal{A}'_2 outputs 1, and otherwise, \mathcal{A}'_2 outputs 0.

We now calculate the probability that \mathcal{A}' correctly guesses which of c_0 and c_1 it received from the challenger, which we know must be less than $\frac{1}{2} + \eta(\kappa + n)$ for some negligible function, since HE is q -IND-CPA, $\kappa + n$ is the security parameter given to HE.KeyGen, and $|m_0| = |m_1| = 2n + N = O(\text{poly}(n)) = O(\text{poly}(n + \kappa))$. If \mathcal{A}' received c_0 , then it acted as KeyGen', whereas if it received c_1 , it acted as AUX.KeyGen. In the former case, the probability that \mathcal{A}' correctly guesses 0 is the probability that \mathcal{A} loses the SymK experiment when it interacts with AUX', $\Pr[\text{SymK}_{\mathcal{A}, \text{AUX}'}^{\text{cpa}}(\kappa) = 0]$. In the latter case, the probability that \mathcal{A}' correctly guesses 1 is the probability that \mathcal{A} wins the

SymK experiment when it interacts with AUX, $\Pr[\text{SymK}_{\mathcal{A}, \text{AUX}}^{\text{cpa}}(\kappa) = 1]$. Thus, since HE is q-IND-CPA, there exists a negligible function η' such that

$$\begin{aligned} \frac{1}{2} \Pr[\text{SymK}_{\mathcal{A}, \text{AUX}'}^{\text{cpa}}(\kappa) = 0] + \frac{1}{2} \Pr[\text{SymK}_{\mathcal{A}, \text{AUX}}^{\text{cpa}}(\kappa) = 1] &\leq \frac{1}{2} + \eta'(\kappa) \\ 1 - \Pr[\text{SymK}_{\mathcal{A}, \text{AUX}'}^{\text{cpa}}(\kappa) = 1] + \Pr[\text{SymK}_{\mathcal{A}, \text{AUX}}^{\text{cpa}}(\kappa) = 1] &\leq 1 + 2\eta'(\kappa) \\ \Pr[\text{SymK}_{\mathcal{A}, \text{AUX}}^{\text{cpa}}(\kappa) = 1] - \Pr[\text{SymK}_{\mathcal{A}, \text{AUX}'}^{\text{cpa}}(\kappa) = 1] &\leq 2\eta'(\kappa). \end{aligned}$$

Setting $\eta = 2\eta'$ completes the proof. \square

The next lemma shows that the output of KeyGen' is actually $(pk, evk, \$)$, which is independent of a, b, k . The proof of Lemma 7.6 is mainly computational, and provides little insight, so we relegate it to App. E.

Lemma 7.6. *Let $N = n|T_1| + \dots + n|T_L|$. For any $a, b \in \{0, 1\}^n$, $\sum_{k \in \{0, 1\}^N} \sigma_{aux}^{a, b, k} = \frac{1}{2^N} \mathbb{I}_{2^N}$.*

To complete the proof of Thm. 7.4, we show that no adversary interacting with $\text{AUX.KeyGen}'$ can win the experiment SymK^{cpa} with probability better than $\frac{1}{2}$.

Lemma 7.7. *For any adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ with access to an encryption oracle,*

$$\Pr[\text{SymK}_{\mathcal{A}, \text{AUX}'}^{\text{cpa}}(\kappa) = 1] = \frac{1}{2}.$$

Proof. Let q be the number of oracle calls made by \mathcal{A}_1 , and write $\mathcal{A}_1 = (\mathcal{A}_1^{(1)}, \dots, \mathcal{A}_1^{(q+1)})$. Let q' be the number of oracle calls made by \mathcal{A}_2 , and write $\mathcal{A}_2 = (\mathcal{A}_2^{(1)}, \dots, \mathcal{A}_2^{(q'+1)})$. If $q \geq n$, then the challenger just outputs \perp , independent of r , so certainly in that case \mathcal{A} cannot win with probability more than $\frac{1}{2}$, so suppose $q < n$. If $q + q' + 1 > n$, then the last $q + q' + 1 - n$ oracle calls made by \mathcal{A}_2 simply return \perp , which \mathcal{A} could simulate without actually making these oracle calls, so suppose without loss of generality that $q + q' + 1 \leq n$.

The output of $\text{AUX}'.\text{KeyGen} = \text{KeyGen}'$ to $\mathcal{A}_1^{(1)}$ is $(pk, evk, \sigma_{aux}^{a, b, k})$, and by Lemma 7.6, for any a, b , this is equal to $(pk, evk, \frac{1}{2^N} \mathbb{I}_{2^N})$. Thus, the interaction of KeyGen' with the experiment is shown in part (a) of Fig. 11. KeyGen' chooses random bits $a_1, b_1, \dots, a_{q+q'+1}, b_{q+q'+1}$, for use in oracle calls and the challenge itself, but these are independent of the information given to \mathcal{A} by KeyGen' . (The other random bits selected by KeyGen' , $a_{q+q'+2}, b_{q+q'+2}, \dots, a_n, b_n$ and the string k , are independent of the interaction with the adversary, so we ignore them.)

It is then easy to see from Fig. 11 that every call to the encryption oracle can be replaced by a channel that discards the input and returns a completely mixed state, since for any input $\rho^{\mathcal{M}}$, the encryption oracle returns

$$\text{Tr}_1 \left(\frac{1}{4} \sum_{a, b \in \{0, 1\}} |a, b\rangle \langle a, b|_1 \otimes X^a Z^b \rho^{\mathcal{M}} Z^b X^a \right) = \frac{1}{4} \sum_{a, b \in \{0, 1\}} X^a Z^b \rho^{\mathcal{M}} Z^b X^a = \frac{1}{2} \mathbb{I}_2.$$




Figure 11: Proof of Lemma 7.7. Part (a) shows how KeyGen' interacts with the experiment. The channel $\$$ outputs a completely mixed state, or equivalently, a uniform random variable. Since the random bits a_i, b_i are independent of the other outputs of KeyGen' , for each i , we can replace each of the oracle calls as well as the challenger with a channel that discards the input and returns a completely mixed state, as shown in part (b). Thus, the experiment is independent of r from the perspective of \mathcal{A} , and so \mathcal{A} can do no better than guessing r .

In other words, we have:

$$\begin{array}{c} \$ \\ \text{---} a, b \\ \text{---} \mathcal{M} \text{---} \text{QEnc}_{a,b} \text{---} C \end{array} \equiv \begin{array}{c} \mathcal{M} \\ \text{---} \text{---} \text{---} \\ \text{---} \$ \text{---} C \end{array}$$

Here $\$$ denotes the channel that outputs a completely mixed state, or equivalently, a uniform random variable.

For the same reason, the call to the challenger $\Xi_{\text{AUX}}^{\text{cpa}, r}$ can also be replaced with the channel that discards the input and returns $\$$, since $\Xi_{\text{AUX}}^{\text{cpa}, r}$ applies a quantum one-time pad using random keys a_{q+1}, b_{q+1} to the input or to $|0\rangle\langle 0|$, and in either case, the resulting state is the completely mixed state. Thus, from the perspective of \mathcal{A} , the experiment is independent of r , as shown in part (b) of Fig. 11. Thus, an adversary cannot win with probability better than $\frac{1}{2}$. \square

Combining Lemma 7.5 and Lemma 7.7 proves Thm. 7.4 immediately.

8 Conclusions and Open Problems

In this work, we have presented three quantum homomorphic encryption schemes. The first, CL, is a stepping stone to the other two, and is homomorphic and compact for the class of stabilizer circuits. The second, EPR, is homomorphic for all quantum circuits, but the compactness property degrades with the number of T-gates. In the third scheme, AUX, the complexity of the evaluation key and the evaluation procedure scale doubly exponentially with the T-depth, so that it is only homomorphic for circuits with constant T-depth, but it is also compact.

The clear central open problem in this work is to come up with a quantum fully homomorphic encryption scheme satisfying Def. 3.8, which must be homomorphic for all quantum circuits and compact. Our schemes EPR and AUX make progress towards this goal from two directions, but still leave open a full solution to this problem.

Our work can be seen as analogous to a number of classical results leading up to fully homomorphic encryption, including classical encryption schemes that were homomorphic for some limited classes of circuits, including limits in the multiplicative depth, as well as quasi-compact homomorphic schemes. In addition, we have attempted, in our security definitions and the theorems in App. B, to set the groundwork for a rigorous treatment of quantum homomorphic encryption, hopefully leading, eventually to quantum fully homomorphic encryption.

Acknowledgements

The authors would like to thank Fang Song for helpful discussions about this paper, in particular regarding security definitions.

A. B. acknowledges support from the Canadian Institute for Advanced Research (CIFAR); A. B. and S. J. acknowledge support from the Natural Sciences and Engineering Research Council of Canada (NSERC). Part of this work was done while the authors were visitors at the Simons Institute for the Theory of Computing.

References

- [ABOE10] D. Aharonov, M. Ben-Or, and E. Eban. Interactive proofs for quantum computations. In *Proceeding of Innovations in Computer Science 2010 (ICS'10)*, pages 453–469, 2010.
- [AMMR13] M. Amy, D. Maslov, M. Mosca, and M. Roetteler. A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 32(6):818–830, June 2013.
- [AMTW00] A. Ambainis, M. Mosca, A. Tapp, and R. De Wolf. Private quantum channels. In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science (FOCS'00)*, pages 547–553, 2000.
- [AS06] P. Arrighi and L. Salvail. Blind quantum computation. *International Journal of Quantum Information*, 4:883–898, 2006.
- [BB84] C.H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing*, pages 175–179, 1984.
- [BCG⁺02] H. Barnum, C. Crépeau, D. Gottesman, A. Smith, and A. Tapp. Authentication of quantum messages. In *Proceedings of the 43rd Annual IEEE Symposium on the Foundations of Computer Science (FOCS'02)*, pages 449–458, 2002.

- [BDF⁺11] D. Boneh, Ö. Dagdelen, M. Fischlin, A. Lehmann, C. Schaffner, and M. Zhandry. Random oracles in a quantum world. In *Advances in Cryptology (ASIACRYPT 2011)*, pages 41–69, 2011.
- [BFK09] A. Broadbent, J. Fitzsimons, and E. Kashefi. Universal blind quantum computation. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS'09)*, pages 517–526, 2009.
- [BGN05] D. Boneh, E.-J. Goh, and K. Nissim. Evaluating 2-DNF formulas on ciphertexts. In *Proceedings of the Second Theory of Cryptography Conference (TCC 2005)*, pages 325–341, 2005.
- [BGS13] A. Broadbent, G. Gutoski, and D. Stebila. Quantum one-time programs. In *Advances in Cryptology (CRYPTO 2013)*, pages 344–360, 2013.
- [BK05] S. Bravyi and A. Kitaev. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Physical Review A*, 71:022316, 2005.
- [BOCG⁺06] M. Ben-Or, C. Crépeau, D. Gottesman, A. Hassidim, and A. Smith. Secure multiparty quantum computation with (only) a strict honest majority. In *Proc. 47th IEEE Symposium on the Foundations of Computer Science (FOCS 2006)*, pages 249–260, 2006.
- [Bro15] A. Broadbent. Delegating private quantum computations. [arXiv:1506.01328\[quant-ph\]](https://arxiv.org/abs/1506.01328), to appear in *Canadian Journal of Physics*, 2015.
- [BV11] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2011)*, pages 97–106, 2011. Full version available at Cryptology ePrint Archive, Report 2011/344.
- [BZ13] D. Boneh and M. Zhandry. Secure signatures and chosen ciphertext security in a quantum computing world. In *Proceedings of CRYPTO 2013*, 2013.
- [Chi05] A. Childs. Secure assisted quantum computation. *Quantum Information and Computation*, 5:456–466, 2005.
- [CLN05] A. M. Childs, D. W. Leung, and M. A. Nielsen. Unified derivations of measurement-based schemes for quantum computation. *Physical Review A*, 71:032318, 2005.
- [DGHV10] M. Dijk, C. Gentry, S. Halevi, and V. Vaikuntanathan. Fully homomorphic encryption over the integers. In *Advances in Cryptology (EUROCRYPT 2010)*, pages 24–43. 2010.
- [DNS10] F. Dupuis, J. B. Nielsen, and L. Salvail. Secure two-party quantum evaluation of unitaries against specious adversaries. In *Proceedings of the 30th International Cryptology Conference (CRYPTO'10)*, pages 685–706, 2010. Detailed version available at [arXiv:1009.2096](https://arxiv.org/abs/1009.2096).
- [DNS12] F. Dupuis, J. B. Nielsen, and L. Salvail. Actively secure two-party evaluation of any quantum operation. In *Advances in Cryptology – Proc. CRYPTO 2012*, volume 7417 of *LNCS*, pages 794–811. Springer, 2012.
- [ElG85] T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In *Advances in Cryptology (CRYPTO '85)*, pages 10–18, 1985.
- [EPR35] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47:777–780, May 1935.
- [FBS⁺14] K. A. G. Fisher, A. Broadbent, L. K. Shalm, Z. Yan, J. Lavoie, R. Prevedel, T. Jennewein, and K. J. Resch. Quantum computing on encrypted data. *Nature communications*, 5, 2014.
- [Gen09a] C. Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. crypto.stanford.edu/craig.
- [Gen09b] C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st annual ACM symposium on Theory of Computing (STOC'09)*, pages 169–178, 2009.

- [GHV10] C. Gentry, S. Halevi, and V. Vaikuntanathan. *i*-hop homomorphic encryption and rerandomizable Yao circuits. In *Advances in Cryptology (CRYPTO 2010)*, pages 155–172, 2010.
- [GKP⁺13] S. Goldwasser, Y. Kalai, R. A. Popa, V. Vaikuntanathan, and N. Zeldovich. Reusable garbled circuits and succinct functional encryption. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, (STOC '13), pages 555–564, 2013.
- [GM84] S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270 – 299, 1984.
- [Got98] D. Gottesman. The Heisenberg representation of quantum computers. In *Group 22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, pages 32–43, 1998.
- [KL08] J. Katz and Y. Lindell. *Introduction to Modern Cryptography: Principles and Protocols*. Chapman & Hall/CRC, 2008.
- [KMM13] V. Kliuchnikov, D. Maslov, and M. Mosca. Asymptotically optimal approximation of single qubit unitaries by Clifford and *T* circuits using a constant number of ancillary qubits. *Phys. Rev. Lett.*, 110:190502, May 2013.
- [Kos07] T. Koshiha. Security notions for quantum public-key cryptography. *IEICE TRANSACTIONS on Fundamentals of Electronics, Communications and Computer Sciences (Japanese Edition)*, J90-A(5):367–375, 2007. English version available as: [arXiv:quant-ph/0702183](https://arxiv.org/abs/quant-ph/0702183).
- [LC97] H.-K. Lo and H. F. Chau. Is quantum bit commitment really possible? *Physical Review Letters*, 78:3410–3413, 1997.
- [May97] D. Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical Review Letters*, 78:3414–3417, 1997.
- [MR09] D. Micciancio and O. Regev. Lattice-based cryptography. In *Post-Quantum Cryptography*. Springer, 2009.
- [OTU00] T. Okamoto, K. Tanaka, and S. Uchiyama. Quantum public-key cryptosystems. In *Advances in Cryptology (CRYPTO 2000)*, pages 147–165, 2000.
- [Pai99] P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Advances in Cryptology (EUROCRYPT '99)*, pages 223–238. 1999.
- [RAD78] R. Rivest, L. Adleman, and M. Dertouzos. On data banks and privacy homomorphisms. In *Foundations of Secure Computation*, pages 169–177, 1978.
- [Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th annual ACM symposium on Theory of computing*, (STOC '05), pages 84–93, 2005.
- [Reg09] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–34:40, September 2009.
- [RFG12] P. P. Rohde, J. F. Fitzsimons, and A. Gilchrist. Quantum walks with encrypted data. *Physical Review Letters*, 109:150501, Oct 2012.
- [Rot11] R. Rothblum. Homomorphic encryption: from private-key to public-key. In *Proceedings of the 8th Theory of Cryptography Conference, (TCC 2011)*, pages 219–234. Springer, 2011.
- [Sel13] P. Selinger. Quantum circuits of T-depth one. *Phys. Rev. A*, 87:042302, Apr 2013.
- [Son14] F. Song. A note on quantum security for post-quantum cryptography. In *Proceedings of the 6th International Conference on Post-Quantum Cryptography (PQCrypto 2014)*, pages 246–265, 2014.
- [SYY99] T. Sander, A. Young, and M. Yung. Non-interactive cryptocomputing for NC^1 . In *Proceedings of the 40th Annual Symposium on the Foundations of Computer Science (FOCS 1999)*, pages 554–566, 1999.

- [TKO⁺14] S.-H. Tan, J. A. Kettlewell, Y. Ouyang, L. Chen, and J. F. Fitzsimons. A quantum approach to homomorphic encryption. <http://arxiv.org/abs/1411.5254>, 2014.
- [Unr12] D. Unruh. Quantum proofs of knowledge. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 135–152. Springer Berlin Heidelberg, 2012.
- [Vai11] V. Vaikuntanathan. Computing blindfolded: New developments in fully homomorphic encryption. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*, (FOCS '11), pages 5–16, 2011.
- [VFPR14] D. Vedran, J. F. Fitzsimons, C. Portmann, and R. Renner. Composable security of delegated quantum computation. In *Proceedings of Advances in Cryptology (ASIACRYPT 2014)*, pages 406–425, 2014.
- [Wat06] J. Watrous. Zero-knowledge against quantum attacks. In *Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing*, STOC '06, pages 296–305, New York, NY, USA, 2006. ACM.
- [Wat13] J. Watrous. Course notes for *Theory of Quantum Information*. Available at: <https://cs.uwaterloo.ca/~watrous/CS766/>, 2013.
- [WZ82] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.
- [XY12] C. Xiang and L. Yang. Indistinguishability and semantic security for quantum encryption scheme. In *Proc. SPIE 8554, Quantum and Nonlinear Optics II*, page 85540G, 2012.
- [YPDF14] L. Yu, C. A. Perez-Delgado, and J. F. Fitzsimons. Limitations on information-theoretically-secure quantum homomorphic encryption. *Physical Review A*, 90(5):050303, 2014.
- [Zha12] M. Zhandry. How to construct quantum random functions. In *Proceedings of the IEEE 53rd Annual Symposium on the Foundations of Computer Science (FOCS 2012)*, pages 679–687, Oct 2012.
- [ZLC00] X. Zhou, D. W. Leung, and I. L. Chuang. Methodology for quantum logic gate construction. *Physical Review A*, 62:052316, 2000.

A Classical Fully Homomorphic Encryption

Here we present the definitions from the full version of [BV11].

Definition A.1. A homomorphic encryption scheme is a 4-tuple (HE.KeyGen, HE.Enc, HE.Dec, HE.Eval) of PPT algorithms such that:

Key Generation. The algorithm $(pk, evk, sk) \leftarrow \text{HE.KeyGen}(1^\kappa)$ takes a unary representation of the security parameter and outputs a public encryption key pk , a public evaluation key evk and a secret decryption key sk .

Encryption. The algorithm $c \leftarrow \text{HE.Enc}_{pk}(\mu)$ takes the public key pk and a single bit message $\mu \in \{0, 1\}$ and outputs a ciphertext c .

Decryption. The algorithm $\mu^* \leftarrow \text{HE.Dec}_{sk}(c)$ takes the secret key sk and a ciphertext c and outputs a message $\mu^* \in \{0, 1\}$.

Homomorphic Evaluation. The algorithm $c_f \leftarrow \text{HE.Eval}_{evk}^f(c_1, \dots, c_\ell)$ takes the evaluation key evk , a classical circuit $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ and a set of ℓ ciphertexts c_1, \dots, c_ℓ and outputs a ciphertext c_f .

We define \mathcal{S} -homomorphic, which is homomorphism with respect to a specified class \mathcal{S} of circuits. This notion is sometimes also referred to as “somewhat homomorphic”.

Definition A.2 (\mathcal{S} -homomorphic). Let $\mathcal{S} = \{\mathcal{S}_\kappa\}_{\kappa \in \mathbb{N}}$ be a class of classical circuits. A scheme HE is \mathcal{S} -homomorphic (or, homomorphic for the class \mathcal{S}) if for any sequence of circuits $\{f_\kappa \in \mathcal{S}_\kappa\}_{\kappa \in \mathbb{N}}$ and respective inputs $\mu_1, \dots, \mu_\ell \in \{0, 1\}$ (where $\ell = \ell(\kappa)$), there exists a negligible function η such that

$$\Pr[\text{HE.Dec}_{sk}(\text{HE.Eval}_{evk}^f(c_1, \dots, c_\ell)) \neq f(\mu_1, \dots, \mu_\ell)] = \eta(\kappa), \quad (6)$$

where $(pk, evk, sk) \leftarrow \text{HE.KeyGen}(1^\kappa)$ and $c_i \leftarrow \text{HE.Enc}_{pk}(\mu_i)$.

Definition A.3 (compactness). A homomorphic scheme HE is compact if there exists a polynomial p such that the circuit complexity of decrypting the output of $\text{HE.Eval}^f(\dots)$ is at most $p(\kappa)$ (regardless of f).

Definition A.4 (fully homomorphic encryption). A scheme HE is fully homomorphic if it is both compact and homomorphic for the class of all arithmetic circuits over \mathbb{F}_2 .

B Equivalence of Definitions for q-IND-CPA

In this section, we prove Thm. 3.5 of Sec. 3.2, restated below:

Theorem 3.5 (Equivalence of q-IND-CPA and q-IND-CPA-mult). Let QHE be a quantum homomorphic encryption scheme. Then QHE is q-IND-CPA if and only if QHE is q-IND-CPA-mult.

In order to prove Thm. 3.5, we will first introduce an intermediate security definition, q-IND-CPA-2, which is like q-IND-CPA-mult when we restrict to $t = 1$. We will show that q-IND-CPA is equivalent to q-IND-CPA-2, and then that q-IND-CPA-2 is equivalent to q-IND-CPA-mult. We note that the proof given in this section is easily modified to a proof of Thm. 3.15 — a similar statement for the symmetric-key setting. The only difference in the symmetric-key case is that adversaries can make calls to an encryption oracle. The main techniques in this section involve constructing an adversary \mathcal{A}' that runs some other adversary \mathcal{A} . If \mathcal{A} is an adversary with oracle access, then another adversary with oracle access, \mathcal{A}' , can easily run \mathcal{A} , so the same ideas go through in an identical manner in the symmetric-key case.

CPA-2 security. The CPA-2 indistinguishability experiment is given below and illustrated in Fig. 12.




Figure 12: The quantum IND-CPA-2 indistinguishability experiment.

The quantum IND-CPA-2 indistinguishability experiment $\text{PubK}_{\mathcal{A}, \text{QHE}}^{\text{cpa-2}}(\kappa)$

1. $\text{KeyGen}(1^\kappa)$ is run to obtain keys (pk, sk, ρ_{evk}) .
2. Adversary \mathcal{A}_1 is given (pk, ρ_{evk}) and outputs a quantum state ρ in $\mathcal{M}_0 \otimes \mathcal{M}_1 \otimes \mathcal{E}$, where $\mathcal{M}_0 \equiv \mathcal{M}_1 \equiv \mathcal{M}$.
3. For $r \in \{0, 1\}$, let $\Xi_{\text{QHE}}^{\text{cpa-2}, r} : D(\mathcal{M}_0 \otimes \mathcal{M}_1) \rightarrow D(\mathcal{C})$ be given by $\Xi_{\text{QHE}}^{\text{cpa-2}, 0}(\rho) = \text{Tr}_{\mathcal{M}_1} \left((\text{Enc}_{pk}^{\mathcal{M}_0} \otimes \mathbb{I}_{\mathcal{M}_1})(\rho) \right)$ and $\Xi_{\text{QHE}}^{\text{cpa-2}, 1}(\rho) = \text{Tr}_{\mathcal{M}_0} \left((\mathbb{I}_{\mathcal{M}_0} \otimes \text{Enc}_{pk}^{\mathcal{M}_1})(\rho) \right)$. A random bit $r \in \{0, 1\}$ is chosen and $\Xi_{\text{QHE}}^{\text{cpa-2}, r} \otimes \mathbb{I}_{\mathcal{E}}$ is applied to ρ (the output being a state in $\mathcal{C} \otimes \mathcal{E}$).
4. Adversary \mathcal{A}_2 obtains the system in $\mathcal{C} \otimes \mathcal{E}$ and outputs a bit r' .
5. The output of the experiment is defined to be 1 if $r' = r$ and 0 otherwise. In case $r = r'$, we say that \mathcal{A} wins the experiment.

Definition B.1 (q-IND-CPA-2). A quantum homomorphic encryption scheme QHE is q-IND-CPA-2 secure if for all quantum polynomial-time adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ there exists a negligible function η such that:

$$\Pr[\text{PubK}_{\mathcal{A}, \text{QHE}}^{\text{cpa-2}}(\kappa) = 1] \leq \frac{1}{2} + \eta(\kappa).$$

Theorem B.2 (Equivalence of q-IND-CPA and q-IND-CPA-2). A quantum homomorphic encryption scheme is q-IND-CPA if and only if it is q-IND-CPA-2.

Proof. It is trivial to see that if a scheme is q-IND-CPA-2 then it is q-IND-CPA.

Suppose a scheme QHE is q-IND-CPA. Then let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be any adversary for the experiment $\text{PubK}_{\mathcal{A}, \text{QHE}}^{\text{cpa-2}, r}$, so \mathcal{A}_1 implements a quantum channel from $D(\mathcal{R}_{\text{evk}})$ to $D(\mathcal{M}_0 \otimes \mathcal{M}_1 \otimes \mathcal{E})$ conditioned on pk , and \mathcal{A}_2 implements a quantum channel on $D(\mathcal{C} \otimes \mathcal{E})$ that outputs a bit. We will use \mathcal{A} to construct an adversary, \mathcal{A}' for the experiment $\text{PubK}_{\mathcal{A}', \text{QHE}}^{\text{cpa}}$ as shown in Fig. 13. We define \mathcal{A}'_1 by $\mathcal{A}'_1(pk, \rho^{\mathcal{R}_{\text{evk}}}) = \text{Tr}_{\mathcal{M}_0}(\mathcal{A}_1(pk, \rho^{\mathcal{R}_{\text{evk}}}))$, and \mathcal{A}'_2 by $\mathcal{A}'_2 = \mathcal{A}_2$.

We now consider the probability that \mathcal{A}' wins the q-IND-CPA experiment: $\Pr[\text{PubK}_{\mathcal{A}', \text{QHE}}^{\text{cpa}}(\kappa) = 1]$. If $r = 0$, the probability that $r' = r$ is $\Pr[\mathcal{A}_2((\text{Enc}(|0\rangle\langle 0|)) \otimes \rho^{\mathcal{E}}) = 0]$. If $r = 1$, the probability that $r' = r$ is $\Pr[\mathcal{A}_2((\text{Enc}_{pk} \otimes \mathbb{I}_{\mathcal{E}})(\rho^{\mathcal{M}_1 \otimes \mathcal{E}})) = 1]$. Thus, the probability that this adversary correctly predicts r is

$$\frac{1}{2} \Pr[\mathcal{A}_2((\text{Enc}(|0\rangle\langle 0|)) \otimes \rho^{\mathcal{E}}) = 0] + \frac{1}{2} \Pr[\mathcal{A}_2(\text{Enc}_{pk} \otimes \mathbb{I}_{\mathcal{E}}(\rho^{\mathcal{M}_1 \otimes \mathcal{E}})) = 1] \leq \frac{1}{2} + \eta'(\kappa), \quad (7)$$

for some negligible function η' , by the fact that QHE is q-IND-CPA.

Consider a slightly different strategy, $\mathcal{A}'' = (\mathcal{A}''_1, \mathcal{A}''_2)$, for the same experiment $\text{PubK}_{\mathcal{A}'', \text{QHE}}^{\text{cpa}}$. This strategy discards the *second* message space, \mathcal{M}_1 , and inputs the first, \mathcal{M}_0 , into $\Xi_{\text{QHE}}^{\text{cpa}, r}$, that is $\mathcal{A}''_1 = \text{Tr}_{\mathcal{M}_1}(\mathcal{A}_1(pk, \rho^{\mathcal{R}_{\text{evk}}}))$. The new adversary \mathcal{A}'' outputs the complement of the output of \mathcal{A} : $\mathcal{A}''_2(\rho^{\mathcal{C}\mathcal{E}}) = \mathcal{A}_2(\rho^{\mathcal{C}\mathcal{E}}) \oplus 1$. We can then see that \mathcal{A}'' correctly predicts r with probability

$$\frac{1}{2} \Pr[\mathcal{A}_2((\text{Enc}(|0\rangle\langle 0|)) \otimes \rho^{\mathcal{E}}) = 1] + \frac{1}{2} \Pr[\mathcal{A}_2(\text{Enc}_{pk} \otimes \mathbb{I}_{\mathcal{E}}(\rho^{\mathcal{M}_0 \otimes \mathcal{E}})) = 0] \leq \frac{1}{2} + \eta''(\kappa). \quad (8)$$

for some negligible function η'' . The addition of (7) and (8), gives:

$$\frac{1}{2} + \frac{1}{2} \Pr[\mathcal{A}_2(\text{Enc}_{pk} \otimes \mathbb{I}_{\mathcal{E}}(\rho^{\mathcal{M}_1 \otimes \mathcal{E}})) = 1] + \frac{1}{2} \Pr[\mathcal{A}_2(\text{Enc}_{pk} \otimes \mathbb{I}_{\mathcal{E}}(\rho^{\mathcal{M}_0 \otimes \mathcal{E}})) = 0] \leq 1 + \eta'(\kappa) + \eta''(\kappa).$$

Since $\eta := \eta' + \eta''$ is still negligible, we conclude that for some negligible function η :

$$\frac{1}{2} \Pr[\mathcal{A}_2(\text{Enc}_{pk} \otimes \mathbb{I}_{\mathcal{E}}(\rho^{\mathcal{M}_1 \otimes \mathcal{E}})) = 1] + \frac{1}{2} \Pr[\mathcal{A}_2(\text{Enc}_{pk} \otimes \mathbb{I}_{\mathcal{E}}(\rho^{\mathcal{M}_0 \otimes \mathcal{E}})) = 0] \leq \frac{1}{2} + \eta(\kappa). \quad \square$$




Figure 13: The adversary \mathcal{A}' described in the proof of Thm. B.2.

Theorem B.3 (Equivalence of q-IND-CPA and q-IND-CPA-mult). A quantum homomorphic encryption scheme is q-IND-CPA if and only if it is q-IND-CPA-mult.

Proof. It is trivial to see that if a scheme is q-IND-CPA-mult, then it is q-IND-CPA.

For the other direction, it is simple to adapt a similar classical proof (see for example [KL08]) to the quantum setting. Suppose QHE is q-IND-CPA, so in particular, it is q-IND-CPA-2. Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an adversary for the q-IND-CPA-mult experiment $\text{PubK}_{\mathcal{A}, \text{QHE}}^{\text{cpa-mult}}$. We will construct an adversary, \mathcal{A}' , for the q-IND-CPA-2 experiment $\text{PubK}_{\mathcal{A}', \text{QHE}}^{\text{cpa-2}}$ from \mathcal{A} .

For any $i \in \{0, \dots, t\}$, define $\Psi_i : D(\mathcal{M}_0^1 \otimes \dots \otimes \mathcal{M}_0^t \otimes \mathcal{M}_1^1 \otimes \dots \otimes \mathcal{M}_1^t \otimes \mathcal{E}) \rightarrow D(\mathcal{C}^1 \otimes \dots \otimes \mathcal{C}^t \otimes \mathcal{E})$ as the channel that applies Enc_{pk} to the systems $\mathcal{M}_0^1, \dots, \mathcal{M}_0^i, \mathcal{M}_1^{i+1}, \dots, \mathcal{M}_1^t$, and traces out the systems $\mathcal{M}_1^1, \dots, \mathcal{M}_1^i, \mathcal{M}_0^{i+1}, \dots, \mathcal{M}_0^t$.

Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a t -message adversary for q-IND-CPA-mult. We define a q-IND-CPA-2 adversary $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2)$ as follows:

$\mathcal{A}'_1(pk, \rho^{\mathcal{R}_{evk}})$: Run $\mathcal{A}_1(pk, \rho^{\mathcal{R}_{evk}})$ to get $\rho \in D(\mathcal{M}_0^1 \otimes \dots \otimes \mathcal{M}_0^t \otimes \mathcal{M}_1^1 \otimes \dots \otimes \mathcal{M}_1^t \otimes \mathcal{E})$. Choose a random $i \in \{1, \dots, t\}$, and apply $\Xi_{\text{QHE}}^{\text{cpa-2}, r}$ to the system $\mathcal{M}_0^i, \mathcal{M}_1^i$. For $j < i$, apply Enc_{pk} to \mathcal{M}_0^j and label the output as \mathcal{C}_j . For $j > i$, apply Enc_{pk} to \mathcal{M}_1^j , and label the output as \mathcal{C}_j . Let $\mathcal{E}' = \left(\bigotimes_{j=1; j \neq i}^t \mathcal{C}_j \right) \otimes \mathcal{E} \otimes \mathbb{C}^{t+1}$. Record i in the last register.

$\mathcal{A}'_2 : D(\mathcal{C} \otimes \mathcal{E}') \rightarrow \{0, 1\}$: Label the output of $\Xi_{\text{QHE}}^{\text{cpa-2}, r}$, \mathcal{C} , as \mathcal{C}_i . Apply \mathcal{A}_2 to the system $\mathcal{C}_1 \otimes \dots \otimes \mathcal{C}_t \otimes \mathcal{E}$, to get a bit r' . Output r' .

We now consider the success probability of \mathcal{A}' on $\text{PubK}_{\mathcal{A}', \text{QHE}}^{\text{cpa-2}}$. Let $\rho = \mathcal{A}_1(pk, \rho^{\mathcal{R}_{evk}})$. We first note that if \mathcal{A}'_1 selects i , then if $r = 0$, the state passed to \mathcal{A}'_2 is $\Psi_i(\rho)$, but if $r = 1$, the state passed to \mathcal{A}'_2 is $\Psi_{i+1}(\rho)$. So if $r = 0$, then the success probability is:

$$\Pr[\text{PubK}_{\mathcal{A}', \text{QHE}}^{\text{cpa-2}}(\kappa) = 1 | r = 0] = \sum_{i=1}^t \frac{1}{t} \Pr[\mathcal{A}'_2(\Psi_i(\rho)) = 0].$$

And if $r = 1$, then the success probability is:

$$\Pr[\text{PubK}_{\mathcal{A}', \text{QHE}}^{\text{cpa-2}}(\kappa) = 1 | r = 1] = \sum_{i=1}^t \frac{1}{t} \Pr[\mathcal{A}'_2(\Psi_{i+1}(\rho)) = 1] = \sum_{i=0}^{t-1} \frac{1}{t} \Pr[\mathcal{A}'_2(\Psi_i(\rho)) = 1].$$

From these two equations, we can compute:

$$\begin{aligned} & \Pr[\text{PubK}_{\mathcal{A}', \text{QHE}}^{\text{cpa-2}}(\kappa) = 1] \\ &= \frac{1}{2} \Pr[\text{PubK}_{\mathcal{A}', \text{QHE}}^{\text{cpa-2}}(\kappa) = 1 | r = 0] + \frac{1}{2} \Pr[\text{PubK}_{\mathcal{A}', \text{QHE}}^{\text{cpa-2}}(\kappa) = 1 | r = 1] \\ &= \frac{1}{2} \left(\frac{1}{t} \Pr[\mathcal{A}'_2(\Psi_0(\rho)) = 1] + \sum_{i=1}^{t-1} \frac{1}{t} (\Pr[\mathcal{A}'_2(\Psi_i(\rho)) = 0] + \Pr[\mathcal{A}'_2(\Psi_i(\rho)) = 1]) + \frac{1}{t} \Pr[\mathcal{A}'_2(\Psi_t(\rho)) = 0] \right) \\ &= \frac{1}{2} \left(\frac{1}{t} \Pr[\mathcal{A}'_2(\Psi_0(\rho)) = 1] + \frac{t-1}{t} + \frac{1}{t} \Pr[\mathcal{A}'_2(\Psi_t(\rho)) = 0] \right). \end{aligned}$$

Since QHE is assumed to be q-IND-CPA, there exists a negligible function η' such that $\Pr[\text{PubK}_{\mathcal{A}', \text{QHE}}^{\text{cpa-2}}(\kappa) = 1] \leq \frac{1}{2} + \eta'(\kappa)$, so we can compute:

$$\frac{1}{t} \Pr[\mathcal{A}'_2(\Psi_0(\rho)) = 1] + \frac{t-1}{t} + \frac{1}{t} \Pr[\mathcal{A}'_2(\Psi_t(\rho)) = 0] \leq 1 + 2\eta'(\kappa). \quad (9)$$

Note that $\Psi_0(\rho) = \left((\text{Enc}_{pk})^{\otimes t} \otimes \mathbb{I}_{\mathcal{E}} \right) (\rho^{\otimes_{j=1}^t \mathcal{M}_0^j} \otimes \mathcal{E})$ and $\Psi_t(\rho) = \left((\text{Enc}_{pk})^{\otimes t} \otimes \mathbb{I}_{\mathcal{E}} \right) (\rho^{\otimes_{j=1}^t \mathcal{M}_1^j} \otimes \mathcal{E})$, so $\frac{1}{2} \Pr[\mathcal{A}'_2(\Psi_0(\rho)) = 1] + \frac{1}{2} \Pr[\mathcal{A}'_2(\Psi_t(\rho)) = 0] = \Pr[\text{PubK}_{\mathcal{A}, \text{QHE}}^{\text{cpa-mult}}(\kappa) = 1]$. From Equation (9), we get:

$$\Pr[\text{PubK}_{\mathcal{A}, \text{QHE}}^{\text{cpa-mult}}(\kappa) = 1] \leq \frac{1}{2} + t \cdot \eta'(\kappa).$$

It must be the case that $t = O(\text{poly}(\kappa))$, since \mathcal{A} is a QPT algorithm, so $t \cdot \eta'(\kappa)$ is negligible in κ . Setting $\eta = t \cdot \eta'$ completes the proof. \square




Figure 14: The strategy described in the proof of Thm. B.3 for the case $t = 2$, when $i = 2$.

C Key Update Rules for Stabilizer Elements

We review here the key update rules for performing stabilizer/Clifford operators on quantum data encrypted with the quantum one-time pad [Got98].

$$\mathsf{X}^{f_{a,i}} \mathsf{Z}^{f_{b,i}} |\psi\rangle \xrightarrow{\mathcal{X}_i} \text{Measurement} = c \quad f_{a,i} \leftarrow f_{a,i}$$

Figure 15: Protocol for measurement on the i^{th} wire: Simply perform the measurement. The resulting bit, c , can be decrypted by applying $\mathsf{X}^{f_{a,i}}$ (The key $f_{b,i}$ is no longer relevant).

$$|0\rangle \xrightarrow{\mathcal{X}_i} \mathsf{X}^0 \mathsf{Z}^0 |0\rangle \quad f_{a,i} \leftarrow 0, \quad f_{b,i} \leftarrow 0$$

Figure 16: Protocol for auxiliary qubit preparation on a new wire, i : Initialize a new wire labelled \mathcal{X}_i and new key-polynomials $f_{i,a} = f_{b,i} = 0$.

$$\mathsf{X}^{f_{a,i}} \mathsf{Z}^{f_{b,i}} |\psi\rangle \xrightarrow{\mathcal{X}_i} \mathsf{X}^{f_{a,i}} \mathsf{Z}^{f_{b,i}} \mathsf{X} |\psi\rangle \quad f_{a,i} \leftarrow f_{a,i}, \quad f_{b,i} \leftarrow f_{b,i}$$

Figure 17: Protocol for an X-gate on the i^{th} wire: Simply apply the X-gate.

$$\mathsf{X}^{f_{a,i}} \mathsf{Z}^{f_{b,i}} |\psi\rangle \xrightarrow{\mathcal{X}_i} \mathsf{X}^{f_{a,i}} \mathsf{Z}^{f_{b,i}} \mathsf{Z} |\psi\rangle \quad f_{a,i} \leftarrow f_{a,i}, \quad f_{b,i} \leftarrow f_{b,i}$$

Figure 18: Protocol for a Z-gate on the i^{th} wire: Simply apply the Z-gate.

$$\mathsf{X}^{f_{a,i}} \mathsf{Z}^{f_{b,i}} |\psi\rangle \xrightarrow{\mathcal{X}_i} \mathsf{X}^{f_{b,i}} \mathsf{Z}^{f_{a,i}} \mathsf{H} |\psi\rangle \quad f_{a,i} \leftarrow f_{b,i}, \quad f_{b,i} \leftarrow f_{a,i}$$

Figure 19: Protocol for an H-gate on the i^{th} wire: Apply the gate and swap the key-polynomials.

$$\mathsf{X}^{f_{a,i}}\mathsf{Z}^{f_{b,i}}|\psi\rangle \xrightarrow{\boxed{\mathsf{P}}_{\mathcal{X}_i}} \mathsf{X}^{f_{a,i}}\mathsf{Z}^{f_{b,i}\oplus f_{a,i}}\mathsf{P}|\psi\rangle \quad f_{a,i} \leftarrow f_{a,i}, \quad f_{b,i} \leftarrow f_{b,i} \oplus f_{a,i}$$

Figure 20: Protocol for a P-gate on the i^{th} wire: Apply the gate and update $f_{b,i}$.

$$(\mathsf{X}^{f_{a,i}}\mathsf{Z}^{f_{b,i}} \otimes \mathsf{X}^{f_{a,j}}\mathsf{Z}^{f_{b,j}})|\psi\rangle \left\{ \begin{array}{c} \text{---} \bullet \text{---} \\ \mathcal{X}_i \uparrow \\ \text{---} \oplus \text{---} \\ \mathcal{X}_j \end{array} \right\} (\mathsf{X}^{f_{a,i}}\mathsf{Z}^{f_{b,i}\oplus f_{b,j}} \otimes \mathsf{X}^{f_{a,i}\oplus f_{a,j}}\mathsf{Z}^{f_{b,j}})\mathsf{CNOT}(|\psi\rangle)$$

$$f_{a,i} \leftarrow f_{a,i}, \quad f_{b,i} \leftarrow f_{b,i} \oplus f_{b,j}, \quad f_{a,j} \leftarrow f_{a,i} \oplus f_{a,j}, \quad f_{b,j} \leftarrow f_{b,j}$$

Figure 21: Protocol for a CNOT-gate with control wire i and target wire j : Apply the gate and update $f_{b,i}$ and $f_{a,j}$.

We remark that an alternative gadget for the X is to update the X -key as $f_{a,i} \rightarrow f_{a,i} \oplus 1$, rather than applying X to the quantum state. A similar alternative holds for the Z -gadget. However, these are the only two gates for which a key update is sufficient to affect the gate. Since we are actually carrying out quantum computations on encrypted quantum data — in contrast to merely simulating a quantum computation — all gates except the Pauli gates require actual quantum operations to be applied during evaluation.

D Correctness of the T-gate Gadget

We give below a step-by-step proof of the correctness of the T-gate protocol from Fig. 8. The basic building block is the circuit identity for an X -teleportation from [ZLC00], which we re-derive here. Also of relevance to this work are the techniques developed by Childs, Leung, and Nielsen [CLN05] to manipulate circuits that produce an output that is correct *up to known Pauli corrections*.

We will make use of the following identities which all hold up to an irrelevant global phase: $\mathsf{XZ} = \mathsf{ZX}$, $\mathsf{PZ} = \mathsf{ZP}$, $\mathsf{PX} = \mathsf{XZP}$, $\mathsf{TZ} = \mathsf{ZT}$, $\mathsf{TX} = \mathsf{XZPT}$, $\mathsf{P}^2 = \mathsf{Z}$ and $\mathsf{P}^{a\oplus b} = \mathsf{Z}^{a\cdot b}\mathsf{P}^{a+b}$ (for $a, b \in \{0, 1\}$).

1. Our first circuit identity (Fig. 22) swaps a qubit $|\psi\rangle$ with the state $|+\rangle$ and is easy to verify.




Figure 22: Circuit identity (easy to verify).

2. We can measure the top qubit in the above circuit and classically control the output correction (Fig. 23). We have thus re-derived the circuit corresponding to the “ X -teleportation” of [ZLC00].




Figure 23: X -teleportation

3. Let the input be $\mathsf{TX}^a\mathsf{Z}^b|\psi\rangle$, and add two gates on the auxiliary wire, P^a and Z^k (Fig. 24). Using the fact that P and Z commute with control, and applying identities given above, we get as output (using $\mathsf{TX} = \mathsf{XZPT}$):

$$\mathsf{P}^a\mathsf{Z}^k\mathsf{X}^c\mathsf{TX}^a\mathsf{Z}^b|\psi\rangle = \mathsf{P}^a\mathsf{Z}^k\mathsf{X}^c\mathsf{X}^a\mathsf{Z}^{a\oplus b}\mathsf{P}^a\mathsf{T}|\psi\rangle. \quad (10)$$

This is equal to (simplifying, then pushing the first P to the end):

$$\begin{aligned}
\mathbf{P}^a \chi^{a \oplus c} \mathbf{Z}^{a \oplus b \oplus k} \mathbf{P}^a \mathbf{T} |\psi\rangle &= \chi^{a \oplus c} \mathbf{Z}^{(a \oplus c)a} \mathbf{P}^a \mathbf{Z}^{a \oplus b \oplus k} \mathbf{P}^a \mathbf{T} |\psi\rangle \\
&= \chi^{a \oplus c} \mathbf{Z}^{a^2 \oplus c \cdot a \oplus a \oplus b \oplus k} \mathbf{P}^{2a} \mathbf{T} |\psi\rangle \\
&= \chi^{a \oplus c} \mathbf{Z}^{c \cdot a \oplus a \oplus b \oplus k} \mathbf{T} |\psi\rangle \quad \text{since } a^2 = a \text{ and } \mathbf{P}^2 = \mathbf{Z}.
\end{aligned}$$




Figure 24: Final circuit for T gate.

E Proof of Lemma 7.6

Lemma 7.6 *Let $N = n|T_1| + \dots + n|T_L|$. For any $a, b \in \{0, 1\}^n$, $\sum_{k \in \{0, 1\}^N} \sigma_{aux}^{a, b, k} = \frac{1}{2^N} \mathbb{I}_{2^N}$.*

Proof. We first note that for any string s of length $|s|$:

$$\begin{aligned}
\sum_{k \in \{0, 1\}^{|s|}} \sigma(s, k) &= \sum_{k \in \{0, 1\}^{|s|}} \bigotimes_{i=1}^{|s|} \mathbf{Z}^{k_i \mathbf{P}^{s_i}} |+\rangle \langle +| \mathbf{P}^{s_i} \mathbf{Z}^{k_i} = \bigotimes_{i=1}^{|s|} \sum_{k \in \{0, 1\}} \mathbf{Z}^k \mathbf{P}^{s_i} |+\rangle \langle +| \mathbf{P}^{s_i} \mathbf{Z}^k \\
&= \bigotimes_{i=1}^{|s|} \sum_{k \in \{0, 1\}} (|0\rangle \langle 0| + |1\rangle \langle 1| + i^{2k+a} |0\rangle \langle 1| + i^{2k-a} |1\rangle \langle 0|) \\
&= \bigotimes_{i=1}^{|s|} (2\mathbb{I}_2 + (i^{2+a} + i^a) |0\rangle \langle 1| + (i^{2-a} + i^{-a}) |1\rangle \langle 0|) \\
&= (2\mathbb{I}_2)^{\otimes |s|} = 2^{|s|} \mathbb{I}_{2^{|s|}}.
\end{aligned}$$

Then it is easy to see that for any $a, b \in \{0, 1\}^n$:

$$\begin{aligned}
\sum_{k \in \{0, 1\}^N} \sigma_{aux}^{a, b, k} &= \sum_{\substack{k^{(1)} \in \{0, 1\}^{n|T_1|}, \dots, \\ k^{(L)} \in \{0, 1\}^{n|T_L|}}} \sigma(s^{(1)}(a, b)^{*n}, k^{(1)}) \otimes \dots \otimes \sigma(s^{(L)}(a, b, k^{(1)}, \dots, k^{(L-1)})^*n, k^{(L)}) \\
&= \sum_{k^{(1)} \in \{0, 1\}^{n|T_1|}} \sigma(s^{(1)}(a, b)^{*n}, k^{(1)}) \otimes \sum_{k^{(2)} \in \{0, 1\}^{n|T_2|}} \sigma(s^{(2)}(a, b, k^{(1)})^*n, k^{(2)}) \otimes \dots \\
&\quad \otimes \sum_{k^{(L)} \in \{0, 1\}^{n|T_L|}} \sigma(s^{(L)}(a, b, k^{(1)}, \dots, k^{(L-1)})^*n, k^{(L)}) \\
&= \sum_{k^{(1)} \in \{0, 1\}^{n|T_1|}} \sigma(s^{(1)}(a, b)^{*n}, k^{(1)}) \otimes \sum_{k^{(2)} \in \{0, 1\}^{n|T_2|}} \sigma(s^{(2)}(a, b, k^{(1)})^*n, k^{(2)}) \otimes \dots \\
&\quad \otimes \sum_{k^{(L-1)} \in \{0, 1\}^{n|T_{L-1}|}} \sigma(s^{(L-1)}(a, b, k^{(1)}, \dots, k^{(L-2)})^*n, k^{(L-1)}) \otimes 2^{n|T_L|} \mathbb{I}_{2^{n|T_L|}} \\
&= 2^{n|T_1| + \dots + n|T_L|} \mathbb{I}_{2^{n|T_1| + \dots + n|T_L|}}.
\end{aligned}$$

□