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.

1 Introduction

An encryption scheme is *homomorphic* over some set of circuits $\mathscr{S}$ if any circuit in $\mathscr{S}$ 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 $\mathsf{C} \in \mathscr{S}$. In *fully homomorphic encryption (FHE)*, $\mathscr{S}$ is the set of all classical circuits. FHE was introduced in 1978 [26], but the existence of such a scheme was an open problem for over 30

[‡]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.

years. Some early public-key encryption schemes were homomorphic over the set of circuits consisting of only additions [18, 23] or only multiplications [12]. 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 [4], or of logarithmic depth [29], until finally in 2009, Gentry established a breakthrough result by giving the first fully homomorphic encryption scheme [15]. Follow-up work showed that FHE could be simplified [11], and based on standard assumptions, such as *learning with errors* [5]. The advent of FHE has unleashed a series of far-reaching consequences, such as delegating computations, and functional encryption [17]. For a survey on FHE, see [32].

A number of works have studied the secure delegation of quantum computation [1, 6–8, 10, 13, 33]. None 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 FHE, even if we allow interaction). Non-interactive approaches are given by [3], [27] and [31]. However, none of these approaches are applicable to universal circuit families. Furthermore, in the case of [3], security is given only in terms of cheat sensitivity, while both [27] and [31] only bound the leakage of their encoding schemes.

Recent work [36] 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 interested in 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 [20, 22, 35].

In terms of QHE schemes, we start by using straightforward techniques to construct a scheme that is homomorphic for Clifford 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 [2]. 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 [5], we get security based on the *learning with errors* (LWE) assumption [24, 25]; this has been equated with worst-case hardness of "short vector problems" on arbitrary lattices [21], 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 $\mathsf{T}\mathsf{X}^a\mathsf{Z}^b = \mathsf{X}^a\mathsf{Z}^{a\oplus b}\mathsf{P}^a\mathsf{T}$. 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 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 are incomparable. 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.

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 (Sec. 6), and AUX (Sec. 7). Further details, including proofs of our main theorems, can be found in the full version [9].

2 Preliminaries and Notation

A negligible function, denoted $\eta(\cdot)$, 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.

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. Two quantum systems, $\mathcal{X}$ and $\mathcal{Y}$, form a composite system by the tensor product, $\mathcal{X} \otimes \mathcal{Y}$. If $\rho \in D(\mathcal{X} \otimes \mathcal{Y})$ is a state on the joint system, we write $\rho^{\mathcal{X}}$ to denote $\mathrm{Tr}_{\mathcal{Y}}(\rho)$. 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 $\Delta(\rho, \sigma) := \mathrm{Tr}\left(\sqrt{(\rho-\sigma)^\dagger(\rho-\sigma)}\right)$.

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 $\rho(X) := \sum_{x \in \Sigma_X} \Pr[X = x]|x\rangle\langle x|$, the density matrix corresponding to X. A *classical-quantum* state is a state of the form $\rho^{\mathcal{M}\mathcal{A}} = \sum_x \Pr[X = x]|x\rangle\langle x|^{\mathcal{M}} \otimes \rho_x^{\mathcal{A}}$.

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* $\Phi : D(\mathcal{A}) \to D(\mathcal{B})$ 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 work with the gate set $\{\mathsf{X},\mathsf{Z},\mathsf{P},\mathsf{CNOT},\mathsf{H}\}$. This gate set applied to arbitrary wires (redundantly) generates the Clifford group, and adding any non-Clifford gate, such as T, gives a generating set for all quantum circuits.

For a single-qubit register $\mathcal{R}$, and $a,b \in \{0,1\}$, we denote by $\mathsf{QEnc}_{a,b} : \mathcal{R} \to \mathcal{R}$ the quantum one-time pad encryption and by $\mathsf{QDec}_{a,b} : \mathcal{R} \to \mathcal{R}$ the quantum one-time pad decryption [2], $\mathsf{QEnc}_{a,b} : \rho \mapsto \mathsf{X}^a\mathsf{Z}^b\rho\mathsf{Z}^b\mathsf{X}^a$ and $\mathsf{QDec}_{a,b} = \mathsf{QEnc}_{a,b}$. It is easy to see that $\mathsf{QDec}_{a,b} \circ \mathsf{QEnc}_{a,b} = \mathbb{I}_{\mathcal{R}}$. 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, $\frac{1}{4}\sum_{a,b}\mathsf{X}^a\mathsf{Z}^b\rho\mathsf{Z}^b\mathsf{X}^a = \frac{\mathbb{I}_2}{2}$.

3 Definitions

We now formally define QHE schemes and their properties. In Sec. 3.1, we define QHE in the public-key setting. Sec. 3.2 carefully defines the security of QHE, giving two definitions for security under chosen plaintext attacks, shown in the full version [9] to be equivalent. Sec. 3.3 defines correctness and compactness for QHE, culminating in a complete definition of quantum fully homomorphic encryption. 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, one of our schemes (AUX) must be used in the symmetric-key setting, defined in Sec. 3.5. We do not address the issue of *circuit privacy* [16], leaving this question for future work.

3.1 Classical and Quantum Homomorphic Encryption

Our schemes rely on a classical fully homomorphic encryption scheme. 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. 1). Fortunately, much of classical fully homomorphic encryption uses lattice-based cryptography, which exploits one of the few conjectured "quantum-safe" assumptions [21]. Among all known solutions, the scheme of [5] appears to be the best for our purposes, as it bases its security on the *learning with errors* (LWE) assumption [24, 25], which has been equated with worst-case hardness of "short vector problems" on arbitrary lattices.

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

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

Although a classical scheme that is q-IND-CPA is also IND-CPA, the converse may not be true. 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 [30]).

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.

Definition 2 (QHE). *A quantum homomorphic encryption scheme is a* 4*-tuple of quantum algorithms* $(\mathsf{QHE.KeyGen}, \mathsf{QHE.Enc}, \mathsf{QHE.Eval}, \mathsf{QHE.Dec})$*:*

Key Generation. $\mathsf{QHE.KeyGen} : 1^\kappa \to (pk, sk, \rho_{evk})$*. 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. $\mathsf{QHE.Enc}_{pk} : D(\mathcal{M}) \to D(\mathcal{C})$*. For every possible* pk*, the quantum channel* Enc_{pk} *maps a state in the message space* $\mathcal{M}$ *to a state (the* cipherstate*) in the cipherspace* $\mathcal{C}$*.*

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

Decryption. $\mathsf{QHE.Dec}_{sk} : D(\mathcal{C}') \to D(\mathcal{M})$*. For every possible* 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. We note that, by taking the evaluation key to be empty, our definitions are trivially applicable to the scenario of quantum public-key encryption (*i.e.* without a homomorphic property).

The CPA indistinguishability experiment is given below and illustrated in Fig. 1. The experiment interacts with an adversary $\mathscr{A}$, which is a pair of polynomial-time quantum algorithms $(\mathscr{A}_1, \mathscr{A}_2)$ (which we also call adversaries).

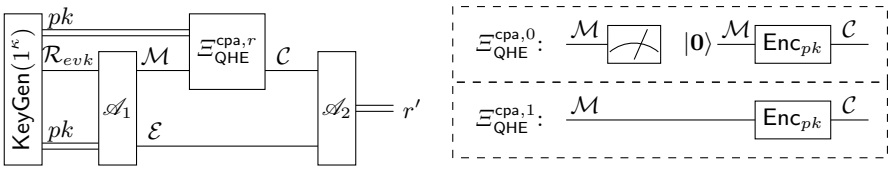

Fig. 1. The quantum CPA indistinguishability experiment.

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

1. $\mathsf{KeyGen}(1^\kappa)$ is run to obtain keys (pk, sk, ρ_{evk}).

2. Adversary $\mathscr{A}_1$ is given (pk, ρ_{evk}) and outputs a quantum state on $\mathcal{M} \otimes \mathcal{E}$.
3. For $r \in \{0,1\}$, let $\Xi_{\mathsf{QHE}}^{\mathsf{cpa},r} : D(\mathcal{M}) \to D(\mathcal{C})$ be: $\Xi_{\mathsf{QHE}}^{\mathsf{cpa},0}(\rho) = \mathsf{QHE.Enc}_{pk}(|\mathbf{0}\rangle\langle\mathbf{0}|)$ and $\Xi_{\mathsf{QHE}}^{\mathsf{cpa},1}(\rho) = \mathsf{QHE.Enc}_{pk}(\rho)$. A random bit $r \in \{0,1\}$ is chosen and $\Xi_{\mathsf{QHE}}^{\mathsf{cpa},r}$ is applied to the state in $\mathcal{M}$ (the output being a state in $\mathcal{C}$).
4. Adversary $\mathscr{A}_2$ 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 $\mathscr{A}$ *wins* the experiment.

Definition 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* $\mathscr{A} = (\mathscr{A}_1, \mathscr{A}_2)$ *there exists a negligible function* η *such that* $\Pr[\mathsf{PubK}_{\mathscr{A},\mathsf{QHE}}^{\mathsf{cpa}}(\kappa) = 1] \leq \frac{1}{2} + \eta(\kappa)$.

In the case of classical cryptosystems, it is known that IND-CPA security, the classical analogue of Def. 1, implies a seemingly stronger security against an adversary who can send multiple messages to a challenger. In the quantum case, we can analogously define an experiment similar to $\mathsf{PubK}_{\mathscr{A},\mathsf{QHE}}^{\mathsf{cpa}}$, but where the adversary prepares a state in $\mathcal{M}^{\otimes t} \otimes \mathcal{M}^{\otimes t}$ and sends it to the challenger, who traces out either the first half or the second half of the system, before applying an encryption map to each of the remaining subspaces. The adversary must then decide which system was traced out. In the full version [9], we give a formal definition of this notion of security, which we call q-IND-CPA-mult, and prove the equivalence of q-IND-CPA and q-IND-CPA-mult. This strengthens our results since security in the most general case (q-IND-CPA-mult) follows from security for the simplest definition (q-IND-CPA).

3.3 Correctness and Compactness of QHE

Next, we give a notion that encapsulates correctness of both encryption and evaluation, with respect to a class $\mathscr{S}$ of quantum circuits. 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 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 4 ($\mathscr{S}$-homomorphic). *Let* $\mathscr{S} = \{\mathscr{S}_\kappa\}_{\kappa\in\mathbb{N}}$ *be a class of quantum circuits. A quantum encryption scheme* QHE *is* $\mathscr{S}$*-homomorphic (or homomorphic for* $\mathscr{S}$*) if for any sequence of circuits* $\{\mathsf{C}_\kappa \in \mathscr{S}_\kappa\}_\kappa$ *with induced channels* $\Phi_{\mathsf{C}_\kappa} : \mathcal{M}^{\otimes n(\kappa)} \to \mathcal{M}^{\otimes m(\kappa)}$*, and input* $\rho \in D(\mathcal{M}^{\otimes n(\kappa)} \otimes \mathcal{E})$*, there exists a negligible function* η *such that for* $(pk, sk, \rho_{evk}) \leftarrow \mathsf{QHE.Keygen}(1^\kappa)$*:*

$$\Delta\left(\mathsf{QHE.Dec}_{sk}^{\otimes m(\kappa)}\left(\mathsf{QHE.Eval}^{\mathsf{C}_\kappa}\left(\rho_{evk}, \mathsf{QHE.Enc}_{pk}^{\otimes n}(\rho)\right)\right), \Phi_{\mathsf{C}_\kappa}(\rho)\right) = \eta(\kappa)\,. \quad (1)$$

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 $\mathsf{QHE.Eval}$ be able to undergo additional homomorphic evaluations; indeed, if the evaluation key ρ_{evk} is quantum, it will in general be "consumed" by the $\mathsf{QHE.Eval}$ process, rendering any future applications of $\mathsf{QHE.Eval}$ impossible.

Analogously to the classical case, we define compactness, which requires that the complexity of $\mathsf{QHE.Dec}$ be independent of the evaluated circuit, ruling out schemes where applying the circuit is delayed until after decryption.

Definition 5 ($\mathscr{S}$-compactness). *Let $\mathscr{S} = \{\mathscr{S}_\kappa\}_{\kappa\in\mathbb{N}}$ be a class of quantum circuits. A quantum encryption scheme* QHE *is* $\mathscr{S}$-compact *if there exists a polynomial p such that for any sequence of circuits $\{\mathsf{C}_\kappa \in \mathscr{S}_\kappa\}_\kappa$, the circuit complexity of applying* $\mathsf{QHE.Dec}$ *to the output of* $\mathsf{QHE.Eval}^{\mathsf{C}_\kappa}$ *is at most $p(\kappa)$.*

If QHE *is $\mathscr{S}$-compact for $\mathscr{S}$ 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 question of quantum fully homomorphic encryption, we have established all the machinery relevant for a formal definition:

Definition 6 (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 [34]) to define the output of $\mathsf{QHE.Eval}$ 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, $\mathsf{QHE.Dec}$ 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 7. *An* indivisible *quantum homomorphic encryption scheme is a QHE scheme with* $\mathsf{QHE.Eval}$ *and* $\mathsf{QHE.Dec}$ *re-defined as:*

Homomorphic Evaluation. $\mathsf{QHE.Eval}^{\mathsf{C}} : D(\mathcal{R}_{evk} \otimes \mathcal{C}^{\otimes n}) \to D(\mathcal{R}_{aux} \otimes \mathcal{C'}^{\otimes m})$. *Compared to* $\mathsf{QHE.Eval}$ *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 $\mathcal{R}_{aux} \otimes \mathcal{C'}^{\otimes m}$ can no longer be considered subsystem-by-subsystem.*

Decryption. $\mathsf{QHE.Dec}_{sk} : D(\mathcal{R}_{aux} \otimes \mathcal{C'}^{\otimes m}) \to D(\mathcal{M}^{\otimes m})$. *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.

Definition 8 ($\mathscr{S}$-compactness for an indivisible scheme). *Fix a class of quantum circuits, $\mathscr{S} = \{\mathscr{S}_\kappa\}_{\kappa\in\mathbb{N}}$. An indivisible QHE scheme* QHE *is $\mathscr{S}$-*compact *if there exists a polynomial p such that for any sequence of circuits $\{\mathsf{C}_\kappa \in \mathscr{S}_\kappa\}_\kappa$ with channels $\Phi_{\mathsf{C}_\kappa} : \mathcal{M}^{\otimes n(\kappa)} \to \mathcal{M}^{\otimes m(\kappa)}$, the circuit complexity of applying* $\mathsf{QHE.Dec}^{\otimes m(\kappa)}$ *to the output of* $\mathsf{QHE.Eval}^{\mathsf{C}_\kappa}$ *is at most $p(\kappa, m(\kappa))$.*

The trivial quantum fully homomorphic encryption scheme, TRIV, is easily phrased as an indivisible scheme. Informally, TRIV is defined by taking $\mathsf{TRIV.KeyGen}$ and $\mathsf{TRIV.Enc}$ from any public-key encryption scheme, letting $\mathsf{TRIV.Eval}^{\mathsf{C}}$ append a description of C to the cipherstate, and $\mathsf{TRIV.Dec}$ decode the cipherstate, and then apply C. Clearly, TRIV is homomorphic, but it is not compact, since $\mathsf{TRIV.Dec}$ must evaluate the quantum circuit C, and so its complexity scales with $G(\mathsf{C})$, 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 [14]). We give a formal definition that quantifies this notion for indivisible quantum homomorphic encryption schemes.

Definition 9 (quasi-compactness). *Let $\mathscr{S} = \{\mathscr{S}_\kappa\}_\kappa$ be the set of all quantum circuits over some fixed universal gate set. For any $f : \mathscr{S} \to \mathbb{R}_{\geq 0}$, an indivisible QHE scheme* QHE *is* f-quasi-compact *if there exists a polynomial p such that for any sequence of circuits $\{\mathsf{C}_\kappa \in \mathscr{S}_\kappa\}_\kappa$ with induced channels $\Phi_{\mathsf{C}_\kappa} : \mathcal{M}^{\otimes n(\kappa)} \to \mathcal{M}^{\otimes m(\kappa)}$, the circuit complexity of decrypting the output of $\mathsf{QHE.Eval}^{\mathsf{C}_\kappa}$ is at most $f(\mathsf{C}_\kappa)p(\kappa, m(\kappa))$.*

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 $G(\mathsf{C})$, the number of gates in C.

3.5 Symmetric-Key Quantum Homomorphic Encryption

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 *symmetric-key* quantum homomorphic encryption. In the case of classical FHE, symmetric-key encryption is known to be *equivalent* to public-key encryption [28]. 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 10. *A* symmetric-key *QHE scheme is a quantum homomorphic encryption scheme with* $\mathsf{QHE.KeyGen}$ *and* $\mathsf{QHE.Enc}$ *re-defined as:*

Key Generation. $\mathsf{QHE.KeyGen} : 1^\kappa \to (sk, \rho_{evk})$. *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. $\mathsf{QHE.Enc}_{sk} : D(\mathcal{M}) \to D(\mathcal{C})$. *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 11. *A* bounded *symmetric-key QHE scheme is a symmetric-key QHE scheme with* $\mathsf{QHE.KeyGen}$*,* $\mathsf{QHE.Enc}$*, and* $\mathsf{QHE.Dec}$ *re-defined as:*

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

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

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

We can define q-IND-CPA security for the symmetric-key setting by allowing the adversary access to an encryption oracle $\mathsf{Enc}_{sk}(\cdot)$. We give details in [9].

4 Main Contributions

We now formally state our main results. Our first theorem, Thm. 1, establishes quantum homomorphic encryption for Clifford circuits.

Theorem 1. *(Clifford scheme,* CL*) Let* $\mathscr{S}$ *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* $\mathscr{S}$*-homomorphic.*

Next, we consider two variants of the scheme given by Thm. 1. Each variant deals with non-Clifford T-gates in a different way. The first scheme, described in Thm. 2 and formally defined in Sec. 6, uses entanglement to implement T-gates, resulting in a QHE scheme in which the complexity of decryption scales with the number of T-gates in the homomorphically evaluated circuit.

Theorem 2. *(entanglement-based scheme,* EPR*) Let* $\mathscr{S}$ *be the set of all quantum circuits over the universal gate set* $\{\mathsf{X}, \mathsf{Z}, \mathsf{P}, \mathsf{H}, \mathsf{CNOT}, \mathsf{T}\}$. *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,* $\mathscr{S}$*-homomorphic and* R^2*-quasi-compact, where* $R(\mathsf{C})$ *is the number of* T*-gates in a circuit* C.

The compactness of the scheme EPR is nontrival 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 QHE scheme that is homomorphic for circuits with constant T-depth, as described in the following theorem:

Theorem 3. *(auxiliary-qubit scheme,* AUX*) Fix a constant* L. *Let* $\mathscr{S}$ *be the set of quantum circuits over the universal gate set* $\{\mathsf{X}, \mathsf{Z}, \mathsf{P}, \mathsf{H}, \mathsf{CNOT}, \mathsf{T}\}$ *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,* $\mathscr{S}$*-homomorphic and compact.*

The QHE scheme in Thm. 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. 2 and 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 Homomorphic Encryption for Clifford Circuits: CL

In this section, we present CL, a compact quantum homomorphic encryption scheme for Clifford circuits. This is a building block for the schemes that follow in Sec. 6 and 7. In the full version [9], we prove that CL is q-IND-CPA secure, and homomorphic for Clifford circuits, hence proving Thm. 1.

By definition, Clifford circuits conjugate Pauli operators to Pauli operators [19]. In other words, for any Clifford C, and any Pauli, Q, there exists a Pauli Q' such that $\mathsf{CQ} = \mathsf{Q}'\mathsf{C}$. Furthermore, applying a random Pauli operator is a perfectly secure symmetric-key quantum encryption scheme: the quantum one-time pad. Thus, 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 $\mathsf{Q}|\psi\rangle$ to get $\mathsf{Q}'(\mathsf{C}|\psi\rangle)$. Now decrypting the state requires applying the Pauli Q'. If Q can be described by the encryption key $(a_1, \ldots, a_n, b_1, \ldots, b_n)$ — that is, $\mathsf{Q} = \mathsf{X}^{a_1}\mathsf{Z}^{b_1} \otimes \cdots \otimes \mathsf{X}^{a_n}\mathsf{Z}^{b_n}$ — then Q' can be described by some key $(a'_1, \ldots, a'_n, b'_1, \ldots, b'_n)$ depending on C and $(a_1, \ldots, a_n, b_1, \ldots, b_n)$. We describe this dependence by a function $f^{\mathsf{C}} : \mathbb{F}_2^{2n} \to \mathbb{F}_2^{2n}$, which we call a *key update rule*. We need only consider key update rules for each gate in our gate set, which consists of the one- and two-qubit gates in $\{\mathsf{X}, \mathsf{Z}, \mathsf{P}, \mathsf{CNOT}, \mathsf{H}\}$. 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 $f_a^{\mathsf{C}}, f_b^{\mathsf{C}} : \mathbb{F}_2^2 \to \mathbb{F}_2$ such that when C is applied to the i^{th} wire, $a'_i = f_a^{\mathsf{C}}(a_i, b_i)$ and $b'_i = f_b^{\mathsf{C}}(a_i, b_i)$:

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

For the 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 the full version [9, App. C]. (We also give key update rules for single-qubit measurement and qubit preparation, so that our scheme is actually homomorphic for stabilizer circuits.) By applying these rules after each gate, we can update the key so that the output is correctly decrypted (since we are actually carrying out computations on encrypted quantum data—in contrast to merely simulating a quantum computation—we note that all gates except the Pauli gates require quantum operations). Such a technique was already used, *e.g.* in [6, 10, 13].

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 $\mathsf{c}_1, \ldots, \mathsf{c}_G$, we apply the gates to the quantum one-time pad encrypted message, and homomorphically evaluate the function $f^{\mathsf{c}_1} \circ \cdots \circ f^{\mathsf{c}_G}$ on the encrypted one-time pad keys $a_1, \ldots, a_n, b_1, \ldots, b_n$, 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 (see [9]) is linear, each $f_{a,i}$ and $f_{b,i}$ is a linear polynomial in $\mathbb{F}_2[a_1, \ldots, a_n, b_1, \ldots, b_n]$ (from the perspective of the evaluation procedure, $a_1, \ldots, a_n, b_1, \ldots, b_n$ 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. We note that since the key update rules (see [9]) 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 FHE scheme that is q-IND-CPA secure (see Def. 1). As noted, such a scheme could be derived from [5]. 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 C is the output space of $\mathsf{HE.Enc}$, and C' is the output space of $\mathsf{HE.Eval}$, then we define $\mathcal{C} = \mathbb{C}^{C \times C} \otimes \mathcal{X}$, where $\mathcal{X} \equiv \mathbb{C}^{\{0,1\}}$, and $\mathcal{C}' = \mathbb{C}^{C' \times C'} \otimes \mathcal{X}$.

Key Generation. $\mathsf{CL.KeyGen}(1^\kappa)$. For key generation, execute $(pk, sk, evk) \leftarrow \mathsf{HE.Keygen}(1^\kappa)$. 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. $\mathsf{CL.Enc}_{pk} : D(\mathcal{M}) \to D(\mathcal{C})$. Encryption is defined as:

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

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

Suppose $\mathsf{C} = \mathsf{c}_1, \ldots, \mathsf{c}_G$ is a Clifford circuit.

1. For all $i \in [n]$, set $f_{a,i} \leftarrow a_i$, $f_{b,i} \leftarrow b_i$.
2. For $j = 1, \ldots, G$ such that c_j is a gate or a measurement:
   (a) Apply the gate c_j to the state: $\rho \leftarrow \mathsf{c}_j \rho \mathsf{c}_j^{-1}$.
   (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 $(f_{a,i}, f_{b,i}) \leftarrow (f_{a,i} \circ f_a^{\mathsf{c}_j}, f_{b,i} \circ f_b^{\mathsf{c}_j})$. If c_j is a CNOT-gate acting on wires i and i', update $(f_{a,i}, f_{a,i'}, f_{b,i}, f_{b,i'})$.
3. Update the classical encryptions by computing

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

4. Output $(c_1, \ldots, c_m, \rho)$.

Decryption. $\mathsf{CL.Dec}_{sk} : D(\mathcal{C}') \to D(\mathcal{M})$. For $\tilde{a}, \tilde{b} \in C'$, decryption is defined:
$\mathsf{CL.Dec}_{sk} : |\tilde{a}\rangle\langle\tilde{a}| \otimes |\tilde{b}\rangle\langle\tilde{b}| \otimes \rho^{\mathcal{X}} \mapsto \mathsf{QDec}_{\mathsf{HE.Dec}_{sk}(\tilde{a}),\mathsf{HE.Dec}_{sk}(\tilde{b})}(\rho^{\mathcal{X}})$.

We prove the homomorphic and security properties of CL in [9].

6 T-gate Computation Using Entanglement: EPR

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 $\mathsf{T}\mathsf{X}^a\mathsf{Z}^b = \mathsf{X}^a\mathsf{Z}^{a\oplus b}\mathsf{P}^a\mathsf{T}$. That is, conditioned on a, the output picks up an undesirable P error, which cannot be corrected by applying Pauli corrections. In [10], 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. As long as the evaluation algorithm does not find out if this correction is being executed or not, security holds. The solution in [10] involves quantum interaction; this was recently improved to a single auxiliary qubit, coupled with classical interaction [6, 13]. As a proof technique (for establishing security), [6, 13] 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. 2. Correctness of Fig. 2 is proven in the full version [9].

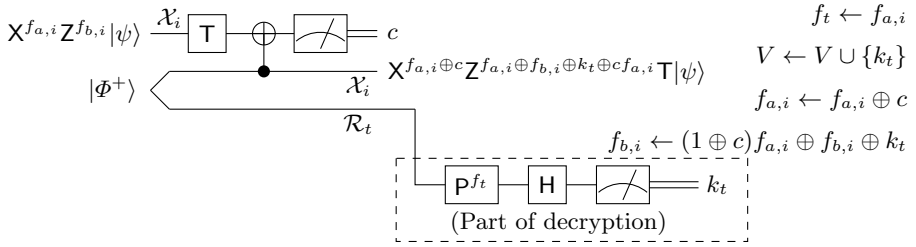

Fig. 2. Evaluation protocol for the t^{th} T-gate, on the i^{th} wire. The key-polynomials $f_{a,i}$ and $f_{b,i}$ are in $\mathbb{F}_2[V]$. After the protocol, V gains a new variable corresponding to the unknown measurement result k_t. The dashed box shows part of the decryption, which happens at some point in the future, after the complete evaluation is finished.

Fig. 2 shows that, using the 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. 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 $\widetilde{f_1}$ that decrypts to $f_1(a_1, b_1, \ldots, a_n, b_n)$. Thus, for this T-gate, the output part of the auxiliary system contains both $\widetilde{f_1}$ and the register $\mathcal{R}_1$. As part of the decryption operation, compute $f_1 \leftarrow \mathsf{HE.Dec}(\widetilde{f_1})$, 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. 2, as well as the linearity of the key update rules, for all t, $f_t \in \mathbb{F}_2[a_1, \ldots, a_n, b_1, \ldots, b_n, k_1, \ldots, k_{t-1}]$ is linear (since c is a known constant), so we can write $f_t = f_t^k + f_t^{ab}$ for $f_t^k \in \mathbb{F}_2[k_1, \ldots, k_{t-1}]$ and $f_t^{ab} \in \mathbb{F}_2[a_1, \ldots, a_n, b_1, \ldots, b_n]$.

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 $O(R^2)$ (see full version [9]).

We now formally define the indivisible QHE 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 $\mathsf{HE.Enc}$ and $\mathcal{X} \equiv \mathbb{C}^{\{0,1\}}$. Since EPR is indivisible, the output space of $\mathsf{EPR.Eval}^{\mathsf{C}}$ has the form $\mathcal{R}_{aux} \otimes \mathcal{C'}^{\otimes m}$. In our case, we have $\mathcal{R}_{aux} = \mathcal{R}_1 \otimes \cdots \otimes \mathcal{R}_R \otimes (\mathbb{C}^{\{0,1\}^{R+1}})^{\otimes R} \otimes (\mathbb{C}^{C'})^{\otimes R}$, where R is the number of T-gates, C' is the output space of $\mathsf{HE.Eval}$, and $\mathcal{R}_t \equiv \mathbb{C}^{\{0,1\}}$. The classical parts of the auxiliary space allow us to output R linear polynomials in $\mathbb{F}_2[k_1, \ldots, k_R]$ corresponding to $\{f_t^k\}_{t=1}^R$, each of which can be represented with $R+1$ bits; as well as R $\mathsf{HE.Eval}$ outputs, corresponding to encryptions of $\{f_t^{ab}(a_1, \ldots, a_n, b_1, \ldots, b_n)\}_{t=1}^R$. Similarly, we have $\mathcal{C}' = (\mathbb{C}^{\{0,1\}^{R+1}})^{\otimes 2} \otimes \mathbb{C}^{C'\times C'} \otimes \mathcal{X}$.

The key generation, $\mathsf{EPR.KeyGen}$, and encryption, $\mathsf{EPR.Enc}$, are defined exactly as $\mathsf{CL.KeyGen}$ and $\mathsf{CL.Enc}$. We now define $\mathsf{EPR.Eval}$ and $\mathsf{EPR.Dec}$.

Evaluation. $\mathsf{EPR.Eval}_{evk}$. As in CL, apply gates in $\{\mathsf{X}, \mathsf{Z}, \mathsf{P}, \mathsf{H}, \mathsf{CNOT}\}$ directly on the encrypted quantum registers. For the T-gate, use the gadget defined in Fig. 2. 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 $\forall\, i \in [n]$, $f_{a,i} \leftarrow a_i$, $f_{b,i} \leftarrow b_i$.
2. Let $\mathsf{g}_1, \ldots, \mathsf{g}_G$ be a topological ordering of the gates in C. For $j = 1, \ldots, G$, evaluate g_j using the appropriate gadget.
3. Let S be the set of output wires. Let $\mathcal{L}$ be the set of labels $\mathcal{L} = \{(a, i), (b, i) : i \in S\} \cup \{1, \ldots, R\}$. 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 $f_\alpha = f^k_\alpha + f^{ab}_\alpha$ for $f^k_\alpha \in \mathbb{F}[k_1, \ldots, k_R]$ and $f^{ab}_\alpha \in \mathbb{F}_2[a_1, \ldots, a_n, b_1, \ldots, b_n]$. Compute $\widetilde{f^{ab}_\alpha} \leftarrow \mathsf{HE.Eval}^{f^{ab}_\alpha}_{evk}(\tilde{a}_1, \ldots, \tilde{a}_n, \tilde{b}_1, \ldots, \tilde{b}_n)$.

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 $\mathcal{R}_1, \ldots, \mathcal{R}_R$ corresponding to auxiliary states created by T-gadgets; the polynomials $\{f^k_\alpha\}_{\alpha \in \mathcal{L}} \subset \mathbb{F}_2[k_1, \ldots, k_R]$ and the homomorphically evaluated polynomials $\{\widetilde{f^{ab}_\alpha}\}_{\alpha \in \mathcal{L}}$.

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

1. For $t = 1, \ldots, R$:
   (a) Decrypt $f^{ab}_t \leftarrow \mathsf{HE.Dec}_{sk}(\widetilde{f^{ab}_t})$.
   (b) Compute $a \leftarrow f^k_t(k_1, \ldots, k_{t-1}) \oplus f^{ab}_t$ 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 $f^{ab}_{a,i} \leftarrow \mathsf{HE.Dec}_{sk}(\widetilde{f^{ab}_{a,i}})$ and $f^{ab}_{b,i} \leftarrow \mathsf{HE.Dec}_{sk}(\widetilde{f^{ab}_{b,i}})$.
   (b) Compute $a_i \leftarrow f^k_{a,i}(k_1, \ldots, k_t) \oplus f^{ab}_{a,i}$ and $b_i \leftarrow f^k_{b,i}(k_1, \ldots, k_t) \oplus f^{ab}_{b,i}$.
3. To each register $\mathcal{X}_i$, apply the map QDec_{a_i,b_i}. Output registers $\mathcal{X}_1, \ldots, \mathcal{X}_m$.

We prove that EPR is homomorphic for all quantum circuits in the universal gate set $\{\mathsf{X}, \mathsf{Z}, \mathsf{P}, \mathsf{CNOT}, \mathsf{H}, \mathsf{T}\}$, R^2-quasi-compact, and q-IND-CPA, in [9].

7 T-gate Computation Using Auxiliary States: AUX

In the previous QHE scheme, we solved the problem of performing the P correction 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, $\mathsf{AUX.Keygen}$ 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 as $|+_{a,k}\rangle := \mathsf{Z}^k\mathsf{P}^a|+\rangle$ (where k is a random bit and a is an encryption key) that are useful for the evaluation of the T-gate (see Fig. 3). 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 $\mathsf{AUX.KeyGen}$ 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

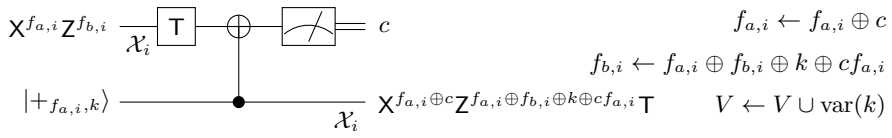

Fig. 3. A T-gadget for the scheme AUX consists of the above circuit and key-update rules. We use $\mathrm{var}(k)$ to denote the set of variables in the polynomial k, which depends on the construction of the auxiliary state $|+_{f_{a,i},k}\rangle$, described below.

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. Fix a constant L. We will now define a scheme $\mathsf{AUX} = \mathsf{AUX}_L$ that is homomorphic for all circuits with T-depth at most L.

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 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:

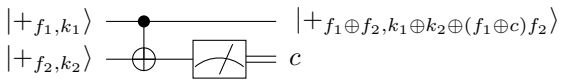

By iterating this procedure, given auxiliary states $|+_{f_1,k_1}\rangle, \ldots, |+_{f_r,k_r}\rangle$, we can construct $|+_{f_1\oplus\cdots\oplus f_r,k}\rangle$, where $k = \bigoplus_{i=1}^{m} k_i \oplus \bigoplus_{i=2}^{r} c_i f_i \oplus \bigoplus_{i=1}^{r}\bigoplus_{j=1}^{i-1} f_i f_j$ for known 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 $|+_{t_1,k_1}\rangle, \ldots, |+_{t_r,k_r}\rangle$, 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 $V \cup \{k_1, \ldots, k_r\}$. 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 the full version [9].

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. 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 $\ell = 1, \ldots, L$.

We now define sets of monomials $T_1, \ldots, T_L$ such that the keys in the ℓ^{th} layer consist of sums of terms from T_ℓ. Let $V_1 := \{a_i, b_i\}_{i\in[n]}$, and define $T_1 \subset \mathbb{F}_2[V_1]$ by $T_1 := \{a_i, b_i\}_{i\in[n]}$. 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, $\{k_{1,i}^{(\ell-1)}, \ldots, k_{|T_{\ell-1}|,i}^{(\ell-1)}\}_{i=1}^n$, 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 [9] for details). Thus, for $\ell > 1$, we can define $T_\ell \subset \mathbb{F}_2[V_\ell]$, where $V_\ell := V_{\ell-1} \cup \{k_{1,i}^{(\ell-1)}, \ldots, k_{|T_{\ell-1}|,i}^{(\ell-1)}\}_{i=1}^n$, by

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

We then provide each of the n wires with an auxiliary state for each term in T_ℓ, for $\ell = 1, \ldots, L$. 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 $a, b \in \{0,1\}^n$, $s^{(1)}(a, b) = (a_1, \ldots, a_n, b_1, \ldots, b_n)$.

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 $k = (k^{(1)}, \ldots, k^{(L)}) \in \{0,1\}^{n|T_1|} \times \cdots \times \{0,1\}^{n|T_L|}$, define

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

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

1. Execute $(pk, sk, evk) \leftarrow$ HE.KeyGen$(1^{\kappa+n})$.

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

Encryption. $\mathsf{AUX.Enc}_{(sk,a,b,k),d} : D(\mathcal{M}) \to D(\mathcal{C})$. 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 $\mathsf{AUX.Enc}$). If $d \leq n$, it applies the quantum one-time pad channel $\mathsf{QEnc}_{a_d,b_d} : D(\mathcal{M}) \to D(\mathcal{C})$. 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. $\mathsf{AUX.Dec}_{(sk,a,b,k),d} : D(\mathcal{C}') \to D(\mathcal{M})$. The decryption is defined the same as $\mathsf{CL.Dec}_{sk}$.

Homomorphic Evaluation. $\mathsf{AUX.Eval}^{\mathsf{C}} : D(\mathcal{R}_{evk} \otimes \mathcal{C}^{\otimes n}) \to D(\mathcal{C}'^{\otimes m})$. For Clifford group gates, we apply the gadgets as in $\mathsf{CL.Eval}$. For T-gates, we apply the gadget in Fig. 3. The full evaluation procedure is as follows:

1. Set $V \leftarrow \{a_i, b_i\}_{i \in [n]}$, and $\forall\, i \in [n]$, $f_{a,i} \leftarrow a_i$, $f_{b,i} \leftarrow b_i$.
2. Let $\mathsf{g}_1, \ldots, \mathsf{g}_G$ be a topological ordering of the gates in C. For $i = 1, \ldots, G$, evaluate 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: $\tilde{a}_i \leftarrow \mathsf{HE.Eval}_{evk}^{f_{a,i}}(\tilde{v} : v \in V)$ and $\tilde{b}_i \leftarrow \mathsf{HE.Eval}_{evk}^{f_{b,i}}(\tilde{v} : v \in V)$.
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 corresponding to the encrypted output qubit i of the circuit.

The correctness of this scheme depends on two facts, which we prove in [9]. First, for every unknown $v \in V$, we have an encrypted copy of $\tilde{v}$, encrypted using $\mathsf{HE.Enc}$. 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. In the full version [9], we prove that AUX requires $O(n^{2^{L-1}+1})$ auxiliary qubits, from which it follows that AUX is homomorphic for quantum circuits with T-depth L. We further show that AUX is q-IND-CPA and compact.

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 $|+_{s,k}\rangle\langle+_{s,k}|$, $\mathsf{HE.Enc}_{pk}(s)$ — 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.

References

1. 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.
2. 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.
3. P. Arrighi and L. Salvail. Blind quantum computation. *International Journal of Quantum Information*, 4:883–898, 2006.
4. 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.
5. 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.
6. A. Broadbent. Delegating private quantum computations. `arXiv:1506.01328[quant-ph]`, to appear in *Canadian Journal of Physics*, 2015.
7. 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.
8. A. Broadbent, G. Gutoski, and D. Stebila. Quantum one-time programs. In *Advances in Cryptology (CRYPTO 2013)*, pages 344–360, 2013.
9. A. Broadbent and S. Jeffery. Quantum homomorphic encryption for circuits of low T-gate complexity (full version). `arXiv:1412.8766[quant-ph]`, 2014.
10. A. Childs. Secure assisted quantum computation. *Quantum Information and Computation*, 5:456–466, 2005.
11. 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.
12. T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In *Advances in Cryptology (CRYPTO '85)*, pages 10–18, 1985.
13. 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.
14. C. Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. `crypto.stanford.edu/craig`.
15. 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.
16. 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.
17. 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.
18. S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270 – 299, 1984.
19. 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.
20. T. Koshiba. 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`.
21. D. Micciancio and O. Regev. Lattice-based cryptography. In *Post-Quantum Cryptography*. Springer, 2009.
22. T. Okamoto, K. Tanaka, and S. Uchiyama. Quantum public-key cryptosystems. In *Advances in Cryptology (CRYPTO 2000)*, pages 147–165, 2000.
23. P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Advances in Cryptology (EUROCRYPT '99)*, pages 223–238. 1999.
24. 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.
25. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–34:40, September 2009.
26. R. Rivest, L. Adleman, and M. Dertouzos. On data banks and privacy homomorphisms. In *Foundations of Secure Computation*, pages 169–177, 1978.
27. P. P. Rohde, J. F. Fitzsimons, and A. Gilchrist. Quantum walks with encrypted data. *Physical Review Letters*, 109:150501, Oct 2012.
28. 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.
29. 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.
30. 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.
31. 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.
32. 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.
33. 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.
34. W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.
35. 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.
36. 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.