

Quantum Fully Homomorphic Encryption With Verification

Gorjan Alagic^{1,2}, Yfke Dulek³, Christian Schaffner³, and Florian Speelman⁴

¹ Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, MD

² National Institute of Standards and Technology, Gaithersburg, MD

³ CWI, QuSoft, and University of Amsterdam

⁴ QMATH, Department of Mathematical Sciences, University of Copenhagen

Abstract. Fully-homomorphic encryption (FHE) enables computation on encrypted data while maintaining secrecy. Recent research has shown that such schemes exist even for quantum computation. Given the numerous applications of classical FHE (zero-knowledge proofs, secure two-party computation, obfuscation, etc.) it is reasonable to hope that quantum FHE (or QFHE) will lead to many new results in the quantum setting. However, a crucial ingredient in almost all applications of FHE is *circuit verification*. Classically, verification is performed by checking a transcript of the homomorphic computation. Quantumly, this strategy is impossible due to no-cloning. This leads to an important open question: can quantum computations be delegated and verified in a non-interactive manner?

In this work, we answer this question in the affirmative, by constructing a scheme for QFHE with verification (vQFHE). Our scheme provides authenticated encryption, and enables arbitrary polynomial-time quantum computations without the need of interaction between client and server. Verification is almost entirely classical; for computations that start and end with classical states, it is completely classical. As a first application, we show how to construct quantum one-time programs from classical one-time programs and vQFHE.

1 Introduction

The 2009 discovery of fully-homomorphic encryption (FHE) in classical cryptography is widely considered to be one of the major breakthroughs of the field. Unlike standard encryption, FHE enables non-interactive computation on encrypted data even by parties that do not hold the decryption key. Crucially, the input, output, and all intermediate states of the computation remain encrypted, and thus hidden from the computing party. While FHE has some obvious applications (e.g., cloud computing), its importance in cryptography stems from its wide-ranging applications to other cryptographic scenarios. For instance, FHE can be used to construct secure two-party computation, efficient zero-knowledge proofs for NP, and indistinguishability obfuscation [3, 14]. In fact, the breadth of its usefulness has led some to dub FHE “the swiss army knife of cryptography” [3].

Recent progress on constructing quantum computers has led to theoretical research on “cloud-based” quantum computing. In such a setting, it is natural to ask whether users can keep their data secret from the server that performs the quantum computation. A recently-constructed quantum fully-homomorphic encryption (QFHE) scheme shows that this can be done in a single round of interaction [12]. This discovery raises an important question: do the numerous classical applications of FHE have suitable quantum analogues? As it turns out, most of the classical applications require an additional property which is simple classically, but non-trivial quantumly. That property is *verification*: the ability of the user to check that the final ciphertext produced by the server is indeed the result of a particular computation, homomorphically applied to the initial user-generated ciphertext. In the classical case, this is a simple matter: the server makes a copy of each intermediate computation step, and provides the user with all these copies. In the quantum case, such a “transcript” would appear to violate no-cloning. The user simply checks a transcript generated by the server. In the quantum case, this would violate no-cloning. In fact, one might suspect that the no-cloning theorem prevents non-interactive quantum verification *in principle*.

In this work, we show that verification of homomorphic quantum computations is in fact possible. We construct a new QFHE scheme which allows the server to generate a “computation log” which can certify to the user that a particular homomorphic quantum computation was performed on the ciphertext. The computation log itself is purely classical, and most (in some cases, all) of the verification can be performed on a classical computer. Unlike in all previous quantum homomorphic schemes, the underlying encryption is now authenticated.

Verification immediately yields new applications of QFHE, e.g., allowing users of a “quantum cloud service” to certify the server’s computations. Verified QFHE (or vQFHE) also leads to a simple construction of quantum one-time programs (qOTPs) [8]. In this construction, the qOTP for a functionality Φ consists of an evaluation key and a classical OTP which performs vQFHE verification for Φ only. Finding other applications of vQFHE (including appropriate analogues of all classical applications) is the subject of ongoing work.

Related Work. Classical FHE was first constructed by Gentry in 2009 [15]. For us, the scheme of Brakerski and Vaikuntanathan [4] is of note: it has decryption in \mathbf{NC}^1 and is believed to be quantum-secure. Quantumly, partially-homomorphic (or partially-compact) schemes were proposed by Broadbent and Jeffery [5]. The first fully-homomorphic (leveled) scheme was constructed by Dulek, Schaffner and Speelman [12]. Recently, Mahadev proposed a scheme, based on classical indistinguishability obfuscation, in which the user is completely classical [17]. A parallel line of work has attempted to produce QFHE with information-theoretic security [23, 19, 21, 18]. There has also been significant research on delegating quantum computation interactively (see, e.g., [1, 7, 10]). Another notable interactive approach is quantum computation on authenticated data (QCAD), which was used to construct quantum one-time programs from classical one-time programs [8] and zero-knowledge proofs for QMA [9].

Summary of Results. Our results concern a new primitive: verified QFHE. A standard QFHE scheme consists of four algorithms: KeyGen, Enc, Eval and Dec [5, 12]. We define vQFHE similarly, with two changes: (i.) Eval provides an extra classical “computation log” output; (ii.) decryption is now called VerDec, and accepts a ciphertext, a circuit description C , and a computation log. Informally, correctness then demands that, for all keys k and circuits C acting on plaintexts,

$$\text{VerDec}_k^C \circ \text{Eval}_{\text{ev}k}^C \circ \text{Enc}_k = \Phi_C. \quad (1)$$

A crucial parameter is the relative difficulty of performing C and VerDec_k^C . In a nontrivial scheme, the latter must be simpler. In our case, C is an arbitrary poly-size quantum circuit and VerDec_k^C is almost entirely classical.

Security of verified QFHE. Informally, security should require that, if a server deviates significantly from the map Eval_k^C in (1), then VerDec_k^C will reject.

1. **Semantic security (SEM-VER).** Consider a QPT adversary \mathcal{A} which manipulates a ciphertext (and side info) and declares a circuit, as in Figure 1 (top). This defines a channel $\Phi_{\mathcal{A}} := \text{VerDec} \circ \mathcal{A} \circ \text{Enc}$. A simulator \mathcal{S} does not receive or output a ciphertext, but does declare a circuit; this defines a channel $\Phi_{\mathcal{S}}$ which first runs \mathcal{S} and then runs a circuit on the plaintext based on the outputs of \mathcal{S} . We say that a vQFHE scheme is semantically secure (SEM-VER) if for all adversaries \mathcal{A} there exists a simulator \mathcal{S} such that the channels $\Phi_{\mathcal{A}}$ and $\Phi_{\mathcal{S}}$ are computationally indistinguishable.
2. **Indistinguishability (IND-VER).** Consider the following security game. Based on a hidden coin flip b , \mathcal{A} participates in one of two protocols. For $b = 0$, this is normal vQFHE. For $b = 1$, this is a modified execution, where we secretly swap out the plaintext $\rho_{\mathcal{A}}$ to a private register (replacing it with a fixed state), apply the desired circuit to $\rho_{\mathcal{A}}$, and then swap $\rho_{\mathcal{A}}$ back in; we then discard this plaintext if VerDec rejects the outputs of \mathcal{A} . Upon receiving the final plaintext of the protocol, \mathcal{A} must guess the bit b . A vQFHE scheme is IND-VER if, for all \mathcal{A} , the success probability is at most $1/2 + \text{negl}(n)$.
3. **New relations between security definitions.** If we restrict SEM-VER to empty circuit case, we recover (the computational version of) the definition of quantum authentication [13, 6]. SEM-VER (resp., IND-VER) generalizes computational semantic security SEM (resp., indistinguishability IND) for quantum encryption [5, 2]. We generalize $\text{SEM} \Leftrightarrow \text{IND}$ [2] as follows.

Theorem 1. *A vQFHE scheme satisfies SEM-VER iff it satisfies IND-VER.*

A scheme for vQFHE for poly-size quantum circuits. Our main result is a vQFHE scheme which admits verification of arbitrary polynomial-size quantum circuits. The verification in our scheme is almost entirely classical. In fact, we can verify classical input/output computations using purely classical verification. The main technical ingredients are (i.) classical FHE with NC^1 decryption [4], (ii.) the trap code for computing on authenticated quantum data [20, 8, 6], and (iii.) the “garden-hose gadgets” from the first QFHE scheme [12]. The scheme is called TrapTP; a brief sketch is as follows.

1. **Key Generation (KeyGen).** We generate keys for the classical FHE scheme, as well as some encrypted auxiliary states (see evaluation below). This procedure requires the generation of single-qubit and two-qubit states from a small fixed set, performing Bell measurements and Pauli gates, and executing the encoding procedure of a quantum error-correcting code on which the trap code is based.
2. **Encryption (Enc).** We encrypt each qubit of the plaintext using the trap code, and encrypt the trap code keys using the FHE scheme. This again requires the ability to perform Paulis, execute an error-correcting encoding, and the generation of basic single-qubit states.
3. **Evaluation (Eval).** Paulis and CNOT are evaluated as in the trap code; keys are updated via FHE evaluation. To measure a qubit, we measure all ciphertext qubits and place the outcomes in the log. To apply P or H, we use encrypted magic states (from the eval key) plus the aforementioned gates. Applying T requires a magic state and an encrypted “garden-hose gadget” (because the T-gate magic state circuit applies a P-gate conditioned on a measurement outcome). In addition to all of the measurement outcomes, the log also contains a transcript of all the classical FHE computations.
4. **Verified decryption (VerDec).** We check the correctness and consistency of the classical FHE transcript, the measurement outcomes, and the claimed circuit. The result of this computation is a set of keys for the trap code, which are correct provided that Eval was performed honestly. We decrypt using these keys and output either a plaintext or reject. In terms of quantum capabilities, decryption requires executing the decoding procedure of the error-correcting code, computational-basis and Hadamard-basis measurements, and Paulis.

Our scheme is *compact*: the number of elementary quantum operations performed by VerDec scales only with the size of the plaintext, and *not* with the size of the circuit performed via Eval. We do require that VerDec performs a classical computation which can scale with the size of the circuit; this is reasonable since VerDec must receive the circuit as input. Like the other currently-known schemes for QFHE, our scheme is leveled, in the sense that pre-generated auxiliary magic states are needed to perform the evaluation procedure.

Theorem 2 (Main result, informal). *Let TrapTP be the scheme outlined above, and let VerDec^{\equiv} be VerDec for the case of verifying the empty circuit.*

1. *The vQFHE scheme TrapTP satisfies IND-VER security.*
2. *The scheme (KeyGen, Enc, VerDec^{\equiv}) is authenticating [13] and IND-CPA [5].*

Application: quantum one-time programs. A one-time program (or OTP) is a device which implements a circuit, but self-destructs after the first use. OTPs are impossible without hardware assumptions, even with quantum states; OTPs that implement quantum circuits (qOTP) can be built from classical OTPs (cOTP) [8]. As a first application of vQFHE, we give another simple construction of qOTPs. Our construction is weaker, since it requires a computational assumption. On the other hand, it is conceptually very simple and serves to demonstrate the

power of verification. In our construction, the qOTP for a quantum circuit C is simply a (vQFHE) encryption of C together with a cOTP for verifying the universal circuit. To use the resulting qOTP, the user attaches their desired input, homomorphically evaluates the universal circuit, and then plugs their computation log into the cOTP to retrieve the final decryption keys.

Preliminaries. Our exposition assumes a working knowledge of basic quantum information and the associated notation. As for the particular notation of quantum gates, the gates (H, P, CNOT) generate the so-called Clifford group (which can also be defined as the normalizer of the Pauli group); it includes the Pauli gates X and Z. In order to implement arbitrary unitary operators, it is sufficient to add the T gate (also known as the $\pi/8$ gate). Finally, we can reach universal quantum computation by adding single-qubit measurements in the computational basis.

We will frequently make use of several standard cryptographic ingredients, as follows. The quantum one-time pad (QOTP) will be used for information-theoretically secret one-time encryption. In its encryption phase, two bits $a, b \in \{0, 1\}$ are selected at random, and the map $X^a Z^b$ is applied to the input, projecting it to the maximally-mixed state. We will also need the computational security notions for quantum secrecy, including indistinguishability (IND, IND-CPA) [5] and semantic security (SEM) [2]. For quantum authentication, we will refer to the security definition of Dupuis, Nielsen and Salvail [13]. We will also make frequent use of the trap code for quantum authentication, described below in Section 3. For a security proof and methods for interactive computation on this code, see [8]. Finally, we will also use classical fully-homomorphic encryption (FHE). In brief, an FHE scheme consists of classical algorithms (KeyGen, Enc, Eval, Dec) for (respectively) generating keys, encrypting plaintexts, homomorphically evaluating circuits on ciphertexts, and decrypting ciphertexts. We will use FHE schemes which are quantum-secure and whose Dec circuits are in NC^1 (see, e.g., [4]).

2 A new primitive: verifiable QFHE

We now define verified quantum fully-homomorphic encryption (or vQFHE), in the symmetric-key setting. The public-key case is a straightforward modification.

Basic definition. The definition has two parameters: the class \mathcal{C} of circuits which the user can verify, and the class \mathcal{V} of circuits which the user needs to perform in order to verify. We are interested in cases where \mathcal{C} is stronger than \mathcal{V} .

Definition 1 (vQFHE). Let \mathcal{C} and \mathcal{V} be (possibly infinite) collections of quantum circuits. A $(\mathcal{C}, \mathcal{V})$ -vQFHE scheme is a set of four QPT algorithms:

- KeyGen : $\{1\}^\kappa \rightarrow \mathcal{K} \times \mathcal{D}(\mathcal{H}_E)$ (security parameter \rightarrow private key, eval key);
- Enc : $\mathcal{K} \times \mathcal{D}(\mathcal{H}_X) \rightarrow \mathcal{D}(\mathcal{H}_C)$ (key, ptext \rightarrow ctext);
- Eval : $\mathcal{C} \times \mathcal{D}(\mathcal{H}_{CE}) \rightarrow \mathcal{L} \times \mathcal{D}(\mathcal{H}_C)$ (circuit, eval key, ctext \rightarrow log, ctext);
- VerDec : $\mathcal{K} \times \mathcal{C} \times \mathcal{L} \times \mathcal{D}(\mathcal{H}_C) \rightarrow \mathcal{D}(\mathcal{H}_X) \times \{\text{acc, rej}\}$

such that (i.) the circuits of VerDec belong to the class \mathcal{V} , and (ii.) for all $(sk, \rho_{\text{evk}}) \leftarrow \text{KeyGen}$, all circuits $c \in \mathcal{C}$, and all $\rho \in \mathfrak{D}(\mathcal{H}_{XR})$,

$$\left\| \text{VerDec}_{sk}(c, \text{Eval}(c, \text{Enc}_k(\rho), \rho_{\text{evk}})) - \Phi_c(\rho) \otimes |\text{acc}\rangle\langle\text{acc}| \right\|_1 \leq \text{negl}(\kappa),$$

where R is a reference and the maps implicitly act on appropriate spaces.

We will refer to condition (ii.) as *correctness*. It is implicit in the definition that the classical registers \mathcal{K}, \mathcal{L} and the quantum registers E, X, C are really infinite families of registers, each consisting of $\text{poly}(\kappa)$ -many (qu)bits. In some later definitions, it will be convenient to use a version of VerDec which also outputs a copy of the (classical) description of the circuit c .

Compactness. We note that there are trivial vQFHE schemes for some choices of $(\mathcal{C}, \mathcal{V})$ (e.g., if $\mathcal{C} \subset \mathcal{V}$, then the user can simply authenticate the ciphertext and then perform the computation during decryption). Earlier work on quantum and classical homomorphic encryption required compactness, meaning that the size of the decrypt circuit should not scale with the size of the homomorphic circuit.

Definition 2 (Compactness of QFHE). *A QFHE scheme S is compact if there exists a polynomial $p(\kappa)$ such that for any circuit C with n_{out} output qubits, and for any input ρ_X , the complexity of applying $S.\text{Dec}$ to $S.\text{Eval}^C(S.\text{Enc}_{sk}(\rho_X), \rho_{\text{evk}})$ is at most $p(n_{\text{out}}, \kappa)$.*

When considering QFHE *with* verification, however, some tension arises. On one hand, trivial schemes like the above still need to be excluded. On the other hand, verifying that a circuit C has been applied requires reading a description of C , which violates Definition 2. We thus require a more careful consideration of the relationship between the desired circuit $C \in \mathcal{C}$ and the verification circuit $V \in \mathcal{V}$. In our work, we will allow the number of classical gates in V to scale with the size of C . We propose a new definition of compactness in this context.

Definition 3 (Compactness of vQFHE (informal)). *A vQFHE scheme S is compact if $S.\text{VerDec}$ is divisible into a classical verification procedure $S.\text{Ver}$ (outputting only an accept/reject flag), followed by a quantum decryption procedure $S.\text{Dec}$. The running time of $S.\text{Ver}$ is allowed to depend on the circuit size, but the running time of $S.\text{Dec}$ is not.*

The procedure $S.\text{Dec}$ is not allowed to receive and use any other information from $S.\text{Ver}$ than whether or not it accepts or rejects. This prevents the classical procedure $S.\text{Ver}$ from de facto performing part of the decryption work (e.g., by computing classical decryption keys). In Section 3, we will see a scheme that does not fulfill compactness for this reason.

Definition 4 (Compactness of vQFHE (formal)). *A vQFHE scheme S is compact if there exists a polynomial p such that $S.\text{VerDec}$ can be written as $S.\text{Dec} \circ S.\text{Ver}$, and the output ciphertext space $\mathfrak{D}(\mathcal{H}_C)$ can be written as a classical-quantum state space $\mathcal{A} \times \mathfrak{D}(\mathcal{H}_B)$, where (i.) $S.\text{Ver} : \mathcal{K} \times \mathcal{C} \times \mathcal{L} \times \mathcal{A} \rightarrow \{\text{acc}, \text{rej}\}$ is*

a classical polynomial-time algorithm, and (ii.) $S.\text{Dec} : \{\text{acc}, \text{rej}\} \times \mathcal{K} \times \mathfrak{D}(\mathcal{H}_C) \rightarrow \mathfrak{D}(\mathcal{H}_X) \times \{\text{acc}, \text{rej}\}$ is a quantum algorithm such that for any circuit C with n_{out} output qubits and for any input ρ_X , $S.\text{Dec}$ runs in time $p(n_{\text{out}}, \kappa)$ on the output of $S.\text{Eval}^C(S.\text{Enc}(\rho_X), \rho_{\text{evk}})$.

Note that in the above definition, the classical registers \mathcal{K} and \mathcal{A} are copied and fed to both $S.\text{Dec}$ and $S.\text{Ver}$.

For privacy, we say that a vQFHE scheme is private if its ciphertexts are indistinguishable under chosen plaintext attack (IND-CPA) [5, 12].

Secure verifiability. In this section, we formalize the concept of verifiability. Informally, one would like the scheme to be such that whenever VerDec accepts, the output can be trusted to be close to the desired output. We will consider two formalizations of this idea: a semantic one, and an indistinguishability-based one.

Our semantic definition will state that every adversary with access to the ciphertext can be simulated by a simulator that only has access to an ideal functionality that simply applies the claimed circuit. It is inspired by quantum authentication [13, 6] and semantic secrecy [2].

The real-world scenario (Figure 1, top) begins with a state $\rho_{XR_1R_2}$ prepared by a QPT (“message generator”) \mathcal{M} . The register X (plaintext) is subsequently encrypted and sent to the adversary \mathcal{A} . The registers R_1 and R_2 contain side information. The adversary acts on the ciphertext and R_1 , producing some output ciphertext $C_{X'}$, a circuit description c , and a computation log \log . These outputs are then sent to the verified decryption function. The output, along with R_2 , is sent to a distinguisher \mathcal{D} , who produces a bit 0 or 1.

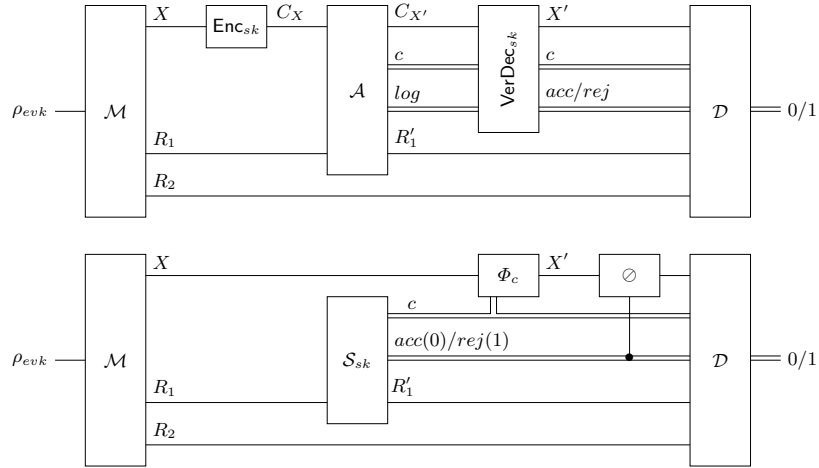


Fig. 1. The real-world (top) and ideal-world (bottom) for SEM-VER.

In the ideal-world scenario (Figure 1, bottom), the plaintext X is not encrypted or sent to the simulator \mathcal{S} . The simulator outputs a circuit c and chooses whether to accept or reject. The channel Φ_c implemented by c is applied to the input register X directly. If reject is chosen, the output register X' is traced out and replaced by the fixed state Ω ; this controlled-channel is denoted $\overline{\text{ctrl}}-\Omega$.

Definition 5 (κ -SEM-VER). A vQFHE scheme $(\text{KeyGen}, \text{Enc}, \text{Eval}, \text{VerDec})$ is semantically κ -verifiable if for any QPT adversary \mathcal{A} , there exists a QPT \mathcal{S} such that for all QPTs \mathcal{M} and \mathcal{D} ,

$$\left| \Pr \left[\mathcal{D} \left(\text{Real}_{sk}^{\mathcal{A}}(\mathcal{M}(\rho_{evk})) \right) = 1 \right] - \Pr \left[\mathcal{D} \left(\text{Ideal}_{sk}^{\mathcal{S}}(\mathcal{M}(\rho_{evk})) \right) = 1 \right] \right| \leq \text{negl}(\kappa),$$

where $\text{Real}_{sk}^{\mathcal{A}} = \text{VerDec}_{sk} \circ \mathcal{A} \circ \text{Enc}_{sk}$ and $\text{Ideal}_{sk}^{\mathcal{S}} = \overline{\text{ctrl}}-\Omega \circ \Phi_c \circ \mathcal{S}_{sk}$, and the probability is taken over $(\rho_{evk}, sk) \leftarrow \text{KeyGen}(1^\kappa)$ and all QPTs above.

Note that the simulator (in the ideal world) gets the secret key sk . We believe that this is necessary, because the actions of an adversary may depend on superficial properties of the ciphertext. In order to successfully simulate this, the simulator needs to be able to generate (authenticated) ciphertexts. He cannot do so with a fresh secret key, because the input plaintext may depend on the correlated evaluation key ρ_{evk} . Fortunately, the simulator does not become too powerful when in possession of the secret key, because he does not receive any relevant plaintexts or ciphertexts to encrypt or decrypt: the input register X is untouchable for the simulator.

Next, we present an alternative definition of verifiability, based on a security game motivated by indistinguishability.

Game 1 For an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$, a scheme S , and a security parameter κ , the $\text{VerGame}_{\mathcal{A}, S}(\kappa)$ game proceeds as depicted in Figure 2.

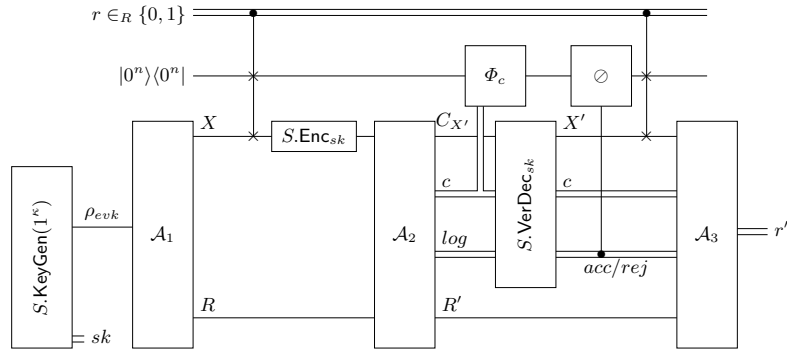


Fig. 2. The indistinguishability game $\text{VerGame}_{\mathcal{A}, S}(\kappa)$, as used in the definition of κ -IND-VER.

The game is played in several rounds. Based on the evaluation key, the adversary first chooses an input (and some side information in R). Based on a random bit b this input is either encrypted and sent to \mathcal{A}_2 (if $b = 0$), or swapped out and replaced by a dummy input $|0^n\rangle\langle 0^n|$ (if $b = 1$). If $b = 1$, the ideal channel Φ_c is applied by the challenger, and the result is swapped back in right before the adversary (in the form of \mathcal{A}_3) has to decide on its output bit b' . If \mathcal{A}_2 causes a reject, the real result is also erased by the channel \circlearrowleft . We say that the adversary *wins* (expressed as $\text{VerGame}_{\mathcal{A},S}(\kappa) = 1$) whenever $b' = b$.

Definition 6 (κ -IND-VER). A vQFHE scheme S has κ -indistinguishable verification if for any QPT adversary \mathcal{A} , $\Pr[\text{VerGame}_{\mathcal{A},S}(\kappa) = 1] \leq \frac{1}{2} + \text{negl}(\kappa)$.

Theorem 3. A vQFHE scheme is κ -IND-VER iff it is κ -SEM-VER.

Proof (sketch). The forward direction is shown by contraposition. Given an adversary \mathcal{A} , define a simulator \mathcal{S} that encrypts a dummy 0-state, then runs \mathcal{A} , and then VerDec. For this simulator, there exist \mathcal{M} and \mathcal{D} such that the difference in acceptance probability between the real and the ideal scenario is nonnegligible. The triple $(\mathcal{M}, \mathcal{A}, \mathcal{D})$ forms an adversary for the VER indistinguishability game.

For the reverse direction, we use the following approach. From an arbitrary adversary \mathcal{A} for the IND-VER indistinguishability game, we define a semantic adversary, message generator, and distinguisher, that together simulate the game for \mathcal{A} . The fact that S is κ -SEM-VER allows us to limit the advantage of the semantic adversary over any simulator, and thereby the winning probability of \mathcal{A} .

For a detailed proof, see the supplementary material. \square

3 TC: A partially-homomorphic scheme with verification

We now present a partially-homomorphic scheme with verification, which will serve as a building block for the fully-homomorphic scheme in Section 4. It is called TC (for “trap code”), and is homomorphic only for CNOT, (classically controlled) Paulis, and measurement in the computational and Hadamard basis. It does not satisfy compactness: as such, it performs worse than the trivial scheme where the client performs the circuit at decryption time. However, TC lays the groundwork for the vQFHE scheme we present in Section 4, and as such is important to understand in detail. It is a variant of the trap-code scheme presented in [8] (which requires classical interaction for T gates), adapted to our vQFHE framework. A variation also appears in [9], and implicitly in [20].

Setup and encryption. Let CSS be a (public) self-dual $[[m, 1, d]]$ CSS code, so that H and CNOT are transversal. CSS can correct d_c errors, where $d = 2d_c + 1$. We choose $m = \text{poly}(d)$ and large enough that $d_c = \kappa$ where κ is the security parameter. The concatenated Steane code satisfies all these requirements.

We generate the keys as follows. Choose a random permutation $\pi \in_R S_{3m}$ of $3m$ letters. Let n be the number of qubits that will be encrypted. For each

$i \in \{1, \dots, n\}$, pick bit strings $x[i] \in_R \{0, 1\}^{3m}$ and $z[i] \in_R \{0, 1\}^{3m}$. The secret key sk is the tuple $(\pi, x[1], z[1], \dots, x[n], z[n])$, and ρ_{evk} is left empty.

Encryption is per qubit: (i.) the state σ is encoded using CSS, (ii.) m computational and m Hadamard ‘traps’ ($|0\rangle$ and $|+\rangle$ states, see [8]) are added, (iii.) the resulting $3m$ qubits are permuted by π , and (iv.) the overall state is encrypted with a quantum one-time pad (QOTP) as dictated by $x = x[i]$ and $z = z[i]$ for the i th qubit. We denote the ciphertext by $\tilde{\sigma}$.

Evaluation. First, consider Pauli gates. By the properties of CSS, applying a logical Pauli is done by applying the same Pauli to all physical qubits. The application of Pauli gates (X and/or Z) to a state encrypted with a quantum one-time pad can be achieved without touching the actual state, by updating the keys to QOTP in the appropriate way. This is a classical task, so we can postpone the application of the Pauli to VerDec (recall it gets the circuit description) without giving up compactness for TC. So, formally, the evaluation procedure for Pauli gates is the identity map. Paulis conditioned on a classical bit b which will be known to VerDec at execution time (e.g., a measurement outcome) can be applied in the same manner.

Next, we consider CNOT. To apply a CNOT to encrypted qubits σ_i and σ_j , we apply CNOT transversally between the $3m$ qubits of $\tilde{\sigma}_i$ and the $3m$ qubits of $\tilde{\sigma}_j$. Ignoring the QOTP for the moment, the effect is a transversal application of CNOT on the physical data qubits (which, by CSS properties, amounts to logical CNOT on $\sigma_i \otimes \sigma_j$), and an application of CNOT between the $2m$ pairs of trap qubits. Since $\text{CNOT}|00\rangle = |00\rangle$ and $\text{CNOT}|++\rangle = |++\rangle$, the traps are unchanged. Note that CNOT commutes with the Paulis that form the QOTP. In particular, for all $a, b, c, d \in \{0, 1\}$, $\text{CNOT}(X_1^a Z_1^b \otimes X_2^c Z_2^d) = (X_1^a Z_1^{b \oplus d} \otimes X_2^{a \oplus c} Z_2^d) \text{CNOT}$. Thus, updating the secret-key bits (a, b, c, d) to $(a, b \oplus d, a \oplus c, d)$ finishes the job. The required key update happens in TC.VerDec (see below).

Next, consider computational-basis measurements. For CSS, logical measurement is performed by measurement of all physical qubits, followed by a classical decoding procedure [8]. In TC.Eval, we measure all $3m$ ciphertext qubits. During TC.VerDec, the contents of the measured qubits (now a classical string $a \in \{0, 1\}^{3m}$) will be interpreted into a logical measurement outcome.

Finally, we handle Hadamard-basis measurements. A transversal application of H to all $3m$ relevant physical qubits precedes the evaluation procedure for the computational basis measurement. Since CSS is self-dual, this applies a logical H. Since $H|0\rangle = |+\rangle$ and $H|+\rangle = |0\rangle$, all computational traps are swapped with the Hadamard traps. This is reflected in the way TC.VerDec checks the traps (see the supplementary material for details). Note that this is a classical procedure (and thus its accept/reject output flag is classical).

Verification and decryption. If a qubit is unmeasured after evaluation (as stated in the circuit), TC.VerDecQubit is applied. This removes the QOTP, undoes the permutation, checks all traps, and decodes the qubit. See the supplementary material for a specification of this algorithm.

If a qubit is measured during evaluation, `TC.VerDec` receives a list \tilde{w} of $3m$ physical measurement outcomes for that qubit. These outcomes are classically processed (removing the QOTP by flipping bits, undoing π , and decoding CSS) to produce the plaintext measurement outcome. Note that we only check the $|0\rangle$ traps in this case. Intuitively, this should not affect security, since any attack that affects only $|+\rangle$ but not $|0\rangle$ will be canceled by computational basis measurement.

The complete procedure `TC.VerDec` updates the QOTP keys according to the gates in the circuit description, and then decrypts all qubits and measurement results as described above (see the supplementary material for details).

Correctness, compactness, and privacy. For honest evaluation, `TC.VerDec` accepts with probability 1. Correctness is straightforward to check by following the description in Section 3. For privacy, note that the final step in the encryption procedure is the application of a (information-theoretically secure) QOTP with fresh, independent keys. If IND-CPA security is desired, one could easily extend TC by using a pseudorandom function for the QOTP, as in [2].

TC is not compact in the sense of Definition 4, however. In order to compute the final decryption keys, the whole gate-by-gate key update procedure needs to be executed, aided by the computation log and information about the circuit. Thus, we cannot break `TC.VerDec` up into two separate functionalities, `Ver` and `Dec`, where `Dec` can successfully retrieve the keys and decrypt the state, based on only the output ciphertext and the secret key.

Security of verification. The trap code is proven secure in its application to one-time programs [8]. Broadbent and Wainwright proved authentication security (with an explicit, efficient simulator) [6]. One can use similar strategies to prove κ -IND-VER for TC. In fact, TC satisfies a stronger notion of verifiability, where the adversary is allowed to submit plaintexts in multiple rounds (letting the choice of the next plaintext depend on the previous ciphertext), which are either all encrypted or all swapped out. Two rounds (κ -IND-VER-2) are sufficient for us; the definitions and proof (see the supplementary material) extend straightforwardly to the general case κ -IND-VER- i for $i \in \mathbb{N}_+$.

Theorem 4. *TC is κ -IND-VER-2 for the above circuit class.*

4 TrapTP: Quantum FHE With Verification

In this section, we introduce our candidate scheme for verifiable quantum fully homomorphic encryption (vQFHE). In this section, we will define the scheme prove correctness, compactness, and privacy. We will show verifiability in Section 5.

Let $\kappa \in \mathbb{N}$ be a security parameter, and let $t, p, h \in \mathbb{N}$ be an upper bound on the number of T, P, and H gates (respectively) that will be in the circuit which is to be homomorphically evaluated. As in Section 3, we fix a self-dual $[[m, 1, d]]$ CSS code CSS which has $m = \text{poly}(d)$ and can correct $d_c := \kappa$ errors

(e.g., the concatenated Steane code). We also fix a classical fully homomorphic public-key encryption scheme HE with decryption in LOGSPACE (see, e.g., [4]). Finally, fix a message authentication code $\text{MAC} = (\text{Tag}, \text{Ver})$ that is existentially unforgeable under adaptive chosen message attacks (EUF-CMA [16]) from a quantum adversary; for example, one may take the standard pseudorandom-function construction with a post-quantum PRF. This defines an authentication procedure $\text{MAC.Sign}_k : m \mapsto (m, \text{MAC.Tag}_k(m))$.

Key generation and encryption. The evaluation key will require a number of auxiliary states, which makes the key generation algorithm `TrapTP.KeyGen` somewhat involved (see Algorithm 1 and Algorithm 2). Note that non-evaluation keys are generated first, and then used to encrypt auxiliary states which are included in the evaluation key (see `TrapTP.Enc` below). Most states are encrypted using the same ‘global’ permutation π , but all qubits in the error-correction gadget (except first and last) are encrypted using independent permutations π_i (see line 15). The T-gate gadgets are prepared by Algorithm 2, making use of garden-hose gadgets from [12].

Algorithm 1. `TrapTP.KeyGen`($1^\kappa, 1^t, 1^p, 1^h$)

```

1:  $k \leftarrow \text{MAC.KeyGen}(1^\kappa)$ 
2:  $\pi \leftarrow_R S_{3m}$  ▷  $S_{3m}$  is the permutation group on  $3m$  elements
3: for  $i = 0, \dots, t$  do
4:    $(sk_i, pk_i, evk_i) \leftarrow \text{HE.KeyGen}(1^\kappa)$ 
5:  $sk \leftarrow (\pi, k, sk_0, \dots, sk_t, pk_0)$ 
6: for  $i = 1, \dots, p$  do ▷ Magic-state generation for P
7:    $\mu_i^P \leftarrow \text{TrapTP.Enc}(sk, P|+)$  ▷ See Algorithm 3 for TrapTP.Enc
8: for  $i = 1, \dots, t$  do ▷ Magic-state generation for T
9:    $\mu_i^T \leftarrow \text{TrapTP.Enc}(sk, T|+)$ 
10: for  $i = 1, \dots, h$  do ▷ Magic-state generation for H
11:    $\mu_i^H \leftarrow \text{TrapTP.Enc}(sk, \frac{1}{\sqrt{2}}(H \otimes \mathbb{I})(|00\rangle + |11\rangle))$ 
12: for  $i = 1, \dots, t$  do ▷ Gadget generation for T
13:    $\pi_i \leftarrow_R S_{3m}$ 
14:    $(g_i, \gamma_i^{\text{in}}, \gamma_i^{\text{mid}}, \gamma_i^{\text{out}}) \leftarrow \text{TrapTP.GadgetGen}(sk_{i-1})$  ▷ See Algorithm 2
15:    $\Gamma_i \leftarrow \text{MAC.Sign}(\text{HE.Enc}_{pk_i}(g_i, \pi_i)) \otimes \text{TrapTP.Enc}((\pi_i, k, sk_0, \dots, sk_t, pk_i), \gamma_i^{\text{mid}}) \otimes$   

    $\text{TrapTP.Enc}(sk, \gamma_i^{\text{in}}, \gamma_i^{\text{out}})$ 
16:  $keys \leftarrow \text{MAC.Sign}(evk_0, \dots, evk_t, pk_0, \dots, pk_t, \text{HE.Enc}_{pk_0}(\pi))$ 
17:  $\rho_{evk} \leftarrow (keys, \mu_0^P, \dots, \mu_p^P, \mu_0^T, \dots, \mu_t^T, \mu_0^H, \dots, \mu_h^H, \Gamma_1, \dots, \Gamma_t)$ 
18: return  $(sk, \rho_{evk})$ 

```

Algorithm 2. `TrapTP.GadgetGen`(sk_i)

```

1:  $g_i \leftarrow g(sk_i)$  ▷ classical description of the garden-hose gadget, see [12], p. 13
2:  $(\gamma_i^{\text{in}}, \gamma_i^{\text{mid}}, \gamma_i^{\text{out}}) \leftarrow$  generate  $|\Phi^+\rangle$  states and arrange them as described by  $g_i$ . Call  

   the first qubit  $\gamma_i^{\text{in}}$  and the last qubit  $\gamma_i^{\text{out}}$ . The rest forms the state  $\gamma_i^{\text{mid}}$ .
3: return  $(g_i, \gamma_i^{\text{in}}, \gamma_i^{\text{mid}}, \gamma_i^{\text{out}})$ 

```

The encryption of a quantum state is similar to TC.Enc , only the keys to the QOTP are now chosen during encryption (rather than during key generation) and appended in encrypted and authenticated form to the ciphertext (see Algorithm 3). Note that the classical secret keys sk_0 through sk_t are not used.

Algorithm 3. $\text{TrapTP.Enc}((\pi, k, sk_0, \dots, sk_t, pk), \sigma)$

- 1: $\tilde{\sigma} \leftarrow \sum_{x, z \in \{0,1\}^{3m}} \left(\text{TC.Enc}((\pi, x, z), \sigma) \otimes \text{MAC.Sign}_k(\text{HE.Enc}_{pk}(x, z)) \right) \triangleright$ Algorithm 13
 - 2: **return** $\tilde{\sigma}$
-

Evaluation. Evaluation of gates is analogous to the interactive evaluation scheme using the trap code [8], except the interactions are replaced by classical homomorphic evaluation. Evaluation of a circuit c is done gate-by-gate, as follows.

In general, we will use the notation $\tilde{\cdot}$ to denote encryptions of classical and quantum states. For example, in the algorithms below, $\tilde{\sigma}$ is the encrypted input qubit for a gate and \tilde{x} and \tilde{z} are classical encryptions of the associated QOTP keys. We will assume throughout that HE.Enc and HE.Eval produce, apart from their actual output, a complete *computation log* describing all randomness used, all computation steps, and all intermediate results.

Measurements. Computational basis measurement is described in Algorithm 4. Recall that $\text{TC.VerDecMeasurement}$ is a completely classical procedure that decodes the list of $3m$ measurement outcomes into the logical outcome and checks the relevant traps. Hadamard-basis measurement is performed similarly, except the qubits of $\tilde{\sigma}$ are measured in the Hadamard basis and $\text{HE.Enc}_{pk}(\times)$ is given as the last argument for the evaluation of $\text{TC.VerDecMeasurement}$.

Algorithm 4. $\text{TrapTP.EvalMeasure}(\tilde{\sigma}, \tilde{x}, \tilde{z}, \tilde{\pi}, pk, evk)$

- 1: $a = (a_1, \dots, a_{3m}) \leftarrow$ measure qubits of $\tilde{\sigma}$ in the computational basis
 - 2: $(\tilde{a}, log_1) \leftarrow \text{HE.Enc}_{pk}(a)$
 - 3: $(\tilde{b}, flag, log_2) \leftarrow \text{HE.Eval}_{evk}^{\text{TC.VerDecMeasurement}}((\tilde{\pi}, \tilde{x}, \tilde{z}), \tilde{a}, \text{HE.Enc}_{pk}(+))$
 - 4: **return** $(\tilde{b}, flag, log_1, log_2) \triangleright b \in \{0, 1\}$ represents the output of the measurement
-

Pauli gates. A logical Pauli-X is performed by (homomorphically) flipping the X-key bits of the QOTP (see Algorithm 5). Since this is a classical operation, the functionality extends straightforwardly to a classically controlled Pauli-X (by specifying an additional bit b encrypted into \tilde{b} that indicates whether or not X should be applied; see Algorithm 6). The (classically controlled) evaluation of a Pauli-Z works the same way, only the relevant bits in \tilde{z} are flipped.

Algorithm 5. $\text{TrapTP.EvalX}(\tilde{\sigma}, \tilde{x}, \tilde{\pi}, pk, evk)$

- 1: $(\tilde{x}, log_1) \leftarrow \text{HE.Eval}_{evk}^{\text{unpermute}}(\tilde{\pi}, \tilde{x})$
 - 2: $(\tilde{x}, log_2) \leftarrow \text{HE.Eval}_{evk}^{\oplus}(\tilde{x}, \text{HE.Enc}_{pk}(1^m 0^{2m})) \triangleright$ this flips the first m bits
 - 3: $(\tilde{x}, log_3) \leftarrow \text{HE.Eval}_{evk}^{\text{permute}}(\tilde{\pi}, \tilde{x})$
 - 4: **return** $(\tilde{\sigma}, \tilde{x}, log_1, log_2, log_3)$
-

Algorithm 6. $\text{TrapTP.EvalCondX}(\tilde{b}, \tilde{\sigma}, \tilde{x}, \tilde{z}, \tilde{\pi}, pk, evk)$


- 1: $(\tilde{x}, \log_1) \leftarrow \text{HE.Eval}_{evk}^{\text{unpermute}}(\tilde{\pi}, \tilde{x})$
 - 2: $\tilde{s} \leftarrow \text{HE.Eval}_{evk}^{y \rightarrow y^{m_0} 2^m}(\tilde{b})$
 - 3: $(\tilde{x}, \log_2) \leftarrow \text{HE.Eval}_{evk}^{\oplus}(\tilde{x}, \tilde{s})$ ▷ this conditionally flips the first m bits
 - 4: $(\tilde{x}, \log_3) \leftarrow \text{HE.Eval}_{evk}^{\text{permute}}(\tilde{\pi}, \tilde{x})$
 - 5: **return** $(\tilde{\sigma}, \tilde{x}, \tilde{z}, \log_1, \log_2, \log_3)$
-

CNOT gates. The evaluation of CNOT in TrapTP is analogous to TC, only the key updates are performed homomorphically during evaluation (see Algorithm 7).

Algorithm 7. $\text{TrapTP.EvalCNOT}(\tilde{\sigma}_1, \tilde{\sigma}_2, \tilde{x}_1, \tilde{x}_2, \tilde{z}_1, \tilde{z}_2, \tilde{\pi}, pk, evk)$


- 1: $(\tilde{\sigma}_1, \tilde{\sigma}_2) \leftarrow$ apply CNOT on all physical qubit pairs of $\tilde{\sigma}_1, \tilde{\sigma}_2$
 - 2: $(\tilde{x}_1, \tilde{x}_2, \tilde{z}_1, \tilde{z}_2, \log_1) \leftarrow \text{HE.Eval}_{evk}^{\text{CNOT-key-update}}(\tilde{x}_1, \tilde{x}_2, \tilde{z}_1, \tilde{z}_2)$ ▷ for commutation rules, see Section 3
 - 3: **return** $(\tilde{\sigma}_1, \tilde{\sigma}_2, \tilde{x}_1, \tilde{x}_2, \tilde{z}_1, \tilde{z}_2, \log_1, \log_2)$
-

Phase gates. Performing a P gate requires homomorphic evaluation of all the above gates: (classically controlled) Paulis, CNOTs, and measurements. We also consume the state μ_i^P (an encryption of the state $P|+\rangle$) for the i th phase gate in the circuit. The circuit below applies P to the data qubit (see, e.g., [8]).




We define TrapTP.EvalP to be the concatenation of the corresponding gate evaluations. The overall computation log is just a concatenation of the logs.

Hadamard gate. The Hadamard gate can be performed using the same ingredients as the phase gate [8]. The i th gate consumes μ_i^H , an encryption of $(H \otimes I)|\Phi^+\rangle$.



The T gate. A magic-state computation of T uses a similar circuit to that for P, using μ_i^T (an encryption of $T|+\rangle$) as a resource for the i th T gate:



The evaluation of this circuit is much more complicated, since it requires the application of a classically-controlled phase correction P. We will accomplish this using the error-correction gadget Γ_i .

First, we remark on some subtleties regarding the encrypted classical information surrounding the gadget. Since the structure of Γ_i depends on the classical secret key sk_{i-1} , the classical information about Γ_i is encrypted under the (independent) public key pk_i (see Algorithm 1). This observation will play a crucial role in our proof that TrapTP satisfies IND-VER, in Section 5.

The usage of two different key sets also means that, at some point during the evaluation of a T gate, all classically encrypted information needs to be reencrypted from the $(i-1)$ st into the i th key set. This can be done because \widetilde{sk}_{i-1} is included in the classical information g_i in Γ_i . The reryption is performed right before the classically-controlled phase gate is applied (see Algorithm 8).

Algorithm 8. $\text{TrapTP.EvalT}(\widetilde{\sigma}, \widetilde{x}, \widetilde{z}, \widetilde{\pi}, \mu_i^\top, \Gamma_i, pk_{i-1}, evk_{i-1}, pk_i, evk_i)$

- 1: $(\widetilde{\sigma}_1, \widetilde{\sigma}_2, \widetilde{x}_1, \widetilde{z}_1, \widetilde{x}_2, \widetilde{z}_2, log_1) \leftarrow \text{TrapTP.EvalCNOT}(\mu_i^\top, \widetilde{\sigma}, \widetilde{x}, \widetilde{z}, \widetilde{\pi}, pk_{i-1}, evk_{i-1})$
 - 2: $(\widetilde{b}, log_2) \leftarrow \text{TrapTP.EvalMeasure}(\widetilde{\sigma}_2, \widetilde{x}_2, \widetilde{z}_2, \widetilde{\pi}, pk_{i-1}, evk_{i-1})$
 - 3: $log_3 \leftarrow$ reencrypt *all* classically encrypted information (except \widetilde{b}) from key set $i-1$ into key set i .
 - 4: $(\widetilde{\sigma}, log_4) \leftarrow \text{TrapTP.EvalCondP}(\widetilde{b}, \widetilde{\sigma}_1, \widetilde{x}_1, \widetilde{z}_1, \Gamma_i, \widetilde{\pi}, pk_i, evk_i)$
 - 5: **return** $(\widetilde{\sigma}, log_1, log_2, log_3, log_4)$
-

Algorithm 9 shows how to use Γ_i to apply logical P on an encrypted quantum state $\widetilde{\sigma}$, conditioned on a classical bit b for which only the encryption \widetilde{b} is available. When TrapTP.EvalCondP is called, b is encrypted under the $(i-1)$ st classical HE-key, while all other classical information (QOTP keys x and z , permutations π and π_i , classical gadget description g_i) is encrypted under the i th key. Note that we can evaluate Bell measurements using only evaluation of CNOT, computational-basis measurements, and H-basis measurements. In particular, no magic states are needed to perform a Bell measurement. After this procedure, the data is in qubit $\widetilde{\gamma}_i^{\text{out}}$. The outcomes a_1, a_2, a of the Bell measurements determine how the keys to the QOTP must be updated.

Algorithm 9. $\text{TrapTP.EvalCondP}(\widetilde{b}, \widetilde{\sigma}, \widetilde{x}, \widetilde{z}, \Gamma_i = (\widetilde{g}_i, \widetilde{\pi}_i, \widetilde{\gamma}_i^{\text{in}}, \widetilde{\gamma}_i^{\text{mid}}, \widetilde{\gamma}_i^{\text{out}}), \widetilde{\pi}, pk_i, evk_i)$

- 1: $(\widetilde{a}_1, \widetilde{a}_2, log_1) \leftarrow$ evaluate Bell measurement between $\widetilde{\sigma}$ and $\widetilde{\gamma}_i^{\text{in}} \triangleright a_1, a_2 \in \{0, 1\}$
 - 2: $(\widetilde{a}, log_2) \leftarrow$ evaluate Bell measurements in $\widetilde{\gamma}_i^{\text{mid}}$ as dictated by the ciphertext \widetilde{b} and the garden-hose protocol for HE.Dec
 - 3: $(\widetilde{x}, \widetilde{z}, log_3) \leftarrow \text{HE.Eval}_{evk_i}^{\text{T-key-update}}(\widetilde{x}, \widetilde{z}, \widetilde{a}_1, \widetilde{a}_2, \widetilde{a}, \widetilde{g}_i)$
 - 4: **return** $(\widetilde{\gamma}_i^{\text{out}}, \widetilde{x}, \widetilde{z}, log_1, log_2, log_3)$
-

Verified Decryption. The decryption procedure (Algorithm 10) consists of two parts. First, we perform several classical checks. This includes MAC-verification of all classically authenticated messages, and checking that the gates listed in the log match the circuit description. We also check the portions of the log which specify

the (purely classical, FHE) steps taken during HE.Enc and HE.Eval; this is the standard transcript-checking procedure for FHE, which we call `TrapTP.CheckLog`. Secondly, we check all unmeasured traps and decode the remaining qubits. We reject if `TrapTP.CheckLog` rejects, or if the traps have been triggered.

Algorithm 10. `TrapTP.VerDec`($sk, \tilde{\sigma}, (\widetilde{x[i]})_i, (\widetilde{z[i]})_i, log, c$)

- 1: Verify classically authenticated messages (in log) using k (contained in sk). If one of these verifications rejects, **reject**.
 - 2: Check whether all claimed gates in log match the structure of c . If not, **return** $(\Omega, |rej\rangle)$. ▷ Recall that Ω is a dummy state.
 - 3: $flag \leftarrow \text{TrapTP.CheckLog}(log)$ If $flag = rej$, **return** $(\Omega, |rej\rangle)$.
 - 4: Check whether the claimed final QOTP keys in the log match \tilde{x} and \tilde{z} . If not, **return** $(\Omega, |rej\rangle)$.
 - 5: **for all** gates G of c **do**
 - 6: **if** G is a measurement **then**
 - 7: $\tilde{x}', \tilde{z}' \leftarrow$ encrypted QOTP keys right before measurement (listed in log)
 - 8: $\tilde{w} \leftarrow$ encrypted measurement outcomes (listed in log)
 - 9: $x', z', w \leftarrow \text{HE.Dec}_{sk_t}(\tilde{x}', \tilde{z}', \tilde{w})$
 - 10: Execute `TC.VerDecMeasurement`((π, x', z'), w , $basis$), where $basis$ is the appropriate basis for the measurement, and store the (classical) outcome.
 - 11: **if** a trap is triggered **then**
 - 12: **return** $(\Omega, |rej\rangle)$.
 - 13: **for all** unmeasured qubits $\tilde{\sigma}_i$ in $\tilde{\sigma}$ **do**
 - 14: $x[i], z[i] \leftarrow \text{HE.Dec}_{sk_t}(x[i], z[i])$
 - 15: $\sigma_i \leftarrow \text{TC.VerDec}_{(\pi, x[i], z[i])}(\tilde{\sigma}_i)$. If `TC.VerDec` rejects, **return** $(\Omega, |rej\rangle)$.
 - 16: $\sigma \leftarrow$ the list of decrypted qubits (and measurement outcomes) that are part of the output of c
 - 17: **return** $(\sigma, |acc\rangle)$
-

4.1 Correctness, compactness, and privacy

If all classical computation was unencrypted, checking correctness of `TrapTP` can be done by inspecting the evaluation procedure for the different types of gates, and comparing them to the trap code construction in [8]. This suffices, since HE and the MAC authentication both satisfy correctness.

Compactness as defined in Definition 4 is also satisfied: verifying the computation log and checking all intermediate measurements (up until line 12 in Algorithm 10) is a completely classical procedure and runs in polynomial time in its input. The rest of `TrapTP.VerDec` (starting from line 13) only uses the secret key and the ciphertext $(\tilde{\sigma}, \tilde{x}, \tilde{z})$ as input, not the log or the circuit description. Thus, we can separate `TrapTP.VerDec` into two algorithms `Ver` and `Dec` as described in Definition 4, by letting the second part (`Dec`, lines 13 to 17) reject whenever the first part (`Ver`, lines 1 to 12) does. It is worth noting that, because the key-update steps are performed homomorphically during the evaluation phase, skipping the classical verification step yields a QFHE scheme without verification that satisfies Definition 2 (and is authenticating). This is not the case

for the scheme TC, where the classical computation is necessary for the correct decryption of the output state.

In terms of privacy, TrapTP satisfies IND-CPA (see Section 2). This is shown by reduction to IND-CPA of HE. This is non-trivial since the structure of the error-correction gadgets depends on the classical secret key. The reduction is done in steps, where first the security of the encryptions under pk_t is applied (no gadget depends on sk_t), after which the quantum part of the gadget I_t (which depends on sk_{t-1}) looks completely mixed from the point of view of the adversary. We then apply indistinguishability of the classical encryptions under pk_{t-1} , and repeat the process. After all classical encryptions of the quantum one-time pad keys are removed, the encryption of a state appears fully mixed. Full details of this proof can be found in Lemma 1 of [12], where IND-CPA security of an encryption function very similar to TrapTP.Enc is proven.

5 Proof of verifiability for TrapTP

In this section, we will prove that TrapTP is κ -IND-VER. By Theorem 3, it then follows that TrapTP is also verifiable in the semantic sense. We will define a slight variation on the VER indistinguishability game, followed by several hybrid schemes (variations of the TrapTP scheme) that fit into this new game. We will argue that for any adversary, changing the game or scheme does not significantly affect the winning probability. After polynomially-many such steps, we will have reduced the adversary to an adversary for the somewhat homomorphic scheme TC, which we already know to be IND-VER. This will complete the argument that TrapTP is IND-VER. The IND-VER game is adjusted as follows.

Definition 7 (Hybrid game $\text{Hyb}_{\mathcal{A},S}(\kappa)$). For an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$, a scheme S , and security parameter κ , $\text{Hyb}_{\mathcal{A},S}(\kappa)$ is the game in Figure 3.

Comparing to Definition 1, we see that three new wires are added: a classical wire from $S.\text{Enc}$ to $S.\text{VerDec}$, and a classical and quantum wire from $S.\text{KeyGen}$ to $S.\text{VerDec}$. We will later adjust TrapTP to use these wires to bypass the adversary; TrapTP as defined in the previous section does not use them. Therefore, for any adversary, $\Pr[\text{VerGame}_{\mathcal{A},\text{TrapTP}}(\kappa) = 1] = \Pr[\text{Hyb}_{\mathcal{A},\text{TrapTP}}(\kappa) = 1]$.

Hybrid 1: Removing Classical MAC. In TrapTP, the initial keys to the QOTP can only become known to VerDec through the adversary. We thus use MAC to make sure these keys cannot be altered. Without this authentication, the adversary could, e.g., homomorphically use $\tilde{\pi}$ to flip only those bits in \tilde{x} that correspond to non-trap qubits, thus applying X to the plaintext. In fact, all classical information in the evaluation key must be authenticated.

In the first hybrid, we argue that the winning probability of a QPT \mathcal{A} in $\text{Hyb}_{\mathcal{A},\text{TrapTP}}(\kappa)$ is at most negligibly higher than in $\text{Hyb}_{\mathcal{A},\text{TrapTP}'}(\kappa)$, where TrapTP' is a modified version of TrapTP where the initial keys are sent directly from KeyGen and Enc to VerDec (via the extra wires above). More precisely, in TrapTP' .KeyGen and TrapTP' .Enc, whenever $\text{MAC}.\text{Sign}(\text{HE}.\text{Enc}(x))$ or

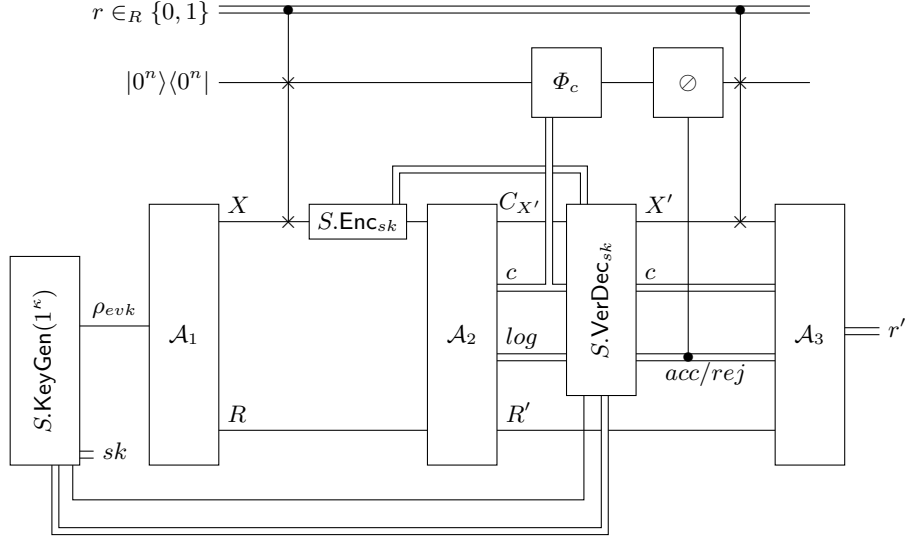


Fig. 3. The hybrid indistinguishability game $\text{Hyb}_{\mathcal{A},S}(\kappa)$, which is a slight variation on $\text{VerGame}_{\mathcal{A},S}(\kappa)$ from Figure 2.

$\text{MAC.Sign}(x)$ is called, the message x is also sent directly to $\text{TrapTP}'.\text{VerDec}$. Moreover, instead of decrypting the classically authenticated messages sent by the adversary, $\text{TrapTP}'.\text{VerDec}$ uses the information it received directly from $\text{TrapTP}'.\text{KeyGen}$ and $\text{TrapTP}'.\text{Enc}$. It still check whether the computation log provided by the adversary contains these values at the appropriate locations and whether the MAC signature is correct. The following fact is then a straightforward consequence of the EUF-CMA property of MAC.

Recall that all adversaries are QPTs, i.e., quantum polynomial-time uniform algorithms. Given two hybrid games H_1, H_2 , and a QPT adversary \mathcal{A} , define

$$\text{AdvHyb}_{H_1}^{H_2}(\mathcal{A}, \kappa) := |\Pr[\text{Hyb}_{\mathcal{A},H_1}(\kappa) = 1] - \Pr[\text{Hyb}_{\mathcal{A},H_2}(\kappa) = 1]|.$$

Lemma 1. For any QPT \mathcal{A} , $\text{AdvHyb}_{\text{TrapTP}}^{\text{TrapTP}'}(\mathcal{A}, \kappa) \leq \text{negl}(\kappa)$.

Hybrid 2: Removing Computation Log. In TrapTP and TrapTP' , the adversary (homomorphically) keeps track of the keys to the QOTP and stores encryptions of all intermediate values in the computation log. Whenever VerDec needs to know the value of a key (for example to check a trap or to decrypt the final output state), the relevant entry in the computation log is decrypted.

In TrapTP' , however, the plaintext initial values to the computation log are available to VerDec , as they are sent through the classical side channels. This means that whenever VerDec needs to know the value of a key, instead of decrypting an entry to the computation log, it can be computed by “shadowing” the computation log in the clear.

For example, suppose the log contains the encryptions \tilde{b}_1, \tilde{b}_2 of two initial bits, and specifies the homomorphic evaluation of XOR, resulting in \tilde{b} where $b = b_1 \oplus b_2$. If one knows the plaintext values b_1 and b_2 , then one can compute $b_1 \oplus b_2$ directly, instead of decrypting the entry \tilde{b} from the computation log.

We now define a second hybrid, TrapTP'' , which differs from TrapTP' exactly like this: VerDec still verifies the authenticated parts of the log, checks whether the computation log matches the structure of c , and checks whether it is syntactically correct. However, instead of decrypting values from the log (as it does in $\text{TrapTP}.\text{VerDec}$, Algorithm 10, on lines 9 and 14), it computes those values from the plaintext initial values, by following the computation steps that are claimed in the log. By correctness of classical FHE, we then have the following.

Lemma 2. *For any QPT \mathcal{A} , $\text{AdvHyb}_{\text{TrapTP}'}^{\text{TrapTP}''}(\mathcal{A}, \kappa) \leq \text{negl}(\kappa)$.*

Proof. Let s be the (plaintext) classical information that forms the input to the classical computations performed by the adversary: initial QOTP keys, secret keys and permutations, measurement results, et cetera. Let f be the function that the adversary computes on it in order to arrive at the final keys and logical measurement results. By correctness of HE, we have that

$$\Pr[\text{HE.Dec}_{sk_t}(\text{HE.Eval}_{evk_0, \dots, evk_t}^f(\text{HE.Enc}_{pk_0}(s))) \neq f(s)] \leq \text{negl}(\kappa).$$

In the above expression, we slightly abuse notation and write $\text{HE.Eval}_{evk_0, \dots, evk_t}$ to include the t reryption steps that are performed during $\text{TrapTP}.\text{Eval}$. As long as the number of T gates, and thus the number of reryptions, is polynomial in κ , the expression holds.

Thus, the probability that $\text{TrapTP}'.\text{VerDec}$ and $\text{TrapTP}''.\text{VerDec}$ use different classical values (decrypting from the log vs. computing from the initial values) is negligible. Since this is the only place where the two schemes differ, the output of the two VerDec functions will be identical, except with negligible probability. Thus \mathcal{A} will either win in both $\text{Hyb}_{\mathcal{A}, \text{TrapTP}'}(\kappa)$ and $\text{Hyb}_{\mathcal{A}, \text{TrapTP}''}(\kappa)$, or lose in both, again except with negligible probability. \square

More Hybrids: Removing Gadgets. We continue by defining a sequence of hybrid schemes based on TrapTP'' . In $4t$ steps, we will move all error-correction functionality from the gadgets to VerDec . This will imply that the adversary has no information about the classical secret keys (which are involved in constructing these gadgets). This will allow us to eventually reduce the security of TrapTP to that of TC.

We remove the gadgets back-to-front, starting with the final gadget. Every gadget is removed in four steps. For all $1 \leq \ell \leq t$, define the hybrids $\text{TrapTP}_1^{(\ell)}$, $\text{TrapTP}_2^{(\ell)}$, $\text{TrapTP}_3^{(\ell)}$, and $\text{TrapTP}_4^{(\ell)}$ (with $\text{TrapTP}_4^{(t+1)} := \text{TrapTP}''$) as follows:

1. $\text{TrapTP}_1^{(\ell)}$ is the same as $\text{TrapTP}_4^{(\ell+1)}$ (or, in the case that $\ell = t$, the same as TrapTP''), except for the generation of the state Γ_ℓ (see Algorithm 1, line 15). In $\text{TrapTP}_1^{(\ell)}$, all classical information encrypted under pk_ℓ is replaced

with encryptions of zeros. In particular, for $i \geq \ell$, line 15 is adapted to

$$\begin{aligned} \Gamma_i &\leftarrow \text{MAC.Sig}(\text{HE.Enc}_{pk_i}(00 \cdots 0)) \\ &\otimes \text{TrapTP}''.\text{Enc}'(sk', \gamma_i^{\text{mid}}) \otimes \text{TrapTP}.\text{Enc}(sk, \gamma_i^{\text{in}} \otimes \gamma_i^{\text{out}}) \end{aligned}$$

where $\text{TrapTP}''.\text{Enc}'$ also appends a signed encryption of zeros, effectively replacing line 1 in Algorithm 3 with

$$\tilde{\sigma} \leftarrow \sum_{x, z \in \{0,1\}^{3m}} \left(\text{TC.Enc}((\pi, x, z), \sigma) \otimes \text{MAC.Sig}_k(\text{HE.Enc}_{pk}(00 \cdots 0)) \right)$$

It is important to note that in both KeyGen and Enc' , the information that is sent to VerDec through the classical side channel is *not* replaced with zeros. Hence, the structural and encryption information about Γ_ℓ is kept from the adversary, and instead is directly sent (only) to the verification procedure. Whenever VerDec needs this information, it is taken directly from this trusted source, and the all-zero string sent by the adversary will be ignored.

2. $\text{TrapTP}_2^{(\ell)}$ is the same as $\text{TrapTP}_1^{(\ell)}$, except that for the ℓ th gadget, the procedure $\text{TrapTP}.\text{PostGadgetGen}$ is called instead of $\text{TrapTP}.\text{GadgetGen}$:

Algorithm 11. $\text{TrapTP}.\text{PostGadgetGen}(sk_i)$

- 1: $g_i \leftarrow 0^{|g(sk_i)|}$
 - 2: $(\gamma_i^{\text{in}}, \gamma_i^{\text{mid}}, \gamma_i^{\text{out}}) \leftarrow$ halves of EPR pairs (send other halves to VerDec)
 - 3: **return** $(g_i, \gamma_i^{\text{in}}, \gamma_i^{\text{mid}}, \gamma_i^{\text{out}})$
-

This algorithm produces a ‘gadget’ in which all qubits are replaced with halves of EPR pairs. These still get encrypted in line 15 of Algorithm 1. All other halves of these EPR pairs are sent to VerDec through the provided quantum channel. $\text{TrapTP}_2^{(\ell)}.\text{VerDec}$ has access to the structural information g_ℓ (as this is sent via the classical side information channel from KeyGen to VerDec) and performs the necessary Bell measurements to recreate γ_ℓ^{in} , γ_ℓ^{mid} and γ_ℓ^{out} after the adversary has interacted with the EPR pair halves. Effectively, this postpones the generation of the gadget structure to decryption time. Of course, the measurement outcomes are taken into account by VerDec when calculating updates to the quantum one-time pad. As can be seen from the description of $\text{TrapTP}_4^{(\ell)}$, all corrections that follow the ℓ th one are unaffected by the fact that the server cannot hold the correct information about these postponed measurements, not even in encrypted form.

3. $\text{TrapTP}_3^{(\ell)}$ is the same as $\text{TrapTP}_2^{(\ell)}$, except that gadget generation for the ℓ th gadget is handled by $\text{TrapTP}.\text{FakeGadgetGen}$ instead of $\text{TrapTP}.\text{PostGadgetGen}$.

Algorithm 12. $\text{TrapTP}.\text{FakeGadgetGen}(sk_i)$

- 1: $g_i \leftarrow 0^{|g(sk_i)|}$
- 2: $(\gamma_i^{\text{in}}, \gamma_i^{\text{mid}}, \gamma_i^{\text{out}}) \leftarrow$ halves of EPR pairs (send other halves to VerDec)
- 3: Send γ_i^{mid} to VerDec as well
- 4: **return** $(g_i, \gamma_i^{\text{in}}, |00 \cdots 0\rangle, \gamma_i^{\text{out}})$

This algorithm prepares, instead of halves of EPR pairs, $|0\rangle$ -states of the appropriate dimension for γ_ℓ^{mid} . (Note that this dimension does not depend on $sk_{\ell-1}$). For γ_ℓ^{in} and γ_ℓ^{out} , halves of EPR pairs are still generated, as in $\text{TrapTP}_2^{(\ell)}$. Via the side channel, the full EPR pairs for γ_ℓ^{mid} are sent to VerDec . As in the previous hybrids, the returned gadget is encrypted in TrapTP.KeyGen .

$\text{TrapTP}_3^{(\ell)}.VerDec$ verifies that the adversary performed the correct Bell measurements on the fake ℓ th gadget by calling TC.VerDec . If this procedure accepts, $\text{TrapTP}_3^{(\ell)}.VerDec$ performs the verified Bell measurements on the halves of the EPR pairs received from $\text{TrapTP}_3^{(\ell)}.KeyGen$ (and subsequently performs the Bell measurements that depend on g_ℓ on the other halves, as in $\text{TrapTP}_2^{(\ell)}$). Effectively, $\text{TrapTP}_3^{(\ell)}.VerDec$ thereby performs a protocol for HE.Dec , removing the phase error in the process.

4. $\text{TrapTP}_4^{(\ell)}$ is the same as $\text{TrapTP}_3^{(\ell)}$, except that VerDec (instead of performing the Bell measurements of the gadget protocol) uses its knowledge of the initial QOTP keys and all intermediate measurement outcomes to compute whether or not a phase correction is necessary after the ℓ th T gate. $\text{TrapTP}_4^{(\ell)}.VerDec$ then performs this phase correction on the EPR half entangled with γ_ℓ^{in} , followed by a Bell measurement with the EPR half entangled with γ_ℓ^{out} .

The first $\ell - 1$ gadgets in $\text{TrapTP}_1^{(\ell)}$ through $\text{TrapTP}_4^{(\ell)}$ are always functional gadgets, as in TrapTP . The last $t - \ell$ gadgets are all completely replaced by dummy states, and their functionality is completely outsourced to VerDec . In four steps described above, the functionality of the ℓ th gadget is also transferred to VerDec . It is important to replace only one gadget at a time, because replacing a real gadget with a fake one breaks the functionality of the gadgets that occur later in the evaluation: the encrypted classical information held by the server does not correspond to the question of whether or not a phase correction is needed. By completely outsourcing the phase correction to VerDec , as is done for all gadgets after the ℓ th one in all $\text{TrapTP}_i^{(\ell)}$ schemes, we ensure that this incorrect classical information does not influence the outcome of the computation. Hence, correctness is maintained throughout the hybrid transformations. We now show that these transformations of the scheme do not significantly affect the adversary's winning probability in the hybrid indistinguishability game.

Lemma 3. *For any QPT \mathcal{A} , there exists a negligible function negl such that for all $1 \leq \ell \leq t$,*

$$\text{AdvHyb}_{\text{TrapTP}_1^{(\ell)}, \text{TrapTP}_4^{(\ell+1)}}(\mathcal{A}, \kappa) \leq \text{negl}(\kappa).$$

Proof (sketch). In $\text{TrapTP}_4^{(\ell+1)}$, no information about $sk_{(\ell)}$ is sent to the adversary. In the original TrapTP scheme, the structure of the quantum state $\Gamma_{\ell+1}$ depended on it, but this structure has been replaced with dummy states in several steps in $\text{TrapTP}_2^{\ell+1}$ through $\text{TrapTP}_4^{\ell+1}$.

This is fortunate, since if absolutely no secret-key information is present, we are able to bound the difference in winning probability between $\text{Hyb}_{\mathcal{A}, \text{TrapTP}_4^{(\ell+1)}}(\kappa)$

and $\text{Hyb}_{\mathcal{A}, \text{TrapTP}_1^\ell}(\kappa)$ by reducing it to the IND-CPA security against quantum adversaries [5] of the classical homomorphic encryption scheme HE.

The proof is closely analogous to the proof of Lemma 1 in [12], and on a high level it works as follows. Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ be a QPT adversary for the game $\text{Hyb}_{\mathcal{A}, \text{TrapTP}_1^\ell}(\kappa)$ or $\text{Hyb}_{\mathcal{A}, \text{TrapTP}_4^{\ell+1}}(\kappa)$ (we do not need to specify for which one, since they both require the same input/output interface). A new quantum adversary \mathcal{A}' for the classical IND-CPA indistinguishability game is defined by having the adversary taking the role of challenger in either the game $\text{Hyb}_{\mathcal{A}, \text{TrapTP}_1^\ell}(\kappa)$ or the game $\text{Hyb}_{\mathcal{A}, \text{TrapTP}_4^{\ell+1}}(\kappa)$. Which game is simulated depends on the coin flip of the challenger for the IND-CPA indistinguishability game, and is unknown to \mathcal{A}' . This situation is achieved by having \mathcal{A}' send any classical plaintext that should be encrypted under pk_ℓ to the challenger, so that either that plaintext is encrypted or a string of zeros is.

Based on the guess of the simulated \mathcal{A} , which \mathcal{A}' can verify to be correct or incorrect in his role of challenger, \mathcal{A}' will guess which of the two games was just simulated. By IND-CPA security of the classical scheme against quantum adversaries, \mathcal{A}' cannot succeed in this guessing game with nonnegligible advantage over random guessing. This means that the winning probability of \mathcal{A} in both games cannot differ by a lot. For details, we refer the reader the proof of Lemma 5, in which a very similar approach is taken.

Technically, the success probability of \mathcal{A}' , and thus the function negl , may depend on ℓ . A standard randomizing argument, as found in e.g. the discussion of hybrid arguments in [16], allows us to get rid of this dependence by defining another adversary \mathcal{A}'' that selects a random value of j , and then bounding the advantage of \mathcal{A}'' by a negligible function that is independent of j . \square

Lemma 4. For $1 \leq \ell \leq t$ and any QPT \mathcal{A} , $\text{AdvHyb}_{\text{TrapTP}_1^\ell}^{\text{TrapTP}_2^\ell}(\mathcal{A}, \kappa) = 0$.

Proof. In TrapTP_1^ℓ , the ℓ th error-correction gadget consists of a number of EPR pairs arranged in a certain order, as described by the garden-hose protocol for HE.Dec. For example, this protocol may dictate that the i th and j th qubit of the gadget must form an EPR pair $|\Phi^+\rangle$ together. This can alternatively be achieved by creating two EPR pairs, placing half of each pair in the i th and j th position of the gadget state, and performing a Bell measurement on the other two halves. This creates a Bell pair $X^a Z^b |\Phi^+\rangle$ in positions i and j , where $a, b \in \{0, 1\}$ describe the outcome of the Bell measurement.

From the point of view of the adversary, it does not matter whether these Bell measurements are performed during **KeyGen**, or whether the halves of EPR pairs are sent to **VerDec** for measurement – because the key to the quantum one-time pad of the ℓ th gadget is not sent to the adversary at all, the same state is created with a completely random Pauli in either case. Of course, the teleportation correction Paulis of the form $X^a Z^b$ need to be taken into account when updating the keys to the quantum one-time pad on the data qubits after the gadget is used. **VerDec** has all the necessary information to do this, because

it observes the measurement outcomes, and computes the key updates itself (instead of decrypting the final keys from the computation log).

Thus, with the extra key update steps in $\text{TrapTP}_2^{(\ell)}. \text{VerDec}$, the inputs to the adversary are exactly the same in the games of $\text{TrapTP}_1^{(\ell)}$ and $\text{TrapTP}_2^{(\ell)}$. \square

Lemma 5. *For any QPT \mathcal{A} , there exists a negligible function negl such that for all $1 \leq \ell \leq t$,*

$$\text{AdvHyb}_{\text{TrapTP}_2^{(\ell)}}^{\text{TrapTP}_3^{(\ell)}}(\mathcal{A}, \kappa) \leq \text{negl}(\kappa).$$

Proof. We show this by reducing the difference in winning probabilities in the statement of the lemma to the IND-VER security of the somewhat homomorphic scheme TC. Intuitively, because TC is IND-VER, if $\text{TrapTP}_2^{(\ell)}$ accepts the adversary's claimed circuit of Bell measurements on the EPR pair halves, the effective map on those EPR pairs is the claimed circuit. Therefore, we might just as well ask VerDec to apply this map, as we do in $\text{TrapTP}_3^{(\ell)}$, to get the same output state. If $\text{TrapTP}_2^{(\ell)}$ rejects the adversary's claimed circuit on those EPR pair halves, then $\text{TrapTP}_3^{(\ell)}$ should reject too. This is why we let the adversary act on an encrypted dummy state of $|0\rangle_s$.

Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ be a set of QPT algorithms on the appropriate registers, so that we can consider it as an adversary for the hybrid indistinguishability game for either $\text{TrapTP}_2^{(\ell)}$ or $\text{TrapTP}_3^{(\ell)}$ (see Definition 7). Note the input/output wires to the adversary in both these games are identical, so we can evaluate $\Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}_2^{(\ell)}}(\kappa) = 1]$ and $\Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}_3^{(\ell)}}(\kappa) = 1]$ for the same \mathcal{A} .

Now define an adversary $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2, \mathcal{A}'_3)$ for the IND-VER game against TC, $\text{VerGame}_{\mathcal{A}', \text{TC}}(\kappa)$, as follows:

1. \mathcal{A}'_1 : Run $\text{TrapTP}_2^{(\ell)}. \text{KeyGen}$ until the start of line 15 in the ℓ th iteration of that loop. Up to this point, $\text{TrapTP}_2^{(\ell)}. \text{KeyGen}$ is identical to $\text{TrapTP}_3^{(\ell)}. \text{KeyGen}$. It has generated real gadgets Γ_1 through $\Gamma_{\ell-1}$, and halves of EPR pairs for γ_ℓ^{in} , γ_ℓ^{mid} and γ_ℓ^{out} . Note furthermore that the permutation π_ℓ is used nowhere. Now send γ_ℓ^{mid} to the challenger via the register X , and everything else (including sk) to \mathcal{A}'_2 via the side register R .

2. \mathcal{A}'_2 : Continue $\text{TrapTP}_2^{(\ell)}. \text{KeyGen}$ using the response from the challenger instead of $\text{TrapTP}. \text{Enc}'(sk', \gamma_\ell^{\text{mid}})$ on line 15 in the ℓ th iteration. Call the result ρ_{evk} . Again, this part of the key generation procedure is identical for $\text{TrapTP}_2^{(\ell)}$ and $\text{TrapTP}_3^{(\ell)}$. Start playing the hybrid indistinguishability game with \mathcal{A} :

- Flip a bit $r \in \{0, 1\}$.
- Send ρ_{evk} to \mathcal{A}'_1 . If $r = 0$, encrypt the response of \mathcal{A}_1 using the secret key sk generated by \mathcal{A}'_1 . Note that for this, the permutation π_ℓ is also not needed. If $r = 1$, encrypt a $|0\rangle$ state of appropriate dimension instead.
- Send the resulting encryption, along with the side info from \mathcal{A}_1 , to \mathcal{A}_2 .
- On the output of \mathcal{A}_2 , start running $\text{TrapTP}_2^{(\ell)}. \text{VerDec}$ until the actions on the ℓ th gadget need to be verified. Since the permutation on the state γ_ℓ^{mid} is unknown to \mathcal{A}'_2 (it was sent to the challenger for encryption), it cannot verify this part of the computation.

- Instead, send the relevant part of the computation log to the challenger for verification, along with the relevant part of the claimed circuit (the Bell measurements on the gadget state), and the relevant qubits, all received from \mathcal{A}_2 , to the challenger for verification and decryption.
- In the meantime, send the rest of the working memory to \mathcal{A}'_3 via register R' .

3. \mathcal{A}'_3 : Continue the simulation of the hybrid game with \mathcal{A} :

- If the challenger rejects, reject and replace the entire quantum state by the fixed dummy state Ω .
- If the challenger accepts, then we know that the challenger applies the claimed subcircuit to the quantum state it did not encrypt (either $|0\rangle$ or γ_ℓ^{mid}), depending on the bit the challenger flipped), and possibly swaps this state back in (again depending on which bit it flipped). Continue the $\text{TrapTP}_2^{(\ell)}. \text{VerDec}$ computation for the rest of the computation log.
- Send the result (the output quantum state, the claimed circuit, and the accept/reject flag) to \mathcal{A}_3 , and call its output bit r' .

Output 0 if $r = r'$, and 1 otherwise. (i.e., output $NEQ(r, r')$)

Recall from Definition 7 that the challenger flips a coin (let us call the outcome $s \in \{0, 1\}$) to decide whether to encrypt the quantum state provided by \mathcal{A}' , or to swap in an all-zero dummy state before encrypting. Keeping this in mind while inspecting the definition of \mathcal{A}' , one can see that whenever $s = 0$, \mathcal{A}' takes the role of challenger in the game $\text{Hyb}_{\mathcal{A}, \text{TrapTP}_2^{(\ell)}}(\kappa)$ with \mathcal{A} , and whenever $s = 1$, they play $\text{Hyb}_{\mathcal{A}, \text{TrapTP}_3^{(\ell)}}(\kappa)$. Now let us consider when the newly defined adversary \mathcal{A}' wins the VER indistinguishability game for TC. If $s = 0$, \mathcal{A}' needs to output a bit $s' = 0$ to win. This happens, by definition of \mathcal{A}' , if and only if \mathcal{A} wins the game $\text{Hyb}_{\mathcal{A}, \text{TrapTP}_2^{(\ell)}}(\kappa)$ (i.e. $r = r'$). On the other hand, if $s = 1$, \mathcal{A}' needs to output a bit $s' = 1$ to win. This happens, by definition of \mathcal{A}' , if and only if \mathcal{A} loses the game $\text{Hyb}_{\mathcal{A}, \text{TrapTP}_3^{(\ell)}}(\kappa)$ (i.e. $r \neq r'$). Thus the winning probability of \mathcal{A}' is:

$$\begin{aligned}
& \Pr[\text{VerGame}_{\mathcal{A}', \text{TC}}(\kappa) = 1] = \\
&= \Pr[s = 0] \cdot \Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}_2^{(\ell)}}(\kappa) = 1] + \Pr[s = 1] \cdot \Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}_3^{(\ell)}}(\kappa) = 0] \\
&= \frac{1}{2} \Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}_2^{(\ell)}}(\kappa) = 1] + \frac{1}{2} \left(1 - \Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}_3^{(\ell)}}(\kappa) = 1] \right) \\
&= \frac{1}{2} + \frac{1}{2} \left(\Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}_2^{(\ell)}}(\kappa) = 1] - \Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}_3^{(\ell)}}(\kappa) = 1] \right)
\end{aligned}$$

From the IND-VER property of TC (see Theorem 4) we know that the above is at most $\frac{1}{2} + \text{negl}(\kappa)$. From this (and a randomizing argument similar to Lemma 3), the statement of the lemma follows directly. \square

Lemma 6. *For any QPT \mathcal{A} , there exists a negligible function negl such that for all $1 \leq \ell \leq t$,*

$$\text{AdvHyb}_{\text{TrapTP}_3^{(\ell)}}^{\text{TrapTP}_4^{(\ell)}}(\mathcal{A}, \kappa) \leq \text{negl}(\kappa).$$

Proof. Let $f(s)$ be the bit that, after the ℓ th T gate, determines whether or not a phase correction is necessary. Here, s is all the relevant starting information (such as quantum one-time pad keys, gadget structure, permutations, and applied circuit), and f is some function that determines the X key on the relevant qubit right before application of the T gate.

In $\text{TrapTP}_3^{(\ell)}$, a phase correction after the ℓ th T gate is applied conditioned on the outcome of

$$\text{HE.Dec}_{sk_{\ell-1}}(\text{HE.Eval}_{evk_0, \dots, evk_{\ell-1}}^f(\text{HE.Enc}_{pk_0}(s))),$$

because the garden-hose computation in the gadget computes the classical decryption. In the above expression, we again slightly abuse notation, as in the proof of Lemma 2, and include decryption steps in $\text{HE.Eval}_{evk_0, \dots, evk_{\ell-1}}$. As long as t is polynomial in κ , we have, by correctness of HE,

$$\Pr[\text{HE.Dec}_{sk_{\ell-1}}(\text{HE.Eval}_{evk_0, \dots, evk_{\ell-1}}^f(\text{HE.Enc}_{pk_0}(s))) \neq f(s)] \leq \text{negl}(\kappa).$$

In $\text{TrapTP}_4^{(\ell)}$, the only difference from $\text{TrapTP}_3^{(\ell)}$ is that, instead of performing the garden-hose computation on the result of the classical homomorphic evaluation procedure, the phase correction is applied directly by VerDec , conditioned on $f(s)$. The probability that in $\text{TrapTP}_4^{(\ell)}$, a phase is applied (or not) when in $\text{TrapTP}_3^{(\ell)}$ it is not (or is), is negligible. The claim follows directly. \square

Final Hybrid: Removing All Classical FHE. In $\text{TrapTP}_4^{(1)}$, all of the error-correction gadgets have been removed from the evaluation key, and the error-correction functionality has been redirected to VerDec completely. Effectively, $\text{TrapTP}_4^{(1)}$. KeyGen samples a permutation π , generates a lot of magic states (for P, H and T) and encrypts them using TC.Enc_π , after which the keys to the quantum one-time pad used in that encryption are homomorphically encrypted under pk_0 . The adversary is allowed to act on those encryptions, but while its homomorphic computations are syntactically checked in the log, VerDec does not decrypt and use the resulting values. This allows us to link $\text{TrapTP}_4^{(1)}$ to a final hybrid, TrapTP^f , where all classical information is replaced with zeros before encrypting.

The proof of the following lemma is analogous to that of Lemma 3, and reduces to the IND-CPA security of the classical scheme HE:

Lemma 7. *For any QPT \mathcal{A} , $\text{AdvHyb}_{\text{TrapTP}_4^{(1)}}^{\text{TrapTP}^f}(\mathcal{A}, \kappa) \leq \text{negl}(\kappa)$.*

Proof of main theorem. Considering TrapTP^f in more detail, we can see that it is actually very similar to TC. This allows us to prove the following lemma, which is the last ingredient for the proof of verifiability of TrapTP .

Lemma 8. *For any QPT \mathcal{A} , $\Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}^f}(\kappa) = 1] \leq \frac{1}{2} + \text{negl}(\kappa)$.*

Proof. To see the similarity with TC, consider the four algorithms of TrapTP^f .

In $\text{TrapTP}^f.\text{KeyGen}$, a permutation π is sampled, and magic states for P, H and T are generated, along with some EPR pair halves (to replace in_i and out_i). For all generated quantum states, random keys for QOTPs are sampled, and the states are encrypted using TC.Enc with these keys as secret keys. No classical FHE is present anymore. Thus, $\text{TrapTP}^f.\text{KeyGen}$ can be viewed as TC.KeyGen , followed by TC.Enc on the magic states and EPR pair halves.

$\text{TrapTP}^f.\text{Enc}$ is identical to TC.Enc , only the keys to the quantum one-time pad are sampled on the fly and sent to $\text{TrapTP}^f.\text{VerDec}$ via a classical side-channel, whereas TC.VerDec receives them as part of the secret key. Since the keys are used exactly once and not used anywhere else besides in Enc and VerDec , this difference does not affect the outcome of the game.

$\text{TrapTP}^f.\text{Eval}$ only requires CNOT, classically controlled Paulis, and computational/Hadamard basis measurements. For the execution of any other gate, it suffices to apply a circuit of those gates to the encrypted data, encrypted magic states and/or encrypted EPR halves.

$\text{TrapTP}^f.\text{VerDec}$ does two things: (i) it syntactically checks the provided computation log, and (ii) it runs TC.VerDec to verify that the evaluation procedure correctly applied the circuit of CNOTs and measurements.

An execution of $\text{Hyb}_{\mathcal{A}, \text{TrapTP}^f}(\kappa)$ for any \mathcal{A} corresponds to the two-round VER indistinguishability game for TC as follows. Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ be a polynomial-time adversary for the game $\text{Hyb}_{\mathcal{A}, \text{TrapTP}^f}(\kappa)$. Define an additional QPT \mathcal{A}_0 that produces magic states and EPR pair halves to the register X_1 . The other halves of the EPR pairs are sent through R , and untouched by \mathcal{A}_1 and \mathcal{A}_2 . The above analysis shows that the adversary $\mathcal{A}' = (\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ can be viewed as an adversary for the VER-2 indistinguishability game $\text{VerGame}_{\mathcal{A}', \text{TC}}^2(\kappa)$ and wins whenever $\text{Hyb}_{\mathcal{A}, \text{TrapTP}^f}(\kappa) = 1$. The other direction does not hold: \mathcal{A} loses the hybrid indistinguishability game if $\text{TrapTP}^f.\text{VerDec}$ rejects check (i), but accepts check (ii) (see above). In this case, \mathcal{A}' would still win the VER-2 indistinguishability game. Hence,

$$\Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}^f}(\kappa) = 1] \leq \Pr[\text{VerGame}_{\mathcal{A}', \text{TC}}^2(\kappa) = 1].$$

Theorem 4 yields $\Pr[\text{VerGame}_{\mathcal{A}', \text{TC}}^2(\kappa) = 1] \leq \frac{1}{2} + \text{negl}(\kappa)$, and the result follows. \square

Theorem 5. *The vQFHE scheme TrapTP satisfies κ -SEM-VER.*

Proof. From Lemmas 1, 2, 3, 4, 5, 6, and 7, we may conclude that if t (the number of T gates in the circuit) is polynomial in κ (the security parameter), then for any polynomial-time adversary \mathcal{A} ,

$$\Pr[\text{VerGame}_{\mathcal{A}, \text{TrapTP}}(\kappa) = 1] - \Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}^f}(\kappa) = 1] \leq \text{negl}(\kappa),$$

since the sum poly-many negligible terms is negligible (it is important to note that there is only a constant number of *different* negligible terms involved). By Lemma 8, which reduces verifiability of TrapTP^f to verifiability of TC,

$\Pr[\text{Hyb}_{\mathcal{A}, \text{TrapTP}^f}(\kappa) = 1] \leq 1/2 + \text{negl}(\kappa)$. It follows that $\Pr[\text{VerGame}_{\mathcal{A}, \text{TrapTP}}(\kappa) = 1] \leq 1/2 + \text{negl}(\kappa)$, i.e., that TrapTP is κ -IND-VER. By Theorem 3, TrapTP is also κ -SEM-VER. \square

6 Application to quantum one-time programs

One-time programs. We now briefly sketch an application of the vQFHE scheme to one-time programs. A classical one-time program (or cOTP) is an idealized object which can be used to execute a function once, but then self-destructs. In the case of a quantum OTP (or qOTP), the program executes a quantum channel Φ . In the usual formalization, Φ has two inputs and is public. One party (the sender) creates the qOTP by fixing one input, and the qOTP is executed by a receiver who selects the other input. To recover the intuitive notion of OTP, choose Φ to be a universal circuit. We will work in the universally-composable (UC) framework, following the approach of [8]. We thus first define the ideal functionality of a qOTP.

Definition 8 (Functionality 3 in [8]). *The ideal functionality $\mathcal{F}_{\Phi}^{\text{OTP}}$ for a channel $\Phi_{XY \rightarrow Z}$ is the following:*

1. **Create:** given register X from sender, store X and send **create** to receiver.
2. **Execute:** given register Y from receiver, send Φ applied to XY to receiver. Delete any trace of this instance.

A qOTP is then a real functionality which “UC-emulates” the ideal functionality [22]. As in [8], we only allow corrupting receivers; unlike [8], we consider computational (rather than statistical) UC security. The achieved result is therefore slightly weaker. The construction within our vQFHE framework is however much simpler, and shows the relative ease with which applications of vQFHE can be constructed.

The construction. Choose a vQFHE scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Eval}, \text{VerDec})$ satisfying SEM-VER. For simplicity, we first describe the classical input/output case, i.e., the circuit begins and ends with full measurement of all qubits. Let C be such a circuit, for the map $\Phi_{XY \rightarrow Z}$. On **Create**, the sender generates keys $(k, \rho_{\text{evk}}) \leftarrow \text{KeyGen}$ and encrypts their input register X using k . The sender also generates a classical OTP for the public, classical function VerDec , choosing the circuit and key inputs to be C and k ; the computation log is left open for the receiver to select. The qOTP is then the triple

$$\Xi_C^X := (\rho_{\text{evk}}, \text{Enc}_k(\rho_X), \text{OTP}_{\text{VerDec}}(C, k)).$$

On **Execute**, the receiver computes as follows. The receiver’s (classical) input Y together with the (public) circuit C defines a homomorphic computation on the ciphertext $\text{Enc}_k(\rho_X)$, which the receiver can perform using Eval and ρ_{evk} . Since C has only classical outputs, the receiver measures the final state completely. At

the end of that computation, the receiver holds the (completely classical) output of the computation log from `Eval`. The receiver plugs the log into $\text{OTP}_{\text{VerDec}}(C, k)$, which produces the decrypted output.

We handle the case of arbitrary circuits C (with quantum input and output) as follows. Following the ideas of [8], we augment the above quantum OTP with two auxiliary quantum states: an “encrypt-through-teleport” gadget σ_{in} and a “decrypt-through-teleport” gadget σ_{out} . These are maximally entangled states with the appropriate map (encrypt or decrypt) applied to one half. The receiver uses teleportation on $\sigma_{Y_1 W_1}^{\text{in}}$ to encrypt their input register Y before evaluating, and places the teleportation measurements into the computation log. After evaluation, the receiver uses $\sigma_{W_2 Y_2}^{\text{out}}$ to teleport the plaintext out, combining the teleportation measurements with the output of $\text{OTP}_{\text{VerDec}}(C, k)$ to compute the final QOTP decryption keys.

Security proof sketch. Starting with a QPT adversary \mathcal{A} which attacks the real functionality, we construct a QPT simulator \mathcal{S} which attacks the ideal functionality (with similar success probability). We split \mathcal{A} into \mathcal{A}_1 (receive input, output the OTP query and side information) and \mathcal{A}_2 (receive result of OTP query and side information, produce final output). The simulator \mathcal{S} will generate its own keys, provide fake gadgets that will trick \mathcal{A} into teleporting its input to \mathcal{S} , who will then use that input on the ideal functionality. Details follow.

The simulator first generates $(k, \rho_{\text{evk}}) \leftarrow \text{KeyGen}$ and encrypts the input X via Enc_k . Instead of the encrypt gadget $\sigma_{Y_1 W_1}^{\text{in}}$, \mathcal{S} provides half of a maximally entangled state in register Y and likewise in register W . The other halves Y'_1 and W'_1 of these entangled states are kept by \mathcal{S} . The same is done in place of the decrypt gadget $\sigma_{W_2 Y_2}^{\text{out}}$, with \mathcal{S} keeping Y'_2 and W'_2 . Then \mathcal{S} runs \mathcal{A}_1 with input $\rho_{\text{evk}}, \text{Enc}_k(\rho_X)$ and registers Y and W . It then executes VerDec_k on the output (i.e., the query) of \mathcal{A}_1 to see if \mathcal{A}_1 correctly followed the `Eval` protocol. If it did not, then \mathcal{S} aborts; otherwise, \mathcal{S} plugs register Y'_1 into the ideal functionality, and then teleports the output into register W'_2 . Before responding to \mathcal{A}_2 , it corrects the one-time pad keys appropriately using its teleportation measurements.

7 Conclusion

In this work, we devised a new quantum-cryptographic primitive: quantum fully-homomorphic encryption with verification (vQFHE). Using the trap code for quantum authentication [8] and the garden-hose gadgets of [12], we constructed a vQFHE scheme `TrapTP` which satisfies (i.) correctness, (ii.) compactness, (iii.) security of verification, (iv.) IND-CPA secrecy, and (v.) authentication. We also outlined a first application of vQFHE, to quantum one-time programs.

We leave open several interesting directions for future research. Foremost is finding more applications of vQFHE. Another interesting question is whether vQFHE schemes exist where verification can be done publicly (i.e., without the decryption key), as is possible classically. Finally, it is unknown whether vQFHE

(or even QFHE) schemes exist with evaluation key that does not scale with the size of the circuit at all.

8 Acknowledgements

This work was completed while GA was a member of the QMATH center at the Department of Mathematical Sciences at the University of Copenhagen. GA and FS acknowledge financial support from the European Research Council (ERC Grant Agreement no 337603), the Danish Council for Independent Research (Sapere Aude), Qubiz - Quantum Innovation Center, and VILLUM FONDEN via the QMATH Centre of Excellence (Grant No. 10059). CS is supported by an NWO VIDI grant.

Bibliography

- [1] Dorit Aharonov, Michael Ben-Or, and Elad Eban. Interactive proofs for quantum computations. *arXiv preprint arXiv:0810.5375*, 2008.
- [2] Gorjan Alagic, Anne Broadbent, Bill Fefferman, Tommaso Gagliardoni, Christian Schaffner, and Michael St. Jules. Computational security for quantum encryption. In *9th International Conference on Information Theoretic Security (ITICS)*, pages 47–71, 2016. doi: 10.1007/978-3-319-49175-2_3.
- [3] Boaz Barak and Zvika Brakerski. Windows on theory: The swiss army knife of cryptography, 2012. URL <https://windowsontheory.org/2012/05/01/the-swiss-army-knife-of-cryptography/>.
- [4] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *52nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 97–106, 2011. doi: 10.1109/FOCS.2011.12.
- [5] Anne Broadbent and Stacey Jeffery. Quantum homomorphic encryption for circuits of low T-gate complexity. In *Advances in Cryptology-CRYPTO 2015*, pages 609–629. Springer, 2015.
- [6] Anne Broadbent and Evelyn Wainwright. Efficient simulation for quantum message authentication. *arXiv preprint arXiv:1607.03075*, 2016.
- [7] Anne Broadbent, Joseph Fitzsimons, and Elham Kashefi. Universal blind quantum computation. In *50th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 517–526. IEEE, 2009.
- [8] Anne Broadbent, Gus Gutoski, and Douglas Stebila. Quantum one-time programs. In *Advances in Cryptology-CRYPTO 2013*, pages 344–360. Springer, 2013.
- [9] Anne Broadbent, Zhengfeng Ji, Fang Song, and John Watrous. Zero-knowledge proof systems for QMA. In *57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 31–40, Oct 2016. doi: 10.1109/FOCS.2016.13.
- [10] Andrea Coladangelo, Alex Grilo, Stacey Jeffery, and Thomas Vidick. Verifier-on-a-leash: new schemes for verifiable delegated quantum computation, with quasilinear resources. *arXiv preprint arXiv:1708.02130*, 2017.

- [11] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A*, 80(1):012304, 2009.
- [12] Yfke Dulek, Christian Schaffner, and Florian Speelman. Quantum homomorphic encryption for polynomial-sized circuits. In *Advances in Cryptology–CRYPTO 2016*, pages 3–32. Springer, 2016. doi: 10.1007/978-3-662-53015-3_1.
- [13] Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Actively secure two-party evaluation of any quantum operation. In *Advances in Cryptology–CRYPTO 2012*, pages 794–811. Springer, 2012.
- [14] S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 40–49, Oct 2013. doi: 10.1109/FOCS.2013.13.
- [15] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 169–178, 2009. doi: 10.1145/1536414.1536440.
- [16] Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography*. CRC press, 2014.
- [17] Urmila Mahadev. Classical homomorphic encryption for quantum circuits. *arXiv preprint arXiv:1708.02130*, 2017.
- [18] M. Newman and Y. Shi. Limitations on Transversal Computation through Quantum Homomorphic Encryption. *ArXiv e-prints*, April 2017.
- [19] Yingkai Ouyang, Si-Hui Tan, and Joseph Fitzsimons. Quantum homomorphic encryption from quantum codes. *arXiv preprint arXiv:1508.00938*, 2015.
- [20] Peter W. Shor and John Preskill. Simple proof of security of the BB84 quantum key distribution protocol. *Phys. Rev. Lett.*, 85:441–444, Jul 2000. doi: 10.1103/PhysRevLett.85.441.
- [21] Si-Hui Tan, Joshua A. Kettlewell, Yingkai Ouyang, Lin Chen, and Joseph Fitzsimons. A quantum approach to homomorphic encryption. *Scientific Reports*, 6:33467, September 2016. doi: 10.1038/srep33467.
- [22] Dominique Unruh. Universally composable quantum multi-party computation. In *Advances in Cryptology–EUROCRYPT 2010*, pages 486–505, 2010. arXiv:0910.2912 [quant-ph].
- [23] Li Yu, Carlos A. Pérez-Delgado, and Joseph F. Fitzsimons. Limitations on information-theoretically-secure quantum homomorphic encryption. *Phys. Rev. A*, 90:050303, 2014. doi: 10.1103/PhysRevA.90.050303.

Supplementary material

A Equivalence of κ -IND-VER and κ -SEM-VER.

In Theorem 3, it was stated that if a scheme is κ -IND-VER, then it is also κ -SEM-VER, and vice versa. We here provide the full proof of Theorem 3.

Proof (of Lemma 3). We first show the forward direction. Suppose a scheme S is *not* κ -SEM-VER. Then there exists a QPT \mathcal{A} such that for all simulators \mathcal{S} , there exist QPTs \mathcal{M} and \mathcal{D} and a polynomial p such that the difference in acceptance probability is at least $1/p(\kappa)$. Choose \mathcal{S} to be

$$\mathcal{S} : (sk, \rho_{R_1}) \mapsto \text{tr}_{X'} \left((\text{VerDec}_{sk} \otimes \mathbb{I}_{R_1}) (\mathcal{A}(\text{Enc}_{sk}(|0^n\rangle\langle 0^n|) \otimes \rho_{R_1})) \right),$$

This simulator encrypts a dummy state and feeds it to the adversary; whatever comes out is then checked. Note that in the accept case, the output is wrong, since the claimed circuit is applied to the dummy state instead of the real input. This does not matter, however, because the simulator throws out the result immediately. Since \mathcal{S} is a possible simulator, we can let \mathcal{M} and \mathcal{D} be as given by the assumption that κ -SEM-VER is false.

This allows us to construct a QPT adversary $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2, \mathcal{A}'_3)$ for the VER indistinguishability game $\text{VerGame}_{\mathcal{A}, S}(\kappa)$ simply by setting $\mathcal{A}'_1 = \mathcal{M}$, $\mathcal{A}'_2 = (\mathcal{A} \otimes \mathbb{I}_{R_2})$, and $\mathcal{A}'_3 = \mathcal{D}$. Informally, the probability that this adversary wins is

$$\Pr[r = 0] \Pr[\mathcal{A}'_3 \text{ guesses } 0 \mid r = 0] + \Pr[r = 1] \Pr[\mathcal{A}'_3 \text{ guesses } 1 \mid r = 1].$$

More precisely, it is

$$\begin{aligned} & \frac{1}{2} \Pr \left[\mathcal{A}'_3 \left((S.\text{VerDec}_{sk} \otimes \mathbb{I}_R) (\mathcal{A}'_2 \left((S.\text{Enc}_{sk} \otimes \mathbb{I}_R) (\mathcal{A}'_1(\rho_{evk})) \right)) \right) = 0 \right] \\ & + \frac{1}{2} \Pr \left[\mathcal{A}'_3 \left((\mathbb{I}_{X'R'} \otimes \Pi_{\text{acc}}) \left((\Phi_c \otimes \mathbb{I}_{FR'}) (\sigma_{XFR'}) \right) + (\mathbb{I}_{X'R'} \otimes \Pi_{\text{rej}}) (\Omega \otimes \sigma_{FR'}) \right) = 1 \right] \end{aligned}$$

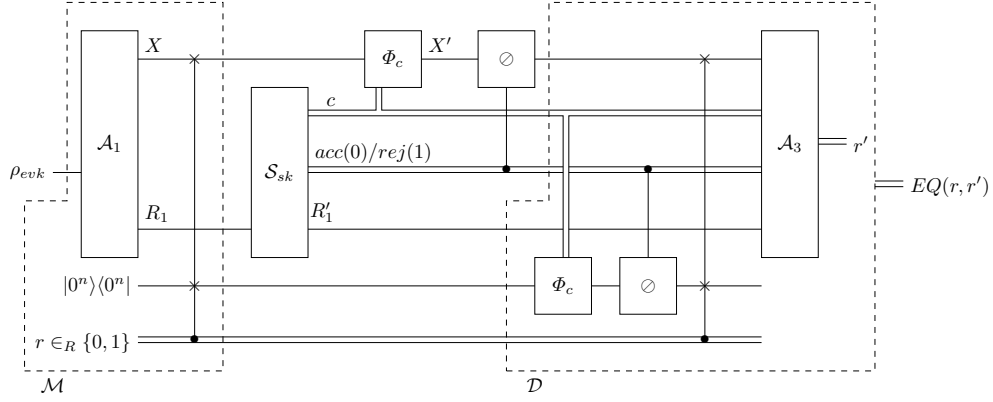
where F is the flag register (accept/reject), and we set $(c, \sigma_{XFR'}) = (\text{tr}_{X'} \circ (\mathbb{I}_X \otimes S.\text{VerDec}_{sk} \otimes \mathbb{I}_{R'}) \circ (\mathbb{I}_X \otimes \mathcal{A}'_2)) (S.\text{Enc}_{sk}(|0^n\rangle\langle 0^n|) \otimes \mathcal{A}'_1(\rho_{evk}))$, and $\Pi_{\text{acc}} = |\text{acc}\rangle\langle \text{acc}|$ and $\Pi_{\text{rej}} = |\text{rej}\rangle\langle \text{rej}|$. This can be seen by following the wires in the indistinguishability game. By our definition of \mathcal{A}' and \mathcal{S} , this is equal to

$$\begin{aligned} & \frac{1}{2} \left(1 - \Pr \left[\mathcal{D} \left((S.\text{VerDec}_{sk} \otimes \mathbb{I}_{R_1 R_2}) \left((\mathcal{A} \otimes \mathbb{I}_{R_2}) \tau_{C_X R} \right) \right) = 1 \right] \right) \\ & + \frac{1}{2} \Pr \left[\mathcal{D} \left((\overline{\text{ctrl}} \circ \Phi_c \circ \mathcal{S}_{sk}) (\mathcal{M}(\rho_{evk})) \right) = 1 \right] \end{aligned}$$

where $\tau_{C_X R} = (S.\text{Enc}_{sk} \otimes \mathbb{I}_{R_1 R_2}) (\mathcal{M}(\rho_{evk}))$. By the assumption that S is not κ -SEM-VER, this is at least $\frac{1}{2} + 1/p(\kappa)$. Hence, this adversary wins the IND-VER indistinguishability game with nonnegligible probability.

For the other direction, suppose that a scheme S is κ -SEM-VER, and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ be an arbitrary QPT adversary for the IND-VER indistinguishability

game for this scheme. By definition of κ -SEM-VER, for \mathcal{A}_2 there exists \mathcal{S} such that for all QPTs \mathcal{M} and \mathcal{D} , the equation from Definition 5 holds with $\mathcal{A} := \mathcal{A}_2$. We choose \mathcal{M} and \mathcal{D} as in the figure below. More precisely, \mathcal{M} does: (i.) run \mathcal{A}_1 on its input (ii.) prepare the state $|0^n\rangle\langle 0^n|$, plus a random bit $r \in_R \{0, 1\}$, and store them in the side information register R_2 , and (iii.) swap the quantum states in X and R_2 conditioned on r . We also choose \mathcal{D} to (i.) run \mathcal{A}_3 on the appropriate input wires, (ii.) either apply Φ_c or \otimes on the quantum state in the register R_2 , conditioned on the accept/reject flag, (iii.) swap those wires back (again, conditioned on r), and finally (iv.) output 1 if \mathcal{A}_3 's output was correct (i.e. equal to r), and 0 otherwise.



Note that these choices ensure that the real channel is an execution of the IND-VER game. In the ideal scenario, \mathcal{A}_3 receives *exactly* the same state in the cases $r = 0$ and $r = 1$. Hence, the best he can do is guess, and the probability that $r' = r$ (and thus that \mathcal{D} outputs 1) is at most $\frac{1}{2}$.

By the assumption that S is κ -SEM-VER, the probability that \mathcal{D} outputs 1 in the real scenario can only be negligibly higher than in the ideal case. As discussed above, the real scenario corresponds exactly to the adversary \mathcal{A} playing the IND-VER game. Therefore, the winning probability for \mathcal{A} (i.e. the probability that $\text{VerGame}_{\mathcal{A}, S}(\kappa) = 1$) is at most negligibly (in κ) higher than $\frac{1}{2}$. \square

B Algorithms for TC

This section formalizes the encryption, verification, and decryption procedures described in Section 3.

Algorithm 13. $\text{TC.Enc}((\pi, x, z), \sigma)$

- 1: $\sigma' \leftarrow \text{CSS.Encode}(\sigma)$
 - 2: $\sigma'' \leftarrow \text{permute}_{\pi}(\sigma' \otimes |0^m\rangle \otimes |+\rangle^m)$
 - 3: $\tilde{\sigma} \leftarrow X^x Z^z \sigma'' X^x Z^z$
 - 4: **return** $\tilde{\sigma}$
-

Algorithm 14. TC.VerDecQubit($(\pi, x, z), \tilde{\sigma}$)

- 1: $\sigma \leftarrow X^x Z^z \tilde{\sigma} X^x Z^z$
- 2: $(\sigma', \text{traps}_X, \text{traps}_Z) \leftarrow \text{permute}_{\pi^{-1}}(\sigma) \triangleright \text{traps}_X$ and traps_Z are quantum states
- 3: Measure traps_X in the computational basis. If not all-zero, **return** $(\Omega, |\text{rej}\rangle)$.
- 4: Measure traps_Z in the Hadamard basis. If not all+, **return** $(\Omega, |\text{rej}\rangle)$.
- 5: $\sigma'' \leftarrow \text{CSS.Decode}(\text{first } m \text{ qubits of } \sigma')$
- 6: **return** $(\sigma'', |\text{acc}\rangle)$

Algorithm 15. TC.VerDecMeasurement($(\pi, x, z), \tilde{w}, \text{basis}$)

- 1: **if** $\text{basis} = +$ **then**
- 2: $w' \leftarrow \text{permute}_{\pi^{-1}}(x \oplus \tilde{w})$
- 3: Check the second m bits of w' . If not all-zero, **return** $(0, \text{rej})$.
- 4: **if** $\text{basis} = \times$ **then**
- 5: $w' \leftarrow \text{permute}_{\pi^{-1}}(z \oplus \tilde{w})$
- 6: Check the third m bits of w' . If not all-zero, **return** $(0, \text{rej})$.
- 7: $w'' \leftarrow \text{CSS.ClassicalDecode}(\text{first } m \text{ bits of } w')$
- 8: **return** (w'', acc)

Algorithm 16. TC.VerDec($(\pi, x, z), \tilde{\sigma}, c$)

- 1: **for all** gates G in c **do**
- 2: **if** $G = X_i$ **then**
- 3: $x[i] \leftarrow x[i] \oplus \text{permute}_{\pi}(1^m 0^{2m})$ \triangleright update keys (see Section 3)
- 4: **else if** $G = Z_i$ **then**
- 5: $z[i] \leftarrow z[i] \oplus \text{permute}_{\pi}(1^m 0^{2m})$ \triangleright update keys (see Section 3)
- 6: **else if** $G = \text{CNOT}$ **then**
- 7: $(x[i], z[i])(x[j], z[j]) \leftarrow (x[i], z[i] \oplus z[j])(x[i] \oplus x[j], z[j])$ \triangleright update keys
- 8: **else if** G is a measurement in basis b on qubit i **then**
- 9: $(a_i, \text{flag}) \leftarrow \text{TC.VerDecMeasurement}((\pi, x[i], z[i]), \tilde{\sigma}_i, b)$
- 10: **if** $\text{flag} = \text{rej}$ **then**
- 11: **return** $(\Omega, |\text{rej}\rangle)$.
- 12: Execute TC.VerDecQubit on all unmeasured qubits. If it rejects, **return** $(\Omega, |\text{rej}\rangle)$.
- 13: $\sigma' \leftarrow$ the list of decrypted qubits (and measurement outcomes a_i).
- 14: $\sigma'' \leftarrow \sigma'$ with all wires that are not part of the output of c traced out.
- 15: **return** $(\sigma'', |\text{acc}\rangle)$

C Security of verification in TC.

The trap code is proven secure in its application to one-time programs [8]. Broadbent and Wainwright proved authentication security (with an explicit, efficient simulator) [6]. One can use similar strategies to [6, 8] to prove κ -IND-VER for TC.

Theorem 6. *For any adversary \mathcal{A} ,*

$$\Pr[\text{VerGame}_{\mathcal{A}, \text{TC}}(\kappa) = 1] \leq \frac{1}{2} + \text{negl}(\kappa),$$

and thus TC is a κ -IND-VER secure (somewhat) homomorphic encryption scheme.

The proof (again following, e.g., [6]) will use the following lemma, the *Pauli Twirl* [11].

Lemma 9 (Pauli Twirl). *Let ρ be an arbitrary n -qubit state. Then for any Pauli operators P, P' it holds that*

$$\frac{1}{4^n} \sum_{a,b \in \{0,1\}^n} (\mathbf{X}^a \mathbf{Z}^b)^\dagger P \mathbf{X}^a \mathbf{Z}^b \rho \mathbf{X}^a \mathbf{Z}^b P'^\dagger (\mathbf{X}^a \mathbf{Z}^b)^\dagger = \begin{cases} P \rho P^\dagger & \text{if } P = P' \\ 0 & \text{otherwise} \end{cases}$$

Proof (Theorem 6). Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ be an adversary for TC, for the $\text{VerGame}_{\mathcal{A},S}(\kappa)$ security game. Let $sk = (\pi, x, z)$ be the uniformly random keys, with $\pi \in S_{3m}$, $x, z \in \{0,1\}^{3mn}$. Let CSS be a $[[m, 1, d]]$ CSS code, that can correct up to d_c errors. We can let d_c be the security parameter κ , and then $d = 2d_c + 1$ and m depends on the exact properties of CSS.

First note that for TC, the circuit c which is output by \mathcal{A}_2 cannot in any way depend on the bit r : All qubits output by TC.Enc are encoded with the quantum one-time pad, and therefore will look completely mixed whether or not the real or dummy input is given to \mathcal{A}_2 . Also, c is measured when it is supplied to VerDec as classical information. Therefore, we can in general view \mathcal{A} as a probabilistic mixture of adversaries for different choices of c . From now on, we assume that \mathcal{A} uses an arbitrary fixed c without loss of generality (since it can always use the circuit c that wins the game with highest probability).

Next, observe that the accept probability of VerDec within the $\text{VerGame}_{\mathcal{A},S}(\kappa)$ game is independent of the random choice r . The decryption procedure only looks at the trap qubits when choosing whether to accept or reject, and so we can imagine delaying undoing the quantum one-time pad on the data qubits until after the accept or reject choice – which cannot depend on r at all since the encrypted data always looks completely mixed.

In the reject case, VerDec outputs a fixed quantum state, and the quantum one-time pad that is applied to the input of TC.Enc will never be revealed. So in that case \mathcal{A}_3 will never be able to do better than a random guess. To prove security, it then suffices to argue that the state \mathcal{A}_2 outputs in the $r = 0$ case is close to the state \mathcal{A}_2 outputs in the $r = 1$ case, conditioned on VerDec accepting.

Now let D be the quantum operations that are performed by the honest evaluation circuit, i.e., the list of CNOT gates applied transversally to the encrypted qubits. Let $B \in \{I, \text{comp}, H\}^n$ describe for each logical qubit whether it is unmeasured, measured in the computational basis, or measured in the Hadamard basis respectively. To simplify notation, we assume without loss of generality that $\mathcal{H}_R = \mathcal{H}_{R'}$ and that \mathcal{A}_2 can be written as a unitary operation that consists of first applying an arbitrary U acting on $\mathcal{H}_R \otimes \mathcal{H}_C$, and then applying the honest actions D .

Define $M_I = I^{\otimes m} \otimes |0\rangle\langle 0|^{\otimes m} \otimes |+\rangle\langle +|^{\otimes m}$ as the projector corresponding to accepting the traps of an unmeasured qubit (after undoing the permutation and

quantum one-time pad). Similarly define $M_{\text{comp}} = I^{\otimes m} \otimes |0\rangle\langle 0|^{\otimes m} \otimes I^{\otimes m}$ and $M_H = I^{\otimes m} \otimes I^{\otimes m} \otimes |+\rangle\langle +|^{\otimes m}$ as the projectors corresponding to accepting the traps of a measured qubit.

The `VerDecMeasurement` function does not undo the Z part of the quantum one-time pad in case of computational-basis measurement (or the X corrections for the Hadamard basis), making it not possible to immediately apply the Pauli twirl (Lemma 9) as with the unmeasured qubits. Consider a scheme where the decryption procedure for the measurement would first undo the entire quantum one-time pad, and only measure the qubits afterward: this scheme would be functionally completely equivalent to the actual TC scheme. (In [8] this property is called the equivalence between decode-then-measure and measure-then-decode.) Also, note that we can write the Pauli operators and the conditional Paulis as occurring after all other gates; because of the commutation rules between the Clifford group and the Pauli group, they will just correspond to different Pauli operations applied later. Since the Pauli key updates only occur on the data qubits and do not change the acceptance probability at all, we can rewrite the decryption procedure as first checking the traps using the keys without the Paulis, and only then apply the (conditional) Paulis after the checks.

Let $\sigma \in \mathcal{H}_{RX}$ be the output state of \mathcal{A}_1 . Write $\Lambda(\cdot) = \text{CSS.Encode}(\cdot) \otimes |0\rangle\langle 0|^{\otimes m} \otimes |+\rangle\langle +|^{\otimes m}$ as the channel representing the part of the encryption which encodes the qubits and appends the traps and define $\rho = \left(\text{id}_R \otimes \bigotimes_{i=1}^n \Lambda\right)(\sigma)$. Now the state $(\mathbb{I}_R \otimes X^x Z^z \pi^{\otimes n})\rho(\mathbb{I}_R \otimes X^x Z^z \pi^{\otimes n})^\dagger$ is the encrypted input to \mathcal{A}_2 . The plaintext that `VerDec` holds right before running `CSS.Decode`, projected to the accepting case, equals

$$\bigotimes_{i \in [n]} M_{B_i} \mathbb{E}_{\pi \in S_{3m}} \left[\mathbb{E}_{x, z \in \{0,1\}^{3mn}} \left[\pi^{\dagger \otimes n} X^{f_c(x)} Z^{f_c(z)} D U X^x Z^z \pi^{\otimes n} \rho \pi^{\dagger \otimes n} X^x Z^z U^\dagger D^\dagger X^{f_c(x)} Z^{f_c(z)} \pi^{\otimes n} \right] \right]$$

where U acts on the reference system R and the ciphertext register, and all other operations only act on the ciphertext register.⁵ The transformation f_c represents the updating of the quantum one-time pad keys as function of the applied circuit – these are the keys that are used by the decryption circuit.

Now define D' as the unitary operation which applies CNOT gates transversally on the $3m + 3m$ qubits when listed in c . By construction of the key-update rules, we have that $X^{f_c(x)} Z^{f_c(z)} \pi^{\dagger \otimes n} D = D' X^x Z^z \pi^{\dagger \otimes n}$. Using that identity and the Pauli twirl (Lemma 9) we decompose U into a probabilistic mixture of Pauli operations:

$$\begin{aligned} & \bigotimes_{i \in [n]} M_{B_i} D' \mathbb{E}_{\pi \in S_{3m}} \left[\mathbb{E}_{x, z \in \{0,1\}^{3mn}} \left[\pi^{\dagger \otimes n} X^x Z^z U X^x Z^z \pi^{\otimes n} \rho \pi^{\dagger \otimes n} X^x Z^z U^\dagger X^x Z^z \pi^{\otimes n} \right] \right] D'^\dagger = \\ & \bigotimes_{i \in [n]} M_{B_i} D' \mathbb{E}_{\pi \in S_{3m}} \left[\mathbb{E}_{P \in \mathcal{P}^{3mn}} \left[|\alpha_P|^2 \pi^{\dagger \otimes n} (P \otimes U_P) \pi^{\otimes n} \rho \pi^{\dagger \otimes n} (P \otimes U_P^\dagger) \pi^{\otimes n} \right] \right] D'^\dagger \end{aligned}$$

⁵ The expectation value is always taken over the uniform distribution, e.g., $\mathbb{E}_{\pi \in S_{3m}}$ is nothing more than a short way of writing $\frac{1}{(3m)!} \sum_{\pi \in S_{3m}}$.

Expressions of this form were carefully analyzed in the earlier trap-code security proofs – we will for completeness finish our security sketch, but see, e.g., [6] for a more precise analysis.

First observe that if these Pauli operators do not change any logical qubit, this expression will be exactly the same as the state that \mathcal{A}_3 receives in the $r = 1$ case – namely the claimed circuit c , as represented by D' , the measurements, and the (conditional) Paulis that will be effectively performed by decryption, applied to the data. Consider what form the Pauli operator P would need to change a specific logical qubit i . First consider an unmeasured qubit i . Because CSS can correct up to d_c errors, only those Paulis that are non-identity on more than d_c qubits will cause the logical qubit after decoding to change. Say without loss of generality that this Pauli operator P_i has an X on at least $d_c/2$ out of the $3m$ physical qubits that encode i . (If the operator consists of more Z components than X components, we could argue using Z instead.) Consider the probability for a randomly chosen $\pi \in S_{3m}$ that all these X do not end up in the positions $m + 1$ to $2m$, i.e., each misses the computational basis traps. For each X, the probability of missing all trap positions, conditioned on no trap being hit yet, is always at most $2/3$. Therefore the probability that all traps are missed is at most $(2/3)^{d_c/2}$. A more careful combinatorial analysis which includes the Z flips improves this to $(2/3)^{d_c}$ [8, 6], but this simple bound suffices for us.

Now, consider the case that i is a qubit on which a computational-basis measurement has been performed, of which only the corresponding traps are checked. For these qubits, the Pauli Z parts of the attack are not detected, but they also do not change the output: Since the data qubits are measured, only the X Paulis will change anything in the data. Therefore, the operator P_i will have to contain at least d_c X Paulis on the $3m$ physical qubits. Now repeating the same argument as for the unmeasured qubits, we see that the probability over a random permutation π that all traps are missed is at most $(2/3)^{d_c}$. The analogous argument works for the Hadamard-basis measurements.

To conclude, the part of the output of \mathcal{A}_2 that has been changed from that what would come out of the honest evaluator, and still is accepted, has norm at most $(2/3)^{d_c}$, both in case $r = 0$ and $r = 1$. This norm gives an upper bound to the trace distance between the states that \mathcal{A}_3 receives in the $r = 0$ case and the $r = 1$ case, since for all lower-weight Pauli attacks these states are exactly the same (by the error-correction property of CSS). The final guessing probability is then bounded as

$$\Pr[\text{VerGame}_{\mathcal{A}, \text{TC}}(\kappa) = 1] \leq \frac{1}{2} + \frac{1}{2} \left(\frac{2}{3}\right)^{d_c}.$$

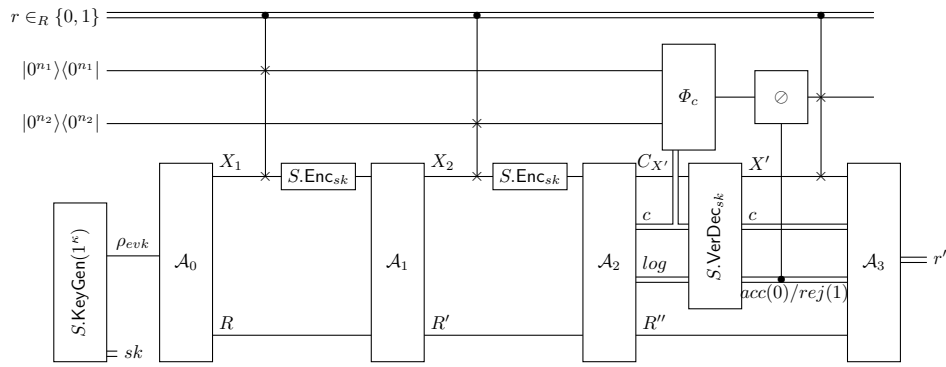
Since we picked the parameters of CSS such that d_c scaled with κ , this completes the proof. \square

D Security of TC with multiple encryptions

In Section 5, we will use the IND-VER property of TC to prove verifiability for our new scheme. In order to achieve this, we will actually need a slightly stronger

notion of verifiability for TC: IND-VER- n , where the adversary is allowed to submit plaintexts in multiple rounds, which are either all encrypted or all swapped out. In this subsection, we show that TC also fulfills this stronger notion. For our purposes in Section 5, it suffices to show that TC is secure against an adversary that is allowed two rounds (IND-VER-2), but the definitions and proof trivially extend to the general case.

Definition 9 (VER-2 indistinguishability game $\text{VerGame}_{\mathcal{A},S}^2(\kappa)$). For an adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$, a scheme S , and a security parameter κ , $\text{VerGame}_{\mathcal{A},S}^2(\kappa)$ is the following game:



Here, n_1 and n_2 are the respective dimensions of the X_1 and X_2 registers.

Definition 10 (κ -IND-VER-2). A vQFHE scheme $S = (\text{KeyGen}, \text{Enc}, \text{Eval}, \text{VerDec})$ has 2-round κ -indistinguishable verification if for any QPT adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$,

$$\Pr[\text{VerGame}_{\mathcal{A},S}^2(\kappa) = 1] \leq \frac{1}{2} + \text{negl}(\kappa).$$

Here, the probability is taken over $\text{KeyGen}(1^\kappa)$, Enc , VerDec , and \mathcal{A} .

Lemma 10. TC is κ -IND-VER-2.

Proof. Let $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ be an arbitrary polynomial-time adversary for the VER-2 indistinguishability game for TC. For notational convenience, write the secret key as $sk = (\pi, x_1, z_1, x_2, z_2)$, where x_1 and z_1 are lists of $3mn_1$ bits, sufficient for encrypting X_1 , and analogously x_2 and z_2 are lists of $3mn_2$ bits.

We now slightly alter the VER-2 game in the following way. In the first encryption step of the game, instead of providing \mathcal{A}_1 with $\text{TC.Enc}_{(\pi, x_1, z_1)}$ applied to the register X_1 , we provide \mathcal{A}_1 with the halves of n_1 EPR pairs, and perform Bell measurements between the other halves and the qubits in X_1 , after they have been CSS-encoded and permuted with traps. Let the outcomes of these measurements be given by $a, b \in \{0, 1\}^{3mn_1}$: a and b describe the effective X and Z Paulis that are applied to X_1 by these teleportation measurements. To undo

these Paulis, we update sk to $(\pi, x_1 \oplus a, z_1 \oplus b, x_2, z_2)$ at this point. Here, \oplus is bitwise addition modulo 2. Since the quantum one-time pad keys x_1 and z_1 are chosen uniformly at random, and are completely hidden from the perspective of the adversary, the new keys $x_1 \oplus a$ and $z_1 \oplus b$ are valid keys that are sampled from the same distribution. Hence, the winning probability of \mathcal{A} is not affected by this change of the game.

A second small change to the game is the following: instead of performing the Bell measurements and the secret-key update immediately, it is done only after \mathcal{A}_1 has provided its query in X_2 . Since these actions happen only on wires which are not accessible to \mathcal{A}_1 and otherwise also not touched in this stage of the game, this change also does not affect the execution or outcome of the game in any way.

We have now arrived at an interesting situation: \mathcal{A}_1 only receives halves of EPR pairs, and so its choice for X_2 or R' is not based on the first ciphertext received from the challenger – that ciphertext will only be generated after execution of \mathcal{A}_1 . We can merge \mathcal{A}_0 and \mathcal{A}_1 into a single QPT algorithm that produces X_1 and X_2 simultaneously. When viewed as such, \mathcal{A} is an adversary for the single-query VER indistinguishability game, and we can conclude that

$$\Pr[\text{VerGame}_{\mathcal{A}, \text{TC}}^2(\kappa) = 1] = \Pr[\text{VerGame}_{\mathcal{A}, \text{TC}}(\kappa) = 1].$$

Since we know that the latter probability is bounded by $\frac{1}{2} + \text{negl}(\kappa)$ from Theorem 4, so is the first. \square