Limitations on Transversal Computation through Quantum Homomorphic Encryption

Michael Newman[1] and Yaoyun Shi[2]

[1]Department of Mathematics
[2]Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor, MI 48109, USA
mgnewman@umich.edu, shiyy@umich.edu

Abstract

Transversality is a simple and effective method for implementing quantum computation fault-tolerantly. However, no quantum error-correcting code (QECC) can transversally implement a quantum universal gate set (Eastin and Knill, *Phys. Rev. Lett.*, 102, 110502). Since reversible classical computation is often a dominating part of useful quantum computation, whether or not it can be implemented transversally is an important open problem. We show that, other than a small set of non-additive codes that we cannot rule out, no binary QECC can transversally implement a classical reversible universal gate set. In particular, no such QECC can implement the Toffoli gate transversally.

We prove our result by constructing an information theoretically secure (but inefficient) quantum homomorphic encryption (ITS-QHE) scheme inspired by Ouyang *et al.* (arXiv:1508.00938). Homomorphic encryption allows the implementation of certain functions directly on encrypted data, i.e. homomorphically. Our scheme builds on almost any QECC, and implements that code's transversal gate set homomorphically. We observe a restriction imposed by Nayak's bound (*FOCS* 1999) on ITS-QHE, implying that any ITS quantum *fully* homomorphic scheme (ITS-QFHE) implementing the full set of classical reversible functions must be highly inefficient. While our scheme incurs exponential overhead, any such QECC implementing Toffoli transversally would still violate this lower bound through our scheme.

1 Introduction

1.1 Restrictions on transversal gates

Transversal gates are surprisingly ubiquitous objects, finding applications in quantum cryptography [29], [25], quantum complexity theory [11], and of course quantum fault-tolerance. The instability of quantum information is well-documented, and quantum error-correcting codes [24] allow the encoding of single qubits into multiple qubit systems so that errors on small subsets of physical qubits can be corrected [21]. Performing computations on these codes carries the risk of propagating errors between different subsystems, unless the code can implement the computation in a way that preserves the subsystem structure. Informally, these types of logical operators that decompose as a product across the subsystems are called *transversal*, and the oft-cited Eastin-Knill theorem [15], [39] limits the ability of quantum codes to prevent this error propagation.

Theorem 1 (Eastin-Knill)**.** *No quantum error-correcting code can implement a quantum universal transversal gate set.*

These transversal gate sets are valuable as most models of fault-tolerant quantum computation implement associated transversal gate sets fault-tolerantly "for free". Incurring comparatively significant overhead, often in the form of magic state distillation [18], [22], gauge fixing [5], [26], or more recently deconstructions of non-transversal gates into fault-tolerant pieces [37], one can fault-tolerantly implement some remaining gate set making the computation space universal. Improving the efficiency of this overhead and designing new fault tolerant architectures to supplement transversal gates is central to quantum fault tolerance.

Implementing fault-tolerant classical reversible computation efficiently would be extremely desirable as many quantum algorithms are primarily classical subroutines with a relatively small number of quantum gates, and there have been several proposals for doing so [31], [13], [23]. For example, factoring a cryptographically large RSA key using Shor's algorithm requires around 3×10^{11} Toffoli gates to perform modular exponentiation alone, and is the dominating portion of the circuit [19]. As Toffoli is universal for classical reversible computation, one might ask if there are any quantum error-correcting codes that can naturally implement Toffoli, and thus classical computations, transversally? We give restrictions on the ability of QECCs to do this.

Theorem 2 (Informal)**.** *Almost no quantum error-correcting code can implement a classical universal transversal gate set. In particular, almost no quantum error-correcting code can implement the Toffoli gate transversally.*

The only exceptions to our theorem are non-additive distance d codes that decompose as d-fold product states in their logical computational basis, where each "subcode" itself fails to be erasure-correcting. Essentially, one can think of these as maximally redundant quantum codes: they are the concatenation of a repetition code with some distance 1 inner code, similar to Shor's stabilizer code written as a 3-fold product of GHZ states. We do not expect that any such code can implement Toffoli transversally, but it remains a case our proof technique cannot rule out. In particular, our proof does apply to all binary additive codes. The result is perhaps slightly surprising since there exist QECCs (e.g. triorthogonal codes) that can implement the CCZ gate transversally [31], and in fact transversal Toffoli gates can map between different quantum Reed-Solomon codes by increasing the degree of the underlying polynomial [13].

1.2 Quantum homomorphic encryption

The main ingredient in our proof is an information-theoretically secure homomorphic encryption scheme. Generally, homomorphic encryption [20] is a means of delegating computation on sensitive data securely. It allows for the encryption of data in such a way that another party can perform meaningful computation on the ciphertext *without* decoding, preserving the security of the underlying plaintext. A scheme is termed *fully* homomorphic encryption (FHE) if it can implement a universal class of functions in some computation space.

Recently, extensions of homomorphic encryption to the quantum setting have been considered. Instead of encrypting classical data and implementing addition and multiplication gates homomorphically, quantum homomorphic encryption aims to encrypt quantum data and implement unitary gates homomorphically. Progress was made in [10], and recently [14] extended this work to a leveled scheme that could homomorphically implement all polynomial-sized quantum circuits.

The aforementioned schemes are only computationally secure, since they use classical FHE as a subroutine. This is no great indictment: FHE is built on the difficulty of certain hard lattice problems that are leading candidates for quantum-secure encryption [12], [33]. However, quantum information often promises information-theoretic security (ITS) guarantees that are impossible classically. Intermediate advances have also been made in this more restrictive setting. One such scheme allows for the implementation of a large class of unitaries homomorphically, but with less stringent ITS guarantees [36].

More recently, [29] proposed a compact ITS-QHE scheme in which the size of the encoding scales polynomially with size of the input for the limited Clifford circuit class. This scheme achieves the strongest notion of imperfect ITS, with the probability of distinguishing between any two ciphertexts exponentially suppressed in the size of the encoding.

The scheme in [29] is based on a "noisy" quantum encoding of the data. They take an encoding circuit for a particular quantum code and replace the ancilla bits of the encoding with uniformly random noise. Their encryption is then choosing a random embedding of this code into yet more uniformly random noise. This scheme links ITS-QHE to transversal gates: the transversal gates for their code are exactly those gates that can be implemented homomorphically.

1.3 Limitations on ITS-QHE

There are fundamental limitations on what ITS homomorphic encryption can do. It is known that for a purely classical scheme, efficient ITS-FHE is impossible, violating lower bounds in the setting of single server private information retrieval [16]. It was further shown that in the best case scenario, when the mutual information between the plaintext and the ciphertext is precisely zero, efficient *quantum* FHE is impossible [38].

This no-go result actually applies to the more restrictive setting of classical data being encrypted into quantum data, while allowing only classical reversible functions to be evaluated homomorphically. Both [38] and [29] ask whether relaxing to imperfect ITS-security might allow for efficient ITS-QFHE. Unfortunately, this is not the case.

Proposition 3. *[Informal] Efficient ITS-QFHE is impossible.*

Concurrent to this work, this proposition was observed in [25]. We provide a precise statement and proof of this restriction in Appendix A. This result can be seen by combining the proof technique in [16] with similar single server private information retrieval bounds in the quantum setting [4]. In essence, the inefficiency of ITS-QFHE follows from viewing ITS-QHE (on *classical data* using a *quantum encoding*) as a certain quantum random access encoding (QRAC) (see [1]) of the function

class we wish to implement homomorphically. Well-known bounds on QRACs [27] place lower bounds on the encoding size of such a scheme, precluding efficiency. Using a variant of the code-based ITS-QHE scheme proposed in [29], we can then argue that (almost) any QECC implementing the Toffoli gate transversally would yield a scheme violating this lower bound.

It is worth noting that similar tasks such as blind quantum computation [9] and computing on encrypted data [17], [8] allow ITS solutions, but they do so at the cost of interactivity between the client and server. We do not allow this interactivity in our definition of homomorphic encryption.

1.4 Comparison to related works

The five works the most closely resemble our results are [15],[39] and [6], which place restrictions on transversal gate sets for QECCs, and [29] and [25], which use similar ITS-QHE constructions. We very roughly summarize these results and compare them to our own.

In [39], Zeng *et al.* were some of the first to place restrictions on quantum universal transversal gate sets for *additive* quantum codes by elucidating the stabilizer group structure. Further work in [2] classified the set of diagonal gates that can implement one and two qubit logical operations in stabilizer codes. Shortly thereafter, [15] showed that for *any* QECC, the transversal gate set must be finite, and so cannot approximate with arbitrary precision the full unitary group. Intuitively, they make a Lie type argument by showing that infinitesimal transversal operations are themselves linear combinations of local error operators. Since these unitaries must act identically on the codespace, it follows that the group of transversal operations must be finite.

More recently, [6] placed restrictions on the more general class of topologically protected logical gates in topological stabilizer codes, which include transversal gates as an optimal subset. They showed that for a topological stabilizer code defined on a d-dimensional lattice, any such gate must lie in the dth level of the Clifford hierarchy. These results were extended in [32] to more general stabilizer subsystem codes, and in Appendix D, we detail how these arguments can be used to rule out classical reversible transversal computation for the subclass of stabilizer codes.

In the direction of ITS-QHE, [29] gave a compact and efficient ITS-QHE scheme for the restricted class of Clifford circuits. Using magic state injection, they complete a universal gate set by adding the T-gate. However, because the client and server cannot communicate during the protocol, they must limit themselves to circuits using a constant number of T-gates. Again, their encryption is a "noisy" encoding of the data into some code followed by a secret embedding into random noise. With this encryption, they are able to generate indistinguishable outputs using only polynomial overhead in the input size. In the more recent work [25], the authors independently observe Proposition 3. They take the more positive approach of arguing what can be done in spite of this limitation, extending the ITS-QHE schematic in [29] to other *particular* error-correcting codes and using code concatenation to achieve security with only polynomial overhead. This achieves ITS-QHE on the larger circuit class IQP^+, which is probably not classically simulable [7].

Because of the stringent lower bounds placed by Nayak, we actually forgo the noisy encoding circuit and embed QECCs directly into random noise after removing a correctable set of qubits. This has the effect of increasing the overhead by an exponential factor in order to achieve security, but thanks to the roomy lower bound, this factor is still too small to allow an ITS-QFHE scheme.

We can argue directly about the security of this scheme using the nonlocality of the quantum information being encoded in almost any QECC. The idea is conceptually simple: in order to obtain encryptions of the data that are both secure and (sufficiently) short, we must inject randomness into the encodings themselves by withholding qubits from the code. While ordinarily this would negatively affect the correctness of homomorphic evaluation, the error-correcting property allows us to inject this randomness while still maintaining perfect recoverability. Then intuitively, spreading

the information across the subsystems limits the complexity of the class of logical operators that don't couple the subsystems, i.e. the *transversal* operators. This differs fundamentally from the approaches in [39] and [15] in that it is a quantitative information-type bound.

It is not without its drawbacks however, as these maximally redundant codes fail to "spread out" the information sufficiently. The prototypical example is Shor's code, which is the concatenation of a bit-flip and phase-flip code. However, we can argue directly using the stabilizer group structure that no such *additive* code can implement Toffoli transversally.

2 Preliminaries

2.1 Quantum Information

We quickly review some standard notation, followed by some less standard tools we will need from quantum information theory. For a more complete view, see [28].

Throughout, we will be working with 2-level qubit quantum systems. We denote by $|\mathcal{H}| = \log(\dim(\mathcal{H}))$, the number of qubits constituting state space $\mathcal{H}$. We define a general quantum state to be a positive semi-definite operator $\rho \in L(\mathcal{H})$ of trace one. We call such a state pure if $\mathrm{rank}(\rho) = 1$, otherwise we call it mixed, and note that such an operator is mixed if and only if $\mathrm{Tr}(\rho^2) < 1$. For any operator $U \in \mathcal{H}_A$, we use the notation U^A to indicate the operator $U_A \otimes I_B \in L(\mathcal{H}_A \otimes \mathcal{H}_\mathcal{B})$. When it is unclear which space a state lives in, we will denote its state space as a superscript (e.g. $|\psi\rangle^A$). We also sometimes adopt the notation that for $\rho \in L(A \otimes B)$, $\rho^A = \mathrm{Tr}_B(\rho)$. By slight abuse of notation, we also adopt the convention that for any permutation $\pi \in S_n$, π can also indicate the unitary permutation operator corresponding to the physical permutation of qubits. We also sometimes omit the dimension of an identity operator I, but usually the dimension is implicitly its trace normalization factor, e.g. I/D acts on a space of dimension D.

The norm $\|\cdot\|_p$ refers the usual Schatten p-norm, so that for any $A \in L(\mathcal{H})$ with singular values $(a_1, \ldots, a_n)$,

$$\|A\|_p = \left(\sum_{i=0}^{n} a_i^p\right)^{1/p} \text{ for } p > 1, \text{ and } \quad \|A\|_1 = \sum_{i=0}^{n} |a_i|.$$

Further recall that

$$\frac{1}{2}\|\rho - \sigma\|_1 = \max_{P \leq I} \mathrm{Tr}(P(\rho - \sigma)),$$

and so we can think of the 1-norm as a means of bounding the ability to distinguish two quantum states, where $\leq$ refers to the positive semidefinite partial ordering. For a collection of quantum states $\{\rho_S\}$ indexed by $S \in \mathcal{S}$, we sometimes write $E_S[\rho_S]$ to denote the expectation over a uniformly random choice of S, $E_S[\rho_S] = \frac{1}{|\mathcal{S}|}\sum_{S\in\mathcal{S}} \rho_S$. We will regularly be referring to several particular gates, and so list them here.

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

$$CX = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \qquad CZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$$

$$\text{Toff} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix} \qquad CCZ = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix}$$

Definition 4. An (n, m, p)-*quantum random access code* (QRAC) is a mapping from an n-bit string x to an m-qubit quantum state ρ_x along with a family of measurements $\{M_i^0, M_i^1\}_{i=1}^n$ satisfying, for all $x \in \{0,1\}^n$, $i \in [n]$,

$$\text{Tr}(M_i^j \rho_x) \geq p \text{ if } x_i = j.$$

More generally, we can consider some protocol M_i for retrieving x_i and we call this protocol the ith query of the QRAC, satisfying $\Pr[M_i(\rho_x) = x_i] \geq p$.

Definition 5. An n-qubit *quantum code* is simply a subspace C of an n-body Hilbert space $\mathcal{H}$ along with a fiducial orthonormal logical basis. Let P_C denote the projection onto C. The code is further called an $[[n, k, d]]$ *quantum error-correcting code* (QECC) if it is a subspace C of dimension 2^k satisfying, for any $\lfloor \frac{d-1}{2} \rfloor$-local operators E_a, E_b,

$$P_C E_a^\dagger E_b P_C = \lambda_{ab} P_C.$$

for some Hermitian (λ_{ab}). We call the E_a, E_b *correctable errors* and we say d is the *distance* of the code.

Operationally, this means that there exists a recovery channel $\mathcal{R}$ by which any $\lfloor \frac{d-1}{2} \rfloor$-local error (acting nontrivially on at most $\lfloor \frac{d-1}{2} \rfloor$ qubits) can be corrected. Recall also that any $[[n, k, d]]$ QECC can correct up to $(d-1)$ errors in known locations. We call these types of errors *erasure errors*, and so call codes satisfying $d \geq 2$ *erasure-correcting codes*.

For simplicity, we will restrict our discussion to QECCs encoding a single qubit, i.e. $[[n, 1, d]]$ QECCs. A quick review of the proof shows that we can make this assumption without loss of generality by arguing against a classical universal transversal gate set on any single encoded logical qubit. By similar reasoning, the proof applies to subsystem codes as well, where we require the logical operators of such a code to be independent of the state of the gauge qubits.

When we have a collection of p logical qubits, *each* encoded into an n-qubit code, we can decompose the np physical qubits into a fixed n-wise partition of p-qubits each so that every partitioning set contains exactly one qubit from each code block. We refer to these partitioning sets as the subsystems of the collective code.

Definition 6. An n-qubit stabilizer group $\mathcal{S}$ is an abelian subgroup of the n-qubit Pauli group not containing $-I$. An n-qubit *additive* (or stabilizer) *code* encoding k logical qubits can be described as the simultaneous $(+1)$-eigenspace of the Pauli operators comprising an n-qubit stabilizer group S with $n-k$ generators. The logical Pauli operators of this code correspond to the normalizer cosets $\mathcal{N}(\mathcal{S})/\mathcal{S}$, and it follows that the distance of the code is the minimal weight operator in $\mathcal{N}(\mathcal{S})/\mathcal{S}$.

Definition 7. For any quantum error correcting code C, we define its logical states $|\psi\rangle_L$ to be the physical encoding of $|\psi\rangle$ in C. We define the logical gate U_L to be a codespace preserving physical gate that satisfies, for all $|\psi\rangle_L \in C$, $U_L|\psi\rangle_L = (U|\psi\rangle)_L$.

The set of *transversal gates* $\mathcal{T}_C$ associated to C are those logical gates that decompose as a product across the subsystems. That is to say, $U_L \in \mathcal{T}_C$ if $U_L = U_1 \otimes \ldots \otimes U_n$, where n is the length of the code, each U_i acts on a single subsystem, and U_L is a codespace preserving map on the code $C^{\otimes r}$ for U an r-qubit gate. We further define a logical gate to be *strongly transversal* if it decomposes as $U_L = U^{\otimes n}$. Following the example of [39], we do not allow coordinate permutations in our definition of transversality.

Definition 8. We say a quantum code $C = \mathrm{Span}_{\mathbb{C}}(|\tilde{0}\rangle, |\tilde{1}\rangle)$ is an *r-fold code* if it can be written as

$$|i\rangle_L = \bigotimes_{j=1}^{r} |\psi_{ij}\rangle.$$

where each vector $|\psi_{ij}\rangle$ does not further decompose as a product state across any bipartition. We additionally assume that $r \leq d$, that $|\psi_{0j}\rangle$ and $|\psi_{1j}\rangle$ occupy the same subsystem, and that $|\psi_{0j}\rangle \perp |\psi_{1j}\rangle$. It then makes sense to refer to $\mathrm{Span}\{|\psi_{0j}\rangle, |\psi_{1j}\rangle\}$ as the *jth subcode*. These assumptions are natural, and we justify them in our discussion.

If the code is additionally an $[[n, 1, d]]$ QECC with $r = d \geq 3$ and each subcode has distance 1, we simply call the resulting code a *maximally redundant code*. Note that any (pure state) code is at least a 1-fold code.

The guiding example is Shor's code, which can be seen as the concatenation of a repetition outer code and a complementary GHZ inner code, neither of which is quantum erasure correcting. In the case that the subcodes are identical, any maximally redundant code is just the concatenation of a repetition code with some distance 1 subcode. Intuitively, these are codes for which you can't erase enough qubits to mix the state while still remaining perfectly correctable: while redundancy can be used in classical error-correcting codes to protect information, quantum error-correcting codes must "spread out" information to protect it. In this sense, these codes are maximally redundant because they "spread out" the information the least. We show that non-additive maximally redundant codes (i.e. maximally redundant codes for which the subcodes are comprised of non-stabilizer subspaces) are the only binary QECCs with the hope of implementing logical Toffoli transversally.

2.2 Homomorphic Encryption

We define an ITS-QHE scheme as three algorithms performed between two parties which we will call Client and Server. We restrict ourselves to the more limited setting of a quantum scheme implementing Boolean functions on classical data using quantum encodings. Of course, any impossibility result then extends to the more difficult task of quantum computations on quantum inputs.

The parameters of such a scheme are given by $(n, m, m', \epsilon, \epsilon')$ and some gate set $\mathcal{F}$. Formally, we define the algorithms of an ITS-QHE scheme as acting on Client's private workspace $\mathcal{C}$, a message space $\mathcal{M}$ sent from Client to Server after encryption, and a message space $\mathcal{M}'$ sent from Server to Client after evaluation.

(*i*) $QHE.Enc(x) = \rho^{\mathcal{CM}}$, in which the client chooses an n-bit input x and encrypts with some private randomness to obtain ρ. We assume that any quantum evaluation key is appended to the encryption. Client then sends the message portion of the encryption $\rho^{\mathcal{M}}$ to the Server. We define m to be the length of this message, the size of the encoding.

(*ii*) $QHE.Eval_f : L(\mathcal{M}) \longrightarrow L(\mathcal{M}')$, in which Server, with description of some circuit f, applies an evaluation map to an encrypted state, possibly consuming an evaluation key in the

process. Server then sends his portion of the state $\sigma^{\mathcal{C}\mathcal{M}'} \coloneqq (I^{\mathcal{C}} \otimes QHE.Eval_f)(\rho^{\mathcal{C}\mathcal{M}})$ back to Client, and we define the length of this message to be m', the size of the evaluated encoding.

(*iii*) $QHE.Dec(\sigma^{\mathcal{C}\mathcal{M}'}) = y$, in which Client decrypts the returned evaluated encoding using her side information $\mathcal{C}$ and recovers some associated plaintext y.

As an encryption scheme, the above should certainly satisfy

$$QHE.Dec(QHE.Enc(x)) = x.$$

and as an ITS-QHE scheme, there are three additional properties the scheme should satisfy as well.

(*i*) ϵ-*Information-theoretic security*: for any inputs x, y, letting $\rho_x, \rho_y \in L(\mathcal{M})$ denote the outputs of $QHE.Enc$ on $\mathcal{M}$ (thinking of these states as mixtures of encryptions under uniformly random choices of secret key),

$$\|\rho_x - \rho_y\|_1 \leq \epsilon.$$

(*ii*) $\mathcal{F}$-*homomorphic*: for any circuit $f \in \mathcal{F}$ and for any input x,

$$\Pr[QHE.Dec(QHE.Eval_f(QHE.Enc(x)))_1 \neq f(x)] \leq \epsilon'$$

where the probability is over the randomness of the protocol and the subscript 1 denotes the first bit of the output. This restriction to the first bit is just to argue directly about Boolean functions. For ease of exposition, we also assume without loss of generality that our protocols are perfectly correct, and allow $\epsilon' = 0$. We call the scheme fully homomorphic if it is homomorphic on the set of all classical Boolean circuits.

(*iii*) *Compactness*: a priori, the server could do nothing except append a description of the circuit f to be run by the decryption function after decrypting. To avoid trivial solutions like this, we demand that the total time-complexity of Client's actions in the protocol do not scale with the complexity of the functions to be evaluated, but only with some fixed function on the size of the input. Intuitively, this captures the motivation behind homomorphic encryption: limiting the computational cost to Client. However, we note that the standard definition of compactness refers to the time-complexity of the decrypt function specifically.

We denote a scheme homomorphic for some class of functions $\mathcal{F}$ and satisfying all of these properties as an $\mathcal{F}$-ITS-QHE scheme. If $\mathcal{F}$ is the set of all Boolean circuits, we denote such a scheme as an ITS-QFHE scheme. We observe that such a scheme must be inefficient. For a precise statement and proof, see Appendix A.

Proposition 9. *The communication cost of ITS-QFHE must be exponential in the size of the input.*

3 A coding based ITS-QHE scheme

We now consider a strategy for implementing compact QHE using quantum codes. This will be a simple "block" embedding encryption scheme homomorphically implementing *quantum* circuits on *classical* input, and is similar to the construction in [29]. We will use the error-correcting property to withhold a correctable set of qubits from the encoding.

Coding QHE Scheme:

Arguments:

$$\begin{array}{rcl} C & = & \text{an } [[n+r,1,d]] \text{ } r\text{-fold QECC with } r < d \\ \vec{i} & \in & \{0,1\}^p \\ m & = & \text{the size of each noise code block} \\ S & \in & [m]^n, \text{ the secret key} \end{array}$$

1. On input $\vec{i} \in \{0,1\}^p$, encode $\vec{i}$ as the pure state $\bigotimes_{\ell=1}^{p} |i_\ell\rangle_L$, for $\{|0\rangle_L, |1\rangle_L\}$ the logical computational basis defining C.

2. Let R be a collection of r subsystems, each of p-qubits, comprised of one subsystem from each subcode. Then form $\gamma^{\vec{i}} = \mathrm{Tr}_R(\bigotimes_\ell |i_\ell\rangle_L)$. Essentially, $\gamma^{\vec{i}}$ is the state of the collection of codewords with each codeword missing one subsystem from each of its subcodes.

3. Initialize n $(p \times m)$ arrays of maximally mixed qubits, and replace the S_j-th column of each array with the j-th subsystem of $\gamma^{\vec{i}}$. This forms the encrypted state.

4. Publish a constant number of labeled encryptions of 0 and 1, to be used as ancilla in homomorphic evaluation.

Figure 1: A description of the encryption procedure for the code based QHE scheme.

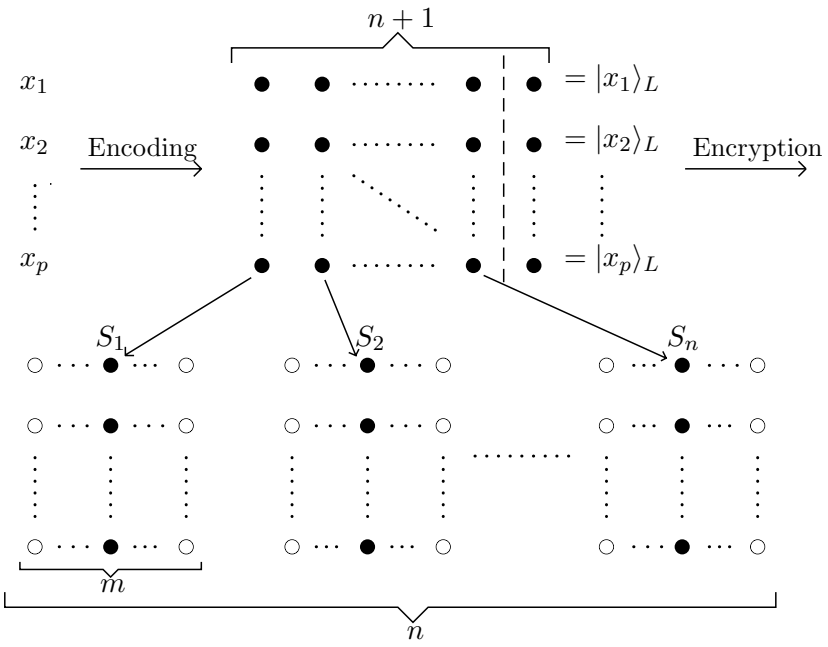

Figure 2: A diagram illustrating the code-based QHE scheme for an $(n+1)$-length 1-fold quantum code while withholding a single subsystem. The $(n+1)$-th subsystem remains in the hands of Client. The arrows connecting the subsystems indicate where each subsystem (i.e. column) is being mapped. The filled dots represent code qubits, while the empty dots represent maximally mixed qubits.

The scheme is detailed in Figures 1 and 2. Using that notation to summarize, our encryption channel $\mathcal{E}$ is defined, for secret key S and input string $\vec{i}$, as $\mathcal{E}(S, \vec{i}) = \gamma_S^{\vec{i}}$. We sometimes use the notation γ_S instead of $\gamma_S^{\vec{i}}$ or γ instead of $\gamma^{\vec{i}}$, omitting $\vec{i}$ when we are unconcerned with the underlying plaintext. The total size of our encrypted input is mnp qubits. In our preceding notation, the described scheme has parameters $(p, mnp, mnp, \epsilon(m,p), 0)$, implementing the set of gates $\mathcal{T}_C$ homomorphically.

Lemma 10. *Let $\mathcal{E}$ be the encryption scheme detailed in Figure 1. Let $\mathcal{T}_C$ denote the group of transversal operators associated to the underlying quantum code C. Then, $\mathcal{E}$ is $\mathcal{T}_C$-homomorphic.*

Proof. Let U_L be the logical operator we wish to apply to some codestate $|\psi\rangle_L$. By definition, $U_L \in \mathcal{T}_C$ implies U_L can be decomposed as a product operator $U_1 \otimes \ldots \otimes U_{n+r}$ where U_i is an operator that acts only on the i-th subsystem of the code. Then, without knowledge of the secret key S, a third party can implement U_L by applying the operator

$$\bigotimes_{i=1}^{n} \bigotimes_{j=1}^{m} U_i$$

where each U_i is an operator local to some subsystem in Server's possession (that is to say, on one of the columns in the corresponding array). Returning the resulting data to a party with the secret key, that party can decrypt to obtain a state of the form $V^R U_L |\psi\rangle_L$, where V^R is supported on the r subsystems that Client has withheld. Since $r < d$, viewing V^R as an erasure error on r subsystems, there exists some recovery channel $\mathcal{R}$ such that $\mathcal{R}(V^R U_L |\psi_L\rangle) = U_L |\psi\rangle_L$. Decoding, we obtain $U|\psi\rangle$ as desired.

□

Note that this scheme is a $\mathcal{T}_C$-homomorphic, non-leveled, and compact QHE scheme, since the recovery and decryption channel do not depend on the complexity of U.

We now aim to compute the security $\epsilon(m,p)$ of the proposed scheme, namely the tradeoff between the size of the input p, the size of the encoding mnp, and the ITS guarantee. To avoid confusion, we point out here that the code size n is a constant, as we are not concatenating to achieve security, just amplifying the size of the noise into which we are embedding.

We want to show that while the scheme is inefficient, its parameters still defeat Nayak's bound. To simplify the security proof, we use the more stringent requirement that the outputs are indistinguishable from uniformly random noise. Here, we will see that the nonlocality of the information stored in QECCs is essential in its allowing us to withhold qubits while still delegating computation to Server. This imposes the requirement of using quantum error-correcting codes, as evidenced by the following observation.

Lemma 11. *Suppose we replace the preceding scheme with one that does not withhold any of the physical qubits comprising the (pure state) code. Then if $m = o(2^p)$, ϵ must be bounded away from zero.*

Proof. Counting the rank of the encrypted state, note that $\text{rank}(\gamma_S) = 2^{np(m-1)}$. Then,

$$\begin{aligned}\text{rank}(E_S[\gamma_S]) &\leq m^n 2^{np(m-1)} \\ &\leq 2^{n(p(m-1)+\log(m))}.\end{aligned}$$

Thus, the fraction of nonzero eigenvalues must be at most $(2^n)^{\log(m)-p}$. Since $\log(m) = o(p)$, the fraction of nonzero eigenvalues goes to zero, and so $\|E_S[\gamma_S] - I/2^{mnp}\|_1$ must be bounded away from zero as claimed. □

4 Security tradeoff for the QHE coding scheme

Our aim is to give (inefficient, but sufficient) security parameters for the coding QHE scheme. We will then argue that if there were a QECC implementing a sufficiently large transversal gate set (such as the set of all classical reversible gates), then it would violate Nayak's bound with these parameters.

We will first need a small lemma on the structure of the partial trace operator. The proof can be found in Appendix B.

Lemma 12. For Hilbert space decomposition $\mathcal{H} = \mathcal{H}_{\bar{\Delta}_1} \otimes \mathcal{H}_\Delta \otimes \mathcal{H}_{\bar{\Delta}_2}$,

$$\text{Tr}\left((\rho^{\bar{\Delta}_1\Delta} \otimes I^{\bar{\Delta}_2})(I^{\bar{\Delta}_1} \otimes \sigma^{\Delta\bar{\Delta}_2})\right) = \text{Tr}\left(\text{Tr}_{\bar{\Delta}_1}(\rho)\text{Tr}_{\bar{\Delta}_2}(\sigma)\right).$$

With this we are ready to prove the security tradeoff between ϵ, p, and m. We adopt the same notation used in the proposed scheme for convenience, and note that we are demanding the stronger condition that outputs are indistinguishable from random noise.

Proposition 13. *For the scheme described in Figure 1, letting $K = 2^p$ be the dimension of any subsystem and for some $c \in (0,1)$, we have*

$$\|\left(I/K^{mn}\right) - E_S[\gamma_S]\|_1 \leq \epsilon(K,m)$$

for $\epsilon(K,m) = \left(\left(\frac{m-1}{m}\right)^n - 1 + K^{-c}\left(\frac{2K}{m}\right)^n\right)^{1/2}$.

Proof. By Cauchy-Schwartz,

$$\begin{aligned}\| (I/K^{mn}) - E_S[\gamma_S]\|_1^2 &\leq K^{mn} \| (I/K^{mn}) - E_S[\gamma_S]\|_2^2 \\ &\leq K^{mn}\mathrm{Tr}(E_S[\gamma_S]^2) - \left(\frac{2}{K^{mn}}\right)\mathrm{Tr}(E_S[\gamma_S]) + \left(\frac{1}{K^{2(mn)}}\right)\mathrm{Tr}(I) \\ &\leq K^{mn}\mathrm{Tr}(E_S[\gamma_S]^2) - 1.\end{aligned}$$

where the third line follows by noting that, as a quantum state, $\mathrm{Tr}(E_S[\gamma_S]) = 1$. We write $|S \cap S'|$ to denote the size of the intersection of S and S' considered as sets. We can then decompose, for $p_\ell = \Pr_{S,S'}[|S \cap S'| = \ell]$,

$$\begin{aligned}K^{mn}\mathrm{Tr}(E_S[\gamma_S]^2) &= \left(\frac{K^{mn}}{m^{2n}}\right)\sum_{S,S'}\mathrm{Tr}(\gamma_S\gamma'_S) \\ (*) \ &= K^{mn}\sum_{\ell=0}^{n} p_\ell \mathrm{Tr}\left(E[(\gamma_S\gamma_{S'}) \,\big|\, |S\cap S'| = \ell]\right).\end{aligned}$$

Note that $p_\ell = \frac{\binom{n}{\ell}(m-1)^{(n-\ell)}}{m^n} \leq \binom{n}{l}/m^\ell$ and that $p_0 = (\frac{m-1}{m})^n$. Furthermore, up to a permutation on the coordinates, we may write for $\dim(I) = K^{mn-2n}$,

$$\begin{aligned}K^{mn}E[(\gamma_S\gamma_{S'}) \,\big|\, |S\cap S'| = 0] &= K^{mn}\mathrm{Tr}\left((\gamma/K^n)\otimes(\gamma/K^n)\otimes\left(I/K^{(mn-2n)}\right)^2\right) \\ &= 1\end{aligned}$$

again by noting that γ is a quantum state of trace one and by multiplicativity of trace over tensor products. Next consider the general case $|S\cap S'| = \ell$. Then up to a permutation on the coordinates and for some $\pi \in S_n$, for Δ the subsystem of the intersection $S \cap S'$,

$$\begin{aligned}K^{mn}\mathrm{Tr}(\gamma_S\gamma_{S'}) &= K^{mn}\mathrm{Tr}\left((I/K^{n-\ell}\otimes\gamma)(\pi\gamma\pi^\dagger\otimes I/K^{n-\ell})\otimes\left(I/K^{(mn-2n+\ell)}\right)^2\right) \\ &= K^\ell\mathrm{Tr}\left((I\otimes\gamma)(\pi\gamma\pi^\dagger\otimes I)\right) \\ &= K^\ell\mathrm{Tr}\left(\mathrm{Tr}_{\bar{\Delta}}(\gamma)\mathrm{Tr}_{\bar{\Delta}}(\pi\gamma\pi^\dagger)\right)\end{aligned}$$

where the final line follows from Lemma 12. Then, because we have withheld a subsystem from each subcode of the underlying QECC, in any row i we have that $\mathrm{Tr}_{\bar{\Delta}}(\gamma^i)$ is mixed. It follows that $\mathrm{Tr}\left(\mathrm{Tr}_{\bar{\Delta}}(\gamma^i)\mathrm{Tr}_{\bar{\Delta}}(\pi\gamma^i\pi^\dagger)\right) < 1$. So by separability across each encoded qubit and again by multiplicativity of trace across tensor products,

$$\mathrm{Tr}\left(\bigotimes_{j=1}^{p}\mathrm{Tr}_{\bar{\Delta}}(\gamma^{i_j})\mathrm{Tr}_{\bar{\Delta}}(\pi\gamma^{i_j}\pi^\dagger)\right) = \prod_{j=1}^{p}\mathrm{Tr}\left(\mathrm{Tr}_{\bar{\Delta}}(\gamma^{i_j})\mathrm{Tr}_{\bar{\Delta}}(\pi\gamma^{i_j}\pi^\dagger)\right).$$

It follows that there exists some $c \in (0,1)$ so that

$$K^{mn}\mathrm{Tr}(\gamma_S\gamma_{S'}) \leq K^{\ell-c}.$$

Putting this all together, we observe that

$$\begin{aligned} K^{mn}\sum_{\ell=1}^{n} p_\ell \mathrm{Tr}\left(E[(\gamma_S\gamma_{S'}) \,\middle|\, |S\cap S'|=\ell]\right) &\leq K^{-c}\sum_{\ell=1}^{n}\binom{n}{\ell}\left(\frac{K}{m}\right)^{\ell} \\ &\leq K^{-c}\left(\left(1+\frac{K}{m}\right)^{n}-1\right) \\ &\leq K^{-c}\left(\frac{2K}{m}\right)^{n} \end{aligned}$$

Including the first term in the sum, we get,

$$(*) \leq \left(\frac{m-1}{m}\right)^{n} + K^{-c}\left(\frac{2K}{m}\right)^{n}$$

and so,

$$\epsilon(K,m) = \left(\left(\frac{m-1}{m}\right)^{n} - 1 + K^{-c}\left(\frac{2K}{m}\right)^{n}\right)^{1/2}$$

as desired.

□

5 Limitations on classical transversal computation

We are left with two competing bounds. On the one hand, it follows from Nayak's bound (Appendix A) that, for any $\mathcal{F}$-ITS-QHE encryption scheme with security ϵ and communication size s,

$$s \geq \log(|\mathcal{F}|)(1-H(\epsilon)).$$

If we choose parameters that do not leak some constant fraction of information about our input, then as $\epsilon \to 0$ we see that for s chosen as some fixed function of the input size, it must be that $s = \Omega(\log(|\mathcal{F}|))$. Using the notation and parameters from the aforementioned coding scheme, this means that $mnp = \Omega(\log(|\mathcal{F}_p|))$ for $\mathcal{F}_p$ the restriction of functions in $\mathcal{F}$ to p-bit inputs. Note that we can assume no ancilla overhead since the constant gets absorbed into this asymptotic bound.

Now by construction of the scheme, $\mathcal{F}$ is the transversal gate set for the underlying choice of quantum error-correcting code. Next, we would like to choose m as a function of K so that $\epsilon \to 0$. For this, it suffices to choose m as a function of K so that

$$\lim_{K\to\infty} K^{-c}\left(\frac{2K}{m}\right)^{n} = 0.$$

Equivalently, we require $m = \omega(K^{1-\left(\frac{c}{n}\right)})$. Then for some $c' < 1$, we can select $m = K^{c'}$ and still have $\epsilon \to 0$. Plugging this back into Nayak's bound, we see that asymptotically

$$K^{c'}\log(K) = \Omega(\log(|\mathcal{F}_p|))$$

for $|\mathcal{F}_p|$ the size of the function class, seen itself as a function returning the number of unique members in the class on p-bit inputs. In particular, $\mathcal{F}_p$ cannot be the set of all Boolean functions, for then $\log(|\mathcal{F}_p|) = K$. This shows that no code satisfying the hypotheses of our scheme can implement Toffoli transversally.

We now justify our earlier assumptions on the structure of candidate r-fold codes. Suppose an r-fold $[[n, 1, d]]$ QECC could implement a logical Toffoli gate transversally. First note that the tensor decomposition between the logical states must align, or else the restriction of logical Toffoli to one element of the product would unitarily map a pure state to a mixed state. Furthermore, we can think of the QECC criterion in Definition 5 as a diagonal and off-diagonal condition: for all $|E| < d$,

$$\langle 0_L|E|0_L\rangle = \langle 1_L|E|1_L\rangle,$$
$$\langle 0_L|E|1_L\rangle = 0.$$

Since the Paulis form an operator basis, we can always assume that E is an element of the Pauli group. Then, for r-fold codes with logical basis states $|i\rangle_L = \bigotimes_{j=1}^{r}|\psi_{ij}\rangle$, this becomes

$$\prod_{k=1}^{r}\langle\psi_{ik}|E_k|\psi_{jk}\rangle = c_E\delta_{ij}$$

where $E = E_1 \otimes \ldots \otimes E_r$. Note then that if $|\psi_{0j}\rangle \not\perp |\psi_{1j}\rangle$, we can trace out the corresponding subsystem and obtain a code with the same correctable error set on the complement of that system. Furthermore, if $r > d$, then we can again trace out any $r - d$ subcode subsystems to obtain a code with the same correctable error set on the complement. Both of these observations follow from noticing that these subcodes must themselves satisfy the diagonal condition,

$$\langle\psi_{0j}|E|\psi_{0j}\rangle = \langle\psi_{1j}|E|\psi_{1j}\rangle.$$

It follows from the security proof that if $r < d$, then the code would satisfy the hypotheses of our scheme and violate the lower bound in Proposition 20. Thus, $r = d$. Furthermore, logical transversal Toffoli on the entire code must restrict (up to global phase) to a logical transversal Toffoli gate on the subcodes, each of which is 1-fold by definition. Thus, each subcode must itself have distance 1. To summarize,

Theorem 14. *If a QECC is not a maximally redundant code, then it does not admit a classical-reversible universal transversal gate set. In particular, no such code can implement the Toffoli gate transversally.*

Note also that for the scheme in Figure 1, for any $m = \omega(K^{1-(\frac{c}{n})})$, $\epsilon(K)$ is negligible in p. Summarizing the parameters of the coding scheme:

Proposition 15. *For any r-fold $[[n, 1, d]]$ quantum error-correcting code C with $r < d$ and with transversal gate set $\mathcal{T}_C$, the described protocol is a compact quantum $\mathcal{T}_C$-homomorphic encryption scheme with security $\epsilon = negl(p)$ for p the input size and with encoding size $m = 2^{pc'}$ for some $c' < 1$.*

While this is highly inefficient, we pause to give some intuition for why it suits our purposes. On the one hand, we can envision trivial "hiding" schemes that have encoding length 2^p in each bit. Nayak's bound allows for higher efficiency, roughly demanding that encodings implementing the set of all classical functions on p bits homomorphically must have length at least $(2^p/p)$ in each bit. Finally our scheme, with encoding length $2^{pc'}$ for some $c' \in (0, 1)$, is just efficient enough to defeat this bound and allow us to argue Theorem 14.

Because these maximally redundant codes have a simple design, if we further assume that they are additive, we can use the additional stabilizer structure to argue directly that they cannot implement logical Toffoli transversally. From this observation, we directly obtain the following.

Corollary 16. *No additive QECC can implement transversal Toffoli.*

For proof, see Appendix C. Note that this also follows from the arguments in Appendix D. Our central result follows.

Corollary 17. *If a QECC is not a non-additive maximally redundant code, then it cannot implement the Toffoli gate transversally.*

Finally, note that by concatenating an $[[n, 1, d]]$ d-fold code with itself, the code remains d-fold while the distance must increase to at least d^2. Furthermore, if such a code implements Toffoli strongly transversally, then so does its concatenation with itself. As a result, we observe the following.

Corollary 18. *No QECC can implement strongly transversal Toffoli.*

6 Discussion

Do there exist non-additive maximally redundant codes that can then implement Toffoli transversally? One can essentially think of these as QECCs formed by concatenating an outer repetition code with a distance 1 inner code that is not a stabilizer subspace. Intuitively, since the inner code is not quantum error-correcting, the code only "spreads out the information in one basis". More precisely, the inner code only satisfies the diagonal QECC criterion. While this is a less restrictive condition, it still must be "complementary" to the outer code, and this allows us to argue impossibility in the additive case. Unfortunately by comparison, the structure of general non-additive codes is less well-understood – in particular, we know of no examples of such a code. We expect that *no QECC can implement Toffoli transversally*, and view this exception as a consequence of the lack of structure on general non-additive codes. We hope to resolve this exception in upcoming work.

The QHE scheme we have detailed is non-leveled and compact, but highly inefficient. An immediate question would be to refine the security proof, which uses too strong a security demand. It would be most interesting to see if a modified approach can achieve *efficient* ITS-QHE for transversal gate sets of general quantum error-correcting codes, where the size of the encoding is some fixed polynomial of the input length. There are certain quantitative properties of "nonlocality" in QECCs (see e.g. [3], [34]) that might be helpful in such an endeavor. Following the outline of [29], we could also expect to extend a scheme built on a code with desirable transversal gates to accommodate a constant number of non-transversal gates. Just as one might tailor a QECC for a specific algorithm that makes heavy use of its transversal gate set, one might also tailor an ITS-QHE scheme to homomorphically implement that algorithm. Furthermore, it would be of theoretical interest to find a protocol matching the lower bound implicit in Proposition 20.

Another interesting open question is to consider leveled ITS-QHE schemes: allow the client some preprocessing to scale with the size of the circuit. Can this relaxation allow more efficient or universal schemes for polynomial sized circuits, mirroring the computational security case? A first step might be to try to apply the techniques of instantaneous nonlocal computation [35] that proved invaluable in the computationally secure scheme. Moreover, through gauge-fixing, we have ways of converting between codes that together form a universal transversal gate set. Its not clear how to implement such a strategy, since the noisy embedding and non-interactivity present barriers to measuring syndromes, but these elements taken together might be useful in extending the current scheme.

Finally, one could ask if there is a correspondence between transversal gates for quantum codes and nontrivial ITS homomorphically-implementable gate sets, based on the "richness" of the function classes they can realize. In particular, [15] asked: what is the maximum size of finite group that can be implemented logically and transversally? Indeed, since the Clifford group on p-qubits is of size at most 2^{2p^2+3p} [30], one could reasonably expect to efficiently implement the Clifford gates homomorphically with information theoretic security, as was done in [29]. We hope that our arguments might extend past classical reversible circuit classes to address this question, although it is unclear how to generalize Nayak's bound to apply to these general finite subgroups of the unitary group.

7 Acknowledgments

The authors would like to thank Cupjin Huang and Fang Zhang for useful discussions, in particular concerning the security proof. They are grateful to Audra McMillan, Anne Broadbent, and Zhengfeng Ji for their comments on an earlier draft of this paper. In particular, we thank Anne for bringing [4] to our attention and pointing out that Proposition 3 follows from the result of that paper and an argument similar to the classical case. They also thank Fernando Pastawksi for pointing us to Lemma 23. This research was supported in part by NSF Awards 1216729, 1318070, and 1526928.

References

[1] A. Ambainis, D. Leung, L. Mancinska, and M. Ozols. Quantum random access codes with shared randomness, October 2008. arXiv:0810.2937.

[2] J. T. Anderson and T. Jochym-O'Connor. Classification of transversal gates in qubit stabilizer codes, 2016. Quantum Information and Computation 16, 0771-0802.

[3] L. Arnaud and N. J. Cerf. Exploring pure quantum states with maximally mixed reductions, January 2013. Phys. Rev. A 87, 012319 (2013).

[4] A. Baumeler and A. Broadbent. Quantum private information retrieval has linear communication complexity, April 2013. Journal of Cryptology. Volume 28, Issue 1, pp 161-175 (2015).

[5] H. Bombin. Gauge color codes: Optimal transversal gates and gauge fixing in topological stabilizer codes, August 2015. New J. Phys. 17 (2015) 083002.

[6] S. Bravyi and R. Koenig. Classification of topologically protected gates for local stabilizer codes, 2013. Phys. Rev. Lett. 110, 170503 (2013).

[7] M. J. Bremner, R. Jozsa, and D. J. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy, August 2010. Proceedings of the Royal Society A, Volume 467, Issue 2126.

[8] A. Broadbent. Delegating private quantum computations, June 2015. Canadian Journal of Physics, 2015, 93(9): 941-946.

[9] A. Broadbent, J. Fitzsimons, and E. Kashefi. Universal blind quantum computation, 2009. Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2009), pp. 517-526.

[10] A. Broadbent and S. Jeffery. Quantum homomorphic encryption for circuits of low T-gate complexity, 2015. In Proceedings of Advances in Cryptology – CRYPTO 2015, pp 609-629.

[11] A. Broadbent, Z. Ji, F. Song, and J. Watrous. Zero-knowledge proof systems for QMA, April 2016. Proceedings of the 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS 2016) pp.31-40.

[12] R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. Cryptology ePrint Archive, Report 2015/313, 2015.

[13] A. W. Cross. Fault-tolerant quantum computer architectures using hierarchies of quantum error-correcting codes, 2008. PhD Thesis.

[14] Y. Dulek, C. Schaffner, and F. Speelman. Quantum homomorphic encryption for polynomial-sized circuits, August 2016. CRYPTO 2016: Advances in Cryptology - CRYPTO 2016, pp 3-32.

[15] B. Eastin and E. Knill. Restrictions on transversal encoded quantum gate sets, July 2009. Phys. Rev. Lett. 102, 110502.

[16] M. Fillinger. Lattice based cryptography and fully homomorphic encryption, 2012. `http://homepages.cwi.nl/~schaffne/courses/reports/MaxFillinger_FHE_2012.pdf`.

[17] K. Fisher, A. Broadbent, L. Shalm, Z. Yan, J. Lavoie, R. Prevedel, T. Jennewein, and K. Resch. Quantum computing on encrypted data, January 2014. Nature Communications 5, Article number: 3074.

[18] A. G. Fowler, S. J. Devitt, and C. Jones. Surface code implementation of block code state distillation, January 2013. Scientific Reports 3, 1939.

[19] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland. Surface codes: Towards practical large-scale quantum computation, August 2012. Phys. Rev. A 86, 032324 (2012).

[20] C. Gentry. A fully homomorphic encryption scheme, 2009. Ph.D. Thesis, Stanford University.

[21] D. Gottesman. Stabilizer codes and quantum error correction, 1997. Caltech Ph.D. Thesis.

[22] C. Jones. Composite toffoli gate with two-round error detection, March 2013. Phys. Rev. A 87, 052334.

[23] C. Jones. Novel constructions for the fault-tolerant toffoli gate, 2013. Phys. Rev. A 87, 022328.

[24] E. Knill and R. Laflamme. A theory of quantum error-correcting codes, April 1996. Phys.Rev.Lett.84:2525-2528.

[25] C.-Y. Lai and K.-M. Chung. On statistically-secure quantum homomorphic encryption. In preparation.

[26] H. P. Nautrup, N. Friis, and H. J. Briegel. Topological code switching in two dimensions, September 2016. arXiv:1609.08062.

[27] A. Nayak. Optimal lower bounds for quantum automata and random access codes, April 1999. FOCS 1999.

[28] M. A. Nielsen and I. L. Chuang. Quantum computation and quantum information, 2011. Cambridge University Press New York, NY.

[29] Y. Ouyang, S.-H. Tan, and J. Fitzsimons. Quantum homomorphic encryption from quantum codes, August 2015. arXiv:1508.00938.

[30] M. Ozols. Notes on the clifford group, July 2008. `http://home.lu.lv/~sd20008/papers/essays/Clifford%20group%20[paper].pdf`.

[31] A. Paetznick and B. W. Reichardt. Universal fault-tolerant quantum computation with only transversal gates and error correction, April 2013. Phys. Rev. Lett. 111, 090505 (2013).

[32] F. Pastawksi and B. Yoshida. Fault-tolerant logical gates in quantum error-correcting codes, 2015. Phys. Rev. A 91, 012305.

[33] C. Peikert. A decade of lattice cryptography, March 2016. Foundations and Trends in Theoretical Computer Science 10(4):283-424.

[34] A. J. Scott. Multipartite entanglement, quantum-error-correcting codes, and entangling power of quantum evolutions, May 2004. Phys. Rev. A 69, 052330.

[35] F. Speelman. Instantaneous non-local computation of low T-depth quantum circuits, November 2015. arXiv:1511.02839.

[36] S.-H. Tan, J. A. Kettlewell, Y. Ouyang, L. Chen, and J. F. Fitzsimons. A quantum approach to homomorphic encryption, 2016. Sci. Rep. 6, 33467.

[37] T. J. Yoder, R. Takagi, and I. L. Chuang. Universal fault-tolerant gates on concatenated stabilizer codes, March 2016. Phys. Rev. X 6, 031039 (2016).

[38] L. Yu, C. A. Perez-Delgado, and J. F. Fitzsimons. Limitations on information theoretically secure quantum homomorphic encryption, June 2014. Phys. Rev. A 90, 050303 (2014).

[39] B. Zeng, A. Cross, and I. L. Chuang. Transversality versus universality for additive quantum codes, September 2011. IEEE Transactions on Information Theory, Volume: 57, Issue: 9, 6272 - 6284.

A A no-go result for ITS-QFHE

We now set out to show that efficient ITS-QFHE is impossible. Define for any $x \in \{0,1\}^n$ the state

$$|s_x\rangle^{\mathcal{CM}} = QHE.Enc(x).$$

Note that for any $|s_x\rangle\langle s_x|^{\mathcal{M}}$, we can always purify to a system of size at most $2|\mathcal{M}|$. So without loss of generality, we may assume that $|\mathcal{C}| = m$, the size of the message sent from Client to Server. Next, by information theoretic security, the state of the encryption on subsystem $\mathcal{M}$ must be almost independent of x. Formally,

$$\|\mathrm{Tr}_{\mathcal{C}}(|s_x\rangle\langle s_x|) - \mathrm{Tr}_{\mathcal{C}}(|s_{0^n}\rangle\langle s_{0^n}|)\|_1 \leq \epsilon$$

for ϵ the security of the scheme. Equivalently, there exists a $V_x^{\mathcal{C}}$ so that, defining $s'_x = (V_x^{\mathcal{C}} \otimes I^{\mathcal{M}})|s_{0^n}\rangle\langle s_{0^n}|(V_x^{\mathcal{C}} \otimes I^{\mathcal{M}})^\dagger$ and $s_x = |s_x\rangle\langle s_x|$,

$$\|s'_x - s_x\|_1 \leq \epsilon.$$

Furthermore, for any $f \in \mathcal{F}$, we have by the homomorphic property that, abbreviating $QHE.Eval_f$ as $f_{ev}(\cdot)$ and $QHE.Dec(\cdot)$ as $D(\cdot)$,

$$\left[I^{\mathcal{C}} \otimes f_{ev}^{\mathcal{M}}(s_x^{\mathcal{CM}})\right]^{\mathcal{CM}'} =: \eta_{f,x},$$

$$D^{\mathcal{CM}'}(\eta_{f,x}) = f(x).$$

But now, defining $\eta'_{f,x}$ by replacing s_x with s'_x in the definition of $\eta_{f,x}$, by contractivity of trace distance we also have

$$\Pr[D^{\mathcal{CM}'}(\eta'_{f,x}) \neq f(x)] \leq \epsilon.$$

To elucidate the underlying QRAC, define the mapping

$$f \mapsto \eta_{f,0^n},$$

and note that $(V_x^{\mathcal{C}} \otimes I^{\mathcal{M}'})\eta_{f,0^n}(V_x^{\mathcal{C}} \otimes I^{\mathcal{M}'})^\dagger = \eta'_{f,x}$. So let $D^{\mathcal{CM}'}\left[(V_x^{\mathcal{C}} \otimes I^{\mathcal{M}'})(\cdot)\right]$ denote the query for index x of f, thinking of f as a 2^n length bit string, with the xth bit defined as $f(x)$. Then we have a $(2^n, m+m', 1-\epsilon)$-QRAC for the set of all Boolean functions, where $m+m'$ is the communication cost of the protocol. We now recall a well-known bound on the efficiency of QRACs [27].

Theorem 19 (Nayak's Bound)**.** *If there exists an (n, m, p)-QRAC, then for $H(\cdot)$ the binary entropy function,*

$$m \geq n(1 - H(p)).$$

So it must be that the total communication cost of the protocol $(m+m') \geq 2^n(1-H(\epsilon))$. For security, allowing $\epsilon \to 0$ and noting that $H(\epsilon) \to 0$, we see that the communication cost $(m+m') = \Omega(2^n)$. Thus, either the size of the encoding or the evaluated ciphertext must be exponentially long in the input, precluding efficiency. In short,

Proposition 20. *The communication cost of ITS-QFHE must be exponential in the size of the input.*

B Proof of Lemma 12

Proof. Expanding in terms of outer products,

$$\begin{aligned}\mathrm{Tr}\left((\rho\otimes I)(I\otimes\sigma)\right) &= \mathrm{Tr}\Bigg(\Big(\sum_{i,i'}\sum_{j,j'}\sum_{k} a_{i,i',j,j'}|i\rangle\langle i|^{\bar{\Delta}_1}\otimes|j\rangle\langle j'|^{\Delta}\otimes|k\rangle\langle k|^{\bar{\Delta}_2}\Big)\cdot \\ &\qquad \Big(\sum_{\ell}\sum_{m,m'}\sum_{n,n'} b_{m,m',n,n'}|\ell\rangle\langle\ell|^{\bar{\Delta}_1}\otimes|m\rangle\langle m'|^{\Delta}\otimes|n\rangle\langle n'|^{\bar{\Delta}_2}\Big)\Bigg) \\ &= \mathrm{Tr}\left(\sum_{i,i'}\sum_{n,n'}\sum_{j,m'}\left(\sum_{j'} a_{i,i',j,j'}b_{j',m',n,n'}\right)|i\rangle\langle i|\otimes|j\rangle\langle m'|\otimes|n\rangle\langle n'|\right) \\ &= \sum_{i}\sum_{n}\sum_{j,j'}\left(a_{i,i,j,j'}b_{j',j,n,n}\right).\end{aligned}$$

On the other hand, we have

$$\begin{aligned}\mathrm{Tr}\left(\mathrm{Tr}_{\bar{\Delta}_1}(\rho)\mathrm{Tr}_{\bar{\Delta}_2}(\sigma)\right) &= \mathrm{Tr}\left(\Big(\sum_{i}\sum_{j,j'} a_{i,i,j,j'}|j\rangle\langle j'|\Big)\Big(\sum_{n}\sum_{m,m'} b_{m,m',n,n}|m\rangle\langle m'|\Big)\right) \\ &= \mathrm{Tr}\left(\sum_{i}\sum_{n}\sum_{j,j'}\Big(\sum_{j'} a_{i,i,j,j'}b_{j',m',n,n}\Big)|j\rangle\langle m'|\right) \\ &= \sum_{i}\sum_{n}\sum_{j,j'}\left(a_{i,i,j,j'}b_{j',j,n,n}\right)\end{aligned}$$

as claimed. □

C Proof of Corollary 16

Proof. By Theorem 14, it suffices to consider maximally redundant codes. So suppose, for the sake of contradiction, that an $[[n,1,d]]$ additive d-fold code could implement Toffoli transversally. Let $[\cdot,\cdot]$ denote the group commutator. We denote by $\bar{\cdot}$ states and operations acting on the subcodes, and $\tilde{\cdot}$ those on the full code. We will assume that each subcode is the same, e.g. $|\tilde{i}\rangle = |\bar{i}\rangle^{\otimes d}$, so that we can speak directly about the inner and outer codes. The general argument follows similarly.

Since the code is additive, the code distance is the minimal weight logical Pauli operator acting on the code. For any $\bar{Z}_L$, by multiplicativity of the inner product over tensor products,

$$\begin{aligned}\frac{1}{2}\left\langle|\tilde{0}\rangle+|\tilde{1}\rangle\,\middle|\bar{Z}_L\middle|\,|\tilde{0}\rangle-|\tilde{1}\rangle\right\rangle &= \frac{1}{2}\left(\langle\tilde{0}|\bar{Z}_L|\tilde{0}\rangle-\langle\tilde{0}|\bar{Z}_L|\tilde{1}\rangle+\langle\tilde{0}|\bar{Z}_L|\tilde{0}\rangle-\langle\tilde{1}|\bar{Z}_L|\tilde{1}\rangle\right) \\ &= \frac{1}{2}\left(\langle\bar{0}|\bar{0}\rangle^{\frac{n}{d}}+\langle\bar{1}|\bar{1}\rangle^{\frac{n}{d}}\right)\neq 0.\end{aligned}$$

Since the outer code has distance d, it follows from the QECC criterion that $\bar{Z}_L$ must have weight at least d. Then $\bar{X}_L$ must have weight 1, since the underlying inner code has distance 1 by assumption. Because the outer classical repetition code factors as a tensor product, transversal $\widetilde{\mathrm{Toff}}_L$ on the outer code must restrict (up to a global phase) to transversal $\overline{\mathrm{Toff}}_L$ on the inner code. Since we're now working with multiqubit gates, let $G_L(i)$ denote the logical gate for G acting on the ith code block. We can compute directly,

$$[\overline{\text{Toff}}_L(1,2,3), \bar{X}_L(1)] = \overline{CX}_L(2,3).$$

Furthermore, because $\overline{\text{Toff}}_L$ and $\bar{X}_L$ are transversal, it follows that $\overline{CX}_L$ has a representative that is also transversal and is supported on the subsystems that support $\bar{X}_L$. By a similar argument

$$[\overline{CX}_L(1,2), \bar{Z}_L(1)] = \bar{Z}_L(2)$$

so that $\bar{Z}_L$ must also be contained in the subsystems supporting $\overline{CX}_L$, and in turn $\bar{X}_L$. As we have already observed, the minimal weight of any representative of $\bar{Z}_L$ must be at least d, a contradiction as $\bar{X}_L$ has a representative of weight 1.

□

D An alternate proof for stabilizer codes

Here we offer an alternate proof limiting universal transversal reversible computation for the subclass of stabilizer codes. The arguments here are based off of results from [6] and [32], which we reproduce for completeness.

Definition 21. The *Clifford hierarchy* $\mathcal{C}$ is a sequence of gate sets $\{\mathcal{C}_k\}_{k\geq 1}$ defined recursively by $\mathcal{C}_k = \{U : U\mathcal{C}_1U^\dagger \subseteq \mathcal{C}_{k-1}\}$, where we define $\mathcal{C}_1$ to be the Pauli group.

Note that $\mathcal{C}_2$ is the Clifford group, and $\mathcal{C}_k$ fails to be a group for $k > 2$. Further note that reversible circuits saturate the Clifford hierarchy (and in fact can lie outside it entirely) by the gate C^kX, the k-controlled bit-flip gate, which lies in $\mathcal{C}_{k+1}$. Toffoli is simply C^2X, and so lies in the third level of the Clifford hierarchy. We next recall the stabilizer cleaning lemma, which can be found in [6].

Lemma 22. Let S be a stabilizer code, and let R be any subset of physical qubits of the code such that any logical operator supported on R acts trivially on S. Then, for any logical operator U_L, there exists a representative of U_L supported on R^c.

We call such subsets R cleanable. Equipped with the cleaning lemma, we can now summarize the following lemma from [32].

Lemma 23. Let S be a stabilizer code and let $\{R_0, \ldots, R_k\}$ be a set of cleanable subsets of the physical qubits comprising S. Let U be a logical operator supported on $\cup_{i=0}^k R_i$ such that U is transversal with respect to the R_i. Then, $U_L \in \mathcal{C}_k$.

Proof. We proceed by induction on k. In the base case, we have a logical operator U supported on cleanable subsets $R_0 \cup R_1$. Let P be any logical Pauli operator cleaned off of R_1, and let $[\cdot,\cdot]$ denote the group commutator. Since in a stabilizer code the logical Pauli operators are transversal, we have $Supp([U,P]) \subseteq R_0$, which by cleanability implies that $[U_L, P_L] = cI_L$. Since this is true for any P_L, it must be that $U_L \in \mathcal{C}_1$.

Similarly, suppose U is supported on $\cup_{i=0}^k R_i$. Then, cleaning any logical Pauli P off of R_k, we see that $Supp([U,P]) \subseteq \cup_{i=0}^{k-1} R_i$. By our inductive hypothesis, $[U_L, P_L] \subseteq \mathcal{C}_{k-1}$, which implies $U_L P_L U_L^\dagger \in \mathcal{C}_{k-1}$ for any logical Pauli P_L. Thus $U_L \in \mathcal{C}_k$, completing the proof. □

This argument generalizes to subsystem codes, and we refer the reader [32] for a more complete description. As a consequence we obtain the following.

Corollary 24. *No erasure-correcting stabilizer code can implement a classical reversible universal transversal gate set.*

Proof. Partition the code block into single subsystem subsets $\{R_1, \ldots, R_n\}$ where n is the length of the code. Then, since the code is erasure-correcting, any logical operator supported on a single subsystem must act trivially on the codespace, and so these subsets are cleanable. By the lemma, any transversal logical gate must lie in $\mathcal{C}_n$. Since reversible circuits saturate $\mathcal{C}$, they cannot be logically transversally implementable. □