

Kent Academic Repository

Full text document (pdf)

Citation for published version

Yu, Li and Pérez-Delgado, Carlos A and Fitzsimons, Joseph F (2014) Limitations on information-theoretic quantum homomorphic encryption. *Physical Review A*, 90 (5). 050303.

DOI

<https://doi.org/10.1103/PhysRevA.90.050303>

Link to record in KAR

<http://kar.kent.ac.uk/58149/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Limitations on information theoretically secure quantum homomorphic encryption

Li Yu,^{1,2} Carlos A. Pérez-Delgado,¹ and Joseph F. Fitzsimons^{1,2,*}

¹*Singapore University of Technology and Design, 20 Dover Drive, Singapore 138682*

²*Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543*

Homomorphic encryption is a form of encryption which allows computation to be carried out on the encrypted data without the need for decryption. The success of quantum approaches to related tasks in a delegated computation setting has raised the question of whether quantum mechanics may be used to achieve information theoretically secure fully homomorphic encryption. Here we show, via an information localisation argument, that deterministic fully homomorphic encryption necessarily incurs exponential overhead if perfect security is required.

The insight that information must be represented and manipulated in accordance with physical laws has led to the blossoming field of quantum information science. The applications of this approach to information processing are diverse, and it has led to discoveries ranging from new algorithms [1, 2] and communications protocols [3, 4] which exploit quantum states for increased efficiency to techniques for enhancing the precision of metrology [5]. Historically, cryptography was one of the first fields for which quantum information processing was shown to offer an advantage over classical processing, when in 1984 Bennett and Brassard introduced a quantum protocol for information theoretically secure key distribution [6]. While for many quantum cryptography has remained synonymous with quantum key distribution, the field has grown substantially, with quantum protocols being discovered which enhance the security with which many cryptographic tasks can be accomplished, including digital signatures [7], anonymous communication [8], private database queries [9], and random number generation [10]. The importance of information theoretically secure cryptography is highlighted by the fact that quantum algorithms offer new attacks against cryptosystems which rely on assumptions of computational intractability for their security [11–13]. Unfortunately, not all cryptographic tasks that we may wish to accomplish admit an information theoretically secure quantum protocol, and indeed a number of no-go theorems have been discovered which show that quantum mechanics alone is insufficient to accomplish certain tasks, such as bit commitment [14] and oblivious transfer [15], with perfect security.

One of the most celebrated results in classical cryptography in recent years has been the discovery of computationally secure protocols for fully homomorphic computation [16–19]. A homomorphic encryption scheme is one which allows data to be encrypted in such a way that certain operations can be performed on the data without decryption. This allows a user to provide encrypted data to a remote server for processing without having to reveal the plaintext. A number of examples of such homomorphic encryption schemes have been known for many years [20, 21], but it was the ground-breaking work of Gentry [16] which for the first time demonstrated a fully homomorphic encryption scheme, one which allowed for arbitrary computations to be performed on the encrypted data, rather than being restricted to some class of non-universal op-

erations. The ability to perform universal computation on encrypted data has greatly increased the utility of homomorphic encryption, and as a result it has become one of the most active areas of modern cryptography.

One drawback of recently discovered classical fully homomorphic encryption schemes is that they derive their security from computational assumptions. At first glance, it is tempting to think that the requirement that encrypted data be manipulable by a third party necessarily precludes information theoretic security. However, the classical one-time pad, in which plaintext bits are XORed with a completely random key, provides an immediate counter-example. Any sequence of bit-flips necessarily commutes with the decryption step, and hence represents a non-trivial set of computations on the plaintext which can be performed directly on the ciphertext. Although this may seem a rather trivial example, the cryptographic community has expended significant effort on the search for information theoretically secure homomorphic encryption systems which support arbitrary algebraic operations (see [22] for a review of recent work in the area). Attempts have also been made to construct fully homomorphic classical encryption schemes with information theoretic security. Current schemes only achieve approximate perfect security, however, and result in encodings which grow exponentially as they approach perfect security [23].

The existence of perfectly secure quantum protocols for blind computation [24–27], and recent experimental demonstrations thereof [28, 29], highlight the possibilities opened by quantum cryptographic techniques in this area. As cryptographic tasks, blind computation and homomorphic encryption are similar in many ways. Both tasks envision a two party scenario, where the first party, Alice, wishes the second party, Bob, to carry out a computation for her, without revealing the input of her computation. In blind computation, however, Alice specifies not only the input data but also the computation to be performed, and the task is to utilise Bob’s resources to perform this computation without revealing either the input or the program. As a result, the current protocols for accomplishing this task are interactive, requiring multiple rounds of communication between Alice and Bob, a significant difference from the setting of homomorphic encryption.

The idea of quantum homomorphic encryption appears in [30], which shows that a perfect, universal, quantum homo-

morphic scheme cannot be constructed using one-time pads and which presents an interactive protocol for achieving similar functionality. Other cryptographic schemes have been proposed that achieve some of the functionality of homomorphic encryption using quantum data [31, 32]. However, these rely on assisted computation, and so require multiple rounds of interaction between Alice and Bob, thus amounting to interactive protocols rather than simply encryption schemes. A quantum homomorphic encryption scheme does exist for a restricted model of quantum computation known as boson scattering, which offers limited information theoretic security [33]. The existence of such schemes raises the question as to whether quantum techniques can be exploited to construct an information theoretically secure fully homomorphic encryption scheme. Here we answer that question in the negative by proving that quantum mechanics does not allow for efficient information theoretically secure fully homomorphic encryption that perfectly conceals the plaintext. To achieve this we first formalise the notion of quantum homomorphic encryption, and then proceed to show via an information localisation argument that any such scheme which perfectly hides Alice's input must necessarily reveal the computation performed, and hence the encoding must be sufficiently long to specify any such computation. For a fully homomorphic encryption scheme this implies that the coding must be exponentially long, and thus rules out the existence of efficient fully homomorphic encryption schemes which perfectly hide Alice's data.

Formally, a classical homomorphic encryption scheme consists of four procedures. The first is a key generation algorithm that generates a classical encryption key, a classical decryption key, and potentially some additional auxiliary key. The second is an encryption algorithm, that encrypts the input using the encryption key. Third is a decryption algorithm that decrypts the output using the decryption key. Finally, there is an evaluation algorithm that performs the computation on the ciphertext without decryption, which may use the auxiliary key. For any permissible logical circuit C , the result of the evaluation algorithm should be such that after decrypting the output, one obtains the result of applying C to the unencrypted input. A fully homomorphic encryption scheme, then, is one in which C can be freely chosen from the set of all classical circuits. Here we shall consider only schemes with perfect completeness, where the evaluation operator must deterministically implement the chosen circuit. We will say that a homomorphic encryption scheme has *perfect information theoretic security* if the ciphertext is a density operator independent of the plaintext.

We will define a *quantum homomorphic encryption* (QHE) scheme using similar criteria as for the classical case, extended to take into account the possibility of entanglement within the protocol. A QHE scheme consists of four components: a key generation protocol which produces a quantum state $|\psi_e\rangle$ used as a key for encryption; an encryption unitary operator U_e which encrypts the input state $|\psi_i\rangle$ using the encryption key state, potentially making use of some ancilla




FIG. 1: A schematic diagram for a general quantum homomorphic encryption scheme with input data $|\psi_i\rangle$ and output $|\psi_o\rangle$. The state $|\psi_e\rangle$ represents the initial state of Alice's key, while U_e and U_d are Alice's encryption and decryption operators. Both parties are also allowed an ancilla system, and access to a shared entanglement resource. Alice's decryption key corresponds to the subsystem she retains after applying U_e to her system. Note that no assumption is made about the dimensionality of subsystems. Time points t_1 and t_2 , used in the proof of Theorem 1, are also shown.

system, and which produces a decryption key in a state ρ_d ; a decryption unitary operator U_d which decrypts the encrypted state using the key state; and a set of evaluation unitary operators $\{U_c\}$, such that after decrypting the output the net effect is equivalent to applying the quantum circuit C directly to the initial input state. Here the decryption key is produced when the encryption unitary is applied. Although this is somewhat more general than the procedure for generating the corresponding classical key, we make this generalization to allow for the possibility of a causal relationship between encryption and decryption keys which, via the no-cloning theorem, may prevent them from existing simultaneously. Note that we have not specified an auxiliary key. This is because, without loss of generality, we can assume that this auxiliary key forms part of the encrypted state. An encryption-evaluation-decryption sequence based on this definition is depicted in Figure 1[38].

As we now prove, for such a scheme to operate deterministically, it is necessary that the dimension of the encrypted state grows as the log of the cardinality of the set of possible choices of C , and hence fully homomorphic encryption with perfect information theoretic security is impossible except when the size of the encoding grows exponentially with the size of the plaintext. To prove this, we begin by proving a modified version of an information localisation theorem due to Griffiths [34].

Lemma 1 (Data Localisation). *Let S be some bipartite quantum system with Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$, initially in state $(|\psi\rangle \otimes |\phi\rangle)_A \otimes |\gamma\rangle_B$, where $|\phi\rangle$ and $|\gamma\rangle$ are fixed states. Let ρ be the state of S after the application of a unitary operator U . Then, if the reduced density operator on system B , $\text{Tr}_A \rho$, is*

independent of the input data state $|\psi\rangle$, there exists a unitary operator $V : \mathcal{H}_A \mapsto \mathcal{H}_A$ such that $\text{Tr}_B \rho = V(|\psi\rangle\langle\psi| \otimes \sigma)V^\dagger$ for some density matrix σ independent of $|\psi\rangle$.

Proof. For simplicity of notation we will define $\rho_A = \text{Tr}_B \rho$ and $\rho_B = \text{Tr}_A \rho$, and use r to denote the rank of ρ_B . We shall further divide the Hilbert space $\mathcal{H}_A = \mathcal{H}_{A_1} \otimes \mathcal{H}_{A_2}$ such that $|\psi\rangle$ is the state of \mathcal{H}_{A_1} and $|\phi\rangle$ is the state of \mathcal{H}_{A_2} .

We begin from the requirement that ρ_B is independent of $|\psi\rangle$. This implies that changing the value of $|\psi\rangle$, while holding $|\phi\rangle$ and $|\gamma\rangle$ constant, will not alter $\text{Tr}_A \rho$. We shall consider the effect on the state of ρ of varying only $|\psi\rangle$. Let an orthonormal basis of \mathcal{H}_{A_1} as $|a^j\rangle$, $j = 1, 2, \dots, d_{A_1}$, where d_{A_1} is the dimension of \mathcal{H}_{A_1} , and let $\{|b^k\rangle : k = 1, 2, \dots, r\}$ be an orthonormal set of eigenstates of ρ_B with corresponding eigenvalues p_k .

For each state $|a^j\rangle$, for $1 \leq j \leq d_{A_1}$, we can expand the state of the system after application of U to yield

$$U(|a^j\rangle_{A_1} \otimes |\phi\rangle_{A_2} \otimes |\gamma\rangle_B) = \sum_{k=1}^r \sqrt{p_k} |\tau^{jk}\rangle_A \otimes |b^k\rangle_B. \quad (1)$$

Note that the possible complex phases have been absorbed into the definition of $|\tau^{jk}\rangle$. Since $|b^k\rangle$ are eigenstates of ρ_B with eigenvalues p_k , the expansion on the right hand side of Eq. (1) is a Schmidt expansion for $U(|a^j\rangle_{A_1} \otimes |\phi\rangle_{A_2} \otimes |\gamma\rangle_B)$, and hence $\{|\tau^{jk}\rangle : k = 1, 2, \dots, r\}$ for any fixed j must be orthonormal. Thus, we have $\langle \tau^{jk} | \tau^{j'k'} \rangle = \delta_{k,k'}$.

Now consider the case where we keep the input on \mathcal{H}_{A_2} and \mathcal{H}_B fixed, while changing the input state on \mathcal{H}_{A_1} to one of the form $|\nu^{jj'}\rangle = (|a^j\rangle + |a^{j'}\rangle)/\sqrt{2}$ for $j \neq j'$. In this case,

$$U(|\nu^{jj'}\rangle_{A_1} \otimes |\phi\rangle_{A_2} \otimes |\gamma\rangle_B) = \sum_{k=1}^r \sqrt{p_k} [(|\tau^{jk}\rangle + |\tau^{j'k}\rangle)/\sqrt{2}]_A \otimes |b^k\rangle_B. \quad (2)$$

Since the output reduced density operator on \mathcal{H}_B is still $\rho_B = \sum_{k=1}^r p_k |b^k\rangle\langle b^k|$, the right hand side of Eq. (2) should be a Schmidt expansion, with the Schmidt coefficients still being $\sqrt{p_k}$. Hence $(|\tau^{jk}\rangle + |\tau^{j'k}\rangle)/\sqrt{2}$ must be already normalised and these states must be orthogonal for different values of k . From this we obtain

$$\begin{aligned} \delta_{k,k'} &= \frac{1}{2} (\langle \tau^{jk} | + \langle \tau^{j'k} |) (\langle \tau^{j'k'} | + \langle \tau^{jk'} |) \\ &= \delta_{k,k'} + \frac{1}{2} (\langle \tau^{jk} | \tau^{j'k'} \rangle + \langle \tau^{j'k} | \tau^{jk'} \rangle), \quad j \neq j', \end{aligned} \quad (3)$$

and hence $\langle \tau^{jk} | \tau^{j'k'} \rangle + \langle \tau^{j'k} | \tau^{jk'} \rangle = 0$ as long as $j \neq j'$.

Similarly, by considering input states on \mathcal{H}_{A_1} of the form $|\eta^{jj'}\rangle = (|a^j\rangle + i|a^{j'}\rangle)/\sqrt{2}$, we obtain $\langle \tau^{jk} | \tau^{j'k'} \rangle - \langle \tau^{j'k} | \tau^{jk'} \rangle = 0$ and hence $\langle \tau^{jk} | \tau^{j'k'} \rangle = 0$ for $j \neq j'$. These criteria can be expressed compactly as $\langle \tau^{jk} | \tau^{j'k'} \rangle = \delta_{j,j'} \delta_{k,k'}$. Hence $\{|\tau^{jk}\rangle\}$ forms an orthonormal set, and it is possible to define the subspaces \mathcal{H}_C and \mathcal{H}_D as having orthonormal bases $\{|c^j\rangle\}$ and $\{|d^k\rangle\}$ such that $\mathcal{H}_A = \mathcal{H}_C \otimes \mathcal{H}_D$, and

$$|\tau^{jk}\rangle = |c^j\rangle \otimes |d^k\rangle, \quad j = 1, 2, \dots, d_{A_1}, \quad k = 1, 2, \dots, r. \quad (4)$$

For a generic input state $|\xi\rangle = |\psi\rangle \otimes |\phi\rangle \otimes |\gamma\rangle$, where $|\psi\rangle = \sum_{j=1}^{d_{A_1}} \alpha_j |a^j\rangle$ we then have

$$\begin{aligned} U|\xi\rangle &= \sum_{j=1}^{d_{A_1}} \alpha_j \sum_{k=1}^r \sqrt{p_k} |\tau^{jk}\rangle_A \otimes |b^k\rangle_B \\ &= \left(\sum_{j=1}^{d_{A_1}} \alpha_j |c^j\rangle \right)_C \otimes \left(\sum_{k=1}^r \sqrt{p_k} |d^k\rangle \otimes |b^k\rangle \right)_{DB} \\ &= |\psi\rangle_C \otimes \left(\sum_{k=1}^r \sqrt{p_k} |d^k\rangle \otimes |b^k\rangle \right)_{DB}. \end{aligned} \quad (5)$$

Now, let $V' : \mathcal{H}_{A_1} \mapsto \mathcal{H}_C$ be an isometry such that

$$V'|a^j\rangle = |c^j\rangle, \quad j = 1, 2, \dots, d_{A_1}, \quad (6)$$

and let V be any extension of V' into a full unitary over \mathcal{H}_A . Then $\text{Tr}_B (U|\xi\rangle\langle\xi|U^\dagger) = V(|\psi\rangle\langle\psi| \otimes \sigma)V^\dagger$, for some density operator σ independent of $|\psi\rangle$, as the lemma requires. \square

Lemma 1 shows that in any quantum homomorphic encryption scheme with perfect information theoretic security, the computation has to occur on Alice's "side". The following theorem formalises this intuition, showing that the encrypted state must contain enough information to identify any operator previously applied to it.

Theorem 1. *Let Q be a quantum homomorphic encryption scheme with perfect information theoretic security with encryption operator U_e and decryption operator U_d and a set of evaluation unitaries. Let $\rho_{a,b}$ ($\rho'_{a,b}$) be the state held by Alice (intended to be the encrypted state plus her decryption key) after applying the evaluation unitary U_c ($U_{c'}$) corresponding to a quantum circuit c (c') (i.e., at time t_2 in Figure 1). Then, if b and b' implement distinct unitary operations, ρ_b and ρ'_b must have orthogonal support.*

Proof. For clarity, we will identify different parts of the encryption, circuit evaluation and decryption process with two parties, Alice and Bob, as depicted in Figure 1. We begin by analysing the state of Alice and Bob's joint system after Alice has sent her encoded data to Bob. This is marked as time t_1 in Figure 1. Let $\rho_{a,1}$ ($\rho_{b,1}$) be the states of Alice's (Bob's) subsystem at this point. From this point forward, all communication flows from Bob to Alice. The requirement that Q be perfectly information theoretically secure implies that $I(\rho_{b,1}; |\psi_i\rangle\langle\psi_i|) = 0$. Hence, by Lemma 1 there exists some unitary operator V such that

$$\rho_{a,1} = V(|\psi_i\rangle\langle\psi_i| \otimes \rho'_{a,1})V^\dagger, \quad (7)$$

for some appropriate $\rho'_{a,1}$.

Now, consider the system after Bob has sent his message back to Alice. This is time t_2 in Figure 1. Due to the previous analysis the state of the system at this point can be written as

$$\rho_{a,2} = (V \otimes I) |\psi_i\rangle\langle\psi_i| \otimes \rho'_{a,1} \otimes \rho_b (V^\dagger \otimes I), \quad (8)$$

where ρ_b represents Bob's message. The density matrix ρ_b cannot in general be assumed to be pure, since Bob could have sent a message that remains entangled to his system. Here V acts only on the part of the system that was in Alice's possession prior to receiving the message from Bob, and the identity operator I acts on Bob's message.

The requirement that the evaluation unitary U_c implements a specific circuit c implies that

$$U_d \rho_{a,2} U_d^\dagger = (W_c |\psi_i\rangle\langle\psi_i| W_c^\dagger) \otimes \rho_{\text{anc}}, \quad (9)$$

where W_c is the unitary operator corresponding to quantum circuit c , and ρ_{anc} is simply some state of the ancilla system. Let $U'_d = U_d(V \otimes I)$, then for all c and all $|\psi_i\rangle$,

$$U'_d (|\psi_i\rangle\langle\psi_i| \otimes \rho'_{a,1} \otimes \rho_b) U'^{\dagger}_d = (W_c |\psi_i\rangle\langle\psi_i| W_c^\dagger) \otimes \rho_{\text{anc}}. \quad (10)$$

As the state $\rho'_{a,1}$ and the operator U'_d are independent of c , in the language of [35] this corresponds to a programmable quantum gate array, where ρ_b acts as a *program* to implement the unitary operator W_c . The *no programming* theorem [35] states that for a programmable quantum gate array to implement two distinct unitary operators, the program states must be orthogonal. Hence if ρ_b and $\rho_{b'}$ correspond to the messages returned from Bob after application of evaluation operators corresponding to two non-equivalent circuits, then ρ_b and $\rho_{b'}$ must have orthogonal support. \square

A direct consequence of this theorem is that for any perfectly information theoretically secure homomorphic scheme (fully homomorphic or otherwise), if a known input state is encrypted, and an evaluation operator from some unknown circuit c is applied, it is always possible to unambiguously determine c from the resulting encrypted state. This mirrors a result obtained for one time programs [36], a similar task in which the secret to be protected is Bob's circuit rather than Alice's input. Further, this property severely compromises the efficiency of any QHE encoding, as we now prove.

Corollary 1. *Let Q be a QHE scheme, with perfect information theoretic security, that corresponds to a permissible set of unitary operations S . Then, the following statements hold:*

1. *The size of the system required to store the encrypted state after the application of an evaluation operator U_c corresponding to an arbitrary operation in S is at least $\log_2 |S|$ qubits.*
2. *If S contains the set of reversible classical operations on n bits, then the size of the encrypted state grows at least exponentially in n .*
3. *If S is a set that is ϵ -approximately universal on n qubits, that is every element of $\text{SU}(2^n)$ can be approximated to an accuracy of ϵ , then the size of the encrypted message grows proportional to $(2^{2n} - 1) \log_2(1/\epsilon)$.*

Proof. The proof of the first part of the corollary follows directly from Theorem 1. Each ρ_b corresponding to an operator in S must have orthogonal support on a distinct subspace. Since each such density operator must have at least unit rank, a system must be at least $|S|$ -dimensional in order to represent every possible ρ_b . The second part of the corollary follows from the fact that there are $(2^n)!$ distinct permutations of the n -bit classical states, and hence any S which contains all such operations must have cardinality at least $\log_2(2^n)! \geq 2^n$. The final part of the corollary, the bound on approximate quantum computation, follows from the fact that an ϵ -net that covers $\text{SU}(d)$ requires $\Omega\left((1/\epsilon)^{d^2-1}\right)$ elements [37]. Hence the cardinality of any set of operators which suffices to approximate an arbitrary element of $\text{SU}(2^n)$ to within an accuracy of ϵ must grow at least as $(1/\epsilon)^{2^{2n}-1}$. \square

From this corollary, it follows that no QHE with perfect information theoretic security can deterministically implement either exact or even approximate universal quantum computation or reversible classical computation without incurring exponential overhead [39]. It should be clear that the first bound in the corollary, from which the others follow, is tight, since it is satisfied by the trivial scheme where the encoding is simply a classical description of the computation to be performed. Hence in order to obtain an information theoretically secure QHE, one must be willing to sacrifice either perfect information security, determinism, or face restriction to a permissible set of circuits which is polynomial in the size of the input. As the results presented here incorporate the classical schemes as a special case, this exponential lower bound for classical reversible computation goes some way towards explaining the scaling found in [23].

The authors thank Joshua Kettlewell, Yingkai Ouyang and Si-Hui Tan for helpful discussions. The authors acknowledge support from Singapore's National Research Foundation and Ministry of Education. This material is based on research funded by the Singapore National Research Foundation under NRF Award NRF-NRFF2013-01.

* Electronic address: joe.fitzsimons@nus.edu.sg

- [1] P. W. Shor, in *Foundations of Computer Science, 1994 Proceedings, 35th Annual Symposium on* (IEEE, 1994), pp. 124–134.
- [2] L. K. Grover, in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing* (ACM, 1996), pp. 212–219.
- [3] C. H. Bennett and S. J. Wiesner, *Physical review letters* **69**, 2881 (1992).
- [4] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, *Physical Review Letters* **70**, 1895 (1993).
- [5] V. Giovannetti, S. Lloyd, and L. Maccone, *Science* **306**, 1330 (2004).
- [6] C. H. Bennett, G. Brassard, et al., in *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing* (New York, 1984), vol. 175, p. 8.
- [7] D. Gottesman and I. Chuang, arXiv preprint quant-ph/0105032 (2001).

- [8] G. Brassard, A. Broadbent, J. Fitzsimons, S. Gambs, and A. Tapp, in *Advances in Cryptology—ASIACRYPT 2007* (Springer, 2007), pp. 460–473.
- [9] V. Giovannetti, S. Lloyd, and L. Maccone, *Physical review letters* **100**, 230502 (2008).
- [10] S. Pironio, A. Acín, S. Massar, A. B. de La Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning, et al., *Nature* **464**, 1021 (2010).
- [11] S. Y. Yan, *Cryptanalytic attacks on RSA* (Springer, 2007).
- [12] D. Boneh and R. J. Lipton, in *Advances in Cryptology—CRYPTO95* (Springer, 1995), pp. 424–437.
- [13] G. Brassard, P. Høyer, and A. Tapp, in *LATIN’98: Theoretical Informatics* (Springer, 1998), pp. 163–169.
- [14] D. Mayers, *Physical review letters* **78**, 3414 (1997).
- [15] H.-K. Lo, *Physical Review A* **56**, 1154 (1997).
- [16] C. Gentry, *Proceedings of the 41st annual ACM Symposium on Theory of Computing (STOC)* pp. 169–178 (2009).
- [17] M. van Dijk, C. Gentry, S. Halevi, and V. Vaikuntanathan, in *Advances in Cryptology—EUROCRYPT 2010, Lecture Notes in Computer Science* (2010), vol. 6110, pp. 24–43.
- [18] N. P. Smart and F. Vercauteren, in *Public Key Cryptography—PKC 2010* (Springer, 2010), pp. 420–443.
- [19] Z. Brakerski and V. Vaikuntanathan, in *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on* (IEEE, 2011), pp. 97–106.
- [20] R. L. Rivest, L. Adleman, and M. L. Dertouzos, *Foundations of secure computation* **4**, 169 (1978).
- [21] R. L. Rivest, L. Adleman, and M. L. Dertouzos, *Foundations of secure computation* **4**, 169 (1978).
- [22] C. Fontaine and F. Galand, *EURASIP Journal on Information Security* **2007** (2007).
- [23] M. Hojsík and V. Půlpánová, in *Topics in Cryptology CT-RSA 2013*, edited by E. Dawson (Springer Berlin Heidelberg, 2013), vol. 7779 of *Lecture Notes in Computer Science*, pp. 375–388, ISBN 978-3-642-36094-7, URL http://dx.doi.org/10.1007/978-3-642-36095-4_24.
- [24] A. Broadbent, J. Fitzsimons, and E. Kashefi, *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2009)* pp. 517–526 (2009).
- [25] J. F. Fitzsimons and E. Kashefi, arXiv preprint arXiv:1203.5217 (2012).
- [26] T. Morimae and K. Fujii, *Physical Review A* **87**, 050301 (2013).
- [27] A. Mantri, C. A. Pérez-Delgado, and J. F. Fitzsimons, *Physical Review Letters* **111**, 230502 (2013).
- [28] S. Barz, E. Kashefi, A. Broadbent, J. F. Fitzsimons, A. Zeilinger, and P. Walther, *Science* **335**, 303 (2012).
- [29] S. Barz, J. F. Fitzsimons, E. Kashefi, and P. Walther, *Nature Physics* (2013).
- [30] M. Liang, *Quantum Inf. Process.* **12**, 3675 (2013).
- [31] A. Childs, *Quantum Information and Computation* **5**, 456 (2005).
- [32] K. Fisher, A. Broadbent, L. Shalm, Z. Yan, J. Lavoie, R. Prevedel, T. Jennewein, and K. Resch, *Nat. Commun.* **5**, 3074 (2014).
- [33] P. P. Rohde, J. F. Fitzsimons, and A. Gilchrist, *Phys. Rev. Lett.* **109**, 150501 (2012).
- [34] R. B. Griffiths, *Phys. Rev. A* **71**, 042337 (2005).
- [35] M. A. Nielsen and I. L. Chuang, *Phys. Rev. Lett.* **79**, 321 (1997).
- [36] A. Broadbent, G. Gutoski, and D. Stebila, in *Advances in Cryptology—CRYPTO 2013* (Springer, 2013), pp. 344–360.
- [37] A. W. Harrow, B. Recht, and I. L. Chuang, *Journal of Mathematical Physics* **43**, 4445 (2002), URL <http://scitation.aip.org/content/aip/journal/jmp/43/9/10.1063/1.1495899>.
- [38] We have assumed that the encryption key $|\psi_e\rangle$ is a pure state, for simplicity. This can be done without lack of generality, since we can always consider the purification of a mixed state.
- [39] Although one might expect that a scheme based on gate teleportation may overcome the limits imposed by Corollary 1, such a scheme must necessarily sacrifice determinism (becoming exponentially unlikely to succeed or having exponentially low fidelity for polynomial encodings) or must be interactive. In either case, such a scheme would fall outside the realm of homomorphic encryption.