Iterated gate-teleportation and blind quantum computation

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*

Blind quantum computation (BQC) allows a user to delegate a computation to an untrusted server while keeping the computation hidden. A number of recent works have sought to establish bounds on the communication requirements necessary to implement blind computation, and a bound based on the no-programming theorem of Nielsen and Chuang has emerged as a natural limiting factor. Here we show that this constraint only holds in limited scenarios, and show how to overcome it using a novel method of iterated gate-teleportations. This technique enables drastic reductions in the communication required for distributed quantum protocols, extending beyond the blind computation setting. Applied to BQC, this technique offers significant efficiency improvements, and in some scenarios offers an exponential reduction in communication requirements.

Blind quantum computation is a cryptographic task whereby a client seeks to hide a delegated computation from the server implementing the computation. A number of protocols for blind computation have been discovered [1–11], and a range of capabilities for the client have also been considered, from the ability to prepare or measure individual single qubit states [1, 5, 6], to the ability to perform universal computation on fixed size systems [2]. Recently work has sought to unify this disparate family of protocols in terms of security definitions [12] and in terms of resource accounting [9, 11].

Recently, Giovannetti et al. proposed a novel cheat-sensitive protocol for blind quantum computation [11]. They also derived a lower bound on communication, and showed their protocol to be optimal with respect to it. Their bound—based on the *no-programming* theorem of Nielsen and Chuang [13]—argues that $\Omega(S\log_2 G)$ qubits must be exchanged between client and server, where S is the total number of gates performed, and G is the cardinality of the gate set. It is tempting to conjecture that the no-programming bound applies to any approach to blind quantum computation [16]. In this paper, however, we show that such an efficiency constraint can be overcome even if the client is only allowed to prepare arbitrary single qubit states. The protocols we construct are not only more efficient than previous blind computation protocols, but require less communication than is required to classically describe the delegated computation.

We introduce three protocols in this Letter. The purpose of the first protocol is to introduce our technique of *iterated gate-teleportation*. This is based on the usual gate teleportation [14], but differs from standard usages in that instead of directly correcting errors induced by teleportation byproducts, we make use of additional gate teleportation steps to correct the state of the system. This change, by itself, would not normally provide any drastic speedups. However, we show that there are certain universal gate-sets that allow for guaranteed correction in very few (constant) number of gate teleportations. This technique provides a general means of increasing the efficiency of delegated computation, independent of whether it is a stand-alone computation or forms part of some cryptographic protocol. We then leverage this new technique for efficiently delegating computation to improve the efficiency of BQC. In Prot. 2 we give a *blind* version of iterated gate-teleportation, and in Prot. 3, we build a complete universal BQC protocol that, using Prot. 2 as a building block, achieves a significant saving in the total communication cost, allowing our protocol to avoid the lower bound on communication of $S\log_2 G$ which results from a naive application of the no-programming theorem.

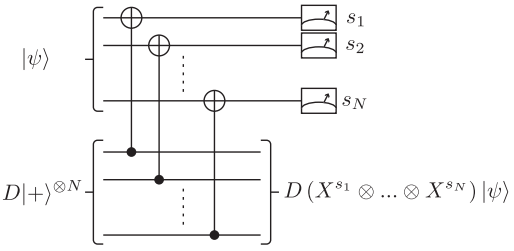

FIG. 1: Gate teleportation procedure. The top set of wires corresponds to register R, while the bottom set correspond to R'.

Consider the set $\mathscr{D}_{m,l}$ of all diagonal operators acting on m qubits of the form $\exp(i\sum_{j\in\{0,1\}^m}\theta_j Z^{j_1}\otimes Z^{j_2}\otimes\ldots\otimes Z^{j_m})$, with $\theta_j\in\{\frac{r\pi}{2^l}|r\in\{0,1,2,3,\ldots,2^l-1\}\}$. Our approach allows Alice to successfully teleport any given operator $D\in\mathscr{D}_{m,l}$ to Bob in at most l steps, each involving the transmission of m qubits. This gives a total cost of $O(ml)$. Compare this to any setting where the no-programming theorem applies, which sets a minimum of $\Omega(l2^m)$ qubits to be transmitted.

We will assume that Bob's system contains two registers R and R'. The multi-qubit gate teleportation circuit we use is depicted in Fig. 1. This procedure is formalised in Prot. 1. Note that at the end of Prot. 1, Bob is in possession of the desired output state, up to a series of Pauli-X corrections, which he can perform himself—in this non-blind version. Before discussing a blind version of Prot. 1 we show that this protocol, if followed by both Alice and Bob, does indeed yield the correct output [17].

We will begin by examining the effect of an iteration of the main loop (Steps 2a through 2d) on an arbitrary input state $|\psi\rangle$ in register R. Each iteration serves to implement a gate teleportation so that an input state $|\psi_\ell\rangle$ is transformed

Protocol 1 Iterated Gate-Teleportation

Alice's input: Gate $D \in \mathscr{D}_{m,l}$ to be teleported.

Bob's input: Initial state $|\psi\rangle$, in register R.

Output: The state $\overline{X}D|\psi\rangle$ in Bob's register R, where $\overline{X}$ is a tensor product of the Pauli-X operator and the identity, known both to Bob and Alice.

Steps:

1. Set $D_1 = D$.
2. For $1 \leq \ell \leq l$
   (a) Alice prepares the state $D_\ell|+\rangle^{\otimes m}$, and sends it to Bob, who stores it in register R'.
   (b) Bob applies the teleportation procedure depicted in Fig. 1 obtaining measurement results $s_1^{(\ell)}, \ldots, s_m^{(\ell)}$. He sends the measurement results to Alice.
   (c) Bob swaps the contents of register R and R'.
   (d) Set $D_{\ell+1} = \overline{X}_\ell D_\ell \overline{X}_\ell D_\ell^\dagger$, where $\overline{X}_\ell = \bigotimes_i^m X^{s_i^{(\ell)}}$.
3. Set $\overline{X} = \left(\prod_{\ell=1}^{l} \overline{X}_\ell\right)$. Bob now has the desired state $\overline{X}D|\psi\rangle$ in register R.

Protocol 2 Blind Iterated Gate-Teleportation

Alice's input: Gate $D \in \mathscr{D}_{m,l}$ to be teleported.

Bob's input: Initial state $|\psi\rangle$, in register R.

Output: The state $\overline{Z}\overline{X}D|\psi\rangle$ in Bob's register R, where $\overline{Z}$ ($\overline{X}$) is a tensor product of the Pauli-Z (Pauli-X) operator and the identity, and $\overline{Z}$ is Alice's encryption key, known only to her.

Steps:

1. Set $D_1 = D$.
2. For $1 \leq \ell \leq l$
   (a) Alice prepares the state $\overline{Z}_\ell \overline{Z}_{\ell-1} D_\ell |+\rangle^{\otimes m}$, where $\overline{Z}_\ell = \bigotimes_{k=1}^m Z^{r_k^{(\ell)}}$, where $r_k^{(\ell)}$, $1 \leq \ell \leq l$ are uniformly random bits, and $\overline{Z}_0 = I$. She transmits it to Bob, who stores it in register R'.
   (b) Bob applies the teleportation procedure depicted in Fig. 1 obtaining measurement results $s_1^{(\ell)}, \ldots, s_m^{(\ell)}$. He sends the measurement results to Alice.
   (c) Bob swaps the contents of register R and R'.
   (d) Set $D_{\ell+1} = \overline{X}_\ell D_\ell \overline{X}_\ell D_\ell^\dagger$, where $\overline{X}_\ell = \bigotimes_i^m X^{s_i^{(\ell)}}$.
3. Set $\overline{X} = \left(\prod_{\ell=1}^{l} \overline{X}_\ell\right)$, and $\overline{Z} = \overline{Z}_l$. Bob now has the desired state $\overline{Z}\overline{X}D|\psi\rangle$ in register R.

to $|\psi_{\ell+1}\rangle = D_\ell \overline{X}_\ell |\psi_\ell\rangle$. Thus,

$$|\psi_l\rangle = \left(\prod_{\ell=1}^{l} D_\ell \overline{X}_\ell\right)|\psi\rangle, \tag{1}$$

where the product operator is used to denote that the left to right ordering is from highest to lowest value of ℓ. Note that if an operator $D \in \mathscr{D}_{m,t}$ then $(\bigotimes_{k=1}^n X^{a_k}) D (\bigotimes_{k=1}^n X^{a_k}) D^\dagger \in \mathscr{D}_{m,t-1}$ for any choice of variables $a_k \in \{0,1\}$. Thus, for any ℓ, we have $D_\ell \in \mathscr{D}_{m,l-\ell+1}$. Since $\mathscr{D}_{m,1}$ corresponds to the set of tensor products of Z and the identity, $D_l \overline{X}_l = \pm \overline{X}_l D_l$. Thus, up to a global phase, we have

$$|\psi_l\rangle = \overline{X}_l D_l \left(\prod_{\ell=1}^{l-1} D_\ell \overline{X}_\ell\right)|\psi_j\rangle, \tag{2}$$

which collapses telescopically, substituting in the definition of $\overline{D}_\ell$, to yield

$$|\psi_l\rangle = \left(\prod_{\ell=1}^{l} \overline{X}_\ell\right) D_1 |\psi_j\rangle. \tag{3}$$

Setting $\overline{X} = \left(\prod_{\ell=1}^{l} \overline{X}_\ell\right)$ completes the proof. Note that, at this stage, Bob can correct his state by applying Pauli-X to his qubits as appropriate, without knowing the teleported gate, and without any further assistance or communication from Alice.

It is worth noting here that the set $\mathscr{D}_{m,l}$ was chosen precisely so that it exhibits the behaviour shown above. If one attempted to use the iterated gate-teleportation approach naively on an arbitrary gate set, one would have to repeat the teleportation process until the current gate teleportation succeeds. For a gate acting on n qubits, this success probability is 2^{-n}, giving an expected communication cost of $\Omega(2^n)$. Using the gate set $\mathscr{D}_{m,l}$, however, guarantees the correct gate to be teleported in at most $O(l)$ steps, independent of n. Setting $l = 3$ will include the $\pi/8$-gate in the set $\mathscr{D}_{m,l}$, making the set, along with Hadamard and CZ operators applied by Bob, universal for quantum computation, as well as making the iterated gate-teleportation succeed in $O(1)$ steps.

Next we present a blind version of Prot. 1. This procedure allows for the same functionality as the previous protocol, enabling Alice and Bob to perform gate teleportation of a gate encoded by Alice, while additionally ensuring that the gate remains unknown to Bob. The procedure for accomplishing this is presented in Prot. 2.

The proof of correctness of Prot. 2 is similar to that of Prot. 1. The only difference is that now each iteration adds a series of random Pauli-Z operators. With each iteration an input state $|\psi_\ell\rangle$ is transformed to $|\psi_{\ell+1}\rangle = \overline{Z}_\ell \overline{Z}_{\ell-1} D_\ell \overline{X}_\ell |\psi_\ell\rangle$. Thus,

$$|\psi_l\rangle = \left(\prod_{\ell=1}^{l} \overline{Z}_\ell \overline{Z}_{\ell-1} D_\ell \overline{X}_\ell\right)|\psi\rangle, \tag{4}$$

where, again, the product operator is used to denote that the left to right ordering is from highest to lowest value of ℓ. As before, we use the property of $\mathscr{D}_{m,1}$, and the definition of D_ℓ to telescopically collapse the last equation to

$$|\psi_l\rangle = \overline{Z}_l \left(\prod_{\ell=1}^{l} \overline{X}_\ell\right) D_1 |\psi_j\rangle. \tag{5}$$

A protocol P with input X is said to be blind while leaking at most $L(X)$[5], if the distribution of classical and quantum

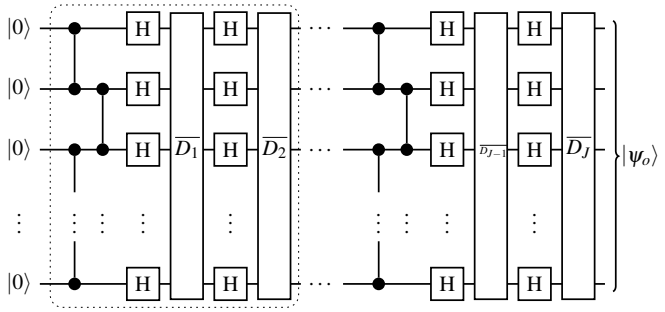

FIG. 2: Blind Quantum Computation with Teleportation Protocol. The dotted-line square shows the repeating pattern of operations.

information obtained by Bob after running P depends only on $L(X)$ [18]. Accordingly, Prot. 2 is blind while leaking at most (l,m). The only information transmitted from Alice to Bob are the set of quantum states $|\phi_\ell\rangle = \overline{Z}_\ell \overline{Z}_{\ell-1} D_\ell |0\rangle^{\otimes m}$. Note that only $|\phi_l\rangle_k$ is a function of $r_k^{(l)}$ for $1 \le k \le m$, and that as these values are unknown to Bob, this state is necessarily the maximally mixed state of m qubits, and hence independent of the values of all other $r_k^{(\ell)}$ when $\ell \neq l$. Hence $|\phi_{l-1}\rangle$ is the only quantum state dependent on $r_k^{(l-1)}$ for $1 \le k \le m$, and must similarly be in a maximally mixed state. We can apply this argument recursively, implying that every state sent from Alice to Bob is in the maximally mixed state due to the unknown values of $\{r_k^\ell | 1 \le k \le m, 1 \le \ell \le l\}$, which serve as Alice's key. As the joint state of all these messages are fixed to the maximally mixed state and hence independent of the computation, the only parameters leaked to Bob are l and m.

We are now in a position to present a complete protocol for universal blind quantum computation. This is presented as a family of protocols which are parameterised by an integer m, corresponding to the size of the client's system. During intermediate steps Alice will instruct Bob to perform m-qubit gates via gate-teleportation, using Prot. 2. The purpose of studying protocols with varying parameter m is the following. When m is small, we obtain a protocol with modest resource requirements on Alice while still outperforming previous protocols. On the other hand, when m is proportional to the number of qubits used in the computation, n, the protocol achieves an exponential (in n) separation in total communication from the naive limit implied by the no-programming theorem.

The protocol proceeds in phases. During the jth *teleportation* phase, Alice will use gate teleportation to send the desired gates to Bob. Without loss of generality we assume that m divides n [19], such that $n = Pm$, for some integer P. Then for each contiguous set of m qubits Alice will teleport an operator $D_{j,p}$, to Bob, for $1 \le p \le P$, who will then apply it to the qubits labelled $(p-1)m+1$ through pm of his current state in memory. Thus, if at the beginning of the jth phase Bob's register is in the state $|\psi_j\rangle$, by the end of the phase it will be in state $\overline{D_j}|\psi_j\rangle$, where $\overline{D_j} = \otimes_{p=1}^P D_{j,p}$.

As the set of gates which may be implemented by gate teleportation do not form a universal gate set, the scheme we present here leverages fixed gates implemented by Bob to bring about universality in a completely blind manner, as follows. Interspersed with the operations that Alice teleports, Bob will also apply the operation $\overline{\mathrm{CZ}} = \prod_{i=1}^{n-1} \mathrm{CZ}(i,i+1)$, where $\mathrm{CZ}(i,i+1)$ is the controlled-Z operator acting on qubits i and $i+1$; as well as the operator $\overline{\mathrm{H}} = \mathrm{H}^{\otimes n}$, where H is the usual Hadamard operator. The order of phases of the protocol is as follows. The first step consists of controlled-Z operators, followed by a Hadamard step, then a teleportation phase, then another Hadamard phase, followed by a second teleportation phase. Then, the pattern repeats itself until J teleportation phases have been achieved. To simplify the analysis, and without loss of generality, we assume J is even. This set of operations forms a universal set of gates for quantum computation for any m, and any $x \ge 2$ (see for example [15]). See Fig. 2 for a schematic diagram and Prot. 3 for formal presentation of the protocol.

The correct operation of the protocol depends on the proper definition of the function f_j used in Step 3c. This function is

Protocol 3 General Iterated Teleportation Blind Quantum Computation

1. Alice chooses a depth J and a set of diagonal operations $\overline{D}_j = \bigotimes_{p=1}^P D_{j,p}$, $D_{j,p} \in \mathscr{D}_{m,x}$, where $P = n/m$ and n is the number of qubits used in the computation, such that her target computation is given by the measurement of $\overline{\mathrm{H}} \overline{D}_J \overline{\mathrm{H}} \overline{D}_{J-1} \overline{\mathrm{H}} \overline{\mathrm{CZ}} \ldots \overline{\mathrm{CZ}} \overline{D}_2 \overline{\mathrm{H}} \overline{D}_1 |+\rangle^{\otimes n}$ in the computational basis.
2. Alice produces the state $\overline{Z}_1 \overline{D}_1 |+\rangle^{\otimes n}$, where $\overline{Z}_1 = \bigotimes_{k=1}^m Z^{r_k^{(1)}}$, where each $r_k^{(1)}$ is chosen uniformly at random from the set $\{0,1\}$, and transmits J and this state to Bob, who stores the quantum state in register R.
3. For $2 \le j \le J$
   (a) If $j \equiv 1 \bmod 2$, then Bob applies $\overline{\mathrm{CZ}}$ to register R.
   (b) Bob applies $\overline{\mathrm{H}}$ to register R.
   (c) For $1 \le p \le P$
      i. Alice calculates the operator $f_{j,p}(D_{j,p})$, where the function $f_{j,p}$ is defined in the main text in Eq. 11.
      ii. Alice and Bob engage in Prot. 2 using $f_{j,p}(D_{j,p})$ as Alice's target gate, and Bob's qubits $(p-1)m$ through pm as the target register.
      iii. Alice keeps a record of the operator $\overline{X}_{j,p}$, the teleportation byproduct resulting from Prot. 2, and $\overline{Z}_{j,p}$ her encryption key.
   (d) Alice calculates the operators
   $$\overline{X}_j = \bigotimes_p^{n/m} \overline{X}_{j,p}, \quad \overline{Z}_j = \bigotimes_p^{n/m} \overline{Z}_{j,p}, \tag{6}$$
   and keeps a record of them.
4. Finally, Bob measures his resulting state in the X basis, and sends the measurement outcomes $m_1 \ldots m_n$ to Alice. Alice computes each output bit for the computation as $o_k = m_k \oplus r_k^J$, where $\overline{Z}_J = \bigotimes_k Z^{r_k^J}$.

meant to correct and remove the X errors and the Z obfuscation operators introduced in previous steps. Before giving a general definition of f_j, lets first consider a simplified version of the protocol where $m = n$ and the phases of controlled-Z operators have been subsumed into the diagonal operator teleportation phases. Hence, the protocol simplifies into a series teleported gate phases followed by a layer of Hadamard gates. The output of the protocol is then given by

$$|\psi_o\rangle = \prod_{j=1}^{J} \left(\overline{Z}_j \overline{X}_j f_j(D_j) \overline{\mathrm{H}}\right) |0\rangle^{\otimes n}, \tag{7}$$

where the product operator is used to denote that the left to right ordering is from highest to lowest value of j.

Because there is a layer of Hadamard gates in between every teleportation stage, the Z operator byproducts are turned into X operators, and *vice versa*, before the the next teleportation. Since Alice can only implement diagonal gates using Prot. 2, she can only correct the X byproducts of the previous stage. She can, however, conjugate her current gate with the previous Z operators, so as to commute that operator forward, so that it can be corrected in the following teleportation stage. In this case, f_j is given by:

$$f_j(D) = \overline{\mathrm{H}}\overline{Z}_{j-1}\overline{\mathrm{H}} D \overline{Z}_{j-2}\overline{\mathrm{H}}\overline{X}_{j-1}\overline{Z}_{j-1}\overline{\mathrm{H}}, \tag{8}$$

where $\overline{Z}_j = \overline{X}_j = I$ for all $j < 1$. It is straightforward to verify from the definition above that f_j maps $\mathscr{D}_{m,l}$ onto itself. Now, substituting into Eq. 7 we get

$$|\psi_o\rangle = \overline{Z}_J \overline{X}_J \prod_{j=1}^{J} \left(D_j \overline{\mathrm{H}}\right) |0\rangle^{\otimes n}. \tag{9}$$

From this state, Alice can get the correct output for computation by having Bob measure in the X basis, and sending her the output. After this, she uses her decryption key.

The analysis of the full protocol is slightly more involved due to the (re-)introduction of the CZ gates, since these affect the propagation of X byproducts. Fortunately, errors propagate only *once*, to the nearest neighbours before they can be corrected. To see this, note that in between every two CZ gate stages, there are two sequences of Hadamard operators followed by gate teleportations. Hence, any error that cannot be fixed in the gate-teleportation phase immediately preceding the CZ stage, can be corrected in the subsequent phase. In order to define the general correction function $f_{j,p}$, let $x_{j,k}$ and $y_{j,k}$ be such that

$$\overline{X}_j = \bigotimes_{k=1}^{n} X^{x_{j,k}}, \quad \text{and} \quad \overline{Z}_j = \bigotimes_{k=1}^{n} Z^{z_{j,k}}. \tag{10}$$

For even j we take $\chi_{j,k} = z_{j-1,k}$ and $\zeta_{j,k} = z_{j-2,k} + x_{j-1,k} + \sum_{t\in\{-1,1\}} \left(z_{j-3,k+t} + x_{j-2,k+t}\right)$, and for odd j we take $\chi_{j,k} = z_{j-1,k} + \sum_{t\in\{-1,1\}} \left(z_{j-2,k+t} + x_{j-1,k+t}\right)$ and $\zeta_{j,k} = z_{j-2,k} + x_{j-1,k}$. Then, finally, for all j and p we define

$$f_{j,p}(D) = \left(\bigotimes_k X^{\chi_{j,k}}\right) D \left(\bigotimes_k Z^{\zeta_{j,k}} X^{\chi_{j,k}}\right), \tag{11}$$

where the tensor products are taken over values of k ranging from $(p-1)m+1$ to pm.

The correctness of the general protocol follows from a similar argument to that of the special case previously considered. The output of the general protocol is given by

$$|\psi_o\rangle = \prod_{j=1}^{J} \left(\overline{Z}_j \overline{X}_j \left(\bigotimes_{p=1}^{n/m} f_{j,p}(D_{j,p})\right) \overline{\mathrm{HCZ}}^j\right) |0\rangle^{\otimes n}, \tag{12}$$

where the product operator is used to denote that the left to right ordering is from highest to lowest value of j. Substituting Eq. 11 into Eq. 12, with some elementary algebra, gives

$$|\psi_o\rangle = \overline{Z}_J \overline{X}_J \prod_{j=1}^{J} \left(\overline{D}_j \overline{\mathrm{HCZ}}^j\right) |0\rangle^{\otimes n}, \tag{13}$$

as required. Furthermore, it follows directly from the blindness of Prot. 2 that Prot. 3 is blind while leaking at most (J, l, m, n).

Setting the parameter m equal to n, the above protocol requires that Alice transmit exactly Jx different n qubit states to Bob, and hence requires only a total of nJx qubits to be sent from Alice to Bob, and nJx classical bits to be sent from Bob to Alice. However, this protocol implements J unknown operations, each of which can be drawn arbitrarily from the set $\mathscr{D}_k$ which has cardinality $(2^k)^{(2^n-1)}$. Any scheme to which the no-programming theorem applies, such as that in [11], would require that at least $Jk(2^n-1)$ qubits or bits be communicated, and hence the protocol presented here is exponentially more efficient than than previous schemes.

However, setting m equal to n requires Alice to prepare large entangled states, which may be undesirable in realistic settings. Our main interest in considering such a setting is to highlight the inapplicability of lower bounds based on the no-programming theorem. On the other hand, setting m to a small constant, gives a protocol that is at least as easily implementable as previous ones in terms of resources needed on Alice's side, while still being universal for quantum computation, and offering a communication advantage. In the intermediate case where m is of order n^α for $\alpha < 1$ the protocol obtained yields an exponential increase over previous schemes in terms of the rate at which gates can be encoded, while retaining the property that Alice's device is insufficient to perform the desired computation alone.

The approach of iterated gate-teleportation introduced in this comment can be applied to other measurement-based blind computation protocols (such as [1], [5] and [9]) to achieve smaller advantages over the no-programming theorem bound. It can be used outside of the blind computation setting, to reduce communication requirements in other delegated computation scenarios via Prot. 1.

This material is based on research funded by the Singapore National Research Foundation under NRF Award NRF-NRFF2013-01. We thank Lorenzo Maccone, Tomoyuki Morimae and Terry Rudolph for helpful comments on the manuscript.

* Electronic address: joseph_fitzsimons@sutd.edu.sg

[1] 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).

[2] D. Aharonov, M. Ben-Or, and E. Eban, in *Proceedings of Innovations in Computer Science 2010* (2010), p. 453.

[3] S. Barz, E. Kashefi, A. Broadbent, J. F. Fitzsimons, A. Zeilinger, and P. Walther, Science **335**, 303 (2012).

[4] T. Morimae and K. Fujii, Nature Communications **3**, 1036 (2012).

[5] J. F. Fitzsimons and E. Kashefi, arXiv preprint arXiv:1203.5217 (2012).

[6] T. Morimae, Phys. Rev. Lett **109** (2012).

[7] T. Morimae and K. Fujii, Physical Review A **87**, 050301 (2013).

[8] S. Barz, J. F. Fitzsimons, E. Kashefi, and P. Walther, Nature Physics (2013).

[9] A. Mantri, C. A. Pérez-Delgado, and J. F. Fitzsimons, Phys. Rev. Lett. **111**, 230502 (2013).

[10] T. Sueki, T. Koshiba, and T. Morimae, Physical Review A **87** (2013).

[11] V. Giovannetti, L. Maccone, T. Morimae, and T. G. Rudolph, Phys. Rev. Lett. **111**, 230501 (2013).

[12] V. Dunjko, J. F. Fitzsimons, C. Portmann, and R. Renner, arXiv preprint arXiv:1301.3662 (2013).

[13] M. A. Nielsen and I. L. Chuang, Phys. Rev. Lett. **79**, 321 (1997).

[14] D. Gottesman and I. L. Chuang, Nature **402**, 390 (1999).

[15] J. Fitzsimons, L. Xiao, S. C. Benjamin, and J. A. Jones, Phys. Rev. Lett. **99**, 030501 (2007).

[16] A slightly weaker bound of $\frac{S \log G}{\log Bd}$ can be shown for any protocol where the client is restricted to preparing or measuring qudits of dimension d in one of B bases, by counting the number of possible branches of the protocol

[17] A minor modification of Prot. 1 involves Alice and Bob halting the protocol as soon as Bob measures all zeroes. In this case, no further corrections are necessary, and the protocol can conclude with the correct output state. In the case where $2^m < l$ this leads to an average-case communication cost which is independent of l.

[18] This definition of blindness is also compatible with the requirements for composable security in the abstract cryptography framework [12].

[19] This may always be done, since Alice can pad her input with unused ancilla qubits.