How to reuse a one-time pad and other notes on authentication, encryption, and protection of quantum information

Jonathan Oppenheim[1,2] and Michał Horodecki[2]

[1]*Racah Institute of Theoretical Physics, Hebrew University of Jerusalem, Givat Ram, Jerusalem 91904, Israel*
[2]*Institute of Theoretical Physics and Astrophysics, University of Gdańsk, Poland*
(Received 28 May 2004; published 5 October 2005)

Quantum information is a valuable resource which can be encrypted in order to protect it. We consider the size of the one-time pad that is needed to protect quantum information in a number of cases. The situation is dramatically different from the classical case: we prove that one can recycle the one-time pad without compromising security. The protocol for recycling relies on detecting whether eavesdropping has occurred, and further relies on the fact that information contained in the encrypted quantum state cannot be fully accessed. We prove the security of recycling rates when authentication of quantum states is accepted, and when it is rejected. We note that recycling schemes respect a general law of cryptography which we introduce relating the size of private keys, sent qubits, and encrypted messages. We discuss applications for encryption of quantum information in light of the resources needed for teleportation. Potential uses include the protection of resources such as entanglement and the memory of quantum computers. We also introduce another application: encrypted secret sharing and find that one can even reuse the private key that is used to encrypt a classical message. In a number of cases, one finds that the amount of private key needed for authentication or protection is smaller than in the general case.

I. INTRODUCTION

When encrypting classical information, the only method that gives unconditional security is the Vernam cipher or one-time pad. Such a private key is a random string of correlated bits shared between two parties, who we shall call Alice and Bob. By XORing [1] a message with the private key, Alice can send a message to Bob which cannot be read by an eavesdropper (Eve). However, this is a rather expensive protocol because classically, Alice and Bob cannot securely increase the size of their private key without meeting. When they have finished using their private key, they may therefore be tempted to reuse it. Of course, reusing the private key is highly insecure, and Eve will be able to exploit redundancies in messages to learn the random string of the key and gain information about the messages. It may be that no one is actually eavesdropping on Alice and Bob, but because they have no way of being certain, they should never reuse the key.

Just as classical information is valuable and may need to be encrypted, quantum information is also of value, and recently it has been proposed that one may also want to encrypt it [2,3]. Protocols have also been introduced to authenticate quantum information [4] (cf. [5] as well as early attempts [6,7]). The case of encrypting classical data that has been encoded in quantum states has also been considered [8].

In [2,3] (cf. also [9]), it was shown that if the source to be encrypted is a quantum message of m qubits, then the size of the private key for perfect encryption needs to be twice as large as in the classical case. That is, a correlated random string of $2m$ bits is necessary and sufficient to encrypt the quantum message. Such a protocol is called a private quantum channel (PQC) and care should be taken to distinguish between the PQC where quantum information is encrypted

(using classical bits), and quantum key distribution [10], where one uses the properties of quantum states to encrypt classical data.

An immediate question which arises in the context of the PQC is what exactly does it mean to encrypt quantum information? In classical encryption, one usually has the situation that Alice knows some message which she wants to send to Bob, and Eve does not know the message and wants to learn it. In the encryption of quantum information, the opposite can be the case: Alice may not know the state being sent, and Eve may know it, and want it in order to use it. Additionally, if the state is unknown to Eve, then the quantum information is in *some* sense, already encrypted. Eve can get no more than one classical bit of information from a single unknown qubit. She cannot access the quantum information. The more likely use of the PQC is therefore when Eve knows something about the state, and Alice and Bob want to make sure that Eve cannot use it. For example, the state may be some valuable resource such as the memory of a quantum computer or entanglement, and Alice and Bob want to *protect* this resource from being used by an adversary.

Perhaps an easier to define scenario exists in the authentication of quantum information. Here, Alice wants to send Bob a state in such a way that Bob can tell if the state has been tampered with. The protocols for a quantum authentication scheme (QAS) are closely related to the PQC, and therefore much of our discussion will be related to them. In fact, any QAS is necessarily a PQC [4].

Notwithstanding the necessity proofs of [2,3], in light of the fact that Eve can only extract one classical bit of information from a qubit, one might wonder in what sense the two classical bits are needed. In fact, we will show that one need not use up this classical key—it can be recycled. When encrypting or authenticating quantum information, one can

break the cardinal rule of cryptography and reuse the one-time pad.

A strong indication that key-recycling of a classical key is possible was given in [5]. There, shared singlets, rather than a classical key, were used to encrypt quantum information. It was then shown that by performing distillation on the singlets, one could reuse them. A qualitative analysis indicated that this was due to the fact that the encrypted information was quantum, rather than the fact that the private key was quantum. It was shown that an attacker's Shannon information about a classical key could be bounded, although it was unclear the extent to which secure recycling was possible, or whether correlations in different messages could be learned even though the key might remain secret.

Key-recycling makes use of one of the fundamental properties of quantum information—it cannot be copied or cloned [11]. What this means is that the encryption is such that Eve cannot make copies of several messages in order to compare them and deduce the values of the private key. Eve could of course steal or tamper with the quantum message, but this will result in a disturbance which can be detected using an error correcting code. This is in stark contrast to the classical case where Alice and Bob have no reliable way of knowing whether someone has successfully eavesdropped. For these reasons, for an authenticated quantum channel using $2m$ $+2s$ bits (where s is a security parameter), $2m+s$ of the bits can be securely recycled when the QAS is accepted. When the QAS is rejected, just under $m+s$ bits can be securely recycled. For the PQC one has a similar rate of recycling. In order to prove such recycling, one needs to show that (i) the probability that Eve gains information about the private key can be made arbitrarily small when the protocol succeeds, and (ii) that Eve cannot exploit correlations in recycled keys in order to gain information about correlations between states sent in successive uses of the channel.

The fact that one can recycle the private key may seem surprising at first, however, a hint that recycling is possible is given by teleportation. Instead of using a private quantum channel, Alice could instead teleport [12] her state to Bob. If the classical channel used in teleportation is authenticated, then teleportation also provides authentication of the quantum state. Both the PQC and teleportation require similar communication resources—both require m uses of a quantum channel to transmit quantum states (either half-singlets in the case of teleportation or the quantum states themselves in the case of the PQC). However, teleportation completely avoids using any private key at all. Teleportation also has the added advantage that there is less danger that the message will be lost in transmission—the quantum state passes directly to Bob, and there is no way to tamper or destroy it. One can therefore ask, what it is that the private key buys us—why is a private key needed at all? In [4] it was noted that the PQC has the advantage that it is noninteractive, i.e., two rounds of classical communication are not needed. The private key therefore buys us a decrease in classical communication. This seems like a rather expensive trade-off, given that classical communication is usually considered to be a cheap resource, while a private key is usually considered valuable (although physically, these resources are incomparable in that neither can be converted into the other).

One might therefore ask whether there are many situations where it is advantageous to use the PQC over teleportation. We will therefore in Sec. III A give examples where the PQC uses as much or less resources than teleportation. These include secure secret sharing, where we find that the related teleportation protocol also requires a private key. The PQC therefore consumes less resources than teleportation (i.e., uses no classical communication). Furthermore, we prove the result that in the case of secure secret sharing using teleportation, the classical message which needs to be encrypted uses a private key which can also be recycled.

We will then discuss other cases where the PQC is useful, namely for protecting quantum resources. For example, one can use quantum data encryption to protect entanglement. We will also analyze protecting the memory of a quantum computer from being stolen and used. We find that often, the attacker will almost always have limited abilities which can be exploited to use a one-time pad which is smaller than the $2m$ bit bound of [2] and [3]. This ability to beat the $2m$ bound in specific cases is in addition to our ability to reuse the pad [13]. For protecting n bits of entanglement, we will see that an n bit reusable pad is sufficient, while for protecting the memory of a quantum computer, it might be possible to use a reusable pad of a size given by the error-correction threshold [14,15].

Although some of the examples we give may be of a practical nature, our primary motivation for studying the PQC is because it is interesting and raises many questions considering the nature of encryption and of quantum information. In particular, it allows us to gain additional insight into teleportation by decoupling the sending of qubits from their encryption. We will also see that even with key-recycling, the PQC obeys a general rule which we prove, regarding the maximum increase of a private key δK as a function of sent qubits δQ and sent private messages δM. Namely

$$\delta K \leq \delta Q - \delta M. \tag{1}$$

We also discuss the thermodynamical nature of such a law.

We will discuss key-recycling in Sec. II and prove that it is secure. Next, in Sec. III we discuss applications of the PQC, including encrypted secret sharing and the protection of entanglement and other resources. We conclude in Sec. IV with a brief discussion.

II. PRIVATE KEY RECYCLING

Let us imagine that Alice wants to send to Bob a state ρ composed of m qubits, and we consider the possibility that the potential adversary Eve may have some prior knowledge of the states that will be sent, or may know from what distribution they arise. To create a PQC, Alice "q-encrypts" the state using her $2m$ bit private key using the method of [2,3]. Essentially, to each qubit, Alice conditionally applies a bit flip in the z direction, and then conditionally applies one in the x direction, using two bits of her private key as the control bits. The private key is a classical bit string randomly chosen from the uniform distribution $\mathcal{K}$. The encrypted state ρ_o is now maximally mixed regardless of what the initial

state was, and can be sent to Bob. Such a procedure gives the private quantum channel (PQC).

Definition 1. A secure private quantum channel (PQC) with error ϵ is a set of classical keys $\mathcal{K}$ and computable superoperators A_k and B_k for each key k such that:

(1) for any ensemble $\mathcal{E} = \{p_i, \rho_i\}$, A_k acts on all m-qubit states of $\mathcal{E}$ and outputs an ensemble $\mathcal{E}' = \{p_i, \rho_{ik}\}$;

(2) For all ρ_i and keys $k \in \mathcal{K}$: $B_k(A_k(\rho_i)) = \rho_i$; and

(3) for all ensembles $\mathcal{E}$ and measurements $\mathcal{M}$ acting on $\mathcal{E}'$, $H(v:i) \leqslant \epsilon$, where $H(v:i)$ is the mutual information between measurement outcomes v and the members of the ensemble i.

In other words, a PQC is an encryption/decryption scheme, such that for all ensembles, the probability that an eavesdropper learns more than ϵ about which state is being sent is small. This definition is somewhat different to that of [2,3] in that we define the PQC with respect to how much information can be gained by an eavesdropper (as opposed to just requiring that PQC produce the maximally mixed state ρ_o). Note that the definition, though based on Shannon information, does not refer to Eve's knowledge of the states. Indeed, in the PQC paradigm, one assumes that Eve may know the states from the very beginning. Note also that such a scheme (unlike teleportation) is only one-way, i.e., Alice just sends the encrypted state to Bob, and no classical communication need be used.

We will discuss key recycling in such a scenario. However, before proceeding to this discussion, we will first consider the case of key recycling in quantum authentication protocols, because there are less subtleties involved.

A. Key recycling in quantum authentication protocols

It usually makes little sense to encrypt quantum information without authenticating it in some way, or protecting the state with some form of error correction. This is because unlike a classical message, any eavesdropping on the channel may damage the state. Since the states may be unknown to Alice, the state cannot be resent. A quantum authentication scheme (QAS) is some encoding that Alice performs using a private key shared with Bob, which enables Bob to tell whether the state has arrived unaltered. Bob will either accept or reject the state, depending on whether he believes he received the correct state. Again, like the PQC, the scheme is one-way.

Definition 2. [4] A one-way quantum authentication scheme (QAS) is a set of classical keys $\mathcal{K}$ and computable super-operators A_k and B_k for each key k such that:

(1) A_k takes an m-qubit state ρ and outputs a system ρ_k of $m+s$ qubits.

(2) B_k takes as input the (possibly altered) state ρ_k' and outputs an m-qubit state ρ', which includes a single qubit which indicates acceptance or rejection (denoted by $|\text{ACC}\rangle$, $|\text{REJ}\rangle$).

A two-way QAS allows for communication between Alice and Bob during the protocol. The scheme is secure if the probability that the protocol is accepted, and that the state is not $|\psi\rangle$, is less than ϵ. That is

Definition 3. [4] A quantum authentications scheme is secure with error ϵ if for all states $|\psi\rangle$ it satisfies:

Completeness: For all keys $k \in \mathcal{K}$: $B_k(A_k(|\psi\rangle\langle\psi|)) = |\psi\rangle\langle\psi| \otimes |\text{ACC}\rangle\langle\text{ACC}|$

Soundness: $\text{Tr}(P\rho') \geqslant 1 - \epsilon$ where $P = |\psi\rangle\langle\psi| \otimes |\text{ACC}\rangle \times \langle\text{ACC}| + I_m \otimes |\text{REJ}\rangle\langle\text{REJ}|$ and ρ' is again the output state.

Here, security is defined for pure states, and can be extended to mixed states by the linearity of quantum mechanics.

We will consider a pessimistic scenario from the point of view of key recycling. Namely, we imagine that Eve may know exactly what state is being sent. However, as in [4] we will assume that Eve has not already managed to acquire part of the state [16]. We will then see that Eve cannot learn too much about the private key because she cannot access all the quantum information.

We will show that almost the entire key can be recycled when the protocol is accepted. When the authentication protocol is rejected, we find that half the key can be recycled. We will show that Alice and Bob can place a sufficient bound on Eve's information (regardless of the initial state and Eve's prior knowledge), that they can perform a process known as privacy amplification [17] to recycle their key. Essentially, they are able to publicly communicate, to distill from their $2n$ bit key, a slightly smaller key of which Eve has an exponentially small probability of knowing anything about. We then show that this recycled key can then be used in another round of the QAS.

To get a true bound we will discuss a general authentication protocol based on stabilizer purity testing codes (error correcting codes) used in Ref. [4]. The protocol of this scheme (SQAS) is as follows.

Protocol 1: Stabilizer-based quantum authentication scheme (SQAS)

(1) Alice and Bob share a secret key x of length $2m$ to be used for q-encryption. For authentication, they additionally agree on some stabilizer purity testing code $\{Q_z\}$ and two secret keys z and y of combined length $2s$.

(2) Alice uses x to encrypt an m qubit state ρ as ρ_o.

(3) Alice encodes ρ_o according to Q_z for the code Q_z and adds syndrome y to produce σ. This requires s additional qubits. She then sends the total state of $n = m + s$ qubits to Bob.

(4) Bob receives the n qubits. Denote the received state by σ'. Bob measures the syndrome y' of the code Q_z on his qubits. Bob compares y to y', and aborts if any error is detected. Bob decodes his n-qubit word according to Q_z, obtaining ρ_o'. Bob q-decrypts ρ_o' using x and obtains ρ'.

Essentially, Alice not only q-encrypts the state ρ using $2m$ bits of classical key, but she also encodes the state using an error-correction protocol based on a set of stabilizer codes and determined by some additional private key. Additional s qubits are used during transmission. The length s is determined by the degree of channel noise or extent of anticipated eavesdropping. Such a protocol can protect against an arbitrarily large amount of eavesdropping. An example of such a scheme was shown in [4] using a particular purity testing code $\{Q_z\}$ which gave a soundness error $2n/s(2^s + 1)$ using an additional key of length $s + \log_2(2^s + 1)$. If the authentication

scheme passes, we will see that they can recycle the key because the probability that Eve has any information about the state or the key can be made arbitrarily small, and Eve is product with the sent states.

The proof of security of Protocol 1 given in [4] is analogous to the Shor-Preskill security proof of BB84 [18]. One essentially shows that such a scheme can be converted into a teleportation protocol without loss of security. Note that the classical communication is taken to be authenticated, since the amount of key required for this is logarithmic in the amount of classical communication, and thus does not contribute to any of the rates of key consumption. We can thus safely ignore this cost in the remainder of the paper.

We now turn to proving security of key-recycling. By security, we mean that if one has two secure QAS with two keys k and l, to authenticate two states, then the scheme is still secure if one can replace the key l with a key recycled from k. Denoting the length of a key by $\|$, we have

Definition 4. A secure key recycling scheme of error ϵ and efficiency k'/k is a map $\Lambda : \mathcal{K} \to \mathcal{K}'$ such that if $\mathcal{E}_k$ and $\mathcal{E}'_l$ are secure QAS on Hilbert spaces $\mathcal{H}$ and $\mathcal{H}'$ then when the QAS are accepted $\mathcal{E}'_{k'} \mathcal{E}_k$ is a secure QAS on $\mathcal{H} \otimes \mathcal{H}'$ with error ϵ when $|l| = |k'|$ with $k' = \Lambda(k)$.

The error ϵ depends not only on the key-recycling scheme but also on the security of the QAS.

The proof in [4] of security of the QAS assumed that Alice and Bob are authenticating pure states (or parts of pure states). That is, it is assumed that the adversary does not already hold part of the state being sent. After the arXiv version of this paper appeared, this condition was removed in the more general and detailed work of [19]. We will also make this assumption, and under it, now prove the following.

Theorem 1. For the QAS of Protocol 1 there exists a secure key recycling scheme with efficiency rate $(2m+s)/(2m+2s)$.

To prove Theorem 1 we first need the following lemma saying that sending half of a singlet does not need encryption by x.

Lemma 1. Consider the maximally entangled state composed of two subsystems $|\psi_+\rangle = \Sigma_i |i\rangle_1 |i\rangle_2$ on $\mathcal{H}_m \otimes \mathcal{H}_m$, with dimension $2^m \times 2^m$ and ρ being the state obtained from tracing out one of the subsystems. Then a secure SQAS using Protocol 1 for ρ with a key of length $2m+2s$ is also secure using a key of length $2s$.

In other words, for half of a maximally entangled state, the key x is not needed, only the key y and z.

Proof. Alice uses Protocol 1 on subsystem 2 except that neither step 2, nor the key x is used. We then note that because

$$I \otimes U |\psi_+\rangle = U^* \otimes I |\psi_+\rangle \qquad (2)$$

Alice can replace step 2 by q-encrypting subsystem 1. Then, due to Eq. (2), this protocol is completely equivalent to Protocol 1 and is a secure QAS. Then, since Alice's actions on subsystem 1 commute with Eve's and Bob's actions on subsystem 2, it cannot matter whether Alice actually performs the q-encryption of subsystem 1. ∎

Now, we prove Theorem 1, using a technique analogous to that of Shor and Preskill.

Proof. Let us consider a variation of Protocol 1 acting on a pure state ψ by assumption (by the linearity of quantum mechanics, our proof will hold if only part of the state is being sent. That is, it holds for any ρ as long as the purification is not with Eve). The modified QAS, Protocol 2, differs in step 2 and the final step from Protocol 1.

Protocol 2: Modified QAS

(1) Alice prepares the state $1/\sqrt{2}^{2m} \Sigma_x |x\rangle$ where the $|x\rangle$ are 2^{2m} orthogonal states. For authentication, Alice and Bob additionally agree on some stabilizer purity testing code $\{Q_z\}$ and two secret keys z and y of combined length $2s$.

(2) Alice q-encrypts ψ conditionally on $|x\rangle$ creating the state $1/\sqrt{2}^{2m} \Sigma_x |x\rangle |\psi_x\rangle$.

(3) Let us call the subsystem which is the mixture of ψ_x's ρ. Alice encodes ρ according to Q_z for the code Q_z with syndrome y to produce σ. This requires s additional qubits. She then sends the total state of $n = m + s$ qubits to Bob.

(4) Bob receives the n qubits. Denote the received state by σ'. Bob measures the syndrome y' of the code Q_z on his qubits. Bob compares y to y', and aborts if any error is detected. Bob decodes his n-qubit word according to Q_z, obtaining ρ'.

(5) If Bob accepts, Alice measures $|x\rangle$ to obtain a random string x.

First we note that because ψ has essentially been q-encrypted by the action's of Alice, ρ is the maximally mixed state ρ_o. Therefore the state $\Sigma_x |x\rangle |\psi_x\rangle$ is a maximally entangled state ψ_+ of $2m$ qubits (actually up to local unitaries, on Alice's side it is a tensor product of half a singlet and the initial state). Therefore, by Lemma 1, if Bob accepts, then Alice and Bob share a state ρ'_{AB} that has an overlap with ψ_+ equal to $1 - \epsilon'$. One can now use standard reasoning originating in [20]: since Alice and Bob share an (almost) pure state, they are decoupled from Eve.

Alice can then measure x without Eve learning what x is. From the point of view of Eve, Protocol 2 is completely equivalent to Protocol 1, and therefore if Bob accepts the authentication, the probability that Eve has obtained more than ϵ information about x is small.

It should be pointed out that Protocol 2 is not equivalent to 1 from the point of view of Bob, since he does not have the key x. However, this is not relevant, since he decides whether to accept the authentication before decoding using x. Therefore neither his actions nor Eve's depend on the modifications that Alice makes. One can in fact show that the preceding protocol is in fact equivalent to teleportation.

Let us now prove that Eve is in a product state with x and the sent state. This will ensure not only reusability of the key in another QAS, but in any other cryptographic task according to [21,22]. This essentially follows from the fact that if the authentication is accepted, the protocol is equivalent to teleportation. Let us consider the total state of the system, before step 5 of Protocol 2. Using the fact that all ancillas and all actions by the parties can be represented as unitaries acting on pure states, we can write the total state of the system in the form

$$\Psi = \frac{1}{2^m 2^s} \sum_{x,y,z} |x\rangle U_B U_E U_A (|\psi_x\rangle|0\rangle_E |0\rangle_{A'} |y,z\rangle_{A''} |y,z\rangle_{B''}) |y,z\rangle_R$$

$$= \sum_i \sqrt{p_i} |\psi_i\rangle |\phi_i\rangle. \tag{3}$$

Here, the key used in the authentication is encoded in the states $|y,z\rangle_{A''}|y,z\rangle_{B''}$ which is labeled by the particular $2s$ bit key $\{y,z\}$ that is used in the QAS. Since this is a classical key, we purify it on an imaginary reference system $|yz\rangle_R$ to which none of the parties have access. The first unitary U_A is the further encoding done by Alice, and it acts on the states ψ_x and the s qubit ancilla $|0\rangle_{A'}$ conditional on the authentication key $|y,z\rangle_{A''}$. The second U_E are the actions of the eavesdropper acting on the $m+s$ qubits sent across the channel, and U_B, the decoding of Bob, acting on the ancilla and sent state ψ_x (conditioned on $|y,z\rangle_{B''}$).

In the second line we have rewritten Ψ in terms of the eigenbasis of maximally entangled states (one of which is the state $\psi_+ \equiv \psi_0 = \Sigma_x |x\rangle|\psi_x\rangle$), correlated with states on the ancillas ϕ_i (which need not be orthogonal). Now, if the QAS is accepted, then we have according to Definition 3 of a secure QAS

$$\mathrm{Tr}(\mathrm{Tr}_{A'A''B''RE}(|\Psi\rangle\langle\Psi|)|\psi_+\rangle\langle\psi_+|) \geq 1 - \epsilon, \tag{4}$$

which implies that with probability $p_0 \geq 1 - \epsilon$, the state of the system is $|\psi_+\rangle \otimes |\phi_0\rangle$ which implies that with arbitrary high probability, Eve is product with the key x and the encoded state ψ_x (even more strongly, all ancillas, as well as the key $\{y,z\}$ are product with x and the sent state ψ). From the composability theorem [21], this implies that the key x can be reused in the second QAS (using the same methodology of [22], since one can also regard this protocol in terms of security of key).

Note that we not only had to prove Eve's lack of knowledge about x, but also that she was decoupled from the sent states. This then gives that the recycled key k' can be used in a secure QAS. The danger is that k' might not be reusable because it is correlated to the sent states. For example, in Sec. II C we will see that with respect to the PQC, one must be concerned with the fact that Eve could learn about correlations between different sent states, even though she learns nothing about the key. However, reusability of k' in the QAS follows from the definition of the QAS. Essentially, since the QAS is not defined in terms of Eve's knowledge, but in terms of Bob verifying he got the correct state, one need not worry whether Eve is correlated with the sent states—we can assume she already knows them.

Let us now turn to the key $\{y,z\}$ which was used in the purity testing code. It is of length $2s$, and we now show that we can perform privacy amplification to obtain a key of length s which can also be reused. In an earlier version of this paper, we showed that the probability that Eve obtains more than an exponentially small amount of information about the key is exponentially small, using results of [23–25], and further, that this part of the key was in a product state with the sent states. Since then, the results of [26] appeared, and we can use their stronger bounds to show that Eve is in a product state with the s of the bits which can be

recycled from the $2s$ key used in the authentication step. Thus we can formally prove that this part of the key can also be reused using the stronger security definition of composability.

Alice and Bob will need to perform privacy amplification [17] on the key $\{y,z\}$ to produce a smaller key of length s, and in order to prove the security of this privacy amplification, we use [26]:

Lemma 2. Given a set of two-universal Hash functions [27] G from J, a random classical distribution with density matrix ρ_J, to a distribution T with range $\{0,1\}^t$ and density matrix $\rho_T(G)$, joint density matrices $\rho_{JE}, \rho_{TE}(G)$, and marginal distribution $\rho_E = \mathrm{tr}_J \rho_{JE} = \mathrm{tr}_T \rho_{TE}(G)$ then $\mathbb{E}_G(\mathrm{tr}|\rho_{TE}(G) - \rho_T(G) \otimes \rho_E|) \leq 2^{1/2\{\log \lambda_{max}(\rho_{JE}) + \log[rank(\rho_E)] + t\}} + 2\epsilon$ where $\mathbb{E}_G$ is the expectation value over G, and λ_{max} is the largest eigenvalue.

Intuitively speaking, the above lemma places a bound on how close the privacy amplified distribution is to being product with an eavesdropper.

In the case under consideration, if the protocol is accepted, then as shown above, with probability $p \geq 1 - \epsilon$ the total state of the system is $|\psi_+\rangle \otimes |\phi_0\rangle$ with $|\psi_+\rangle$ with Alice and Bob, and $|\phi_0\rangle$ on $A'A''B''R$ and E. With high probability, Eve is completely in a product state with the m qubit singlet, and the total state is

$$|\psi_+\rangle_{AB} \otimes |\phi_0\rangle_{A'A''B''ER}. \tag{5}$$

Upon acceptance of the protocol, we can thus simulate the situation by giving the entire A' to Eve, since this can only further degrade the security. Then, with probability $p \geq 1 - \epsilon$, Eve's state ρ_E has maximal rank s, since she is effectively only acting on the s qubit ancilla. The largest eigenvalue λ_{max} of the density matrix on $A'A''B''E$ is no greater than $1/2^{2s}$ (since the probability of any y,z is this large). Thus by Lemma 2, $\mathbb{E}_G(\mathrm{tr}|\rho_{TE} - \rho_T \otimes \rho_E|)$ will be small if we choose the recycled key of length $t = s - 2$. Here, $\rho_T(G)$ is the density matrix of the privacy amplified key $\{y,z\}$. This is precisely the condition needed for composability of a key found in [22], and ensures that just under half of the $2s$ bit key used in the authentication process can be recycled. The total efficiency of the recycling is then a rate of $(2m+s)/(2m+2s)$. ∎

Here some remarks concerning our use of composability are in order. Our primary goal is to prove that the key can be recycled. In the case of x we had two problems we fought with: not only (1) Eve should have small knowledge about the recycled key, but also (2) Eve should not be correlated with the total system: message plus recycled key. This is precisely a place where recycling of the classical one-time pad fails: knowing that the key was used once more, Eve will get to know information about correlations between subsequent messages.

In our case we assume from the very beginning that Eve knows the message. Then, correlations of Eve with system "recycled key plus message" would be dangerous because Eve would then get to know the key, and in some further round, she would break authentication (i.e., tamper with the sent state, without causing Bob's rejection). Thus both in the

cases of key x and z, y we showed that the recycled keys are not correlated with the states subjected to authentication.

Let us now consider the case that the QAS is rejected. Typically, Bob will reject the authentication if the noise level or eavesdropping is higher than anticipated. That is, if Alice uses an error correction code which is not large enough. In such a case, we will show

Theorem 2. For the QAS of Protocol 1, if the protocol is rejected, there exists a recycling scheme with efficiency rate $(m+s)(2m+2s)$.

In an earlier version of this paper, our recycling scheme in the case of rejecting the QAS yielded s bits of key. We obtain a better rate of recycling by simply giving Eve all the $m+s$ qubits. Then, the maximal rank of Eve's state is thus $m+s$, and the maximum eigenvalue of the total state (authentication key plus Eve's state) is bounded by $1/2^{2m+2s}$. Thus by use of Lemma 2 we can reuse the key if we hash the total key to a size just under $m+s$. ∎

An important aspect of the key recycling scheme is that it does not particularly take away from the advantage of the QAS over teleportation. The QAS has the advantage that it is noninteractive (only one-way communication is needed, but at the price of a private key). Key recycling will require a small amount of back communication (Bob must tell Eve whether he accepts or rejects the state). However, the back-communication and key recycling can be done at the convenience of the two parties. That is if Alice needs to send states so that Bob can use them right away, then this can be done, and later, during a break in the transmission, they can engage in key recycling. Therefore a QAS plus key recycling still has an advantage over teleportation in terms of interaction.

B. Slippery slope towards teleportation

Having determined that the key can be recycled in QAS, one must then wonder if it is needed at all. Indeed, this is the case. At the end of Protocol 2, Alice after measuring what x is, can just tell Bob the result publicly. This is no less secure than Protocol 2 since the stabilizer code ensures that Eve did not touch the state, and therefore, by the time she learns what x is, it is too late. In fact, the keys z and y in Protocol 2 can also be replaced by public communication from Alice to Bob (as long as the classical public communication channel is authenticated). This results in Protocol 3.

Protocol 3: Modified QAS

(1) Alice chooses random strings x, y, and z. Alice and Bob additionally agree on some set of stabilizer purity testing codes $\{Q_z\}$.

(2) Alice q-encrypts ψ conditionally on x creating the state ψ_x.

(3) Let us call the subsystem which is the mixture of ψ_x's ρ. Alice encodes ρ according to Q_z for the code Q_z with syndrome y to produce σ. This requires s additional qubits. She then sends the total state of $n=m+s$ qubits to Bob.

(4) Bob receives the n qubits. Denote the received state by σ'. He indicates receipt to Alice.

(5) Alice tells Bob z, and Bob measures the syndrome of Q_z on σ' obtaining the result y'. Alice and Bob compare y to y', and abort if any error is detected. Bob decodes his n-qubit word according to Q_z, obtaining ρ'.

(6) If the protocol is accepted, Alice tells Bob x and Bob decrypts ρ' to obtain ψ'.

That this protocol is secure can be seen by noting that it is essentially teleportation (as was Protocol 2). In teleportation, Alice sends n half-singlets to Bob, and then they measure the syndrome y of the random code Q_z. If they both get the same y, they can presume that the singlets are pure, and they begin teleportation. The preceding protocol differs from teleportation, only in that it is as if Alice has made the Bell measurement to start teleportation, and measured the syndrome of Q_z, before sending the state to Bob. Since such measurements do not change the density matrix, they clearly cannot reduce security.

This procedure has the advantage over the QAS in that less key is needed—not a surprise given that much of it can be reused in the QAS. Of course, it is interactive, while for the QAS with recycling, the interaction can be performed at any time, well after Bob has used his received state. This in some sense highlights the role of the secret key x used in the QAS. Its purpose is not as much for encryption (since Alice could use a random string instead and later tell Bob), but rather, it serves a role of communication. In other words, it allows one to have a noninteractive protocol. In fact, there is still interaction—the key was distributed at some point in the past, and might be recycled at some point in the future—but this interaction can occur at more convenient times. The price (the secret key), is of course rather high.

C. Key recycling for the private quantum channel

Having discussed key recycling for the QAS, we now turn to key-recycling for the PQC. We consider here a PQC such as that of [2,3]—namely, for each qubit of message, one uses 2 bits of private key, with Alice performing one of four operations on each qubit conditional on the key. We will now explain why to get key recycling for the PQC, one must actually use authentication or purity testing. One therefore must actually do key recycling of the QAS.

At first, it might seem that for the PQC one can easily recycle half the key, based simply on the Holevo bound, and the bound of Lemma 2. Eve will be ignorant of half the key, no matter what states are being sent and how much she knows about the states that are being sent.

There is, however, a problem with this. Eve may not be interested in learning the key or the identity of a particular state, but may instead be interested in learning about correlations between the various quantum states being sent between Alice and Bob. Let us imagine that Eve does not try to learn anything about the state ρ encrypted with k but instead, just steals it. Then, when Alice and Bob send a new state ρ' encrypted with k' she steals that as well. Now she may not know much about these two states, but she will learn something by having both of them. The reason is that k', although unknown to Eve, is in fact correlated with k. This means that Eve will be able to learn something about the correlations between ρ and ρ', even though she has no knowledge about k and k'.

The amount of information that Eve can learn about the key depends not only on how much she knows *a priori* about

the states being sent, but also on which states being sent. For example, for mixed states, from the Holevo bound one has that Eve will learn less than m bits of information about the key. Even for pure states, Eve will in general learn much less. Consider, for example, the state

$$|\pi/4\rangle = \cos(\pi/8)|0\rangle + \sin(\pi/8)|1\rangle. \quad (6)$$

It is easy to verify that this state gets encrypted into one of four nonorthogonal states $|0\rangle$, $|1\rangle$, $|0+1\rangle$, or $|0-1\rangle$, conditional on the private key. These are the same states used in BB84 [10].

Now, it may be that Eve did not steal the state, or was unsuccessful in her eavesdropping. It therefore seems a pity for Alice and Bob to throw out k if this is indeed the case. Classically, Alice and Bob have no way to determine whether eavesdropping was successful, but this is not the case for the encryption of quantum information. Alice and Bob can indeed discover whether Eve tampered with their state using a QAS or a purity-testing protocol to ensure that they still possess the original state ρ. Then, because of the no-cloning theorem [11], they can be sure that Eve did not steal or tamper with ρ.

While an authentication scheme such as that of Protocol 1 will be enough to enable key recycling, one can also have Alice publicly announce the authentication strings y and z, exactly as was done for Protocol 3. However, since these have to be done over an authenticated classical channel, it is doubtful that this would be less resource-intensive.

In summary, if one wishes to perform key recycling of a PQC, one needs to augment the PQC with some authentication. The amount of key recycling is then given by Theorem 1. Since any authentication scheme is also a PQC [4], the theorem also guarantees security of the PQC. As was noted in our proof of Theorem 1, Eve will be virtually product with the encrypted states ψ_x. It is for this reason that Eve is not able to gain information about correlations between the various states being sent.

D. Basic law of privacy

It might make one uncomfortable that the length K of the private key need not decrease by 1 for each qubit of encrypted message sent. We know that classically, one cannot increase privacy through communication. However, we know that due to quantum key distribution, the size of a private key can increase for each qubit sent. Quantum communication therefore allows one to increase the size of a private key by 1 per qubit sent. If the number of sent qubits is δQ, and the size of an encrypted message which gets sent (whether quantum or classical) is δM, then we note the following general law. For any communication between two parties

$$\delta K \leq \delta Q - \delta M, \quad (7)$$

where δK is the change in length of the private key. Such a law, while probably known on some intuitive level, has never been stated or proven to the best of our knowledge. We note that our recycling protocol respects such a law, as does teleportation.

Proof. Assume for contradiction that the basic law is violated. Then, in the case where the additional message δM is classical, we imagine that the message sent is used to create an additional key of length δM (since the private communication can always be communication of the private key). Thus a violation of the basic law would imply $\delta K \geq \delta Q$. That the latter inequality cannot be true follows from [28], where the theory of privacy was recast in terms of entanglement theory, and thus it was shown that the relative entropy of entanglement,

$$E_r(\rho_{AB}) \equiv \min_{\sigma_{AB} \in sep} \operatorname{Tr}(\rho_{AB} \log \rho_{AB} - \rho_{AB} \log \sigma_{AB}) \quad (8)$$

(where the minimum is taken over separable states), is an upper bound on the key rate. Since E_r is less than E_c (the entanglement cost) of a state (which is by definition Q), it follows that $\delta K \leq \delta Q$. This gives the desired contradiction.

If the additional message δM is quantum, then the quantum channel must work for all states (since by definition of a quantum message the sent states are unknown and decoupled from the environment). We could thus send half of a singlet for each qubit of message, and then use it to create one bit of key. We would then again have $\delta K \geq \delta Q$. ∎

It is actually quite surprising that [28] is needed for what appears to be such a basic and simple result. One imagines that $\delta K \leq \delta Q$ would follow from the Holevo bound, but because a key is a strictly weaker resource than communication, and can be distilled using many rounds of public communication, standard techniques appear not to work.

The laws of privacy bear a resemblance to thermodynamics—not surprising given their information-theoretic nature. Alice and Bob each possess a random string of maximal entropy. K therefore, represents a decrease in the total entropy of Alice and Bob's string, due to correlations—it is a *negentropy*. In the absence of eavesdropping, both teleportation and our recycling protocol have $\delta K = 0$, and are therefore optimal. They are, in a sense, isentropic processes. On the other hand, eavesdropping disturbs the system and results in an increase of entropy between Alice and Bob. The size of the private key (or negentropy), goes down. This is like the second law of thermodynamics. The leakage of information to Eve is physically very closely related to the processes that produce an increase of entropy in thermodynamics: Indeed, the process that conserves energy and increases entropy is exactly eavesdropping, i.e., it is pure decoherence (if we will not consider just coarse graining). In the case of no eavesdropping, then the bound of Eq. (7) gets saturated, and you have a law which looks very similar to the first law of thermodynamics. In that sense, one should think of the sent messages as work which is being performed.

III. APPLICATIONS

As noted in the Introduction, if two parties are interested in sending an unknown state between them, they are likely to chose teleportation over the PQC. To teleport an unknown state, Alice sends a number of halves of singlets to Bob, and then they use distillation to ensure that they are indeed sharing pure singlets (since the sent entanglement may get cor-

rupted due to noise, or the actions of an eavesdropper). After distillation, Alice makes a measurement on her half of the singlets, and sends the results to Bob, who can then transform his half of the singlets into the state which is to be transmitted. Since the outcomes of Alice's measurements are completely random, Eve gets no information about the state which is being sent. Teleportation therefore automatically ensures encryption. If the classical channel is authenticated, teleportation also ensures authentication of the quantum information. Additionally, the unknown state is never held by Eve, and so it cannot get destroyed or corrupted.

Both teleportation and the PQC therefore need one usage of the quantum channel for every qubit which is to be sent and encrypted. The amount of classical communication needed for teleportation n qubits is one round of $O(\log n)$ classical bits (c-bits) from Bob to Alice, and $2n$ c-bits from Alice to Bob. The first round is used in the distillation process (Bob needs to tell Alice which singlets he tested for purity, and how noisy the singlets are), and the second round is needed for Alice to tell Bob the result of her measurement. The PQC does not need these two rounds of classical communication, but it does need the private key. It therefore seems that this saving in classical communication for the PQC comes at an expensive price (the private key).

One might argue that the private key can be recycled, and therefore, it does not represent much of a cost. However, some of the private key will be lost in the case of heavy eavesdropping or noise. Furthermore, recycling the key also requires interactive communication. Bob must inform Alice how much noise is in the check-bits and which privacy amplification function they will use. This communication can occur at a more convenient time (after Bob has decrypted the states), but it is still needed.

We therefore inquire into situations where one might prefer the PQC over teleportation. One situation where the two are equivalent is if one wishes to store the quantum information in one's own laboratory, but does not trust other people in the laboratory. One could then q-encrypt the states and keep the key somewhere safe. One could also use teleportation, and keep the results of the measurement locked away, but preparing singlets to teleport a state to oneself seems rather silly. Here, we show other applications of the PQC. In Sec. III A we consider a situation which we call secure secret sharing which requires either a modified version of the PQC or of teleportation. In this case, both protocols require the same amount of resources—namely, a private key and 2 cbits of communication. We also find the interesting result that in this case, the private key, which is used to encrypt classical information, can also be recycled. We will present an asymmetric version of secure secret sharing (one party is more trusted) in which the teleportation protocol is more resource intensive. We also explore two other examples where one would not use teleportation, namely, the protection of quantum resources such as entanglement or the memory of a quantum computer. For protecting entanglement, the private key is half the size of the general case. These examples are discussed in Sec. III B. It is the protection of quantum information where the PQC may find its most useful application.

A. Encrypted quantum secret sharing

Let us first consider the following scenario which we will call encrypted quantum secret sharing. We imagine a scenario similar to original secret sharing [29]: Alice wants to send a quantum state to Bob and Claire, but she does not want either one to be able to use the state separately. Alice desires that they cooperate in order to decrypt the state. However, we also imagine that both Bob and Claire are extremely adversarial and may try to eavesdrop on each other's communication channels in order to learn the state. Or perhaps, there is a concern that another party may eavesdrop on the channel.

Consider the following modification of the QAS (e.g. Protocol 1). Alice shares a classical private key X with Bob (to be used for q-encrypting), but rather than q-encrypting the state using this key, she instead q-encrypts the quantum state using another random string J and then applies the authentication protocol using key S (shared with Claire). She then sends the encrypted quantum state to Claire. Since Bob does not know J, he will gain nothing by eavesdropping on the quantum channel between Alice and Claire. Once Claire receives the quantum state, she checks to see that the state has arrived intact using S. If the authentication test fails, they abort. If the test succeeds, Alice classically encrypts the random string J using X and sends the encrypted string to Bob. This prevents Claire from learning the key used to encrypt the state she has, while Bob can decrypt J using X. Claire now possesses the encrypted quantum state, while Bob possess the key, and they will not be able to decrypt it unless they cooperate. The total protocol has consumed $2n$ bits of shared private key and $2m$ bits of classical communication between Alice and Bob, and n uses of the quantum channel from Alice to Claire. The state can only be decrypted if Bob and Claire get together and cooperate.

The protocol involving teleportation requires the same amount of resources. Alice sends half-singlets to Claire, and they use a purification testing protocol over an authenticated classical channel to ensure that they indeed hold singlets. Alice then makes the joint measurement on her half-singlets and the state to be shared. Alice then encrypts the results of the measurements using the key shared between herself and Bob, and sends it to him.

Key recycling can be done in both these cases, although it is a bit more problematic if one is concerned not only about a party learning the key, but also correlations between different states. The transmitted states will need two layers of authentication. Essentially Claire needs to be sure that either the singlets or the transmitted state is authentic, and then Bob would need to test for authenticity when they decrypt. This would need to be done in Claire's laboratory and so the situation may be a bit awkward for this to occur with Claire hovering over his shoulder (however, one may assume that they are polite enough to each other in person to allow this to happen). Alternatively, Claire could test for authenticity in the presence of Bob, assuming that Bob is allowed to test Claire's apparatus. Key recycling makes more sense in the case where one is concerned about eavesdropping from some other party, in which case only one layer of authentication is needed, since Alice can trust Claire's acceptance of the authentication.

This recycling in the case of encrypted quantum secret sharing is perhaps more surprising than recyclability in the original PQC protocol, since in this case, the information being encrypted is classical. In the case of teleportation, it is the result of Alice's measurements which, although random, could be used by Claire to learn what state she has. The proof of this recyclability is almost identical to Theorem 1. The only difference is that here, one does not need to show that Eve must be product with the message (J in the case of QAS-type protocol, and the result of Bell measurements in the case of the teleportation protocol). Indeed, she will not be, as both strings are classical. However, we do not need this to be the case because there are no correlations from one round to another—the messages are completely random strings.

One can envision situations where the PQC is less resource intensive than teleportation. Consider the following: Alice wants to send a state to Bob, but because quantum channels and labs are expensive, Bob has neither of these. The only option then is for Alice to send the state to Claire, a technician in a large company which owns all the quantum channels, and which also has storage equipment for keeping quantum states. Claire is, however, not trusted.

Alice can, however, encrypt the quantum state using the key shared between her and Bob, and send the state to Claire. Bob can then go over to Claire's laboratory any time he chooses, and decrypt the state (we once again imagine that he is able to test the equipment first). The teleportation protocol, however, uses more resources. The half-singlets are sent to Claire, and the results of the measurements sent to Bob. The measurement results that are sent to Bob have to be encrypted, to ensure that Claire does not eavesdrop on Bob's classical channel. Therefore the teleportation protocol also needs a private key. It is therefore more expensive because it needs two way communication, while the PQC does not.

B. Protecting entanglement

We now turn to the issue of using PQC to protect various resources, and we demonstrate a scenario where only n bits of private key are necessary and sufficient rather than $2n$. Let us imagine the following: there is a teleportation device [12] set-up between Alice (who is on Earth) and Claire (a robot) who is on the planet Venus. The more practical reader can instead imagine a device which can be used to perform quantum key distribution between Earth and a satellite, and we note that related devices for satellite deployment are already being tested [30].[1] Now, the teleportation device between Alice and Claire is essentially made up of many entangled states shared between Alice and Claire which we take to be the singlet

$$\psi_- = |01\rangle_{AC} - |10\rangle_{AC}. \tag{9}$$

We assume that these singlets have already been authenticated, and so Alice possesses many singlet-halves which

---

[1]It is usually simpler to perform quantum key distribution using sent photons, but one can also use singlets as described in the protocol of [20].

look locally like maximally mixed states. She may decide that she will not use the device, but instead wants to send her states to Bob for his use. They may have two concerns: (1) to ensure that an adversary Eve has no idea what quantum states are being sent, or (2) they may not care what Eve knows but want to ensure that Eve will not steal the states for her own use (since she may want to use the device to also teleport something to Venus).

Here, we consider the latter case (which we shall refer to as quantum data protection), and assume that Eve is also located on Earth. One way to ensure that the singlets will never fall into the wrong hands is to have an authentication protocol set up between Claire, and the authorized party on Earth. However, this is only useful if one can be sure that Claire's singlets will also never fall into the wrong hands. It may be that Eve has a friend on Venus who might take control of the teleporter.

Alice would therefore want to perform some sort of protection protocol on her half of the singlets, to ensure that if Eve steals them, they will be of no use. In this case, we do not need the full encryption scheme of [2,3] which uses $2n$ bits of private key, but can instead use a scheme using only an n bit private key. We now show that n bits are necessary and sufficient to protect entanglement.

To show sufficiency, consider the following protocol. Alice and Bob share a private key k and Alice divides her states into large blocks of size n and performs the bit flip operation (σ_x) on each of her qubits conditional on each bit of the private key. The total state of each qubit is then

$$\rho_{ABC} = \frac{1}{2}(|00\rangle\langle00| \otimes |\psi_-\rangle\langle\psi_-| + |11\rangle\langle11| \otimes |\psi_+\rangle\langle\psi_+|), \tag{10}$$

where ψ_+ is the Bell-state $|00\rangle_{AC} + |11\rangle_{AC}$ shared between Alice and Claire, and the states $|00\rangle$ and $|11\rangle$ are the classically correlated private key shared between Alice and Bob. An adversary who does not have access to the key receives the state

$$\rho_{CE} = \frac{1}{2}|\psi_-\rangle\langle\psi_-| + \frac{1}{2}|\psi_+\rangle\langle\psi_+|. \tag{11}$$

It is immediate to find using, e.g., the partial transpose test that this state is completely unentangled, and is of no use whatsoever to an adversary wishing to use the singlets.

To show that n bits of private key are necessary, we note that we would not be able to use less than n bits of randomness because any mixture of a singlet with another pure state is unentangled unless $p=1/2$. More explicitly, the encryption amounts to applying some unitaries U_i with probabilities p_i to Alice's halves of singlets. The amount of used bits is then given by the entropy $H(\{p_i\})$. To get a bound for the entropy, we note that Alice wants the state to be separable, while the separable state satisfies $S(\varrho_{AB}) \geq S(\varrho_A)$ [31]. In our case ϱ_{AB} is the state obtained from singlets after applying the random unitaries. We have $S(\varrho_A)=n$, where n is the amount of singlets while $S(\varrho_{AB}) \leq H(\{p_i\})$. If the state is to be separable, $H(\{p_i\})$ must be then no smaller than n.

After protecting the singlet, Alice can send her *n* qubits to Bob. Bob can then test the received qubits to ensure that Eve has not stolen any, by simply testing for purity (for example, by using the stabilizer codes used in Protocol 1). In the above scenario, we are not encrypting quantum information but rather protecting it—we assume that Eve already knows that Alice is trying to send the half-singlets to Bob.

We can use key-recycling also in this case. If, for example, we use the additional layer of authentication, and the test passes, then Alice and Bob reuse the key. Here, they can reuse the key, even though it is only *n* bits long to begin with. This is because the bound of Lemma 2 ensures that Eve can learn nothing about the key. The half singlet states are already maximally mixed, and so there is no measurement that Eve can make which would allow her to guess the private key.

Finally, one can also use quantum data protection even when not sending data to another party. One might imagine that Alice is concerned that someone in her laboratory (Eve) might steal the entanglement and use it for some unauthorized purpose. Alice could use quantum authentication, which would let her know that the entanglement has been tampered with, however, by then it is too late, and Eve has possession of the half-singlets. Again, setting up an authentication protocol is no good if Eve has a friend on Venus who might also steal the other half of the singlets. Alice can therefore use the protocol above using her own private key. This private key can easily be stored somewhere safe, as opposed to the quantum states which presumably must be stored in a prominent place in a laboratory, as the states need to constantly be protected against decoherence. Note that in the case of storage, a teleportation protocol makes little sense.

There are other examples where one can use the PQC to protect quantum data. For example, one might protect the memory of a quantum computer, especially if one is running a long factoring algorithm, and one does not trust other people in the laboratory In such a case, one might make do with a key which is only as long as the threshold, beyond which error correction is impossible [14,15]. That is, one introduces errors conditional on a key, such that the computation can no longer be performed. These errors can be undone if one has the key, and the computation can proceed, but without the key, the computer will not run. Whether an adversary might still be able to get some information by performing some measurement on the quantum memory in such a scenario is not clear.

IV. DISCUSSION

From a conceptual point of view, the private quantum channel and the authenticated quantum channel are interesting because it allows us to decouple the sending of a quantum state, and the encryption/authentication of a quantum state which are automatically coupled in teleportation. Since teleportation does not require a private key, we have inquired into the necessity of the private key used in the QAS. Indeed, we find that in some sense, the private key is not needed, as one can keep reusing the same private key, adding only a small amount of additional private key each time. The role of the private key was seen to be more significant in terms of communication—it allows the QAS to be noninteractive. It would be interesting to better understand the role of the private key. If it can be recycled, one wonders the extent to which it is needed at all.

Encryption of quantum information is different from encryption in the classical case. One can have the exact opposite scenario: Instead of Alice knowing the state, and Eve trying to learn it, one will likely be concerned about the case of Alice not knowing the state, and Eve knowing something about it. We therefore introduced the notion of *protecting* quantum resources. Namely, the private quantum channel allows us to ensure that certain resources cannot be used by an adversary. This is in contrast to the usual role of encryption, which seeks to prevent an adversary from knowing a message. It was found that for protecting some resources (such as entanglement), a shorter private key was needed. Other applications, such as an encrypted version of quantum secret sharing were also introduced.

We also introduced a basic law of privacy which governs how the size of the private key changes as a function of sent qubits and encrypted messages. We believe it would be interesting to explore this law further, as well as the analogy between eavesdropping and thermodynamics.

[1] An XOR flips a target bit, depending on whether a control bit is a 0 or 1. In this case, the target bit is the message to be encrypted, and the control bit is the secret key.

[2] P. Boykin and V. Roychowdhury, Phys. Rev. A **67**, 042317 (2003).

[3] M. Mosca, A. Tapp, and R. de Wolf, e-print quant-ph/0003101.

[4] H. Barnum, C. Crepeau, D. Gottesman, A. Smith, and A. Tapp Proc. 43rd Annual IEEE Symposium on the Foundations of Computer Science (IEEE, New York, 2002).

[5] D. Leung, Quantum Inf. Comput. **2**, 13 (2001).

[6] M. Curty, D. J. Santos, E. Perez, and P. Garcia-Fernandez, Phys. Rev. A **66**, 022301 (2002).

[7] Howard Barnum, e-print quant-ph/0103123.

[8] D. Gottesman, Quantum Inf. Comput. **3**, 581 (2003).

[9] H.-H. Lo, e-print quant-ph/9912011.

[10] C. H. Bennett and G. Brassard, in *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing* (IEEE Computer Society Press, New York, 1984), pp. 175–179.

[11] W. Wooters and W. Zurek, Nature (London) **299**, 802 (1982).

[12] C. H. Bennett, G. Brassard, C. Crepeau, R. Jozsa, A. Peres, and W. K. Wootters, Phys. Rev. Lett. **70**, 1895 (1983).

[13] One might think that if one can recycle the key, that the size of the key does not matter. However, we will see that this is not the case, since encryption must be done in large blocks before each round of recycling. Furthermore, in some cases, the key-recycling protocol will be aborted, in which case, the key is discarded.

[14] P. Shor, in *Proceedings of the 29th Symposium on Theory of Computer Science* (IEEE Computer Society, Washington, DC, 1996), pp. 56–65.

[15] D. Aharonov and M. Ben-Or, in *Proceedings of the 37th Symposium on the Foundations of Computer Science* (IEEE Computer Society, Washington, DC, 1997), pp. 175–188; e-print quant-ph/9611025.

[16] We thank Debbie Leung for emphasizing this point.

[17] C. Bennett, G. Brassard, and J.-M. Robert, in *Proceedings of Crypto 85, Lecture Notes in Computer Science* (Springer-Verlag, Berlin, 1986), Vol. 218, pp. 468–476.

[18] P. W. Shor and J. Preskill, Phys. Rev. Lett. **85**, 441 (2000); e-print quant-ph/0003004.

[19] P. Hayden, D. Leung, and D. Mayers, QIP 2004, http://www.iqc.ca/conferences/qip/presentations/leung.pdf.

[20] A. K. Ekert, Phys. Rev. Lett. **67**, 661 (1991).

[21] M. Ben-Or and D. Mayers, presented at MSRI 2002, http://www.msri.org/publications/ln/msri/2002/qip/mayers/1/index.html; e-print quant-ph/0409062.

[22] M. Ben-Or, M. Horodecki, D. W. Leung, D. Mayers, and J. Oppenheim, in *Theory of Cryptography: Second Theory of Cryptography Conference, TCC 2005* (2005), pp. 386–406; e-print quant-ph/0409078.

[23] A. Ambainis, A. Nayak, A. Ta-Shma, and U. Vazirani, in *Proceedings of the 31st ACM Symposium on Theory of Computing* (1999); e-print quant-ph/9804043.

[24] A. Nayak, in *Proceedings of the 40th Annual Symposium of Computer Science* (1999), pp. 369–377; e-print quant-ph/9904093.

[25] R. Konig, U. Maurer, and R. Renner, IEEE Trans. Inf. Theory (to be published).

[26] R. Renner and R. König, in *Theory of Cryptography: Second Theory of Cryptography Conference TCC 2005*, edited by J. Killian (2005), p. 407; e-print quant-ph/0403133.

[27] A random function G from X to X' is called *two-universal* if for any distinct $x_1, x_2 \in X$ the probability that $G(x_1) = G(x_2)$ is at most $1/|X'|$.

[28] K. Horodecki, M. Horodecki, P. Horodecki, and J. Oppenheim, Phys. Rev. Lett. **94**, 160502 (2005); **94**, 200501 (2005).

[29] R. Cleve, D. Gottesman, and H.-K. Lo, Phys. Rev. Lett. **83**, 64 (1999).

[30] See for example, work along the lines of R. J. Hughes, J. Nordholt, D. Derkacs, and C. Peterson, e-print quant-ph/0206092.

[31] R. Horodecki and P. Horodecki, Phys. Lett. A **194**, 147 (1994).