




Ivan Damgaard



Louis Salvail



Thomas Pedersen



1 Introduction

It is well known that only assuming a quantum channel and an authenticated classical channel, Quantum Key Distribution (QKD) can be used to generate an unconditionally secure shared key between two parties. If we want to use this key for encrypting classical messages, the simplest way is to use it as a one-time pad. This way, an m -bit key can be used to encrypt no more than m message bits, since re-using the key would not be secure (without extra assumptions like in the bounded storage model[19,10,13]).

However, if we allow the same communication model for message transmission as for key exchange — which seems quite natural — an obvious question is whether we might gain something by using the quantum channel to transmit ciphertexts. The reason why this might be a good idea is that the ciphertext is now a quantum state, and so by the laws of quantum mechanics, the adversary cannot avoid affecting the ciphertext when trying to eavesdrop. We may therefore hope being able to detect — at least with some probability — whether the adversary has interacted with the ciphertext. Clearly, if we know he has not, we can re-use the entire key. Even if he has, we may still be able to bound the amount of information he can obtain on the key, and hence we can still re-use part of the key. Note that the authenticated classical channel is needed in such a scheme, in order for the receiver to tell the sender whether the ciphertext arrived safely, and possibly also to exchange information needed to extract the part of the key that can be re-used. Such a system is called a Quantum Key-Recycling Scheme (QKRS).

A possible objection against QKRS is that since it requires interaction, we might as well use QKD to generate new key bits whenever needed. However, in the model where the authenticated classical channel is given as a black-box (i.e., not implemented via a shared key) QKD requires at least three messages: the quantum channel must be used, and the authenticated channel must be used in both directions, since otherwise the adversary could impersonate one of the honest parties. Further, each move requires a substantial amount of communication (if N qubits were transmitted then the two classical moves require more than N classical bits each). Finally, N is typically larger than the length of the secret-key produced. Hence, if we can build a QKRS scheme that is efficient, particularly in terms of how much key material can be re-used, this may be an advantage over straightforward use of QKD.

From a more theoretical point of view, our work can be seen as a study of the recycling capabilities of quantum ciphers in general. In particular, how many key bits can be recycled, and how much feedback information must go from receiver to sender in order to guarantee the security of the recycled key? How do these capabilities differ from those of classical ciphers? In this paper we give precise answers to these questions.

The idea behind a QKRS originates from Bennett and Brassard during the early days of quantum cryptography[4]. Although they did not provide any fully satisfying solution or security proof, their approach to the problem is similar to our. More recently Leung studied recycling of quantum keys in a model where Alice and Bob are allowed three moves of interaction[12]. In this model, however, quantum key distribution can be applied. Leung also suggested that classical keys can be recycled when no eavesdropping is detected. In [16], a QKRS was proposed based on quantum authentication codes[2]. The key-recycling capabilities of their scheme can be described in terms of 2 parameters: the message length m and the security parameter ℓ . The scheme uses $2m + 2\ell$ bits of key, and is based on quantum authentication schemes that, as shown in [2], must always encrypt the message. The receiver first checks the authenticity of the received quantum state and then sends the result to the sender on the authenticated channel. Even when the receiver accepts, the adversary may still have obtained a small amount of information on the key. The receiver therefore also sends a universal hash function, and privacy amplification is used to extract from the original key a secure key of length $2m + \ell$. If the receiver rejects then a secure key of length $m + \ell$ can be extracted.

In this paper, we propose a QKRS for encrypting classical messages. Our QKRS is based on a new technique where we append a k -bit classical authentication tag to the message, and then encrypt the $n = m + \ell$ -bit plaintext using the W_n -quantum cipher introduced in [8]. The authentication is based on universal hashing using an m -bit key. The cipher uses $2n = 2(m + \ell)$ bits of key, where $m + \ell$ bits are used as a one-time pad, and $m + \ell$ bits are used to select in which basis to send the result, out of a set of $2^{m+\ell}$ so called mutually unbiased bases. Thus, the entire key of the QKRS consists of $3m + 2\ell$ bits. The receiver decrypts and checks the authentication tag. If the tag is correct, we can show that the adversary has exponentially small information about the key, and the entire key can therefore be recycled. If the tag is incorrect, we can still identify $2m + \ell$ bits of the key, about which the adversary has no information, and they can therefore be re-used. Since this subset of bits is always the same, the receiver only needs to tell the sender whether he accepts or not.

Being able to recycle the entire key in case the receiver accepts is of course optimal. On the other hand, we can show that any QKRS must discard at least $m - 1$ bits of key in case the receiver rejects. Since m can be chosen to be much larger than ℓ , discarding $m + \ell$ bits, as we do, is almost optimal.

In comparison with earlier works, our technique completely eliminates the use of privacy amplification, and hence reduces the communication on the authenticated channel to a single bit. Moreover, we can recycle the entire key when the receiver accepts the authentication tag. Hence, in scenarios where interference from the adversary is not too frequent, our keys can last much longer than with previous schemes, even though we initially start with a longer key.

Our results differ from those of [16], since quantum authentication based QKRS do not guarantee the privacy of the authentication tag. Therefore, part of the key must be discarded even if the receiver accepts. Instead of quantum authentication, we use classical Wegman and Carter authentication codes[6] and a quantum encryption of classical messages[8] applied to both the message and the tag. This construction allows to recycle the entire authentication key securely.

The scheme we introduce can also be used as an authentication code for quantum messages. However, it requires a longer secret-key than the scheme in [2], but allows for recycling the authentication key entirely upon acceptance.

Our QKRS is composable since the security is expressed in terms of distance from uniform. The secret-keys and plaintexts are private when, from the adversary's point of view, they look like uniformly distributed random variables. This has been shown to provide *universal composability* in the quantum world[17].

We end this introduction with some remarks on the authenticated classical channel. Having such a channel given for free as a black-box may not be a realistic assumption, but it is well known that it can be implemented assuming the players initially have a (short) shared key.¹ In this model, the distinction between QKD and QKRS is not as clear as before, since we now assume an initial shared key for both primitives. Indeed, our QKRS can be seen as an alternative way to do QKD: we can form a message as the concatenation of new random key bits to be output and a short key for implementing the next usage of the authenticated channel. Having sent enough messages of this form successfully, we can generate a much larger number of secure key bits than we started from. Note that this is harder to achieve when using the earlier QKRS scheme since bits of the original key are lost even in successful transmissions.

2 Preliminaries

2.1 Density Operators and Distance Measures

We denote by $\mathcal{S}(\mathcal{H})$ the set of density operators on Hilbert space \mathcal{H} (i.e. positive operators σ such that $\text{tr}(\sigma) = 1$). In the following, \mathcal{H}_n denotes the 2^n -dimensional Hilbert space over \mathbb{C} , $\mathbb{1}_n$ denotes the $2^n \times 2^n$ identity operator, and $\mathbb{I}_n = 2^{-n} \mathbb{1}_n$ denotes the completely mixed state. The trace-norm distance between two quantum states $\rho, \sigma \in \mathcal{S}(\mathcal{H})$ is defined as:

$$D(\rho, \sigma) = \frac{1}{2} \text{tr} (|\rho - \sigma|),$$

where the right-hand side denotes half the sum over the absolute value of all eigenvalues of $\rho - \sigma$. The trace-norm distance is a metric over the set of density operators in $\mathcal{S}(\mathcal{H})$. In the following, we use the same notation as [17]. Let (Ω, P) be a discrete probability space. A *random state* ρ is a function from Ω to $\mathcal{S}(\mathcal{H})$. This means that to $\omega \in \Omega$ corresponds the mixed state $\rho(\omega)$. To an observer ignorant of the randomness $\omega \in \Omega$, the density operator described by ρ is given by

$$[\rho] = \sum_{\omega \in \Omega} P(\omega) \rho(\omega).$$

For any event \mathcal{E} , the density operator described by ρ conditioned on \mathcal{E} is given by

$$[\rho|\mathcal{E}] = \frac{1}{\text{Pr}(\mathcal{E})} \sum_{\omega \in \mathcal{E}} P(\omega) \rho(\omega).$$

Classical random variables can also be represented as random states. Let X be a random variable with range \mathcal{X} and let \mathcal{H} be a $\#\mathcal{X}$ -dimensional Hilbert space with orthonormal basis $\{|x\rangle\}_{x \in \mathcal{X}}$. The random state corresponding to X is denoted by $\{X\} = |X\rangle\langle X|$ and $[\{X\}] = \sum_{x \in \mathcal{X}} P(x)|x\rangle\langle x|$ denotes its associated density operator. Let $\rho \otimes \{X\}$ be a random state with a classical part $\{X\}$. The corresponding density operator is given by

$$[\rho \otimes \{X\}] = \sum_{x \in \mathcal{X}} P(x)[\rho|X=x] \otimes |x\rangle\langle x|.$$

If X is independent of ρ then $[\rho \otimes \{X\}] = [\rho] \otimes [\{X\}]$. Let X be a classical random variable with range \mathcal{X} and let ρ be a random state. The *distance to uniform of X given ρ* is defined by

$$d(X|\rho) = D([\{X\}] \otimes \rho, [\{U\}] \otimes [\rho]), \quad (1)$$

where U is a random variable uniformly distributed over \mathcal{X} .

2.2 Quantum Ciphers

A quantum encryption scheme for classical messages is the central part of any QKRS. Such schemes were introduced in [1], and further studied in [8], where their performances were analyzed against known-plaintext attacks. We adopt a similar definition here except that we allow for the encryption to provide only statistical instead of perfect privacy. As in [1,8], we model encryption under key k by an appropriate unitary operator E_k acting upon the message and a possible ancilla of any size initially in state $|0\rangle$. Decryption is simply done by applying the inverse unitary.

For convenience we will use the notation

$$\rho_x = \sum_{k \in \{0,1\}^n} 2^{-n} E_k |x\rangle\langle x| \otimes |0\rangle\langle 0| E_k^\dagger,$$

for the equal mixture of a plaintext $x \in \{0,1\}^m$ encrypted under all possible keys with uniform probability. A quantum cipher is private if, given a cipherstate, almost no information can be extracted about the plaintext.

Definition 1. For a non-negative function $\epsilon(n)$, a $\epsilon(n)$ -private (n, m) -quantum cipher is described by a set of 2^n unitary encryption operators $\{E_k\}_{k \in \{0,1\}^n}$, acting on a set of m -bit plaintexts and an arbitrary ancilla initially in state $|0\rangle$ such that,

$$(\forall x, x' \in \{0, 1\}^m)[D(\rho_x, \rho_{x'}) < \epsilon(n)].$$

If $\epsilon(n)$ is a negligible function of n we say that the scheme is statistically private.

The total mixture of ciphertexts associated with an ϵ -private (n, m) -quantum cipher with encryption operators $\{E_k\}_{k \in \{0,1\}^n}$ is

$$\xi = \sum_{k \in \{0,1\}^n} 2^{-n} \sum_{x \in \{0,1\}^m} 2^{-m} E_k |x\rangle\langle x| \otimes |0\rangle\langle 0| E_k^\dagger. \quad (2)$$

The next technical Lemma states that the total mixture of any ϵ -private quantum cipher is ϵ -close to any plaintext encryption under a random and private key.

Lemma 1. *Any ϵ -private (n, m) -quantum cipher satisfy that for all $x \in \{0, 1\}^m$, $D(\xi, \rho_x) < \epsilon$.*

Proof. Simply observe that,

$$D(\xi, \rho_x) = D\left(2^{-m} \sum_{y \in \{0,1\}^m} \rho_y, \rho_x\right) \leq \sum_{y \in \{0,1\}^m} 2^{-m} D(\rho_y, \rho_x) < \epsilon,$$

from the convexity of $D(\cdot, \cdot)$ and the ϵ -privacy of the quantum cipher. \square

2.3 Mutually Unbiased Bases

A set $\mathcal{B}_n = \{B_1, \dots, B_{2^t}\}$ of 2^t orthonormal bases in a Hilbert space of dimension 2^n is said to be *mutually unbiased* (we abbreviate mutually unbiased bases set as MUBS) if for all $|u\rangle \in B_i$ and $|v\rangle \in B_j$ for $i \neq j$, we have $|\langle u|v\rangle| = 2^{-n/2}$. Wootters and Fields[20] have shown that there are MUBSS of up to $2^n + 1$ bases in a Hilbert space of dimension 2^n , and such sets are *maximum*. They also give a construction for a maximal MUBS in Hilbert spaces of prime-power dimensions. For $\mathcal{B}_n = \{B_b\}_{b \in \{0,1\}^t}$ a MUBS, $w \in \{0,1\}^n$, and $b \in \{0,1\}^t$, we denote by $|v_w^{(b)}\rangle$ the w -th state in basis $B_b \in \mathcal{B}_n$.

Lawrence, Brukner, and Zeilinger[11] introduced an alternative construction for maximal MUBSS based on algebra in the Pauli group. Their construction plays an important role in the security analysis of our QKRS. The method for constructing a maximal MUBS in \mathcal{H}_n relies on a special partitioning of all Pauli operators in \mathcal{H}_n . These operators live in a vector space of dimension 4^n . Let $\Sigma = \{\sigma_x, \sigma_y, \sigma_z, \sigma_1\}$ (where $\sigma_1 = \mathbb{1}_1$) be the set of Pauli operators in \mathcal{H}_1 . This set forms a basis for all one-qubit operators. A basis for operators on n qubits is constructed as follows for $i \in \{0, \dots, 4^n - 1\}$:

$$O_i = \sigma_{\mu(1,i)}^1 \sigma_{\mu(2,i)}^2 \cdots \sigma_{\mu(n,i)}^n = \prod_{k=1}^n \sigma_{\mu(k,i)}^k, \quad (3)$$

such that $\sigma_{\mu(k,i)}^k$ is an operator in Σ acting only on the k -th qubit. We use the convention $O_0 = \mathbb{1}_n$. The action of O_i on the k -th qubit is $\sigma_{\mu(k,i)}$ where $\mu(k,i) \in \{x, y, z, \mathbb{1}\}$. The basis described in (3) is orthogonal, $\text{tr}(O_i O_j) = 2^n \delta_{i,j}$ where $i = j$ means that $\mu(k,i) = \mu(k,j)$ for any qubit k . Every Pauli operator O_i is such that $O_i^2 = \mathbb{1}_n$. Apart from the identity $\mathbb{1}_n$, all O_i 's are traceless and have eigenvalues ± 1 .

In [11], it is first shown how to partition the set of $4^n - 1$ non-trivial Pauli operators $\{O_i\}_{i=1}^{4^n-1}$ into $2^n + 1$ subsets, each containing $2^n - 1$ commuting members. Second, each such partitioning is shown to define a maximal MUBS. Let us denote by $P_\beta^b = |v_\beta^{(b)}\rangle\langle v_\beta^{(b)}|$ the projector on the β -th vector in basis B_b . Saying that $\mathcal{B}_n = \{B_i\}_i$ is a MUBS means that $\text{tr}(P_\alpha^a P_\beta^b) = 2^{-n}$ when $a \neq b$ and $\text{tr}(P_\beta^b P_{\beta'}^b) = \delta_{\beta,\beta'}$. Let $(\varepsilon_{b,\beta})_{b,\beta}$ be a $2^n \times 2^n$ matrix consisting of orthogonal rows, one of which is all $+1$, and the remaining ones all contain as many $+1$ as -1 . The b -th partition contains Pauli operators $\{O_\beta^b\}_{\beta=1}^{2^n-1}$ such that

$$O_\beta^b = \sum_{\alpha=1}^{2^n} \varepsilon_{\beta,\alpha} P_\alpha^b. \quad (4)$$

In the following, $(\varepsilon_{\beta,\alpha})_{\beta,\alpha}$ will always denote the operator $2^{n/2}H^{\otimes n}$ where $H^{\otimes n}$ is the n -qubit Hadamard transform (i.e. $\varepsilon_{\beta,\alpha} = (-1)^{\beta \cdot \alpha}$).

The number of partitions $\{O_{\beta}^b\}_{\beta}$ defined by (4) is $2^n + 1$ when constructed from a maximal MUBS. Each partition contains $2^n - 1$ operators after discarding the identity (they all contain the identity). Each of these operators is traceless and has ± 1 eigenvalues as for the Pauli operators. It is easy to verify that for $a \neq b$,

$$\mathrm{tr}(O_{\alpha}^a O_{\beta}^b) = \sum_{\mu,\nu} \varepsilon_{\alpha,\mu} \varepsilon_{\beta,\nu} \mathrm{tr}(P_{\mu}^a P_{\nu}^b) = 0. \quad (5)$$

Moreover,

$$\mathrm{tr}(O_{\beta}^b O_{\beta'}^b) = \sum_{\mu,\nu} \varepsilon_{\beta,\mu} \varepsilon_{\beta',\nu} \mathrm{tr}(P_{\mu}^b P_{\nu}^b) = \sum_{\mu} \varepsilon_{\beta,\mu} \varepsilon_{\beta',\mu} = 2^n \delta_{\beta,\beta'}. \quad (6)$$

It follows from (5) and (6) that all operators in (4) are unitarily equivalent to Pauli operators. This essentially shows that partitioning the Pauli operators the way we want is always possible.

It remains to argue that any such partitioning defines a maximal MUBS. Notice that partition $\{O_1^b, \dots, O_{2^n-1}^b\}$ (i.e. without the identity O_0^b) defines a unique basis $\{P_\beta^b\}_\beta$ where

$$P_\beta^b = 2^{-n} \sum_{\mu} \varepsilon_{\mu,\beta} O_\mu^b. \quad (7)$$

It is not difficult to verify that $\text{tr}(P_\beta^b P_{\beta'}^b) = \delta_{\beta,\beta'}$ and for $a \neq b$, $\text{tr}(P_\beta^b P_\alpha^a) = 2^{-n}$ thus leading to a maximal MUBS.

In other words, there is a one-to-one correspondence between maximal MUBSs and the partitionings $\{\{O_\beta^b\}_\beta\}_b$ of the $4^n - 1$ Pauli operators (except the identity), acting on n qubits, into $2^n + 1$ partitions $\{O_\beta^b\}_\beta$ of $2^n - 1$ commuting members. Each partition is a subgroup of the n -qubit Pauli group and is generated by n of these operators. Any Pauli operator commutes with all other operators in the partition in which it is, and anti-commutes with exactly half of the operators, including the identity, in all other partitions. See [11] for more details.

2.4 The W_n -Cipher

In [8], quantum ciphers based on MUBSS were introduced and studied with respect to their secret-key uncertainty against known-plaintext attacks. Our QKRS, presented in Sect. 5.1, uses one of these ciphers, the W_n -cipher, as its main building block. The W_n -cipher is a $(2n, n)$ -quantum cipher, that is, it encrypts n -bit classical messages with the help of a $2n$ -bit secret-key. The W_n -cipher enjoys perfect privacy when the secret-key is perfectly private. It is easy to verify that the cipher is ϵ -private if the secret-key is only ϵ -close to uniform[17].

Let $\mathcal{B}_n = \{B_b\}_{b \in \{0,1\}^n}$ be a MUBS of cardinality 2^n for \mathcal{H}_n . Remember that $|v_w^{(b)}\rangle$ denotes the w -th basis state in basis $B_b \in \mathcal{B}$. The secret-key k for the W_n -cipher is conveniently written as $k = (z, b)$ where $z, b \in_R \{0,1\}^n$. Encryption according secret-key $k = (z, b)$ of message $x \in \{0,1\}^n$ consists in preparing the following state:

$$E_k|x\rangle = E_{(z,b)}|x\rangle = \left|v_{x \oplus z}^{(b)}\right\rangle \in B_b.$$

In other words, the encryption process first one-time pad message x with key z before mapping the resulting state to basis B_b . Encryption and decryption can be performed efficiently on a quantum computer[20,8].

3 Key-Recycling Schemes

A QKRS is an encryption scheme with authentication. In addition, there are two key-recycling mechanisms, $R_{\text{ok}}^{n,s}$ and $R_{\text{no}}^{n,t}$, allowing one to recycle part of the secret-key shared between Alice and Bob in case where the authentication succeeds and fails respectively. We model the recycling mechanism by privacy amplification. That is, $R_{\text{ok}}^{n,s}$ and $R_{\text{no}}^{n,t}$ are classes of hashing functions mapping the current key $k \in \{0, 1\}^n$ into a recycled key \hat{k} of length s and t respectively. In order to apply privacy amplification, an *authentic classical feedback channel* is necessary for announcing Bob's random recycling function $R \in_R R_{\text{ok}}^{n,s}$ or $R \in_R R_{\text{no}}^{n,t}$ depending on the outcome of authentication. Alice and Bob then compute $\hat{k} = R(k)$ as their recycled secret-key. We do not allow further interaction between Alice and Bob since otherwise quantum key distribution could take place between them allowing not only to recycle their secret-key but even to increase its length. Key-recycling should be inherently non-interactive from Bob to Alice since the authentication outcome should anyway be made available to Alice. For simplicity, we assume that the classical feedback channel between Bob and Alice is authenticated. In general, a small secret key could be used for providing classical message-authentication on the feedback channel.

Definition 2. A (n, m, s, t) -QKRS is defined by a pair $(\mathfrak{C}^{m,n}, (R_{\text{ok}}^{n,s}, R_{\text{no}}^{n,t}))$ where

- $\mathfrak{C}^{m,n}$ is a (m, n) -quantum cipher, and
- $(R_{\text{ok}}^{n,s}, R_{\text{no}}^{n,t})$ is a key-recycling mechanism.

In this paper, the privacy of the recycled key is characterized by its distance from uniform. In [17], it is shown that when the distance is negligible, the key behaves as a perfectly private key except with negligible probability. It follows that the application is composable provided the adversary is static[14,17,3].

For a QKRS to be secure, we require that even knowing the plaintext, the function R , and the authentication outcome, the adversary's view about the recycled key is at negligible distance from uniform. This should hold except for a negligible number of functions in $\mathbf{R}_{\text{ok}}^{n,s}$ and $\mathbf{R}_{\text{no}}^{n,t}$. Security against known plaintext attacks is an important property of good key-recycling mechanisms. Otherwise, extra conditions on the *a posteriori* probability distribution over plaintexts have to be enforced. In particular a recycled key could be compromised if a previous plaintext gets revealed to the adversary.

The adversary's view typically changes depending on whether the authentication succeeds or fails. Let \mathcal{A}_{ok} (resp. \mathcal{A}_{no}) be the event consisting in a successful (resp. unsuccessful) authentication. Conditioned on \mathcal{A}_{ok} , the adversary should have access only to very limited amount of information about the secret-key. The better the authentication scheme is, the more key material the recycling mechanism can handle. When \mathcal{A}_{no} occurs, however, the adversary may hold the entire cipherstate. Let K be the random variable for the secret-key. Let $\rho(x)$ be the random state corresponding to the adversary's view on an encryption of classical message x using a random key. We denote by $[\rho_{ok}(x)] = [\rho(x)|\mathcal{A}_{ok}]$ and $[\rho_{no}(x)] = [\rho(x)|\mathcal{A}_{no}]$ the random state $\rho(x)$ conditioned on the event \mathcal{A}_{ok} and \mathcal{A}_{no} respectively.

Definition 3. A key-recycling mechanism $(R_{\text{ok}}^{n,s}, R_{\text{no}}^{n,t})$ is $(\delta_{\text{ok}}, \delta_{\text{no}})$ -indistinguishable if for all $x \in \{0, 1\}^m$:

1. $d(R(K)|\rho_{\text{ok}}(x) \otimes \{R\}) \leq \delta_{\text{ok}}$ (where $R \in_R R_{\text{ok}}^{n,s}$), and
2. $d(R(K)|\rho_{\text{no}}(x) \otimes \{R\}) \leq \delta_{\text{no}}$ (where $R \in_R R_{\text{no}}^{n,t}$).

For $\delta_{\text{ok}}, \delta_{\text{no}}$ negligible functions of n , we say that the key-recycling mechanism is statistically indistinguishable. The class of key-recycling functions $R_{\text{ok}}^{n,s}$ or $R_{\text{no}}^{n,t}$ is said to be δ -indistinguishable if condition 1 or 2 respectively holds relative to δ .

Finally, a QKRS is secure if it is a private encryption scheme together with a statistically indistinguishable key-recycling mechanism. In general,

Definition 4. A (n, m, s, t) -QKRS defined by $(\mathfrak{E}^{m,n}, (\mathbf{R}_{\text{ok}}^{n,s}, \mathbf{R}_{\text{no}}^{n,t}))$ is $(\epsilon, \delta_{\text{ok}}, \delta_{\text{no}})$ -secure if

1. $\mathfrak{E}^{m,n}$ is ϵ -private,
2. when no eavesdropping occurs the key-recycling mechanism $\mathbf{R}_{\text{ok}}^{n,s}$ is used, and
3. $(\mathbf{R}_{\text{ok}}^{n,s}, \mathbf{R}_{\text{no}}^{n,t})$ is a $(\delta_{\text{ok}}, \delta_{\text{no}})$ -indistinguishable key-recycling mechanism.

If the scheme is such that $\epsilon, \delta_{\text{ok}}$, and δ_{no} are all negligible functions of n then we say that the scheme is statistically secure.

The efficiency of a QKRS is characterized by n, s and t . When authentication succeeds $n - s$ bits of secret-key must be thrown away while, when authentication fails, $n - t$ have to be discarded. Clearly, any purely classical key-recycling scheme must have $s, t \leq n - m$. This does not have to hold for quantum schemes. However, we show next that quantum schemes suffer the same restrictions as classical ciphers when authentication fails.

4 Upper Bound on Key-Recycling

In this section, we show that any statistically secure QKRS must discard as many key-bits as the length of the plaintext (minus one bit) when the authentication fails. In other words, when authentication fails no QKRS does better than the classical one-time-pad.

When authentication fails, the adversary may have kept the entire ciphertext and may know the plaintext $x \in \{0, 1\}^m$. On the other hand, condition 2 in Definition 3 requires that the key-recycling mechanism satisfies $d(R(K)|\rho_{\text{no}}(x) \otimes \{R\}) \leq \delta(n)$ where $\delta(n)$ is negligible and $R \in_R \mathcal{R}_{\text{no}}^{n,t}$. Using (1), it follows that

$$D([\{R(K)\} \otimes \rho_{\text{no}}(x) \otimes \{R\}], [\{U\}] \otimes [\rho_{\text{no}}(x) \otimes \{R\}]) \leq \delta(n). \quad (8)$$

The density operator $\rho_{\text{no}}(\hat{k}, x, R) = [\rho_{\text{no}}(x) | R(K) = \hat{k}]$ corresponds to the adversary's view when the plaintext is x , the recycled key is $\hat{k} \in \{0, 1\}^t$, and the privacy amplification function is $R \in \mathbf{R}_{\text{no}}^{n,t}$. We have that,

$$\rho_{\text{no}}(\hat{k}, x, R) = \sum_{k: R(k) = \hat{k}} \frac{1}{\#R^{-1}(\hat{k})} E_k |x\rangle\langle x| \otimes |0\rangle\langle 0| E_k^\dagger. \quad (9)$$

For convenience, we define $\rho_{\text{no}}(\hat{k}, x) = \frac{1}{\#\mathbb{R}_{\text{no}}^{n,t}} \sum_{R \in \mathbb{R}_{\text{no}}^{n,t}} \rho_{\text{no}}(\hat{k}, x, R) \otimes |R\rangle\langle R|$. If a key-recycling scheme is statistically indistinguishable then for a negligible function $\delta(n)$,

$$\delta(n) \geq d(R(K) | \rho_{\text{no}}(x) \otimes \{R\}) \quad (10)$$

$$= D \left(\sum_{\hat{k}} p_{\hat{K}}(\hat{k}) \left| \hat{k} \right\rangle\left\langle \hat{k} \right| \otimes \rho_{\text{no}}(\hat{k}, x), \mathbb{I}_t \otimes \sum_{\hat{k}} p_{\hat{K}}(\hat{k}) \rho_{\text{no}}(\hat{k}, x) \right) \quad (11)$$

$$\geq \frac{1}{\#\mathbb{R}_{\text{no}}^{n,t}} \sum_R \sum_{\hat{k}} 2^{-n} \#R^{-1}(\hat{k}) D(\rho_{\text{no}}(\hat{k}, x, R), \rho_x), \quad (12)$$

where (10) follows by definition of statistical indistinguishability, and (11) is obtained using (8) and (9). The last step follows from the fact that $D(\rho, \sigma) = \max_{\{E_m\}_m} D(p(m), q(m))$ where the maximum is computed over all POVMs $\{E_m\}_m$ and $p(m) = \text{tr}(\rho E_m)$, $q(m) = \text{tr}(\sigma E_m)$ are probability distributions for the outcomes of $\{E_m\}_m$ when applied to ρ and σ respectively (see for example Theorem 9.1 in [15]). In order to get (12) from (11) one only has to consider a POVM that first measures R and \hat{k} before performing the POVM $\{E'_m\}_m$ (depending on R and \hat{k}) on the residual state that satisfies $D(\rho_{\text{no}}(\hat{k}, x, R), \rho_x) = d(p'(m), q'(m))$.

It can be shown that for $t \geq n - m + 2$, (12) implies the existence of $R \in \mathbb{R}_{\text{no}}^{n,t}$ and $\hat{k}_0 \in \{0, 1\}^t$ such that $\#R^{-1}(\hat{k}_0) \leq 2^{m-1}$ and $D(\rho_{\text{no}}(\hat{k}_0, x, R), \rho_x) \leq c$ for any constant $0 < c \leq 1$. Moreover, since the cipher is statistically private, there exists a negligible function $\epsilon(n)$ such that,

$$D(\rho_{\text{no}}(\hat{k}_0, x, R), \xi) \leq D(\xi, \rho_x) + D(\rho_x, \rho_{\text{no}}(\hat{k}_0, x, R)) \leq \epsilon(n) + c. \quad (13)$$

On the other hand, an argument along the lines of the proof of Lemma IV.3.2 in [5] allows us to conclude that when $\#R^{-1}(\hat{k}_0) \leq 2^{m-1}$, $D(\rho_{\text{no}}(\hat{k}_0, x, R), \xi) \geq 1/2$ which contradicts (13) when $c < 1/2$ and $\epsilon(n)$ is negligible (see Lemma 3 in [9]). Next Theorem, proven in [9], follows:

Theorem 1 (Key-Recycling Bound). *Any statistically secure (n, m, s, t) -QKRS is such that $t \leq n - m + 1$.*

We believe that a more careful analysis would show that statistically secure (n, m, s, t) -QKRS must satisfy $t \leq n - m$. Theorem 1 implies that in order to recycle more secret-key bits than any classical scheme, quantum ciphers must provide authentication. It is only when the authentication succeeds that a QKRS may perform better than classical ones.

5 A Near Optimal Quantum Key-Recycling Scheme

We introduce a QKRS, called W_nC_m , that recycles an almost optimal amount of key material. Moreover, the key-recycling mechanism does not use privacy amplification. Deterministic functions are sufficient to guarantee the statistical indistinguishability of the recycled key. The scheme is introduced in Sect. 5.1. In Sect. 5.2 we present an EPR-version of the scheme and we prove it secure. In Sect. 5.3 we reduce the security of W_nC_m to that of the EPR-version.

5.1 The Scheme

The W_nC_m -cipher encrypts a message together with its Wegman-Carter one-time authentication tag[6] using the W_n -cipher[8]. We need an authentication code constructed from XOR-universal classes of hash-functions:

Definition 5 ([6]). *An XOR-universal family of hash-functions is a set of functions $H_{m,\mu} = \{h_u : \{0,1\}^m \rightarrow \{0,1\}^\mu\}_u$ such that for all $a \neq b \in \{0,1\}^m$ and all $x \in \{0,1\}^\mu$, $\#\{h \in H_{m,\mu} | h(a) \oplus h(b) = x\} = \frac{\#H_{m,\mu}}{2^\mu}$.*

There exists an XOR-universal class of hash-functions $H_{m,\mu}^\oplus$ (for any $m \geq \mu$) that requires only m bits to specify and such that picking a function at random can be done efficiently.

For the transmission of m -bit messages, W_nC_m requires Alice and Bob to share a secret-key of size $N = 2n + m$ bits where $n = m + \ell(m)$, and $\ell(m) \in \Omega(m)$ is the size of the Wegman-Carter authentication tag. We denote secret-key k by the triplet: $k = (z, b, u)$ where $z, b \in \{0, 1\}^n$ is the key for the W_n -cipher and $u \in \{0, 1\}^m$ is the description of a random function $h_u \in H_{m, \ell(m)}^\oplus$. Encrypting message $x \in \{0, 1\}^m$ is performed by first computing the Wegman-Carter one-time authentication tag $h_u(x)$. The message $(x, h_u(x)) \in \{0, 1\}^n$ is then encrypted using the W_n -cipher with secret-key (z, b) . Bob decrypts the W_n -cipher and verifies that a message of the form $(x, h_u(x))$ is obtained. Bob announces to Alice the outcome of the authentication using the authenticated feedback channel. When it is successful, Alice and Bob recycle the whole secret-key. If the authentication fails then Alice and Bob throw away the one-time-pad z . The remaining part (b, u) is entirely recycled. In other words, $R_{\text{ok}}^{N, s}$ is the identity with $s = N$ and $R_{\text{no}}^{N, t}$ is deterministic with $t = N - n = N - m - \ell(m)$.

It is almost straightforward to show that our key-recycling function is perfectly indistinguishable when authentication fails.

Lemma 2. *Let $N = 2n + m$ where $n = m + \ell(m), \ell(m) > 0$ be the key-length used in W_nC_m and let $R(z, b, u) = (b, u)$ for $z, b \in \{0, 1\}^n$ and $u \in \{0, 1\}^m$. The key-recycling mechanism $R_{\text{no}}^{N, N-n} = \{R\}$ is 0-indistinguishable.*

Proof. Since $\rho_{\text{no}}((b, u), x, R) = \mathbb{I}_n = \rho_{\text{no}}((b', u'), x, R)$ for all $(b, u), (b', u')$, and x , it easily follows that $d(R(K)|\rho_{\text{no}}(x) \otimes \{R\}) = 0$. \square

Since W_nC_m encrypts m -bit messages and recycles $N - n$ bits of key, the scheme is sub-optimal according Theorem 1. In the next sections, we see that W_nC_m remains statistically secure for any $\ell(m) \in \Omega(m)$. It follows that although sub-optimal, W_nC_m is *nearly* optimal.

Private-Key: $(z, b, u) \in_R \{0, 1\}^{2n+m}$ where $n = m + \ell(m)$.

1. Alice creates the message $c = (x, h_u(x))$ where $h_u \in H_{m, \ell(m)}^\oplus$. She then encrypts this message with key (z, b) according to the W_n -cipher.
2. Bob decodes the received W_n -cipher with key (z, b) and gets $c' = (x', t')$. He then verifies the authentication tag $t' = h_u(x')$. Bob sends the result of the test to Alice through a classical authentic channel.
3. [**Key-Recycling**] If Bob accepts then Alice and Bob recycle the entire key (b, z, u) . If Bob rejects then Alice and Bob recycle (b, u) and throw away $z \in \{0, 1\}^n$.

Fig. 1. The W_nC_m

5.2 An EPR Variant of W_nC_m

We establish the security of the key-recycling mechanism in W_nC_m when the authentication is successful. We prove this case using a Shor-Preskill argument [18] similar to the ones invoked in [16] and [2] for key-recycling and quantum authentication respectively.

We first define a variant of W_nC_m , called $\text{EPR-}W_nC_m$, using EPR-pairs and having access to an additional authenticated and private classical channel. The key-recycling mechanism of $\text{EPR-}W_nC_m$ can be proven secure more easily since it has access to more powerful resources. Second, we show that the security of W_nC_m follows from the security of $\text{EPR-}W_nC_m$.

In $\text{EPR-W}_n\mathcal{C}_m$, Alice and Bob initially share an n -bit key b , and an m -bit key u . They agree on 2^n mutually unbiased bases in \mathcal{H}_n , and a family of XOR-universal hash-functions $H_{m,\mu}^\oplus = \{h_u\}_{u \in \{0,1\}^m}$. As for $\text{W}_n\mathcal{C}_m$, the key b is used to select in which of the bases of the MUBS the encryption will take place. The key u indicates the selection of the hash-function for authentication. The key z in $\text{EPR-W}_n\mathcal{C}_m$ is not shared beforehand but will be implicitly generated by measuring the shared EPR-pairs. This corresponds to refreshing z before each round of $\text{EPR-W}_n\mathcal{C}_m$.

In order for Alice to send classical message $x \in \{0,1\}^m$ to Bob, Alice and Bob proceeds as described in Fig. 2. The key-recycling mechanism of $\text{EPR-}\mathcal{W}_n\mathcal{C}_m$ only takes place when authentication succeeds. The quantum transmission in $\mathcal{W}_n\mathcal{C}_m$ is replaced by transmitting half of a maximally entangled state consisting of n EPR-pairs.

$$|\Psi\rangle = \sum_{x \in \{0,1\}^n} 2^{-n/2} |x\rangle^A |x\rangle^B = \sum_{x \in \{0,1\}^n} 2^{-n/2} |\xi_x^{(b)}\rangle^A |v_x^{(b)}\rangle^B, \quad (14)$$

for some orthonormal basis $\{|\xi_x^{(b)}\rangle\}_x$.

Private-Key: $(b, u) \in_R \{0, 1\}^{n+m}$.

1. Alice prepares the n EPR-pairs in state $|\Psi\rangle^{AB}$.
2. Alice sends the B -register to Bob.
3. Bob acknowledges receiving the state using the classical authentic feedback channel.
4. Alice measures her A -register in basis $\{|\xi_c^{(b)}\rangle\}_{c \in \{0,1\}^n}$ (See (14)). On classical outcome c , she computes $z := c \oplus (x, h_u(x))$.
5. Alice sends z to Bob through the additional private and authenticated classical channel.
6. Bob measures his B -register in the b -th basis of the MUBS, gets outcome c' , and computes $(x', t') = c' \oplus z$. Bob verifies that $t' = h_u(x')$ and announces the result to Alice through the classical authenticated feedback channel.
7. If Bob accepts, Alice and Bob recycle the whole key (b, u) .

Fig. 2. The $\text{EPR-}W_n\mathcal{C}_m$ -cipher using an extra private and authentic classical channel

Any trace-preserving operator the adversary can apply to Bob's half EPR-pairs can be described in terms of the 4^n Pauli operators,

$$\hat{\rho} = \mathcal{E}(|\Psi\rangle\langle\Psi|) = \sum_{i=0}^{4^n-1} \sum_{j=0}^{4^n-1} c_i c_j (\mathbb{1}_n \otimes O_i) |\Psi\rangle\langle\Psi| (\mathbb{1}_n \otimes O_j)^\dagger, \quad (15)$$

where $O_0 = \mathbb{1}_n$. We can split (15) into the case where the error leaves the state untouched, and the case where the state is changed

$$\hat{\rho} = |c_0|^2 |\Psi\rangle\langle\Psi| + (1 - |c_0|^2) \rho_E^{b,u}, \quad (16)$$

where $\rho_E^{b,u} = \sum_{(i,j) \neq (0,0)} \frac{c_i c_j}{(1 - |c_0|^2)} (\mathbb{1}_n \otimes O_i) |\Psi\rangle\langle\Psi| (\mathbb{1}_n \otimes O_j)^\dagger$, and $|c_0|^2$ is the probability that the state is left unchanged by \mathcal{E} .

The idea behind the security of the key-recycling mechanism is that an eavesdropper, performing any non-trivial action upon Bob's system, will fail authentication with high probability. Any eavesdropping strategy that remains undetected with a *not too small* probability is such that $|c_0|^2$ is at negligible distance from 1. This means that the ciphertext will be left untouched with probability essentially 1. In other words the probability of being detected is closely related to $1 - |c_0|^2$.

The probability that Bob will accept the authentication tag, when Alice and Bob share key (b, u) can be expressed by the observable projecting onto the space of states where Alice has her untouched EPR-halves, and Bob has anything that passes the authentication test:

$$\Pi_{\text{Acc}}^{b,u} = \sum_{z \in \{0,1\}^n} \sum_{\hat{x} \in \{0,1\}^m} \left| \xi_{e_{z,u}(x)}^{(b)} \right\rangle \left\langle \xi_{e_{z,u}(x)}^{(b)} \right| \otimes \left| v_{e_{z,u}(\hat{x})}^{(b)} \right\rangle \left\langle v_{e_{z,u}(\hat{x})}^{(b)} \right|, \quad (17)$$

where $e_{z,u}(x) = z \oplus (x, h_u(x))$. The probability that Bob will accept the authentication, when using key (b, u) , is $p_{\text{Acc}}^{b,u} = \text{tr}(\Pi_{\text{Acc}}^{b,u} \hat{\rho})$.

As mentioned in Sect. 2.3, all $4^n - 1$ Pauli operators (excluding the identity) are partitioned into $2^n + 1$ sets, each containing $2^n - 1$ commuting members. Each operator, O_i , appearing in (15), will be in one of the $2^n + 1$ partitions (i.e. which each forms a subgroup). In the partition or basis where an error operator O_i belongs, its action will leave all cipherstates unchanged. For each other 2^n basis b , O_i will anti-commute with exactly half the operators (including the identity). This means that in basis b , the action of O_i permutes the basis vectors. Since this permutation is independent of the authentication code, we can show that the probability for O_i to remain undetected is negligible when the class of Wegman-Carter authentication functions is XOR-universal. Let $\hat{\rho}_{\text{Acc}}^{b,u}$ be the normalized state conditioned on \mathcal{A}_{ok} defined as,

$$\hat{\rho}_{\text{Acc}}^{b,u} = \frac{\Pi_{\text{Acc}}^{b,u} \hat{\rho} \Pi_{\text{Acc}}^{b,u}}{\text{tr} \left(\Pi_{\text{Acc}}^{b,u} \hat{\rho} \right)}. \quad (18)$$

We are going to estimate the average fidelity² of $\hat{\rho}_{\text{Acc}}^{b,u}$ to the ideal state $|\Psi\rangle\langle\Psi|$. To do so we split $\hat{\rho}$ according to (16) and use the concavity of the fidelity, $F(\hat{\rho}_{\text{Acc}}^{b,u}, |\Psi\rangle\langle\Psi|) \geq \frac{|c_0|^2}{p_{\text{Acc}}^{b,u}}$. Applying (16) to $p_{\text{Acc}}^{b,u}$, gives us

$$F(\hat{\rho}_{\text{Acc}}^{b,u}, |\Psi\rangle\langle\Psi|) \geq \frac{|c_0|^2}{|c_0|^2 + (1 - |c_0|^2) \text{tr} \left(\Pi_{\text{Acc}}^{b,u} \rho_E^{b,u} \right)}.$$

To lower bound the average fidelity, $\sum_{b,u} 2^{-n-m} F(\hat{\rho}_{\text{Acc}}^{b,u}, |\Psi\rangle\langle\Psi|)$. We split the sum into keys (bases and authentication keys) for which $\text{tr}(\Pi_{\text{Acc}}^{b,u} \rho_E^{b,u})$ is small, and keys for which this probability is large. We know from the previous argument, that the probability of accepting a non-trivial error will be small in most bases, and indeed the terms with $\text{tr}(\Pi_{\text{Acc}}^{b,u} \rho_E^{b,u})$ negligible compared to $|c_0|^2$ give the main contribution to the fidelity.

² Where the fidelity $F(\hat{\rho}_{\text{Acc}}^{b,u}, |\Psi\rangle\langle\Psi|) = \langle\Psi|\hat{\rho}_{\text{Acc}}^{b,u}|\Psi\rangle$.

In summary, an undetected attack is almost always trivial since it corresponds to the case where no eavesdropping occurred. Next Theorem, proven in [9], gives the desired result.

Theorem 2. *For all adversary strategies for which $p_{Acc} \geq 2^{-(n-m-2)/2+1}$,*

$$\sum_{b \in \{0,1\}^n} \sum_{u \in \{0,1\}^m} 2^{-n-m} F\left(\hat{\rho}_{Acc}^{b,u}, |\Psi\rangle\langle\Psi|\right) \geq 1 - 2^{-\frac{n-m-2}{4}+1},$$

provided n is sufficiently large.

Let $\rho_{\text{ok}}^{\text{epr}}(x)$ be the random state corresponding to the adversary's view in $\text{EPR-}\mathcal{W}_n\mathcal{C}_m$ given \mathcal{A}_{ok} . Let $K = (B, U, Z)$ be the random variable describing the key $(b, u) \in \{0, 1\}^n \times \{0, 1\}^m$, and $z \in \{0, 1\}^n$ computed from the measurement outcome. Using the same line of arguments as [3] (for completeness, the proof can be found in [9]), Theorem 2 implies that:

Theorem 3. *For all adversary strategies for which $p_{\text{Acc}} \geq 2^{-(n-m-2)/2+1}$,*

$$d(K | \rho_{\text{ok}}^{\text{epr}}(x) \otimes \{R\}) \leq 2^{-\frac{(n-m-2)}{8}+1},$$

provided n is sufficiently large.

5.3 Back to W_nC_m

We now show that Theorem 3 also applies to W_nC_m . Similarly to other Shor-Preskill arguments[18,2,16], we transform $EPR-W_nC_m$ into W_nC_m by simple modifications leaving the adversary's view unchanged.

In Step 4 of $EPR-W_nC_m$, Alice measures her part of the entangled pair in order to extract $c \in \{0, 1\}^n$. Instead, she could have measured already in Step 1 since the measurement commutes with everything the adversary and Bob do up to Step 4. Measuring half the EPR-pairs immediately after creating them is equivalent to Alice preparing $c \in_R \{0, 1\}^n$ before sending $|v_c^{(b)}\rangle$ in Step 2.

Instead of picking $c \in_R \{0, 1\}^n$ in Step 1, Alice could choose $z \in_R \{0, 1\}^n$ at random before sending $|v_{z \oplus (x, h_u(x))}^{(b)}\rangle$ to Bob. All these modifications change nothing to the adversary's view.

Now, sending z through the private and authenticated classical channel in Step 5 becomes unnecessary if Alice and Bob share z before the start of the protocol (thus making z part of the key). We have now removed the need for the private and authenticated classical channel.

The resulting protocol is such that Bob first acknowledges receiving the cipher, then measures it, and finally replies with either accept or reject. The acknowledgment of Step 3 is unnecessary and can safely be postponed to Bob's announcement in Step 6. The $\text{EPR-}W_nC_m$ -cipher has now been fully converted into the W_nC_m -cipher without interfering with the eavesdropper's view. It follows directly that Theorem 3 also applies to W_nC_m .

Theorem 3 shows that one use of the W_nC_m -cipher leaves the secret-key at negligible distance from uniform when it was initially 0-indistinguishable. In general, if a random variable K is at distance no more than ϵ from uniform then K behaves exactly like a uniform random variable except with probability at most ϵ [17]. Our main result follows:

Theorem 4 (Main Result). *Let $n = m + \ell(m)$. For all adversary strategies the W_nC_m -cipher used with an initial ϵ -indistinguishable private-key satisfies,*

1. *either $d(K|\rho_{\text{ok}}(x) \otimes \{R\}) \leq \epsilon + 2^{-\frac{\ell(m)-2}{8}+1}$ or $p_{\text{Acc}} \leq 2^{-(\ell(m)-2)/2+1}$,*
2. *$d(K|\rho_{\text{no}}(x) \otimes \{R\}) \leq \epsilon$,*

provided n is sufficiently large.

In other words, the key-recycling mechanism is statistically indistinguishable when $\ell(m) \in \Omega(m)$. It follows that, when starting from a statistically indistinguishable secret-key, key-recycling can take place exponentially many times without compromising the statistical indistinguishability of the resulting key. As mentioned in Sect. 3, Theorem 4 and the discussion in [17] imply that the W_nC_m -cipher is universally composable against static adversaries.

6 Conclusion and Open Questions

We have shown that the $W_n C_m$ -cipher is an almost optimal key-recycling cipher with one-bit feedback. There are many possible improvements of our scheme. In this paper, we assume noiseless quantum communication. This is of course an unrealistic assumption. Our scheme can easily be made resistant to noise by encoding the quantum cipher using a quantum error-correcting code. Since a quantum error-correcting code is also a secret-sharing[7], it can be shown that when authentication succeeds almost no information about the cipherstate is available to the eavesdropper. On the other hand, if the eavesdropper gains information about the cipherstate then authentication will fail similarly to the case where no error-correction is used.

It would be interesting to show that the key recycling bound(i.e. Theorem 1) can be improved to $t \leq n - m$ (instead of $n - m + 1$) as for classical schemes. It is an open question whether there exists a QKRS achieving this upper bound.

It is also possible to allow for more key-recycling mechanisms associated to different output values for the authentication process. Such a generalized scheme would allow to recycle key-material as a function of the adversary's available information but would require more than one-bit feedback.

It is easy to see that the W_nC_m -cipher can be used as a re-usable quantum authentication scheme when authentication succeeds. Our construction (using MUBSS) is different than the ones based on purity testing codes[2] and may be of independent interest.

THE
END