On the Key-Uncertainty of Quantum Ciphers and the Computational Security of One-way Quantum Transmission

Ivan Damgård, Thomas Pedersen*, and Louis Salvail*

BRICS**, FICS***, Dept. of Computer Science, University of Århus,
{ivan|pede|salvail}@brics.dk

Abstract. We consider the scenario where Alice wants to send a secret (classical) n-bit message to Bob using a classical key, and where only one-way transmission from Alice to Bob is possible. In this case, quantum communication cannot help to obtain perfect secrecy with key length smaller then n. We study the question of whether there might still be fundamental differences between the case where quantum as opposed to classical communication is used. In this direction, we show that there exist ciphers with perfect security producing quantum ciphertext where, even if an adversary knows the plaintext and applies an optimal measurement on the ciphertext, his Shannon uncertainty about the key used is almost maximal. This is in contrast to the classical case where the adversary always learns n bits of information on the key in a known plaintext attack. We also show that there is a limit to how different the classical and quantum cases can be: the most probable key, given matching plain- and ciphertexts, has the same probability in both the quantum and the classical cases. We suggest an application of our results in the case where only a short secret key is available and the message is much longer. Namely, one can use a pseudorandom generator to produce from the short key a stream of keys for a quantum cipher, using each of them to encrypt an n-bit block of the message. Our results suggest that an adversary with bounded resources in a known plaintext attack may potentially be in a much harder situation against quantum stream-ciphers than against any classical stream-cipher with the same parameters.

1 Introduction

In this paper, we consider the scenario where Alice wants to send a secret (classical) n-bit message to Bob using an m-bit classical shared key, and where only one-way transmission from Alice to Bob is possible (or at least where interaction is only available with a prohibitively long delay). If interaction had been available, we could have achieved (almost) perfect secrecy using standard quantum

* Part of this research was funded by European project PROSECCO.
** Funded by the Danish National Research Foundation.
*** FICS, Foundations in Cryptography and Security, funded by the Danish Natural Sciences Research Council.

key exchange, even if $m < n$. But with only one-way communication, we need $m \geq n$ even with quantum communication [1].

We study the question of whether there might still be some fundamental differences between the case where quantum as opposed to classical communication is used. In this direction, we present two examples of cryptosystems with perfect security producing n-bit quantum ciphertexts, and with key length $m = n + 1$, respectively $m = 2n$. We show that given plaintext and ciphertext, and even when applying an optimal measurement to the ciphertext, the adversary can learn no more than $n/2$, respectively 1 bit of Shannon information on the key. This should be compared to the fact that for a classical cipher with perfect security, the adversary always learns n bits of information on the key. While proving these results, we develop a method which may be of independent interest, for estimating the maximal amount of Shannon information that a measurement can extract from a mixture. We note that the first example can be implemented without quantum memory, it only requires technology similar to what is needed for quantum key exchange, and is therefore within reach of current technology. The second example can be implemented with a circuit of $O(n^3)$ gates out of which only $O(n^2)$ are elementary quantum gates.

We also discuss the composition of ciphers, i.e., what happens to the uncertainty of keys when the same quantum cipher is used to encrypt several blocks of data using independent keys. This requires some care, it is well known that cryptographic constructions do not always compose nicely in the quantum case. For composition of our ciphers, however, we shows that the adversary's uncertainty about the keys grows linearly with the number of blocks encrypted, and in some cases it can be shown to grow exactly as one would expect classically.

On the other hand, we show that there is a limit to how different the quantum and classical cases can be. Namely, the most probable key (i.e. the min-entropy of the key), given matching plain- and ciphertexts, has the same probability in both cases.

On the technical side, a main observation underlying our results on Shannon key-uncertainty is that our method for estimating the optimal measurement w.r.t. Shannon entropy can be combined with known results on so called entropic uncertainty relations [6, 4, 8] and mutually unbiased bases [9]. We note that somewhat related techniques are used in concurrent independent work by DiVincenzo et al. [3] to handle a different, non-cryptographic scenario.

While we believe the above results are interesting, and perhaps even somewhat surprising from an information theoretic point of view, they have limited practical significance if perfect security is the goal: a key must never be reused, and so we do not really have to care whether the adversary learns information about it when it is used.

However, there is a different potential application of our results to the case where only a short secret key is available, and where no upper bound on the message length is known a priori. In such a case, only computational security is possible and the standard classical way to encrypt is to use a stream-cipher: using a pseudorandom generator, we expand the key into a long random looking keystream, which is then combined with the plaintext to form the ciphertext. The simplest way of doing such a combination is to take the bit-wise XOR of key and plaintext streams. In a known plaintext attack, an adversary will then be able to learn full information on a part of the keystream and can try to analyze it to find the key or guess other parts of the keystream better than at random. In general, any cipher with perfect secrecy, n-bit plain- and ciphertext and m-bit keys can be used: we simply take the next m bits from the keystream and use these as key in the cipher to encrypt the next n bits of the plaintext. It is easy to see that for any classical cipher, if the adversary knows some n-bit block of plaintext and also the matching ciphertext, then he learns n bit of Shannon information on the keystream.

If instead we use quantum communication and one of our quantum ciphers mentioned above, intuition suggests that an adversary with limited resources is in a more difficult situation when doing a known plaintext attack: if measuring the state representing the ciphertext only reveals a small amount of information on the corresponding part of the keystream, then the adversary will need much more known plaintext than in the classical case before being able to cryptanalyze the keystream.

Care has to be taken in making this statement more precise: our results on key uncertainty tell us what happens when keys are random, whereas in this application they are pseudorandom. It is conceivable that the adversary could design a measurement revealing more information by exploiting the fact that the keystream is not truly random. This, however, is equivalent to cryptanalyzing the generator using a *quantum* computation, and is likely to be technologically much harder than implementing the quantum ciphers. In particular, unless the generator is very poorly designed, it will require keeping a coherent state much larger than what is required for encryption and decryption – simply because one will need to involve many bits from the keystream simultaneously in order to distinguish it efficiently from random. Thus, an adversary limited to measurements involving only a *small* number of qubits will simply have to make many such measurements, hoping to gather enough classical information on the keystream to cryptanalyze it. Our results apply to this situation: first, since the adversary makes many measurements, we should worry about what he learns on average, so Shannon information is the appropriate measure. Second, even though the keystream is only pseudorandom, it may be genuinely random when considering only a small part of it (see Maurer and Massey [5]).

In Sect. 9, we prove a lower bound on the amount of known plaintext the adversary would need in order to obtain a given amount of information on the keystream, for a particular type of keystream generator and assuming the size of coherent states the adversary can handle is limited. We believe that quantum communication helps even for more general adversaries and generators. However, quantifying this advantage is an open problem. We stress that our main goal here is merely to point out the potential for improved security against a bounded adversary.

2 Preliminaries

We assume the reader is familiar with the standard notions of Shannon entropy $H(\cdot)$ of a probability distribution, conditional entropy, etc. A related notion that also measures "how uniform" a distribution is, is the so called *min-entropy*. Given a probability distribution $\{p_1, ..., p_n\}$, the min-entropy is defined as

$$H_\infty(p_1, ..., p_n) = -\log_2(max\{p_1, ..., p_n\}) \tag{1}$$

As usual, $H_\infty(X)$ for random variable X is the min-entropy of its distribution. Min-entropy is directly related to the "best guess" probability: if we want to guess which value random variable X will take, the best strategy is to guess at a value with maximal probability, and then we will be correct with probability $2^{-H_\infty(X)}$. Given the value of another random variable Y, we can define $H_\infty(X|Y = y)$ simply as the min-entropy of the distribution of X given that $Y = y$, and similarly to Shannon entropy, we can define $H_\infty(X|Y) = \sum_y Pr(Y = y) \cdot H_\infty(X|Y = y)$.

The min-entropy can be thought of as a worst-case measure, which is more relevant when you have access to only one sample of some random experiment, whereas Shannon entropy measures what happens on average over several experiments. To illustrate the difference, consider the two distributions $(1/2, 1/2)$ and $(1/2, 1/4, 1/4)$. They both have min-entropy 1, even though it intuitively seems there should be more uncertainty in the second case, indeed the Shannon entropies are 1 and 1.5. In fact, we always have $H(X) \geq H_\infty(X)$, with equality if X is uniformly distributed.

3 Classical Ciphers

Consider a classical cryptosystem with n-bit plain and ciphertexts, m-bit keys and perfect secrecy (assuming, of course, that keys are used only once). We identify the cryptosystem with its encryption function $E(\cdot, \cdot)$. We call this an (m, n)-cipher for short.

Definition 1. *Consider an (m, n)-cipher E. We define the* Shannon key-uncertainty *of E to be the amount of Shannon entropy that remains on an m-bit key given n-bit blocks of plain- and ciphertexts, i.e. $H(K|P, C)$, where K, P, C are random variables corresponding to the random choices of key, plaintext and ciphertext blocks for E, and where the key is uniformly chosen. The* min-entropy key-uncertainty *of E is defined similarly, but w.r.t. min-entropy, as $H_\infty(K|P, C)$.*

From the definition, it may seem that the key uncertainties depend on the distribution of the plaintext. Fortunately, this is not the case. The key-uncertainty in the classical case is easy to compute, using the following slight generalization of the classical perfect security result by Shannon:

Proposition 1. *Let E be a cipher with perfect security, and with plaintext, ciphertext and keyspace $\mathcal{P}, \mathcal{C}, \mathcal{K}$, where $|\mathcal{P}| = |\mathcal{C}|$. Furthermore, assume that keys* *are chosen uniformly. For any such cipher, it holds that the distribution of the key, given any pair of matching ciphertext and plaintext is uniform over a set of* $|\mathcal{K}|/|\mathcal{P}|$ *keys.*

Proof. By perfect security, we must have $|\mathcal{K}| \geq |\mathcal{P}|$. Now, let us represent the cipher in a table as follows: we index rows by keys and columns by plaintexts, and we fill each entry in the table with the ciphertext resulting from the key and plaintext on the relevant row and column. Then, since correct decryption must be possible and $|\mathcal{P}| = |\mathcal{C}|$, each ciphertext appears exactly once in each row. Fix any ciphertext c, and let t_c be the number of times c appears in, say, the first column. Since the probability distribution of the ciphertext must be the same no matter the plaintext, c must appear t_c times in every column. Since it also appears in every row, it follows that the length of a column satisfies $|\mathcal{K}| = t_c|\mathcal{P}|$. So $t_c = |\mathcal{K}|/|\mathcal{P}|$ is the same for every c. If we know a matching plaintext/ciphertext pair, we are given some c and a column, and all we know is that the key corresponds to one of the t_c possible rows. The proposition follows. □

Corollary 1. *For any classical* (m, n)*-cipher, both the Shannon- and min-entropy key-uncertainty is* $m - n$ *bits.*

This result shows that there is no room for improvement in classical schemes: the natural constraints on (m, n)-ciphers imply that the key-uncertainty is always the same, once we fix m and n. As we shall see, this is not true for quantum ciphers. Although they cannot do better in terms of min-entropy key uncertainty, they can when it comes to Shannon key-uncertainty.

4 Quantum Ciphers and Min-Entropy Key-Uncertainty

In this section, we consider quantum ciphers which encrypt classical messages using classical keys and produce quantum ciphers.

We model both the encryption and decryption processes by unitary operations on the plaintext possibly together with an ancilla. This is the same model as used in [1], with the restriction that we only encrypt classical messages.

Definition 2 (**(m, n)-quantum cipher**)**.** *A general* (m, n)-quantum cipher *is a tuple* $(\mathcal{P}, \mathcal{E})$*, such that*

- $\mathcal{P} \subseteq \mathcal{H}$ *is a finite set of orthonormal pure-states (plaintexts) in the Hilbert space* $\mathcal{H}$*, and* $\|\mathcal{P}\| = N$ *and* $N = 2^n$*.*
- $\mathcal{E} = \{\mathsf{E}_k : \mathcal{H} \to \mathcal{H} |\ k = 1, \ldots, M\}$ *is a set of unitary operators (encryptions), and* $M = 2^m$*. Decryption using key* k *is performed using* $\mathsf{E}_k^\dagger$*.*

And the following properties hold:

- *Key hiding:* $(\forall k, k' \in \{1, \ldots, M\})$,

$$\sum_{a \in \mathcal{P}} \frac{1}{N} \mathsf{E}_k |a\rangle|0\rangle\langle 0|\langle a| \mathsf{E}_k^\dagger = \sum_{a \in \mathcal{P}} \frac{1}{N} \mathsf{E}_{k'} |a\rangle|0\rangle\langle 0|\langle a| \mathsf{E}_{k'}^\dagger. \tag{2}$$

– *Data hiding:* $(\forall|a\rangle, |b\rangle \in \mathcal{P})$,

$$\sum_{k=1}^{M} \frac{1}{M} \mathsf{E}_k |a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger = \sum_{k=1}^{M} \frac{1}{M} \mathsf{E}_k |b\rangle|0\rangle\langle 0|\langle b|\mathsf{E}_k^\dagger. \tag{3}$$

The key and data hiding properties guarantee that an adversary cannot gain any information about the key and message respectively when an arbitrary ciphertext is seen. In [1], it was shown that data hiding implies that $m \geq n$.

The key hiding property states that an adversary with no information on the message encrypted expects to see the same ensemble no matter what key was used. We denote this ensemble

$$\rho = \sum_{a \in \mathcal{P}} \frac{1}{N} \mathsf{E}_k |a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger, \tag{4}$$

for any $k \in \{1, 2, \ldots, M\}$. As motivation for the key-hiding property, we mention that it is always satisfied if ciphertexts are as short as possible $(dim(\mathcal{H}) = 2^n)$. On the other hand, if the key-hiding property does not hold then the cipher-state on its own reveals information about the secret-key. This is certainly an unnecessary weakness that one should avoid when designing ciphers.

The data hiding property states that the adversary expects to see the same ensemble no matter what message was encrypted. We denote this ensemble

$$\sigma = \sum_{k=1}^{M} \frac{1}{M} \mathsf{E}_k |a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger, \tag{5}$$

for any $a \in \mathcal{P}$. We first prove that $\rho = \sigma$.

Lemma 1. $\rho = \sigma$.

Proof. Define the state

$$\xi = \sum_{k=1}^{M} \sum_{a \in \mathcal{P}} \frac{1}{MN} \mathsf{E}_k |a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger. \tag{6}$$

Observe that

$$\xi = \sum_{k=1}^{M} \sum_{a \in \mathcal{P}} \frac{1}{MN} \mathsf{E}_k |a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger = \sum_{k=1}^{M} \frac{1}{M} \rho = \rho. \tag{7}$$

Similarly, when switching the sums in (6), we get $\xi = \sigma$. We conclude that $\rho = \sigma$. $\square$

We are now ready to prove that for any (m, n)-quantum cipher there exists a measurement that returns the secret key with probability 2^{n-m} given any plaintext and its associated cipher-state. In other words and similarly to the classical case, the min-entropy key-uncertainty of any (m, n)-quantum cipher is at most $m - n$.

Theorem 1 (Min-entropy key uncertainty). *Let $(\mathcal{P}, \mathcal{E})$ be an (m,n)-quantum cipher, encoding the set $\mathcal{P}$. Then*

$$(\forall a \in \mathcal{P})(\exists\ POVM\ \{M_i\}_{i=1}^{M})(\forall k \in \{1,\dots,M\})[\mathrm{tr}(M_k \mathcal{E}_k(|a\rangle\langle a|)) = 2^{n-m}]. \quad (8)$$

Proof. Let $|a\rangle \in \mathcal{P}$ be given. Consider the set $\mathcal{M} = \{M_k = \frac{N}{M}\mathsf{E}_k|a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger \mid k = 1,\dots,M\}$. Lemma 1 gives

$$\sum_{k=1}^{M} M_k = \sum_{k=1}^{M} \frac{N}{M}\mathsf{E}_k|a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger = N\sigma = N\rho. \quad (9)$$

Since the plaintexts are orthogonal quantum states, and since unitary operators preserve angles, we have that $N\sum_{a\in\mathcal{P}} \frac{1}{N}\mathsf{E}_k|a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger$ is the eigen decomposition of $N\rho$, and that 1 is the only eigenvalue. Therefore there exists a positive operator P such that $N\rho + P = \mathbb{I}$, and thus

$$\sum_{k=1}^{M} M_k + P = N\rho + P = \mathbb{I}, \quad (10)$$

and $\mathcal{M} \cup \{P\}$ (and therefore also $\mathcal{M}$) is a valid POVM.

The probability of identifying the key with the measurement $\mathcal{M}$ is

$$\begin{aligned}\mathrm{tr}(M_k\mathsf{E}_k|a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger) &= \mathrm{tr}(\frac{N}{M}\mathsf{E}_k|a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger\mathsf{E}_k|a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger) \\ &= \frac{N}{M}\mathrm{tr}(\mathsf{E}_k|a\rangle|0\rangle\langle 0|\langle a|\mathsf{E}_k^\dagger) \\ &= 2^{n-m},\end{aligned} \quad (11)$$

which proves the theorem. □

5 Some Example Quantum Ciphers

In this section, we suggest a general method for designing quantum ciphers that can do better in terms of Shannon key-uncertainty than any classical cipher with the same parameters. The properties of our ciphers are analyzed in the next section.

The first example is extremely simple:

Definition 3. *The H_n cipher is an $(n+1,n)$-quantum cipher. Given message $b_1, b_2, \dots, b_n$ and key $c, k_1, \dots, k_n$, it outputs the following n q-bit state as ciphertext:*

$$(H^{\otimes n})^c(X^{k_1} \otimes X^{k_2} \otimes \dots \otimes X^{k_n}|b_1 b_2 \dots b_n\rangle), \quad (12)$$

where X is the bit-flip operator and H is the Hadamard transform. That is, we use the last n bits of key as a one-time pad, and the first key bit determines whether or not we do a Hadamard transform on all n resulting q-bits.

Decryption is trivial by observing that the operator $(X^{k_1} \otimes X^{k_2} \otimes \cdots \otimes X^{k_n})(H^{\otimes n})^c$ is the inverse of the encryption operator. It is also easy to see that the data hiding property is satisfied: if $c, k_1, \ldots, k_n$ are uniformly random, then the encryption of any message produces the complete mixture (in fact this would be the case, already if only $k_1, \ldots, k_n$ were uniformly random).

This cipher can be described from a more general point of view: let $\mathcal{B} = \{B_0, \ldots, B_{2^t-1}\}$ be a set of 2^t orthonormal bases for the Hilbert space of dimension 2^n. We require that the bases do not overlap, i.e., no unit vector occurs in more than one basis. For instance $\mathcal{B}$ could consist of the computational basis and the diagonal basis (i.e. $\{H^{\otimes n}|x\rangle | x \in \{0,1\}^n\}$). Let U_i be the unitary operator that performs a basis shift from the computational basis to the basis B_i. Finally, let $[k_1, \ldots, k_t]$ be the number with binary representation $k_1, \ldots, k_t$. Then we can define an $(n+t, n)$-cipher $C_{\mathcal{B}}$ which on input a key $c_1, \ldots, c_t, k_1, \ldots, k_n$ and a plaintext $b_1, \ldots, b_n$ outputs

$$U_{[c_1,\ldots,c_t]}(X^{k_1} \otimes X^{k_2} \otimes \ldots \otimes X^{k_n} |b_1 b_2 \ldots b_n\rangle). \tag{13}$$

The H_n-cipher above is a special case with $U_0 = Id, U_1 = H^{\otimes n}$. Using arguments similar to the above, it is easy to see that

Lemma 2. *For any set of orthonormal non-overlapping bases $\mathcal{B}$, $C_{\mathcal{B}}$ is a quantum cipher satisfying the data hiding and unique decryption properties.*

The lemma holds even if $\mathcal{B}$ contains only the computational basis, in which case $C_{\mathcal{B}}$ is equivalent to the classical one-time pad. The point of having several bases is that if they are well chosen, this may create additional confusion for the adversary, so that he will not learn full information on the key, even knowing the plaintext. We shall see this below.

For now, we note that Wootters and Fields have shown that in a Hilbert space of dimension 2^n, there exists $2^n + 1$ orthonormal bases that are *mutually unbiased*, i.e., the inner product between any pair of vectors from different bases has norm $2^{-n/2}$. Using, say, the first 2^n of these bases, we get immediately from the construction above a $(2n, n)$ cipher:

Definition 4. *The W_n-cipher is the cipher $C_{\mathcal{B}}$ obtained from the above construction when $\mathcal{B}$ is the set of 2^n mutually unbiased bases obtained from [9].*

5.1 Efficient Encoding/Decoding

In this section we look at how to implement W_n efficiently. In [9], a construction for $2^n + 1$ mutually unbiased bases in the space of n qubits is given. In the following, we denote by $v_{\boldsymbol{s}}^{(\boldsymbol{r})}$ with $\boldsymbol{s}, \boldsymbol{r} \in \{0,1\}^n$ the $\boldsymbol{s}$-th vector in the $\boldsymbol{r}$-th mutually unbiased basis. We write $v_{\boldsymbol{s}}^{(\boldsymbol{r})}$ in the computational basis as,

$$|v_{\boldsymbol{s}}^{(\boldsymbol{r})}\rangle = \sum_{\boldsymbol{l} \in \{0,1\}^n} \left(v_{\boldsymbol{s}}^{(\boldsymbol{r})}\right)_{\boldsymbol{l}} |\boldsymbol{l}\rangle, \tag{14}$$

where $\sum_{\boldsymbol{l}} |(v_{\boldsymbol{s}}^{(\boldsymbol{r})})_{\boldsymbol{l}}|^2 = 1$. Wootters and Field[9] have shown that 2^n mutually unbiased bases are obtained whenever

$$\left(v_{\boldsymbol{s}}^{(\boldsymbol{r})}\right)_{\boldsymbol{l}} = \frac{1}{\sqrt{2^n}} i^{\boldsymbol{l}^T(\boldsymbol{r}\cdot\boldsymbol{\alpha})\boldsymbol{l}} (-1)^{\boldsymbol{s}\cdot\boldsymbol{l}}, \tag{15}$$

for $\boldsymbol{\alpha}$ a vector of n matrices each of dimensions $n \times n$ with elements in $\{0, 1\}$. The arithmetic in the exponent of i should be carried out over the integers (or equivalently $\bmod 4$). The elements of $\boldsymbol{\alpha}$ are defined by

$$f_i f_j = \sum_{m=1}^{n} \boldsymbol{\alpha}_{i,j}^{(m)} f_m, \tag{16}$$

where $\{f_i\}_{i=1}^n$ is a basis for $GF(2^n)$ when seen as a vector space. Therefore, $\boldsymbol{\alpha}$ can be computed on a classical computer (and on a quantum one) in $O(n^3)$.

Let $c = c_1, \ldots, c_n$ and $k = k_1, \ldots, k_n$ be the $2n$ bits of key with c defining one out of 2^n mutually unbiased basis and k defining the key for the one-time-pad encoding. The circuit for encrypting classical message a starts by computing:

$$|\psi_a^k\rangle = H^{\otimes n} X^{\otimes k} |a\rangle = H^{\otimes n} |a \oplus k\rangle = 2^{-n/2} \sum_{\boldsymbol{l}} (-1)^{(a \oplus k)\cdot \boldsymbol{l}} |\boldsymbol{l}\rangle. \tag{17}$$

The state (17) differs from (14) only with respect to the phase factor $i^{\boldsymbol{l}^T(\boldsymbol{r}\cdot\boldsymbol{\alpha})\boldsymbol{l}}$ in front of each $|\boldsymbol{l}\rangle$ with $\boldsymbol{r} = c$. Transforming (17) into (14) (i.e. that is transforming $|\psi_a^k\rangle \mapsto |v_{k\oplus a}^{(c)}\rangle$) can easily be achieved using a few controlled operations as described in App. A. The complexity of the quantum encryption circuit is $O(n^3)$ out of which only $O(n^2)$ are quantum gates. The decryption circuit is the same as for the encryption except that it is run in reverse order. A similar encryption/decryption circuit can easily be implemented for any $C_{\mathcal{B}}$-cipher where $\mathcal{B}$ is a set of mutually unbiased bases.

6 Optimal measurements w.r.t. Shannon Entropy

Our ultimate goal is to estimate the Shannon key-uncertainty of an (m, n)-quantum cipher, i.e., the amount of entropy that remains on the key after making an optimal measurement on a ciphertext where the plaintext is given. But actually, this scenario is quite general and not tied to the cryptographic application: what we want to answer is: given a (pure) state chosen uniformly from a given set of states, how much Shannon entropy must (at least) remain on the choice of state after having made a measurement that is optimal w.r.t. minimizing the entropy?

So what we should consider is the following experiment: choose a key $k \in \mathcal{K}$ uniformly. Encrypt a given plaintext p under key k to get state $|c_k\rangle$ (we assume here for simplicity that this is a pure state). Perform some measurement (that may depend on p) and get outcome u. Letting random variables K, U correspond to the choices of key and outcome, we want to estimate

$$H(K|U) = \sum_u Pr(U = u)H(K|U = u). \tag{18}$$

Now, $H(K|U = u)$ is simply the Shannon entropy of the probability distribution $\{Pr(K = k|U = u)|k \in \mathcal{K}\}$. By the standard formula for conditional probabilities, we have

$$Pr(K = k|U = u) = \frac{Pr(U = u|K = k)Pr(K = k)}{Pr(U = u)}. \tag{19}$$

Note that neither $Pr(U = u)$, nor $Pr(K = k)$ depend on the particular value of k (since keys are chosen uniformly).

The measurement in question can be modeled as a POVM, which without loss of generality can be assumed to contain only elements of the form $a_u|u\rangle\langle u|$, i.e., a constant times a projection determined by a unit vector $|u\rangle$. This is because the elements of any POVM can be split in a sum of scaled projections, leading to a measurement with more outcomes which cannot yield less information than the original one. It follows immediately that

$$Pr(U = u|K = k) = |a_u|^2|\langle u|c_k\rangle|^2. \tag{20}$$

Note that also the factor $|a_u|^2$ does not depend on k. Then by (19) and (20), we get

$$1 = \sum_{l\in\mathcal{K}} Pr(K = l|U = u) = \frac{|a_u|^2 Pr(K = l)}{Pr(U = u)} \sum_{l\in\mathcal{K}} |\langle u|c_l\rangle|^2. \tag{21}$$

Which means that we have

$$Pr(K = k|U = u) = \frac{|\langle u|c_k\rangle|^2}{\sum_{l\in\mathcal{K}} |\langle u|c_l\rangle|^2}. \tag{22}$$

In other words, $H(K|U = u)$ can be computed as follows: compute the set of values $\{|\langle u|c_k\rangle|^2|k \in \mathcal{K}\}$, multiply by a normalization factor so that the resulting probabilities sum to 1, and compute the entropy of the distribution obtained. We call the resulting entropy $H[|u\rangle, S_K]$, where S_K is the set of states that may occur $\{|c_k\rangle|k \in \mathcal{K}\}$. This is to emphasize that $H[|u\rangle, S_K]$ can be computed only from $|u\rangle$ and S_K, we do not need any information about other elements in the measurement. From (18) and $H(K|U = u) = H[|u\rangle, S_K]$ follows immediately

Lemma 3. *With notation as above, we have:*

$$H(K|U) \geq min_{|u\rangle}\{H[|u\rangle, S_K]\}, \tag{23}$$

where $|u\rangle$ runs over all unit vectors in the space we work in.

This bound is not necessarily tight, but it will be, exactly if it is possible to construct a POVM consisting only of (scaled) projections $a_u|u\rangle\langle u|$, that minimize $H[|u\rangle, S_K]$. In general, it may not be easy to solve the minimization problem suggested by the lemma, particularly if S_K is large and lives in many dimensions. But in some cases, the problem is tractable, as we shall see.

7 The Shannon Key-Uncertainty of Quantum Ciphers

In this section, we study the cipher $C_{\mathcal{B}}$ constructed from a set of 2^t orthonormal bases $\mathcal{B}$ as defined in Sect. 5. For this, we first need a detour: each basis in our set defines a projective measurement. Measuring a state $|u\rangle$ in basis $B_i \in \mathcal{B}$ produces a result, whose probability distribution depends on $|u\rangle$ and B_i. Let $H[|u\rangle, B_i]$ be the entropy of this distribution. We define the Minimal Entropy Sum (MES) of $\mathcal{B}$ as follows:

$$MES(\mathcal{B}) = min_{|u\rangle}\{\sum_{i=0}^{2^t-1} H[|u\rangle, B_i]\}, \tag{24}$$

where $|u\rangle$ runs over all unit vectors in our space. Lower bounds on the minimal entropy sum for particular choices of $\mathcal{B}$ have been studied in several papers, under the name of entropic uncertainty relations [6, 8, 4]. This is motivated by the fact that if the sum is large, then it is impossible to simultaneously have small entropy on the results of all involved measurements. One can think of this as a "modern" version of Heisenberg's uncertainty relations. It turns out that the key uncertainty of $C_{\mathcal{B}}$ is directly linked to $MES(\mathcal{B})$:

Lemma 4. *The Shannon key uncertainty of the cipher $C_{\mathcal{B}}$ (with 2^t bases) is at least $MES(\mathcal{B})/2^t + t$.*

Proof. We may use Lemma 3, where the set of states S_K in our case consists of all basis states belonging to any of the bases in $\mathcal{B}$. To compute $H[|u\rangle, S_K]$, we need to consider the inner products of unit vector $|u\rangle$ with all vectors in S_K. In our case, this is simply the coordinates of $|u\rangle$ in each of the 2^t bases, so clearly the norm squares of the inner products sum to 2^t. Let z_{ij} be the i'th vector in the j'th basis from $\mathcal{B}$. We have,

$$\begin{aligned}
H[|u\rangle, S_K] &= \sum_{j=0}^{2^t-1}\sum_{i=0}^{2^n-1} \frac{1}{2^t}|\langle u|z_{ij}\rangle|^2 \log(2^t|\langle u|z_{ij}\rangle|^{-2}) \\
&= \sum_{j=0}^{2^t-1}\sum_{i=0}^{2^n-1} \frac{1}{2^t}|\langle u|z_{ij}\rangle|^2 \log(|\langle u|z_{ij}\rangle|^{-2}) + \sum_{j=0}^{2^t-1}\sum_{i=0}^{2^n-1} \frac{1}{2^t}|\langle u|z_{ij}\rangle|^2 \log(2^t) \\
&= \frac{1}{2^t}\sum_{j=0}^{2^t-1}\sum_{i=0}^{2^n-1} |\langle u|z_{ij}\rangle|^2 \log(|\langle u|z_{ij}\rangle|^{-2}) + t\frac{1}{2^t}\sum_{j=0}^{2^t-1}\sum_{i=0}^{2^n-1} |\langle u|z_{ij}\rangle|^2 \\
&= \frac{1}{2^t}\sum_{j=0}^{2^t-1} H[|u\rangle, B_j] + t \geq \frac{1}{2^t} MES(\mathcal{B}) + t.
\end{aligned} \tag{25}$$

The lemma follows. □

We warn the reader against confusion about the role of $|u\rangle$ and $\mathcal{B}$ at this point. When we estimate the key uncertainty of $C_{\mathcal{B}}$, we are analyzing a POVM, where $|u\rangle$ is one of the unit vectors defining the POVM. But when we do the proof of the above lemma and use the entities $H[|u\rangle, B_j]$, we think instead of $|u\rangle$ as the vector being measured according to basis B_j. There is no contradiction, however, since what matters in both cases is the inner products of $|u\rangle$ with the vectors in the bases in $\mathcal{B}$. We are now in a position to give results for our two concrete ciphers H_n and W_n defined earlier.

Theorem 2. *The H_n-cipher has Shannon key-uncertainty $n/2+1$ bits.*

Proof. The main result of [6] states that when $\mathcal{B}$ is a set of two mutually unbiased bases in a Hilbert space of dimension 2^n then $MES(\mathcal{B}) \geq n$. Using Lemma 4, it follows that H_n has Shannon key-uncertainty at least $n/2+1$. Moreover, there exists measurements (i.e. for example the Von Neumann measurement in either the rectilinear or Hadamard basis) achieving $n/2+1$ bit of Shannon key-uncertainty. The result follows. □

For the case of W_n, we can use a result by Larsen[4]. He considers the probability distributions induced by measuring a state $|u\rangle$ in $N+1$ mutually unbiased bases, for a space of dimension N. Let the set of bases be $B_1, \ldots, B_{N+1}$, and let $\pi_{|u\rangle,i}$ be the collision probability for the i'th distribution, i.e., the sum of the squares of all probabilities in the distribution. Then Larsen's result (actually a special case of it) says that

$$\sum_{i=1}^{N+1} \pi_{|u\rangle,i} = 2 \tag{26}$$

In our case, $N = 2^n$. However, to apply this to our cipher W_n, we would like to look at a set of only 2^n bases and we want a bound on the sum of the entropies $H[|u\rangle, B_i]$ and not the sum of the collision probabilities. This can be solved following a line of arguments from Sánchez-Ruiz[8]. Using Jensen's inequality, we obtain the following:

$$\begin{aligned}
\sum_{i=1}^{N} H[|u\rangle, B_i] &\geq -\sum_{i=1}^{N} \log \pi_{|u\rangle,i} \\
&\geq -N \log \left(\frac{1}{N} \sum_{i=1}^{N} \pi_{|u\rangle,i} \right) \\
&= -N \log \left(\frac{1}{N} \left(-\pi_{|u\rangle,N+1} + \sum_{i=1}^{N+1} \pi_{|u\rangle,i} \right) \right) \\
&= N \log \left(\frac{N}{2 - \pi_{|u\rangle,N+1}} \right) \geq N \log \left(\frac{N}{2 - 1/N} \right).
\end{aligned} \tag{27}$$

Together with Lemma 4, we get:

Theorem 3. *The W_n-cipher has Shannon key-uncertainty greater than $2n-1$ bits.*

Unlike for H_n (i.e. Theorem 2), Theorem 3 only provides a lower bound for the key uncertainty of W_n.

Let $\mathcal{B}$ be any set of 2^t mutually unbiased bases living in a Hilbert space of dimension 2^n. The largest value we could hope for $MES(\mathcal{B})$ is $(2^t-1)n$ bits, since this value is exactly matched when the state measured is a state that belongs to a basis in $\mathcal{B}$. It is natural to define $\Delta(n,t)$ as the distance between $MES(\mathcal{B})$ and the the maximum possible value:

$$\Delta(n,t) = (2^t - 1)n - MES(\mathcal{B}).$$

Given what we know already, it seems reasonable to conjecture that $\Delta(n,t)$ is, in some sense, small: we know that $\Delta(n,1) = 0$ and also that $\Delta(n,n) \leq (2^n - 1)n - 2^n(n-1) = 2^n - n$. Let us consider the following conjecture:

Conjecture 1. For any set $\mathcal{B}$ containing 2^n mutually unbiased bases in a Hilbert space of dimension 2^n, it holds that $\frac{\Delta(n,n)}{2^n} \in o(1)$ (i.e. note that we know the fraction is strictly smaller than 1).

In this case, we easily conclude that cipher W_n has almost full Shannon key-uncertainty:

Lemma 5. *Under Conjecture 1, W_n has Shannon key-uncertainty at least $2n - o(1)$ bits.*

Proof. From Lemma 4, the Shannon key-uncertainty of W_n is at least $n + MES(\mathcal{B})/2^n$. Conjecture 1 leads to $MES(\mathcal{B})/2^n = ((2^n-1)n - \Delta(n,n))/2^n = n - o(1)$. The result follows. □

The H_n and W_n-ciphers represent two extremes, using the minimal non-trivial number of bases, respectively as many of the known mutually unbiased bases as we can address with an integral number of key bits. It is not hard to define example ciphers that are "in between" and prove results on their key-uncertainty using the same techniques as for W_n. However, what can be derived from Larsen's result using the above line of argument (i.e. Equation 27) becomes weaker as one considers a smaller number of bases.

8 Composing Ciphers

What happens to the key uncertainty if we use a quantum cipher twice to encrypt two plaintext blocks, using independently chosen keys? Intuition based on classical behavior suggests that the key uncertainty should now be twice that of a single application of the cipher, since the keys are independent. But in the quantum case, this requires proof: the adversary will be measuring a product state composed of of two ciphertext blocks. If the adversary was to measure each block individually then clearly the key uncertainty would be twice the key uncertainty of a single block. However, coherent measurements involving both blocks simultaneously may provide more information on the key than what is achievable by measuring the blocks individually.

In the following, we consider composition of the cipher $C_{\mathcal{B}}$ with itself, where $\mathcal{B}$ consists of 2^t bases for a space of dimension 2^n. This is a $(2(t+n), 2n)$-cipher which we call $C^2_{\mathcal{B}}$. Say $\mathcal{B}$ consists of the bases $\mathcal{B} = \{B_0, ..., B_{2^t-1}\}$. Let us consider the tensor product of two Hilbert spaces of dimension 2^n each. Then $B_i \otimes B_j$ denotes the basis of this tensor product space that one obtains by taking all pairwise tensor products of the 2^n basis vectors in each of B_i and B_j. We will let $\mathcal{B} \otimes \mathcal{B}$ denote the set of all 2^{2t} bases that can be formed this way. Since each such basis consists of 2^{2n} basis vectors, $\mathcal{B} \otimes \mathcal{B}$ can also be thought of as a collection of 2^{2t+2n} pure states.

On the adversary's point of view, determining the two $t+n$-bit keys from two ciphertext blocks is equivalent to the following experiment: choose uniformly a state in $\mathcal{B} \otimes \mathcal{B}$, now the adversary wants to make a measurement that minimizes the uncertainty about the state that was picked.

To study this question, we split $\mathcal{B} \otimes \mathcal{B}$ in subsets: let $\mathcal{B}_i$ be the set of 2^t bases defined by

$$\mathcal{B}_i = \{B_j \otimes B_{j+i \mod 2^t} |\ j = 0, 1, ..., 2^t - 1\} \tag{28}$$

It is now easy to see that $\mathcal{B} \otimes \mathcal{B}$ is the disjoint union of the $\mathcal{B}_i$'s, for $i = 0, 1, ..., 2^t - 1$.

Now, the choice of a state in $\mathcal{B} \otimes \mathcal{B}$ can be rephrased as follows: choose i uniformly from $[0..2^t - 1]$, and then choose a state uniformly from $\mathcal{B}_i$. Let I, J be random variables representing these choices, and let U be the random variable representing the adversary's measurement result. Standard properties of Shannon entropy give:

$$H(I, J|\ U) = H(I|\ U) + H(J|\ I, U).$$

It is straightforward to see that a uniform mixture over all 2^{t+2n} states in $\mathcal{B}_i$ is in fact the complete mixture, and so has the same density matrix for any i, hence no measurement can reveal information on I and we have $H(I|\ U) = t$. We define $M_2(\mathcal{B}) = min_i\{MES(\mathcal{B}_i)\}$. Then, using exactly the same line of argument as for Lemma 4, one finds that for each particular value of i, we have $H(J|\ I = i, U) \geq t + MES(\mathcal{B}_i)/2^t$ and hence $H(J|\ I, U) \geq t + M_2(\mathcal{B})/2^t$. Putting things together gives,

Lemma 6. *$C^2_{\mathcal{B}}$ has Shannon key-uncertainty at least $2t + M_2(\mathcal{B})/2^t$.*

Considering composition of $C_{\mathcal{B}}$ v times with itself, denoted $C^v_{\mathcal{B}}$, the techniques above extend in a straightforward way. In particular, we end up defining a minimum $M_v(\mathcal{B})$ over entropy sums for a generalization of the $\mathcal{B}_i$'s. This leads to,

Lemma 7. *$C^v_{\mathcal{B}}$ has Shannon key-uncertainty at least $vt + M_v(\mathcal{B})/2^t$.*

Note that by the construction defined in (28), each $\mathcal{B}_i$ is a set of mutually unbiased bases, and this holds also for any of the v-wise generalizations. In the special case of H_n, we have $t = 1$, and each $\mathcal{B}_i$ (as well as its v-wise generalization) contains 2 mutually unbiased bases. Lemma 7 together with the result of [6] (i.e. which in our notation reads $M_v(\mathcal{B}) = vn$) immediately implies,

Theorem 4. *The cipher H_n^v has Shannon key uncertainty $v(n/2+1)$ bits.*

We do not know of any strong results on the minimal entropy sum for any set of mutually unbiased bases except when its cardinality is 2[6] or is close to the dimension of the space[4, 8]. Therefore, we cannot prove a good lower bound on the Shannon key-uncertainty for the composition of W_n. Already for W_n^2, we need to consider a set of 2^n mutually unbiased bases living in a space of dimension 2^{2n}. Using the notation of the previous section, we need to bound $\Delta(2n, n)$, or more generally $\Delta(vn, n)$.

While $\Delta(vn, n) = 0$ may be too much to hope for, it seems reasonable to conjecture a result similar to the one we know for $\Delta(n, n)$:

Conjecture 2. For any set $\mathcal{B}$ of 2^n mutually unbiased bases living in a Hilbert space of dimension 2^{vn}, it holds that $\Delta(vn, n) \leq 2^n - vn$.

We then have,

Lemma 8. *Under Conjecture 2, W_n^v has Shannon key-uncertainty at least $2vn-1$ bits.*

9 Application to Stream-Ciphers

We can use the quantum ciphers we just described to build a (computationally secure) quantum stream-cipher using a short key K of length independent from the message length. In fact, any (m, n)-cipher and classical pseudorandom generator can be used: we seed the generator with key K, and use its output as a keystream. To encrypt, we simply take the next m bits from the keystream and use these as key in the cipher to encrypt the next n bits of the plaintext.

Since an (m, n)-cipher has perfect security, this construction would have perfect security as well if the keystream was genuinely random. By a standard reduction, this implies that breaking it is at least as hard as distinguishing the output of the generator from a truly random string.

All this is true whether we use a classical or an (m, n)-quantum cipher. However, by our results on Shannon key-uncertainty, the adversary is in a potentially much harder situation in the quantum case. For intuition on this, we refer to the discussion in the introduction. As a more concrete illustration, we consider the following scenario:

1. We have a pseudorandom generator G, expanding a k-bit seed K into an N-bit sequence $G(K)$. Furthermore, any subset containing at most e bits of $G(K)$ is uniformly random. Finally, no polynomial time (in k) classical algorithm can with non-negligible advantage distinguish $G(K)$ from a truly random sequence when given any piece of data that is generated from $G(K)$ and contains at most t bits of Shannon information on $G(K)$. Both e and t are assumed to be polynomial in k.
2. Coherent measurements simultaneously involving μ qubits or more are not possible to implement in practice. However, technology has advanced so that the W_n-cipher can be implemented for some $n << \mu$.

3. We will consider an adversary that first obtains some amount of known plaintext. Given the plaintext, he decides on a number of complete measurements that he executes on parts of the ciphertext (under the constraints of assumption 2). For simplicity we assume that each measurement involves an integral number of n-bit ciphertext blocks.[1] Finally he executes any polynomial time classical algorithm to analyze the results.

The first assumption can be justified using a result by Maurer and Massey [5] on locally random pseudorandom generators. Their result asserts that there exists pseudorandom generators satisfying the assumption that any e bits are genuinely random, provided $e \leq k/\log_2 N$. Their generators may not behave well against attacks having access to more than e bits of the sequence, but one can always xor the output from their generator with the output of a more conventional one using an independent key. This will preserve the local randomness.

Note that the size of k does not influence the size of the quantum computer required for the honest party to encrypt or decrypt. The third assumption essentially says that we do not expect that results of (incomplete) measurements obtained on one part of the ciphertext will help significantly in designing measurements on other parts. This is justified, as long as not too many measurements are performed: as long as results from previous measurements contain less than t bits of information on the keystream, then by assumption 1, these results might (from the adversary's point of view) as well have been generated from measuring a random source, and so they do not help in designing the next measurement. This assumption can therefore be dropped in a more careful analysis since it esssentially follows from assumptions 1 and 2. For simplicity, we choose to make it explicit.

Lemma 9. *Assume we apply the W_n-cipher for stream encryption using a pseudorandom generator and with an adversary as defined by assumptions 1,2, and 3 above. Suppose we choose $e = 2\mu$ and $k \geq 2\mu \log_2 N$. Then, assuming Conjecture 2, the adversary will need to obtain tn bits of known plaintext, in order to distinguish the case of a real encryption from the case where the keystream is random.*

Proof. Assume the PRG satisfies assumption 1 which is possible since $k \geq e \log_2 N$. By assumption 2, any attack that measures several blocks of ciphertext in one coherent measurement can handle at most $\mu = e/2$ qubits at any one time. By construction, this ciphertext was created using less than e bits of the keystream, which is random by assumption 1. Therefore, the measurement will give the same result as when attacking the composition $W_n^{v/n}$ since the measurement involves $v \leq \mu$ qubits (since different blocks of the keystream are independent if the stream is truly random) and by assumption 3. Hence, by Lemma 8 and under Conjecture 2, the adversary learns less than 1 bit of information on the key stream from each measurement. Now, if the adversary has

[1] This assumption can be dropped so that we can still prove Lemma 9 using a more complicated argument and provided the local randomness of the generator is expanded from e to $n^2 e$

T bits of known plaintext, and hence measures T ciphertext bits, the maximal number of measurements that can take place is T/n so he needs to have $T/n > t$ in order for the classical distinguisher to work, by assumption 1. The lemma follows. □

This lemma essentially says that for a generator with the right properties, and for an adversary constrained as we have assumed, quantum communication allows using the generator securely to encrypt tn bits, rather than the t bits we would have in the classical case. Depending on how close the actual key uncertainty of compositions of W_n is to the maximal value, the number of required plaintext bits can be much larger. The best we can hope for would be if $\Delta(vn, n) = 0$ for all n, v, in which case the adversary would need $t2^n$ plaintext bits.

A similar result can be shown without assuming any conjecture for the H_n cipher. In this case, we gain essentially a factor 2 in plaintext size over the classical case.

Of course, these results do not allow to handle adversaries as general as we would like, our constraints are different from just assuming the adversary is quantum polynomial time. Nevertheless, we believe that the scenario we have described can be reasonable with technology available in the foreseeable future. Moreover, it seems to us that quantum communication should help even for more general adversaries and generators. Quantifying this advantage is an open problem.

10 Conclusion and Open Problems

We have seen that, despite the fact that quantum communication cannot help to provide perfect security with shorter keys when only one-way communication is used, there are fundamental differences between classical and quantum ciphers with perfect security, in particular the Shannon key uncertainty can be much larger in the quantum case. However, the min-entropy key-uncertainty is the same in the two cases. It is an open question whether encryption performed by general quantum operations allows for quantum ciphers to have more min-entropy key-uncertainty than classical ones.

We have also seen an application of the results on Shannon key uncertainty to some example quantum ciphers that could be used to construct a quantum stream-cipher where, under a known plaintext attack, a resource-bounded adversary would be in a potentially much worse situation than with any classical stream-cipher with the same parameters.

For the ciphers we presented, the Shannon key-uncertainty is known exactly for the H_n-cipher but not for the W_n-cipher. It is an interesting open question to determine it. More generally, are Conjectures 1 and 2 true?

Acknowledgements

We are grateful to Renato Renner for pointing out a mistake in the proof of the Shannon key-uncertainty for the composition of cipher $C_{\mathcal{B}}$ appearing in the proceedings of Eurocrypt 2004 (i.e. the proof of Theorem 4 in[2] is wrong!).

References

1. A. Ambainis, M. Mosca, A. Tapp and R. de Wolf, *Private Quantum Channels*, Proceedings of the 41st Annual Symposium on Foundations of Computer Science, 2000, pp. 547–553.
2. I. Damgård, T. Pedersen, and L. Salvail, *On the Key-Uncertainty of Quantum Ciphers and the Computational Security of One-Way Quantum Transmission,* Proceedings of Eurocrypt'04, LLNCS 3027, Springer-Verlag, 2004,pp. 91–108.
3. D. DiVincenzo, M. Horodecki, D. Leung, J. Smolin and B. Terhal, *Locking Classical Correlation in Quantum States*, Phys. Rev. Letters,vol. 92, 067902, 2004.
4. U. Larsen, *Superspace Geometry: the exact uncertainty relationship between complementary aspects*, J.Phys. A: Math. Gen. 23 (1990), pp. 1041–1061.
5. U. Maurer and J. Massey, *Local Randomness in Pseudorandom Sequences*, Journal of Cryptology, vol. 4, 1991, pp. 135–149.
6. H. Maassen and J. B. M. Uffink, *Generalized Entropic Uncertainty Relations*, Phys. Rev. Letters, vol. 60, 1988, pp. 1103–1106.
7. M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, 2000.
8. J. Sánchez-Ruiz, *Improved bounds in the entropic uncertainty and certainty relations for complementary observables*, Physics Letters A 201, 1995, pp. 125–131.
9. W.K. Wootters and B.D. Fields, *Optimal state-determination by mutually unbiased measurements*, Annals of Physics 191, pp. 363–381.

A Encryption Circuit for the $\boldsymbol{W_n}$-Cipher

The circuit depicted in Fig. 2 implements the encryption of any plaintext $a = a_1, \ldots, a_n \in \{0,1\}^n$ according the secret key $(c,k) \in \{0,1\}^{2n}$. It uses three sub-circuits (1), (2), and (3) as defined in Fig. 1.

$\mathcal{A}$, given c and $\boldsymbol{\alpha}$, produces the matrix $c \cdot \boldsymbol{\alpha}$ in the register denoted A. Notice that circuit $\mathcal{A}$ is a classical circuit. It can be implemented with $O(n^3)$ classical gates. The sub-circuit (2) accepts as input $\hat{\boldsymbol{\alpha}} = c \cdot \boldsymbol{\alpha}$ together with $\boldsymbol{l}$, computes $d = \boldsymbol{l}^T \hat{\boldsymbol{\alpha}} \boldsymbol{l} \in [0, \ldots, 3]$, and stores the result in a 2-qubit register I. In (3), an overall phase factor i^d is computed in front of the computational basis element $|\boldsymbol{l}\rangle$. The last gates allow to reset registers I and A making sure registers containing the encrypted data are separable from the other registers. It is straightforward to verify that registers initially in state $|a_1\rangle \otimes \ldots \otimes |a_n\rangle$ ends up in state $|v^{(c)}_{k \oplus a}\rangle$ as required. The overall complexity is $O(n^2)$ quantum gates since (3) requires only $O(n^2)$ CNOT's which is of the same complexity as super-gate (2). In conclusion, the total numbers of gates is $O(n^3)$ out of which $O(n^2)$ are quantum.

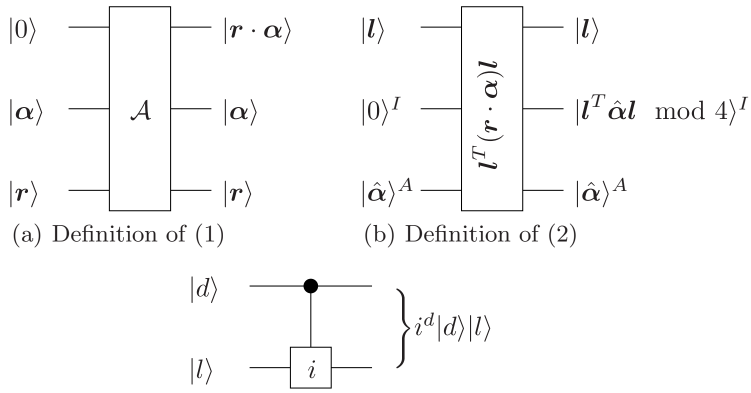

(a) Definition of (1) (b) Definition of (2)

(c) Definition of (3)

Fig. 1. Sub-circuits to the encryption circuit of Fig. 2.

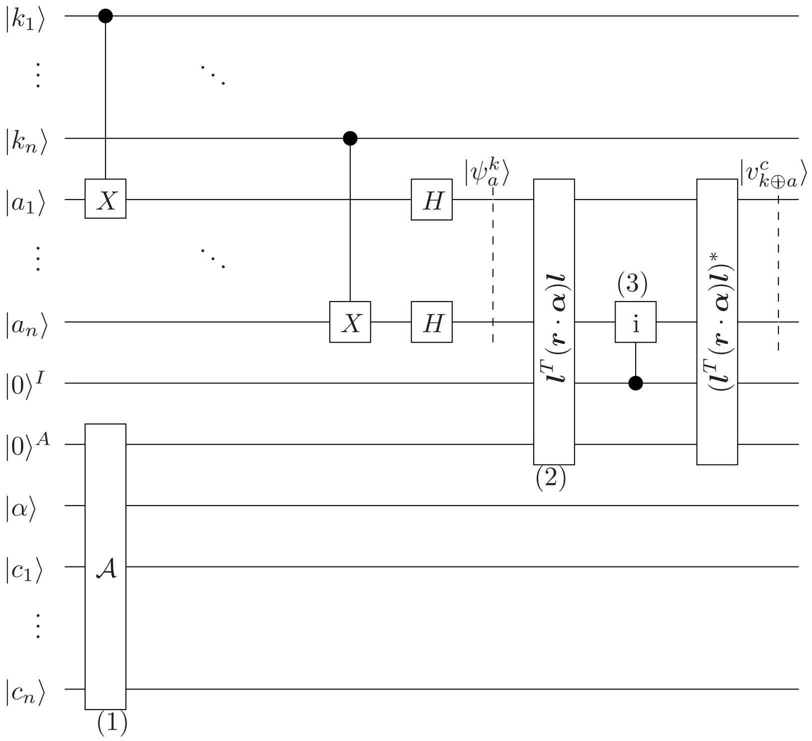

Fig. 2. Encoding circuit for cipher W_n.