Single-Shot Security for One-Time Memories in the Isolated Qubits Model

Yi-Kai Liu

Applied and Computational Mathematics Division,
National Institute of Standards and Technology (NIST),
Gaithersburg, MD, USA
`yi-kai.liu@nist.gov`

Abstract. One-time memories (OTM's) are simple, tamper-resistant cryptographic devices, which can be used to implement sophisticated functionalities such as one-time programs. Can one construct OTM's whose security follows from some physical principle? This is not possible in a fully-classical world, or in a fully-quantum world, but there is evidence that OTM's can be built using “isolated qubits” — qubits that cannot be entangled, but can be accessed using adaptive sequences of single-qubit measurements.

Here we present new constructions for OTM's using isolated qubits, which improve on previous work in several respects: they achieve a stronger “single-shot” security guarantee, which is stated in terms of the (smoothed) min-entropy; they are proven secure against adversaries who can perform arbitrary local operations and classical communication (LOCC); and they are efficiently implementable.

These results use Wiesner's idea of conjugate coding, combined with error-correcting codes that approach the capacity of the q -ary symmetric channel, and a high-order entropic uncertainty relation, which was originally developed for cryptography in the bounded quantum storage model.

Keywords: Quantum cryptography, information theory, local operations and classical communication (LOCC), oblivious transfer, one-time programs.

1 Introduction

One-time memories (OTM's) are a simple type of tamper-resistant cryptographic hardware. An OTM has the following behavior: a user Alice can write two messages s and t into the OTM, and then give the OTM to another user Bob; Bob can then choose to read either s or t from the OTM, but he can only learn one of the two messages, not both. A single OTM is not especially exciting by itself, but when many OTM's are combined in an appropriate way, they can be used to implement *one-time programs*, which are a powerful form of secure computation [3,4,5,6]. (Roughly speaking, a one-time program is a program that can be run exactly once, on an input chosen by the user. After running once,

the program “self-destructs,” and it never reveals any information other than the output of the computation.)

Can one construct OTM’s whose security follows from some physical principle? At first glance, the answer seems to be “no.” OTM’s cannot exist in a fully classical world, because information can always be copied without destroying it. One might hope to build OTM’s in a quantum world, where the no-cloning principle limits an adversary’s ability to copy an unknown quantum state. However, this is also impossible, because an OTM can be used to perform oblivious transfer with information-theoretic security, which is ruled out by various “no-go” theorems [7,8,9,10].

One way around these no-go theorems is to try to construct protocols that are secure against restricted classes of quantum adversaries, e.g., adversaries who can only perform k -local measurements [11], or adversaries who only have bounded or noisy quantum storage [12,13,14,15,16,17]. More recently, Liu has proposed a construction for OTM’s in the *isolated qubits model* [1], where the adversary is only allowed to perform local operations and classical communication (LOCC). That is, the adversary can perform single-qubit quantum operations, including single-qubit measurements, and can make adaptive choices based on the classical information returned by these measurements; but the adversary cannot perform entangling operations on sets of two or more qubits. (Honest parties are also restricted to LOCC operations.) The isolated qubits model is motivated by recent experimental work using solid-state qubits, such as nitrogen vacancy (NV) centers; see [1] for a more complete discussion of this model, and [18] for earlier work on implementing quantum money using NV centers.¹

In this paper we show a new construction and security analysis for OTM’s in the isolated qubits model, which improves on the results of [1] in several respects. First, we show a stronger “single-shot” security guarantee, which is stated in terms of the (smoothed) min-entropy [19,20]. This shows that a constant fraction of the message bits remain hidden from the adversary. This stronger statement is necessary for most cryptographic applications; note that the previous results of [1] were not sufficient, as they used the Shannon entropy.

Second, we prove security against general LOCC adversaries, who can perform arbitrary measurements (including weak measurements), and can measure each qubit multiple times. This improves on the results of [1], which only showed security against 1-pass LOCC adversaries that use 2-outcome measurements. Our new security proof is based solely on the definition of the isolated qubits model, without any additional assumptions.

Third, we show a construction of OTM’s that is efficiently implementable, i.e., programming and reading out the OTM can be done in polynomial time. This improves on the construction in [1], which was primarily an information-theoretic

¹ Note that the devices constructed in [1], and in this paper, are more precisely described as *leaky* OTM’s, because they can leak additional information to the adversary. It is not known whether such leaky OTM’s are sufficient to construct one-time programs as defined in [3]. We will discuss this issue in Section 1.2; for now, we will simply refer to our devices as OTM’s.

result, using random error-correcting codes that did not allow efficient decoding. (In fact, our new construction is quite flexible, and does not depend heavily on the choice of a particular error-correcting code. Our OTM's can be constructed using any code that satisfies two simple requirements: the code must be linear over $GF(2)$, and it must approach the capacity of the q -ary symmetric channel. We show one such code in this paper; several more sophisticated constructions are known [22,23,24].)

We will describe our OTM construction in the following section. Here, we briefly comment on some related work. Note that OTM's cannot make use of standard techniques such as privacy amplification. This is because OTM's are non-interactive and asynchronous: all of the communication between Alice and Bob occurs at the beginning, while the adversary can wait until later to attack the OTM. (To do privacy amplification, Alice would have to first force the adversary to take some action, and then send one more message to Bob. This trick is very natural in protocols for quantum key distribution and oblivious transfer, but it is clearly impossible in the case of an OTM.) As we will see below, the security of our OTM's follows from rather different arguments. (A similar issue was studied recently in [17], albeit with a weaker, non-adaptive adversary.)

In addition, it is a long-standing open problem to prove strong upper-bounds on the power of LOCC operations. Previous results in this area include demonstrations of “nonlocality without entanglement” [25] (see [26] for a recent survey), and constructions of data-hiding states [27,28,29,30]. Our OTM's are not directly comparable to these earlier results, as the security requirements for our OTM's are quite different.

1.1 Our Construction

We now describe our OTM construction, which is based on Wiesner's idea of conjugate coding [21]. Our OTM will store two messages $s, t \in \{0, 1\}^\ell$, and will use $n \lg q$ qubits, where q is a (large) power of 2. Let $C : \{0, 1\}^\ell \rightarrow \{0, 1\}^{n \lg q}$ be any error-correcting code that satisfies the following two requirements: C is linear over $GF(2)$, and C approaches the capacity of the q -ary symmetric channel $\mathcal{E}_q$ with error probability $p_e := \frac{1}{2} - \frac{1}{2q}$ (where the channel treats each block of $\lg q$ bits as a single q -ary symbol). Note that, when q is large, the capacity of the channel $\mathcal{E}_q$ is roughly $1 - p_e$, which is roughly $\frac{1}{2}$, so we have $n \lg q \approx 2\ell$.

Given two messages s and t , let $C(s)$ and $C(t)$ be the corresponding codewords, and view each codeword as n blocks consisting of $\lg q$ bits. We prepare the qubits in the OTM as follows. For each $i = 1, 2, \dots, n$,

- Let $\gamma_i \in \{0, 1\}$ be the outcome of a fair and independent coin toss.
- If $\gamma_i = 0$, prepare the i 'th block of qubits in the standard basis state corresponding to the i 'th block of $C(s)$.
- If $\gamma_i = 1$, prepare the i 'th block of qubits in the Hadamard basis state corresponding to the i 'th block of $C(t)$.

To recover the first message s , we measure every qubit in the standard basis, which yields a string of measurement outcomes $z \in \{0, 1\}^{n \lg q}$, and then we

run the decoding algorithm for C . To recover the second message t , we measure every qubit in the Hadamard basis, then follow the same procedure. It is easy to see that all of these procedures require only single-qubit state preparations and single-qubit measurements, which are allowed in the isolated qubits model.²

(We remark that this OTM construction uses blocks of qubits, rather than individual qubits as in [21] and [1]. That is, we set q large, instead of using $q = 2$. This difference seems to help our security proof, although it is not clear whether it affects the actual security of the scheme.)

We now sketch the proofs of correctness and security for this OTM. With regard to correctness, note that an honest player who wanted to learn s will obtain measurement outcomes that have the same distribution as the output of the q -ary symmetric channel $\mathcal{E}_q$ acting on $C(s)$; hence the decoding algorithm will return s . A similar argument holds for t .

To prove security, we consider adversaries that make *separable* measurements (which include LOCC measurements as a special case). The basic idea is to consider the distribution of the messages s and t , conditioned on one particular measurement outcome z obtained by the adversary. Since the adversary is separable, the corresponding POVM element M_z will be a tensor product of single-qubit operators $\bigotimes_{a=1}^{n \lg q} R_a$ (up to normalization). Now, one can imagine a fictional adversary that measures the qubits one at a time, and happens to observe this same string of single-qubit measurement outcomes $R_1, R_2, \dots, R_{n \lg q}$. This event leads to the same conditional distribution of s and t . But the fictional adversary is easier to analyze, because it is non-adaptive, it measures each qubit only once, and the measurements can be done in arbitrary order.

Now, our proof will be based on the following intuition. In order to learn both messages s and t , the adversary will want to determine the basis choices $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$, so that he will know which blocks of qubits should be measured in the standard basis, and which blocks of qubits should be measured in the Hadamard basis. The choice of the code C is crucial to prevent the adversary from doing this; for instance, if the adversary could predict some of the bits in the codewords $C(s)$ and $C(t)$, he could then measure the corresponding qubits, and gain some information about which bases were used to prepare them. (Note moreover that the adversary has full knowledge of C , before he measures any of the qubits.) We will argue that certain properties of the code C prevent the adversary from learning these basis choices γ perfectly, and that this in turn limits the adversary's knowledge of the messages s and t .

Since C is a linear code over $GF(2)$, it has a generator matrix G , which has rank ℓ . Thus there must exist a subset of ℓ bits of the codeword $C(s)$ that look uniformly random, assuming the message s was chosen uniformly at random; and a similar statement holds for $C(t)$. Now, let A be the subset of ℓ qubits that encode these bits of $C(s)$ and $C(t)$. We can imagine that the fictional adversary happens to measure these qubits *first*. Therefore, during these first ℓ steps, the fictional adversary learns nothing about which bases had been used to prepare

² We note in passing that Winter's "gentle measurement lemma" [31] does not imply an attack on this OTM using LOCC operations; see the full paper [2] for details.

the state, i.e., the basis choices γ are independent of the fictional adversary's measurement outcomes.

One can then show that the conditional distribution of s and t after these first ℓ steps of the fictional adversary is related to the distribution of measurement outcomes when the state $\bigotimes_{a \in A} R_a$ is measured in a random basis. This kind of situation has been studied previously, in connection with cryptography in the bounded quantum storage model. In particular, we can use a high-order entropic uncertainty relation from [16] to show a lower-bound on the smoothed min-entropy of this distribution. We then use trivial bounds to analyze the remaining $n \lg q - \ell$ steps of the fictional adversary. Roughly speaking, we get a bound of the form:

$$H_{\infty}^{\varepsilon}(S, T|Z) \gtrsim \frac{1}{2} \ell, \quad (1)$$

for any separable adversary (where Z denotes the adversary's measurement outcome). Thus, while the OTM may leak some information, it still hides a constant fraction of the bits of the messages s and t . For more details, see Section 3.

Finally, we show one construction of a code C that satisfies the above requirements and is efficiently decodable. The basic idea is to fix some $q_0 < q$, first encode the messages s and t using a random linear code $C_0 : \{0, 1\}^{\ell} \rightarrow \{0, 1\}^{n \lg q_0}$, then encode each block of $\lg q_0$ bits using a fixed linear code $C_1 : \{0, 1\}^{\lg q_0} \rightarrow \{0, 1\}^{\lg q}$. The code C_1 is used to detect the errors made by the q -ary symmetric channel; these corrupted blocks of bits are then treated as erasures, and we can decode C_0 by solving a linear system of equations, which can be done efficiently. Moreover, choosing C_0 to be a random linear encode ensures that, with high probability, C approaches the capacity of the q -ary symmetric channel. For more details, see Section 4.

1.2 Outlook

The results of this paper can be summarized as follows: we construct OTM's based on conjugate coding, which achieve a fairly strong ("single-shot") notion of security, are secure against general LOCC adversaries, and can be implemented efficiently. These results are a substantial improvement on previous work [1].

We view these results as a first step in a broader research program that aims to develop practical implementations of isolated qubits, one-time memories, and ultimately one-time programs. We now comment briefly on some different aspects of this program.

Experimental realization of isolated qubits is quite challenging, though there has been recent progress in this direction [39,40]. Broadly speaking, isolated qubits seem to be at an intermediate level of difficulty, somewhere between photonic quantum key distribution (which already exists as a commercial product), and large-scale quantum computers (which are still many years in the future).

Working with quantum devices in the lab also raises the question of fault-tolerance: can our OTM's be made robust against minor imperfections in the qubits? We believe this can be done, by slightly modifying our OTM construction: we would use a slightly noisier channel to describe the imperfect

measurements made by an honest user, and we would choose the error-correcting code C accordingly. The proof of security would still hold against LOCC adversaries who can make perfect measurements. There is plenty of “slack” in the security bounds, to allow this modification to the OTM’s.

In addition, one may wonder whether our OTM’s are secure against so-called “ k -local” adversaries [11], which can perform entangled measurements on small numbers of qubits (thus going outside the isolated qubits model). There is some reason to be optimistic about this: while we have mainly discussed separable adversaries in this paper, our security proof actually works for a larger set of adversaries, who can generate entanglement among some of the qubits, but are still separable across the partition defined by the subset A (as described in the proof). Also, from a physical point of view, k -local adversaries are quite natural. In particular, even when one can perform entangling operations on pairs of qubits, it may be hard to entangle large numbers of qubits, due to error accumulation.

Finally, let us turn to the construction of one-time programs. Because our OTM’s leak some information, it is not clear whether they are sufficient to construct one-time programs. There are a couple of approaches to this problem. On one hand, one can try to strengthen the security proof, perhaps by proving constraints on the *types* of information that an LOCC adversary can extract from the OTM. We conjecture that, when our OTM’s are used to build one-time programs as in [3], the specific information that is relevant to the security of the one-time program does in fact remain hidden from an LOCC adversary.

On the other hand, one can try to strengthen the OTM constructions, in order to eliminate the leakage. As noted previously, standard privacy amplification (e.g., postprocessing using a randomness extractor) does not work in this setting, because the adversary also knows the seed for the extractor. However, there are other ways of solving this problem, for instance by assuming the availability of a random oracle, or by using something similar to leakage-resilient encryption [32,33] (but with a different notion of leakage, where the “leakage function” is restricted to use only LOCC operations, but is allowed access to side-information).

2 Preliminaries

2.1 Notation

For any natural number n , let $[n]$ denote the set $\{1, 2, \dots, n\}$. Let $\lg(x) = \log_2(x)$ denote the logarithm with base 2.

For any random variable X , let P_X be the probability density function of X , that is, $P_X(x) = \Pr[X = x]$. Likewise, define $P_{X|Y}(x|y) = \Pr[X = x|Y = y]$, etc. For any event $\mathcal{E}$, define $P_{\mathcal{E}X}$ to be the probability density function of X smoothed by $\mathcal{E}$, that is $P_{\mathcal{E}X}(x) = \Pr[X = x \text{ and } \mathcal{E} \text{ occurs}]$.

We say that C is a binary code with codeword length n and message length k if C is a subset of $\{0, 1\}^n$ with cardinality 2^k . We say that C has minimum distance $d = \min_{x, y \in C} d_H(x, y)$, where $d_H(\cdot, \cdot)$ denotes the Hamming distance.

We say that C is a binary linear code if C is a linear subspace of $GF(2)^n$. (Note, $GF(2)$ and $\{0, 1\}$ denote the same set, but we will write $GF(2)$ in situations where we use arithmetic operations.) In this case, there exists a matrix $G \in GF(2)^{k \times n}$, such that the map $x \mapsto x^T G$ is a bijection from $GF(2)^k$ to the code subspace C . We will overload the notation and use C to denote the map $x \mapsto x^T G$; then the codewords consist of the strings $C(x)$ for all $x \in GF(2)^k$.

2.2 The q -ary Symmetric Channel

The q -ary symmetric channel with error probability p_e acts as follows: given an input $x \in GF(q)$, it returns an output $y \in GF(q)$, with conditional probabilities $\Pr(y|x) = 1 - p_e$ (if $y = x$) and $\Pr(y|x) = p_e/(q - 1)$ (if $y \neq x$). The capacity of this channel, measured in q -ary symbols per channel use, is given by [23]:

$$\begin{aligned} L(p_e) &= 1 + (1 - p_e) \log_q(1 - p_e) + p_e \log_q(p_e) - p_e \log_q(q - 1) \\ &= 1 - \frac{h_2(p_e)}{\lg q} - p_e \frac{\lg(q - 1)}{\lg q} \geq 1 - \frac{1}{\lg q} - p_e, \end{aligned} \quad (2)$$

where $h_2(\cdot)$ is the binary entropy function.

2.3 LOCC Adversaries and Separable Measurements

An LOCC adversary is an adversary that uses only local operations and classical communication (LOCC). Here, “local operations” consist of quantum operations on single qubits, and “classical communication” refers to the adversary’s ability to choose each single-qubit operation adaptively, depending on classical information, such as measurement outcomes, that were obtained from previous single-qubit operations. However, the adversary is not allowed to make adaptive choices that depend on quantum information, or perform entangling operations on multiple qubits.

Formally, an LOCC adversary can be described as follows. Consider a system of n qubits. The adversary makes a sequence of steps, labelled by $i = 1, 2, 3, \dots$. At step i , the adversary chooses one of the qubits $q_i \in [n]$, and performs a general quantum measurement $\mathcal{M}_i$ on that qubit; this returns a measurement outcome, which is described by a classical random variable Z_i . The adversary’s choices of q_i and $\mathcal{M}_i$ can depend on $Z_1, Z_2, \dots, Z_{i-1}$. Also, note that the adversary can perform weak measurements, and can measure the same qubit multiple times. Finally the adversary discards the qubits, and outputs the sequence of measurement outcomes $Z_1, Z_2, Z_3, \dots$.

A POVM measurement $\mathcal{M} = \{M_z \mid z = 1, 2, 3, \dots\}$ is called *separable* if every POVM element M_z can be written as a tensor product of single-qubit operators. It is easy to see that any LOCC adversary can be simulated by a separable measurement, i.e., for any LOCC adversary $\mathcal{A}$, there exists a separable POVM measurement $\mathcal{M}$, such that for every quantum state ρ , the output of $\mathcal{M}$ acting on ρ has the same distribution as the output of $\mathcal{A}$ acting on ρ [38].

2.4 Leaky OTM's

We will use the following definition of a leaky OTM [1].

Definition 1. Fix some class of adversary strategies $\mathbb{M}$, some leakage parameter $\delta \in [0, 1]$, and some failure probability $\varepsilon \in [0, 1]$. A leaky one-time memory (leaky OTM) with parameters $(\mathbb{M}, \delta, \varepsilon)$ is a device that has the following behavior. Suppose that the device is programmed with two messages s and t chosen uniformly at random in $\{0, 1\}^\ell$; and let S and T be the random variables containing these messages. Then:

1. *Correctness:* There exists an honest strategy $\mathcal{M}^{(1)} \in \mathbb{M}$ that interacts with the device and recovers the message s with probability $\geq 1 - \varepsilon$. Likewise, there exists an honest strategy $\mathcal{M}^{(2)} \in \mathbb{M}$ that recovers the message t with probability $\geq 1 - \varepsilon$.
2. *Leaky security:* For every strategy $\mathcal{M} \in \mathbb{M}$, if Z is the random variable containing the classical information output by $\mathcal{M}$, then $H_\infty^\varepsilon(S, T|Z) \geq (1 - \delta)\ell$.

Here H_∞^ε is the smoothed conditional min-entropy, which is defined as follows [19,20]:

$$H_\infty^\varepsilon(X|Y) = \max_{\mathcal{E}: \Pr(\mathcal{E}) \geq 1 - \varepsilon} \min_{x, y} \left[-\lg[P_{\mathcal{E}X|Y}(x|y)] \right], \quad (3)$$

where the maximization is over all events $\mathcal{E}$ (defined by the conditional probabilities $P_{\mathcal{E}|XY}$) such that $\Pr(\mathcal{E}) \geq 1 - \varepsilon$. Observe that a lower-bound of the form $H_\infty^\varepsilon(X|Y) \geq h$ implies that there exists an event $\mathcal{E}$ with $\Pr(\mathcal{E}) \geq 1 - \varepsilon$ such that, for all x and y , $\Pr[\mathcal{E}, X = x|Y = y] \leq 2^{-h}$.

The definition of a leaky OTM is weaker than that of an ideal OTM in two important respects: it assumes that the messages s and t are chosen uniformly at random, independent of all other variables; and it allows the adversary to obtain partial information about both s and t , so long as the adversary still has $(1 - \delta)k$ bits of uncertainty (as measured by the smoothed min-entropy). We suspect that this *definition* of a leaky OTM is not strong enough to construct one-time programs (although we conjecture that our actual *constructions* of OTM's in Sections 3 and 4 are, in fact, strong enough for this purpose).

2.5 Uncertainty Relations for the Min-entropy

We will use an uncertainty relation from [16], with a slight modification to describe quantum systems that consist of many non-identical subsystems:

Theorem 1. Consider a quantum system with Hilbert space $\bigotimes_{i=1}^{\ell_0} \mathbb{C}^{d_i}$, i.e., the system can be viewed as a collection of ℓ_0 subsystems, where the i 'th subsystem has Hilbert space dimension d_i .

For each $i \in [\ell_0]$, let B_i be a finite collection of orthonormal bases for $\mathbb{C}^{d_i}$, and suppose that these bases satisfy the following uncertainty relation: for every quantum state ρ on $\mathbb{C}^{d_i}$, $|B_i|^{-1} \sum_{\omega \in B_i} H(P_\omega) \geq h_i$, where P_ω is the distribution of measurement outcomes when ρ is measured in basis ω .

Now let ρ be any quantum state over $\bigotimes_{i=1}^{\ell_0} \mathbb{C}^{d_i}$, let $\Theta = (\Theta_1, \dots, \Theta_{\ell_0})$ be chosen uniformly at random from $B_1 \times \dots \times B_{\ell_0}$, and let $X = (X_1, \dots, X_{\ell_0})$ be the measurement outcome when ρ is measured in basis Θ (i.e., each X_i is the outcome of measuring subsystem i in basis Θ_i).

Then, for any $\tau > 0$, and any $\lambda_1, \dots, \lambda_{\ell_0} \in (0, \frac{1}{2})$, we have:

$$H_{\infty}^{\varepsilon}(X|\Theta) \geq -\tau + \sum_{i=1}^{\ell_0} (h_i - \lambda_i), \quad (4)$$

where $\varepsilon \leq \exp(-2\tau^2/c)$, and $c = \sum_{i=1}^{\ell_0} 16(\lg \frac{|B_i|d_i}{\lambda_i})^2$.

The proof is essentially the same as in [16]; it uses a martingale argument and Azuma's inequality, but it allows the martingale to have different increments at each step.

In addition, we will use the following chain rule for the smoothed min-entropy [20]:

$$H_{\infty}^{\varepsilon+\varepsilon'}(X|Y) > H_{\infty}^{\varepsilon}(X, Y) - H_0(Y) - \lg(\frac{1}{\varepsilon'}). \quad (5)$$

3 One-Time Memories

We now show the correctness and security of the OTM construction described in Section 1.1. Recall that this OTM uses $n \lg q$ qubits, stores two messages of length ℓ , and uses an error-correcting code C . We will show how to set n and q , and how to choose the code C .

Let us introduce some notation. We view the code C as a function $C : \{0, 1\}^{\ell} \rightarrow \{0, 1\}^{n \lg q}$. We view each codeword $x \in \{0, 1\}^{n \lg q}$ as a sequence of n blocks, where each block is a binary string of length $\lg q$. We write the codeword as $x = (x_{ij})_{i \in [n], j \in [\lg q]}$, and we write the i 'th block as $x_i = (x_{ij})_{j \in [\lg q]}$. Finally, let H be the Hadamard gate acting on a single qubit.

We now prepare the qubits in the OTM as follows. For each $i = 1, 2, \dots, n$,

- Let $\gamma_i \in \{0, 1\}$ be the outcome of a fair and independent coin toss.
- If $\gamma_i = 0$, prepare the i 'th block of qubits in the state $|C(s)_i\rangle$.
- If $\gamma_i = 1$, prepare the i 'th block of qubits in the state $H^{\otimes(\lg q)}|C(t)_i\rangle$.

To recover the first message s , we measure every qubit in the standard basis, which yields a string of measurement outcomes $z \in \{0, 1\}^{n \lg q}$, and then we run the decoding algorithm for C . To recover the second message t , we measure every qubit in the Hadamard basis, obtain a string of measurement outcomes z , and again run the decoding algorithm for C .

We will prove the following general theorem, which works for any code C that satisfies certain properties:

Theorem 2. *Let $q \geq 2$ be any power of 2. Let $\mathcal{E}_q$ be the q -ary symmetric channel with error probability $p_e = (1/2) - (1/2q)$. Let $\ell \geq 1$ and $n \geq 1$, and let $C : \{0, 1\}^{\ell} \rightarrow \{0, 1\}^{n \lg q}$ be any error-correcting code that satisfies the following two requirements:*

1. C can transmit information reliably over the channel $\mathcal{E}_q$ (where the channel treats each block of $\lg q$ bits as a single q -ary symbol).
2. C is a linear code over $GF(2)$.

Then the above OTM stores two messages $s, t \in \{0, 1\}^\ell$, and has the following properties:

1. The OTM behaves correctly for honest parties.
2. For any small constants $0 < \lambda \ll \frac{1}{2}$, $0 < \tau_0 \ll 1$, and $0 < \delta \ll 1$, the following statement holds. Suppose the messages s and t are chosen independently and uniformly at random in $\{0, 1\}^\ell$. For any separable adversary,³ we have the following security bound:

$$\begin{aligned}
 H_\infty^{\delta+\varepsilon}(S, T|Z) &\geq \left(\left(\frac{1}{2} - \lambda \right) - 4\tau_0 \left(1 + \frac{1}{\sqrt{\lg q}} \left(1 + \lg \frac{1}{\lambda} \right) \right) + \left(2 - \frac{1}{\alpha} \right) \right) \cdot \ell - \lg \frac{1}{\delta} \\
 &\gtrsim \left(\frac{1}{2} + \left(2 - \frac{1}{\alpha} \right) \right) \cdot \ell.
 \end{aligned} \tag{6}$$

Here S and T are the random variables describing the two messages, Z is the random variable representing the adversary's measurement outcome, we have $\varepsilon \leq \exp(-2\tau_0^2 \ell / \lg q)$, and $\alpha = \ell / (n \lg q)$ is the rate of the code C .

Note that, to get a strong security bound, one must use a code C whose rate α is large. It is useful to ask, then, how large α can be. Let L_q denote the capacity of the channel $\mathcal{E}_q$, measured in q -ary symbols per channel use. Using a good code C , we can hope to have rate $\alpha \approx L_q$. Moreover, L_q is lower-bounded by:

$$L_q \geq 1 - \frac{1}{\lg q} - p_e = \frac{1}{2} - \frac{1}{\lg q} + \frac{1}{2q} \approx \frac{1}{2}, \tag{7}$$

which is nearly tight when q is large. So we can hope to have $\alpha \approx \frac{1}{2}$, in which case our security bound becomes:

$$H_\infty^{\delta+\varepsilon}(S, T|Z) \gtrsim \frac{1}{2} \ell. \tag{8}$$

3.1 Correctness for Honest Parties

We first show the “correctness” part of Theorem 2. Without loss of generality, suppose we want to recover the first message s . (A similar argument applies if we want to recover the second message t .) Let $z \in \{0, 1\}^{n \lg q}$ be the string of measurement outcomes obtained by measuring each qubit in the standard basis. Observe that z is the output of a q -ary symmetric channel $\mathcal{E}_q$ with error probability $p_e = (1/2) - (1/2q)$, acting on the string $C(s) \in \{0, 1\}^{n \lg q}$ (viewed as a sequence of n symbols in $GF(q)$). Since the code C can transmit information reliably over this channel, it follows that we can recover s .

³ Note that this includes LOCC adversaries as a special case.

3.2 Security against Separable Adversaries

We now show the “security” part of Theorem 2. Let us first introduce some notation (see Figure 1). Suppose the OTM is programmed with two messages s and t that are chosen independently and uniformly at random in $\{0, 1\}^\ell$. Let S and T be the random variables representing these messages. Let Γ be the random variable representing the coin flips $\gamma = (\gamma_1, \dots, \gamma_n)$ used in programming the OTM. C denotes the error-correcting code, which maps $\{0, 1\}^\ell$ to $\{0, 1\}^{n \lg q}$. “Select” is an operation that maps $\{0, 1\}^{n \lg q} \times \{0, 1\}^{n \lg q}$ to $\{0, 1\}^{n \lg q}$, depending on the value of Γ , as follows:

$$\text{Select}(x, y)_{i,j} = \begin{cases} x_{i,j} & \text{if } \Gamma_i = 0, \\ y_{i,j} & \text{if } \Gamma_i = 1, \end{cases} \quad \text{for all } i \in [n], j \in [\lg q]. \quad (9)$$

“Select” outputs a string of $n \lg q$ classical bits, which are converted into $n \lg q$ qubits (in the standard basis states $|0\rangle$ and $|1\rangle$). H denotes a Hadamard gate controlled by the value of Γ ; that is, for each $i \in [n]$ and $j \in [\lg q]$, if $\Gamma_i = 1$, then H is applied to the (i, j) ’th qubit.

Fix any separable adversary $\mathcal{A}$, let L be the number of possible outcomes that can be observed by the adversary, and let $\mathcal{M} = \{M_z \mid z \in [L]\}$ be the separable POVM measurement performed by the adversary. Let Z be the random variable representing the adversary’s output; so Z takes values in $[L]$.

Fig. 1. OTM with separable adversary $\mathcal{A}$, and “fictional” adversary $\mathcal{A}^{(f)}$. In the proof, we will analyze the distributions of S and T conditioned on the events $Z = z$ and $Q = 0$.

Fix some small constant $\delta > 0$. We say that a measurement outcome $z \in [L]$ is “negligible” if $\Pr[Z = z] \leq (\delta/2^{n \lg q}) \text{tr}(M_z)$. Note that the probability of observing any of these “negligible” measurement outcomes is small:

$$\Pr[Z \text{ is “negligible”}] = \sum_{z \text{ “negl.”}} \Pr[Z = z] \leq (\delta/2^{n \lg q}) \sum_{z \text{ “negl.”}} \text{tr}(M_z) \leq \delta. \quad (10)$$

The proof will proceed as follows: for all messages $s, t \in \{0, 1\}^\ell$, and for all measurement outcomes $z \in [L]$ that are not “negligible,” we will upper-bound $\Pr[S = s, T = t \mid Z = z]$. This will imply a lower-bound on $H_\infty^\delta(S, T \mid Z)$, which is what we desire.

A Fictional Adversary. We begin by fixing some measurement outcome $z \in [L]$ that is not “negligible.” Since the adversary performed a separable measurement, we can write the corresponding POVM element M_z as a tensor product of single-qubit operators. In particular, we can write $M_z = \text{tr}(M_z) \bigotimes_{i=1}^n \bigotimes_{j=1}^{\lg q} R_{ij}$, where each R_{ij} is a single-qubit operator, positive semidefinite, with trace 1.

We now construct a fictional adversary $\mathcal{A}^{(f)}$, which we will use in the proof. The fictional adversary acts in the following way: for each qubit $(i, j) \in [n] \times [\lg q]$, it performs the POVM measurement $\{R_{ij}, I - R_{ij}\}$ on qubit (i, j) , which yields a binary measurement outcome Q_{ij} (where $Q_{ij} = 0$ corresponds to the POVM element R_{ij} , and $Q_{ij} = 1$ corresponds to $I - R_{ij}$). Let us write the vector of measurement outcomes as $Q = (Q_{ij})_{i \in [n], j \in [\lg q]}$, which takes values in $\{0, 1\}^{n \lg q}$. Let 0 denote the vector $(0, 0, \dots, 0) \in \{0, 1\}^{n \lg q}$.

Intuitively, the event $Q = 0$ (in an experiment using the fictional adversary) corresponds to the event $Z = z$ (in an experiment using the real adversary). More precisely, for any $s, t \in \{0, 1\}^\ell$, we have

$$\begin{aligned} P_{ST|Z}(s, t|z) &= \frac{P_{Z|ST}(z|s, t)P_{ST}(s, t)}{P_Z(z)} \\ &= \frac{P_{Q|ST}(0|s, t) \text{tr}(M_z)P_{ST}(s, t)}{P_Q(0) \text{tr}(M_z)} = P_{ST|Q}(s, t|0). \end{aligned} \quad (11)$$

We will proceed by upper-bounding $P_{ST|Q}(s, t|0)$ (with the fictional adversary); this will imply an upper-bound on $P_{ST|Z}(s, t|z)$ (with the real adversary).

Properties of the Codewords $C(S)$ and $C(T)$. Recall that the messages S and T are independently and uniformly distributed in $GF(2)^\ell$. Now consider the codewords $C(S)$ and $C(T)$. We claim that there exists a subset of ℓ coordinates of $C(S)$ and $C(T)$ that are independently and uniformly distributed in $GF(2)^\ell$.

To see this, recall that C is a linear code over $GF(2)$. Hence the encoding operation $C : GF(2)^\ell \rightarrow GF(2)^{n \lg q}$ can be written in the form $C(x) = x^T G$ for some matrix $G \in GF(2)^{\ell \times n \lg q}$. Since the codewords $C(x)$ are all distinct, the matrix G must have row-rank ℓ . Hence the column-rank of G must also be ℓ , so there exists a subset of ℓ columns of G that are linearly independent over $GF(2)$. Let us denote this subset by $A \subset [n] \times [\lg q]$, $|A| = \ell$.

Now look at those coordinates of $C(S)$ and $C(T)$ that correspond to the subset A ; we write these as $C(S)_A = (C(S)_{ij})_{(i,j) \in A}$ and $(C(T)_{ij})_{(i,j) \in A}$. It follows that $C(S)_A$ and $C(T)_A$ are independently and uniformly distributed in $GF(2)^\ell$.

Behavior of the Fictional Adversary on the Subset of Qubits A . We now analyze the behavior of the fictional adversary on those qubits belonging to the subset A . Without loss of generality, we can assume that the fictional adversary measures the qubits in the subset A first, and then measures the remaining qubits in the subset $([n] \times [\lg q]) \setminus A$. (This follows because the fictional adversary is *non-adaptive*, in that it makes all its decisions about what measurements to perform, before seeing any of the results of the measurements; and because all

of the measurements commute with one another, since each measurement only involves a single qubit.)

For convenience, let $B = ([n] \times [\lg q]) \setminus A$. Let $Q_A = (Q_{ij})_{(i,j) \in A}$ denote the measurement outcomes of the qubits in the subset A , and let $Q_B = (Q_{ij})_{(i,j) \in B}$ denote the measurement outcomes of the qubits in the subset B .

We claim that the OTM's coin tosses Γ , conditioned on the event $Q_A = 0$, are still uniformly distributed in $\{0, 1\}^n$. This is a fairly straightforward calculation; see the full paper [2] for details.

Using the Uncertainty Relation. We will upper-bound these probabilities $P_{ST|\Gamma Q_A}(s, t|\gamma, 0)$, using an entropic uncertainty relation. The basic idea is to consider another experiment, where one runs the OTM and the fictional adversary “backwards” in time. This experiment can be analyzed using the uncertainty relation in Theorem 1 (originally due to [16]).

We now describe this new experiment (see Figure 2). One prepares the quantum state $\bigotimes_{(i,j) \in A} R_{ij}$, one chooses a uniformly random sequence of measurement bases $\Theta = (\Theta_1, \dots, \Theta_n)$ (where $\Theta_i = 0$ denotes the standard basis and $\Theta_i = 1$ denotes the Hadamard basis), and then one measures each qubit $(i, j) \in A$ in the basis Θ_i to get a measurement outcome X_{ij} (which can be either 0 or 1).

Intuitively, the state $\bigotimes_{(i,j) \in A} R_{ij}$ corresponds to the fictional adversary's measurement outcome $Q_A = 0$, the random bases Θ correspond to the OTM's coin flips Γ , and the measurement outcomes X correspond to those bits $C(S)_A$ and $C(T)_A$ used in the OTM. (Note that the OTM's coin flips Γ are uniformly distributed, even when one conditions on the event $Q_A = 0$, as shown in the previous section.)

Fig. 2. In order to understand the behavior of the fictional adversary, conditioned on the event $Q_A = 0$, we consider an analogous experiment, where the state $\bigotimes_{(i,j) \in A} R_{ij}$ is measured in a random basis. We will analyze this using an entropic uncertainty relation.

To make this intuition precise, we will first show that:

$$H_{\infty}^{\varepsilon}(S, T|\Gamma, Q_A = 0) = H_{\infty}^{\varepsilon}(X|\Theta) + \ell. \quad (12)$$

(See the full version [2] for details.) Then note that conditioning on Γ can only reduce the entropy, hence we have:

$$H_{\infty}^{\varepsilon}(S, T|Q_A = 0) \geq H_{\infty}^{\varepsilon}(X|\Theta) + \ell. \quad (13)$$

We then use Theorem 1 to show a lower-bound on $H_{\infty}^{\varepsilon}(X|\Theta)$; see the full paper [2] for details.

Combining All the Pieces. The fictional adversary’s complete sequence of measurement outcomes is denoted by $Q = (Q_A, Q_B)$. So far we have analyzed the adversary’s actions on those qubits belonging to the subset A , and we have shown a lower-bound on $H_\infty^\varepsilon(S, T|Q_A = 0)$. Now, we will show a lower-bound on $H_\infty^\varepsilon(S, T|Q = 0)$. To do this, we bound the adversary’s actions on the subset B in a more-or-less trivial way, using the fact that $\Pr[Q = 0] = \Pr[Z = z]/\text{tr}(M_z) \geq \delta/2^{n \lg q}$, since z was assumed to be “non-negligible.”

We will then consider the real adversary, and show a lower-bound on $H_\infty^{\delta+\varepsilon}(S, T|Z)$. Here we use the following identity that relates the real adversary and the fictional adversary (see equation (11)):

$$H_\infty^\varepsilon(S, T|Z = z) = H_\infty^\varepsilon(S, T|Q = 0). \quad (14)$$

Finally we combine these results to prove the theorem; see the full paper [2] for details.

4 Efficient Implementations of One-Time Memories

In the previous section, we showed that one-time memories can be constructed from any code that approaches the capacity of the q -ary symmetric channel, and is linear over $GF(2)$. In this section, we will construct codes that have these properties, and moreover can be encoded and decoded efficiently. Using these codes, we will get efficient implementations of one-time memories.

Fig. 3. Efficient codes for the q -ary symmetric channel, based on erasure coding and error detection

There are several known constructions for codes that approach the capacity of the q -ary symmetric channel, and are efficiently decodable [22,23,24]. To illustrate how these techniques can be applied in our setting, we will describe one simple approach, which is based on erasure coding and error detection [23]. (See Figure 3.)

The basic idea is to take the message s , encode it using a code C_0 that outputs a string of q_0 -ary symbols (where $q_0 < q$), and then encode each q_0 -ary symbol using a code C_1 that outputs a q -ary symbol. The code C_1 is used to detect errors made by the q -ary symmetric channel; once detected, these errors can be

treated as erasures. The code C_0 is then used to correct these erasures, which is relatively straightforward. For instance, we can choose C_0 to be a random linear code; then we can decode in the presence of erasure errors by solving a linear system of equations, which we can do efficiently.

We now describe the construction in detail. Let $k \geq 2$ be an integer, let $p_e \in (0, 1)$, and choose any small constants $0 < \varepsilon \ll 1$, $0 < \delta \ll 1$ and $0 < \theta \ll 1$. Define:

$$n = \left\lfloor \frac{k}{1 - p_e - \theta} \right\rfloor, \quad (15)$$

$$q = 2^c, \quad c = \lg q = \left\lfloor \frac{2}{\delta} \right\rfloor \lceil \varepsilon n + \lg(np_e) \rceil, \quad (16)$$

$$q_0 = 2^{c_0}, \quad c_0 = \lg q_0 = \left\lceil \frac{2}{\delta} - 2 \right\rceil \lceil \varepsilon n + \lg(np_e) \rceil. \quad (17)$$

Note that our setting is slightly unusual, in that we will be constructing codes for the q -ary symmetric channel where q is not fixed. In particular, $\lg q$ (the number of bits used to describe each q -ary symbol) grows polynomially with the codeword length n , which is proportional to the message length k .

We will construct a code $C : \{0, 1\}^{k \lg q_0} \rightarrow \{0, 1\}^{n \lg q}$ as follows:

1. Choose a uniformly random matrix $G_0 \in GF(2)^{k \lg q_0 \times n \lg q_0}$, and define a code $C_0 : \{0, 1\}^{k \lg q_0} \rightarrow \{0, 1\}^{n \lg q_0}$ by setting $C_0(s) = s^T G_0$.
2. Fix any full-rank matrix $G_1 \in GF(2)^{\lg q_0 \times \lg q}$, and define a code $C_1 : \{0, 1\}^{\lg q_0} \rightarrow \{0, 1\}^{\lg q}$ by setting $C_1(v) = v^T G_1$.
3. Define $C(s) = C_1 \circ C_0(s)$, where we view $C_0(s) \in \{0, 1\}^{n \lg q_0}$ as a sequence of n blocks of $\lg q_0$ bits, and C_1 acts separately on each of these blocks. Equivalently, we can write $C(s) = s^T G_0 (\bigoplus_{i=1}^n G_1)$, where $\bigoplus_{i=1}^n G_1$ denotes a direct sum of n copies of the matrix G_1 .

We use the following decoding algorithm:

1. Given a string $z \in \{0, 1\}^{n \lg q}$, write it as a sequence of n blocks of $\lg q$ bits: $z = (z_{ij})_{i \in [n], j \in [\lg q]}$.
2. For each $i \in [n]$, try to decode the q -ary symbol $z_i \in \{0, 1\}^{\lg q}$, i.e., try to find some $v \in \{0, 1\}^{\lg q_0}$ such that $C_1(v) = z_i$. Let b_i be the result (or set $b_i = *$ if z_i lies outside the image of C_1). Thus we get a string $b = (b_1, b_2, \dots, b_n) \in (\{0, 1\}^{\lg q_0} \cup \{*\})^n$.
3. Try to decode the string b , treating the $*$ symbols as erasures, i.e., try to find some $a \in \{0, 1\}^{k \lg q_0}$ such that, for all $i \in [n]$ such that $b_i \neq *$, and for all $j \in [\lg q]$, $C_0(a)_{ij} = b_{ij}$. If a solution exists, output it; if there are multiple solutions, choose any one of them and output it; otherwise, abort.

Finally, we introduce some more notation. Let us choose a message (represented by a random variable S) uniformly at random in $\{0, 1\}^{k \lg q_0}$. Let $\mathcal{E}_q$ be the q -ary symmetric channel with error probability p_e . We take the message S , encode it using the code C , transmit it through the channel $\mathcal{E}_q$, then run the decoding algorithm, and get an estimate of the original message; call this $\hat{S}$.

We prove the following statement (see the full paper [2] for details):

Theorem 3. *Let $k \geq 2$ be an integer, let $p_e \in (0, 1)$, and choose any small constants $0 < \varepsilon \ll 1$, $0 < \delta \ll 1$ and $0 < \theta \ll 1$. Let us construct the code $C : \{0, 1\}^{k \lg q_0} \rightarrow \{0, 1\}^{n \lg q}$ as described above. Then C has the following properties:*

1. *With high probability (over the choice of the random matrix G_0), C can transmit information reliably over the q -ary symmetric channel $\mathcal{E}_q$ with error probability p_e .*

More precisely, choose any small constant τ such that $0 < \tau < \theta$, and choose any large constant $\lambda \gg 1$. Then, with probability $\geq 1 - \frac{1}{\lambda}$ (over the choice of G_0), the code C can transmit information over the channel $\mathcal{E}_q$, and the probability of decoding failure is bounded by:

$$\Pr[\hat{S} \neq S] \leq \lambda(e^{-2\tau^2 n} + 2^{-\varepsilon n} + 2^{(-n\theta + n\tau + 1) \lg q_0}) \leq e^{-\Omega(n)}. \quad (18)$$

2. *C is a linear code over $GF(2)$.*
3. *C has rate $\alpha := \frac{k \lg q_0}{n \lg q} \geq (1 - p_e - \theta)(1 - \delta)$. (Note that this approaches the capacity of the channel $\mathcal{E}_q$, as shown in equation (2), when q is large.)*
4. *The encoding and decoding algorithms for C run in time polynomial in $n \lg q$. (Also note that $\lg q$ grows at most linearly with n , and n is proportional to k .)*

Finally, we can take the code C constructed above (for $p_e = \frac{1}{2}$), and combine it with the OTM construction of Theorem 2, to get the following result:

Corollary 1. *For any $k \geq 2$, and for any small constant $0 < \mu \ll 1$, there exists an OTM construction that stores two messages $s, t \in \{0, 1\}^\ell$, where $\ell = \Theta(k^2)$, and has the following properties:*

1. *The OTM behaves correctly for honest parties.*
2. *The OTM can be implemented in time polynomial in k .*
3. *Let $0 < \delta \ll 1$ be any small constant. Suppose the messages s and t are chosen independently and uniformly at random in $\{0, 1\}^\ell$. For any separable adversary,⁴ we have the following security bound:*

$$H_\infty^{\delta+\varepsilon}(S, T|Z) \geq (\tfrac{1}{2} - \mu) \ell - \lg \tfrac{1}{\delta}. \quad (19)$$

Here S and T are the random variables describing the two messages, Z is the random variable representing the adversary's measurement outcome, and we have $\varepsilon \leq \exp(-\Omega(k))$.

Acknowledgements. It is a pleasure to thank Serge Fehr, Stephen Jordan, Maris Ozols, Rene Peralta, Eren Sasoglu, Christian Schaffner, Barbara Terhal, Alexander Vardy, and several anonymous reviewers, for helpful suggestions about this work. Some of these discussions took place at the Schloss Dagstuhl – Leibniz Center for Informatics. This paper is a contribution of NIST, an agency of the US government, and is not subject to US copyright.

⁴ Note that this includes LOCC adversaries as a special case.

References

1. Liu, Y.-K.: Building one-time memories from isolated qubits. In: 5th Conference on Innovations in Theoretical Computer Science (ITCS 2014), pp. 269–286 (2014)
2. Liu, Y.-K.: Single-shot security for one-time memories in the isolated qubits model. ArXiv:1402.0049
3. Goldwasser, S., Kalai, Y.T., Rothblum, G.N.: One-Time Programs. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 39–56. Springer, Heidelberg (2008)
4. Goyal, V., Ishai, Y., Sahai, A., Venkatesan, R., Wadia, A.: Founding Cryptography on Tamper-Proof Hardware Tokens. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 308–326. Springer, Heidelberg (2010)
5. Bellare, M., Hoang, V.T., Rogaway, P.: Adaptively Secure Garbling with Applications to One-Time Programs and Secure Outsourcing. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 134–153. Springer, Heidelberg (2012)
6. Broadbent, A., Gutoski, G., Stebila, D.: Quantum one-time programs. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 344–360. Springer, Heidelberg (2013)
7. Lo, H.-K., Chau, H.F.: Is quantum bit commitment really possible? Phys. Rev. Lett. 78, 3410 (1997)
8. Lo, H.-K.: Insecurity of quantum secure computations. Phys. Rev. A 56(2), 1154–1162 (1997)
9. Mayers, D.: Unconditionally secure quantum bit commitment is impossible. Phys. Rev. Lett. 78, 3414–3417 (1997)
10. Buhrman, H., Christandl, M., Schaffner, C.: Complete Insecurity of Quantum Protocols for Classical Two-Party Computation. Phys. Rev. Lett. 109, 160501 (2012)
11. Salvail, L.: Quantum Bit Commitment from a Physical Assumption. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 338–353. Springer, Heidelberg (1998)
12. Damgaard, I., Fehr, S., Salvail, L., Schaffner, C.: Cryptography In the Bounded Quantum-Storage Model. In: FOCS 2005, pp. 449–458 (2005)
13. Koenig, R., Terhal, B.M.: The Bounded Storage Model in the Presence of a Quantum Adversary. IEEE Trans. Inf. Th. 54(2) (2008)
14. Damgård, I.B., Fehr, S., Salvail, L., Schaffner, C.: Secure Identification and QKD in the Bounded-Quantum-Storage Model. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 342–359. Springer, Heidelberg (2007)
15. Wehner, S., Schaffner, C., Terhal, B.: Cryptography from Noisy Storage. Phys. Rev. Lett. 100, 220502 (2008)
16. Damgård, I.B., Fehr, S., Renner, R., Salvail, L., Schaffner, C.: A Tight High-Order Entropic Quantum Uncertainty Relation with Applications. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 360–378. Springer, Heidelberg (2007)
17. Bouman, N.J., Fehr, S., González-Guillén, C., Schaffner, C.: An all-but-one entropic uncertainty relation, and application to password-based identification. In: Kawano, Y. (ed.) TQC 2012. LNCS, vol. 7582, pp. 29–44. Springer, Heidelberg (2012)
18. Pastawski, F., Yao, N.Y., Jiang, L., Lukin, M.D., Cirac, J.I.: Unforgeable Noise-Tolerant Quantum Tokens. Proc. Nat. Acad. Sci. 109, 16079–16082 (2012)
19. Renner, R.: Security of Quantum Key Distribution. PhD thesis, ETH Zurich (2005)
20. Renner, R., Wolf, S.: Simple and Tight Bounds for Information Reconciliation and Privacy Amplification. In: Roy, B. (ed.) ASIACRYPT 2005. LNCS, vol. 3788, pp. 199–216. Springer, Heidelberg (2005)
21. Wiesner, S.: Conjugate coding. ACM SIGACT News 15(1), 78–88 (1983); original manuscript written circa 1970

22. Bleichenbacher, D., Kiayias, A., Yung, M.: Decoding of Interleaved Reed Solomon Codes over Noisy Data. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) ICALP 2003. LNCS, vol. 2719, pp. 97–108. Springer, Heidelberg (2003)
23. Shokrollahi, A.: Capacity-approaching codes on the q -ary symmetric channel for large q . In: ITW 2004, pp. 204–208 (2004)
24. Brown, A., Minder, L., Shokrollahi, M.A.: Improved Decoding of Interleaved AG Codes. In: Smart, N.P. (ed.) Cryptography and Coding 2005. LNCS, vol. 3796, pp. 37–46. Springer, Heidelberg (2005)
25. Bennett, C.H., DiVincenzo, D.P., Fuchs, C.A., Mor, T., Rains, E., Shor, P.W., Smolin, J.A., Wootters, W.K.: Quantum nonlocality without entanglement. *Phys. Rev. A* 59, 1070–1091 (1999)
26. Childs, A.M., Leung, D., Mancinska, L., Ozols, M.: A framework for bounding nonlocality of state discrimination. *Comm. Math. Phys.* 323, 1121–1153 (2013)
27. DiVincenzo, D.P., Leung, D.W., Terhal, B.M.: Quantum Data Hiding. *IEEE Trans. Inf. Theory* 48(3), 580–599 (2002)
28. Egging, T., Werner, R.F.: Hiding Classical Data in Multipartite Quantum States. *Phys. Rev. Lett.* 89, 097905 (2002)
29. DiVincenzo, D.P., Hayden, P., Terhal, B.M.: Hiding Quantum Data. *Found. Phys.* 33(11), 1629–1647 (2003)
30. Hayden, P., Leung, D., Smith, G.: Multiparty data hiding of quantum information. *Phys. Rev. A* 71, 062339 (2005)
31. Winter, A.: Coding theorem and strong converse for quantum channels. *IEEE Trans. Inform. Theory* 45(7), 2481–2485 (1999)
32. Akavia, A., Goldwasser, S., Vaikuntanathan, V.: Simultaneous Hardcore Bits and Cryptography against Memory Attacks. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 474–495. Springer, Heidelberg (2009)
33. Naor, M., Segev, G.: Public-Key Cryptosystems Resilient to Key Leakage. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 18–35. Springer, Heidelberg (2009)
34. Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information. Cambridge University Press (2000)
35. Maassen, H., Uffink, J.: Generalized Entropic Uncertainty Relations. *Phys. Rev. Lett.* 60, 1103 (1988)
36. Wehner, S., Winter, A.: Entropic uncertainty relations - A survey. *New J. Phys.* 12, 025009 (2010)
37. Dubhashi, D.P., Panconesi, A.: Concentration of Measure for the Analysis of Randomized Algorithms. Cambridge University Press (2009)
38. Horodecki, R., Horodecki, P., Horodecki, M., Horodecki, K.: Quantum Entanglement. *Rev. Mod. Phys.* 81, 865–942 (2009)
39. Saeedi, K., et al.: Room-Temperature Quantum Bit Storage Exceeding 39 Minutes Using Ionized Donors in Silicon-28. *Science* 342(6160), 830–833 (2013)
40. Dreau, A., et al.: Single-Shot Readout of Multiple Nuclear Spin Qubits in Diamond under Ambient Conditions. *Phys. Rev. Lett.* 110, 060502 (2013)