Secure Identification and QKD in the Bounded-Quantum-Storage Model

Ivan B. Damgård[1], Serge Fehr[2,⋆], Louis Salvail[1,⋆⋆], and Christian Schaffner[2,⋆⋆⋆]

[1] BRICS[†], FICS, Aarhus University, Denmark
{ivan|salvail}@brics.dk
[2] CWI[‡] Amsterdam, The Netherlands
{fehr|c.schaffner}@cwi.nl

Abstract. We consider the problem of secure identification: user U proves to server S that he knows an agreed (possibly low-entropy) password w, while giving away as little information on w as possible, namely the adversary can exclude at most one possible password for each execution of the scheme. We propose a solution in the bounded-quantum-storage model, where U and S may exchange qubits, and a dishonest party is assumed to have limited quantum memory. No other restriction is posed upon the adversary. An improved version of the proposed identification scheme is also secure against a man-in-the-middle attack, but requires U and S to additionally share a high-entropy key k. However, security is still guaranteed if one party loses k to the attacker but notices the loss. In both versions of the scheme, the honest participants need no quantum memory, and noise and imperfect quantum sources can be tolerated. The schemes compose sequentially, and w and k can securely be re-used. A small modification to the identification scheme results in a quantum-key-distribution (QKD) scheme, secure in the bounded-quantum-storage model, with the same re-usability properties of the keys, and without assuming authenticated channels. This is in sharp contrast to known QKD schemes (with unbounded adversary) without authenticated channels, where authentication keys must be updated, and unsuccessful executions can cause the parties to run out of keys.

1 Introduction

Secure Identification. Consider two parties, a *user* U and a *server* S, which share a common secret-key (or password or Personal Identification Number PIN)

⋆ Supported by the Dutch Organization for Scientific Research (NWO).

⋆⋆ QUSEP, Quantum Security in Practice, funded by the Danish Natural Science Research Council.

⋆⋆⋆ Supported by the European project SECOQC.

† Basic Research in Computer Science (www.brics.dk), and Foundations in Cryptography and Security, funded by the Danish Natural Sciences Research Council.

‡ Centrum voor Wiskunde en Informatica, the national research institute for mathematics and computer science in the Netherlands.

w. In order to obtain some service from S, U needs to convince S that he is the legitimate user U by "proving" that he knows w. In practice—think of how you prove to the ATM that you know your PIN—such a proof is often done simply by announcing w to S. This indeed guarantees that a dishonest user U^* who does not know w cannot identify himself as U, but of course incurs the risk that U might reveal w to a malicious server S^* who may now impersonate U. Thus, from a secure identification scheme we also require that a dishonest server S^* obtains (essentially) no information on w.

There exist various approaches to obtain secure identification schemes, depending on the setting and the exact security requirements. For instance zero-knowledge proofs (and some weaker versions), as initiated by Feige, Fiat and Shamir[12, 11], allow for secure identification. In a more sophisticated model, where we allow the common key w to be of low entropy and additionally consider a man-in-the-middle attack, we can use techniques from password-based key-agreement (like [14, 13]) to obtain secure identification schemes. Common to these approaches is that security relies on the assumption that some computational problem (like factoring or computing discrete logs) is hard and that the attacker has limited computing power.

Our Contribution. In this work, we take a new approach: we consider quantum communication, and we develop two identification schemes which are information-theoretically secure under the *sole* assumption that the attacker can only reliably store quantum states of limited size. This model was first considered in [4]. On the other hand, the honest participants only need to send qubits and measure them immediately upon arrival, no quantum storage or quantum computation is required. Furthermore, our identification schemes are robust to both noisy quantum channels and imperfect quantum sources. Our schemes can therefore be implemented in practice using off-the-shelf technology.

The first scheme is secure against dishonest users and servers but not against a man-in-the-middle attack. It allows the common secret-key w to be non-uniform and of low entropy, like a human-memorizable password. Only a user knowing w can succeed in convincing the server. In any execution of this scheme, a dishonest user or server cannot learn more on w than excluding one possibility, which is unavoidable. This is sometimes referred to as *password-based* identification. The second scheme requires in addition to w a uniformly distributed high-entropy common secret-key k, but is additionally secure against a man-in-the-middle attack. Furthermore, security against a dishonest user or server holds as for the first scheme even if the dishonest party knows k (but not w). This implies that k can for instance be stored on a smartcard, and security of the scheme is still guaranteed even if the smartcard gets stolen, assuming that the affected party notices the theft and thus does not engage in the scheme anymore. Both schemes compose sequentially, and w (and k) may be safely re-used super-polynomially many times, even if the identification fails (due to an attack, or due to a technical failure).

A small modification of the second identification scheme results in a quantum-key-distribution (QKD) scheme secure against bounded-quantum-memory adversaries. The advantage of the proposed new QKD scheme is that no authenticated channel is needed and the attacker can *not* force the parties to run out of authentication keys. The honest parties merely need to share a password w and a high-entropy secret-key k, which they can safely re-use (super-polynomially many times), independent of whether QKD succeeds or fails. Furthermore, like for the identification scheme, losing k does not compromise security as long as the loss is noticed by the corresponding party. One may think of this as a quantum version of password-based authenticated key exchange. The properties of our solution are in sharp contrast to all known QKD schemes without authenticated channels (which do not pose any restrictions on the attacker). In these schemes, an attacker can force parties to run out of authentication keys by making the QKD execution fail (e.g. by blocking some messages). Worse, even if the QKD execution fails only due to technical problems, the parties can still run out of authentication keys after a short while, since they cannot exclude that an eavesdropper was in fact present. This problem is an important drawback of QKD implementations, especially of those susceptible to single (or few) point(s) of failure[9].

Other Approaches. We briefly discuss how our identification schemes compare with other approaches. We have already given some indication on how to construct *computationally* secure identification schemes. This approach typically allows for very practical schemes, but requires some unproven complexity assumption. Another interesting difference between the two approaches: whereas for (known) computationally-secure password-based identification schemes the underlying computational hardness assumption needs to hold indefinitely, the restriction on the attacker's quantum memory in our approach only needs to hold *during* the execution of the identification scheme, actually only at one single point during the execution. In other words, having a super-quantum-storage-device at home in the basement only helps you cheat at the ATM if you can communicate with it on-line quantumly – in contrast to a computational solution, where an off-line super-computer in the basement can make a crucial difference.

Furthermore, obtaining a satisfactory identification scheme requires *some* restriction on the adversary, even in the quantum setting: considering only passive attacks, Lo[15] showed that for an unrestricted adversary, no password-based quantum identification scheme exists. In fact, Lo's impossibility result only applies if the user U is guaranteed not to learn anything about the outcome of the identification procedure. We can argue, however, that a different impossibility result holds even without Lo's restriction: We first show that secure computation of a classical AND gate (in which both players learn the output) can be reduced to a password-based identification scheme. The reduction works as follows. Let w_0, w_0' and w_1 be three distinct elements from $\mathcal{W}$. If Alice has private input $x_A = 0$ then she sets $w_A = w_0$ and if $x_A = 1$ then she sets $w_A = w_1$, and if Bob has private input $x_B = 0$ then he sets $w_B = w_0'$ and if $x_B = 1$ then he sets $w_B = w_1$. Then, Alice and Bob run the identification scheme on inputs w_A and w_B, and if the identification is rejected, the output is set to 0 while if it is accepted, the output is set to 1. Security of the identification scheme is easily seen to imply security of the AND computation. Now, the secure computation of an AND gate—with statistical security and using quantum communication—can be shown to require a superpolynomial number of rounds if the adversary is unbounded[18]. Therefore, the same must hold for a secure password-based identification scheme.[1]

Another alternative approach is the classical bounded-storage model[17, 2, 1]. In contrast to our approach, only classical communication is used, and it is assumed that the attacker's *classical* memory is bounded. Unlike in the quantum case where we do not need to require the honest players to have any quantum memory, the classical bounded-storage model requires honest parties to have a certain amount of memory which is related to the allowed memory size of the adversary: if two legitimate users need n bits of memory in an identification protocol meeting our security criterion, then an adversary must be bounded in memory to $O(n^2)$ bits. The reason is that given a secure password-based identification scheme, one can construct (in a black-box manner) a key-distribution scheme that produces a one-bit key on which the adversary has an (average) entropy of $\frac{1}{2}$. On the other hand it is known that in any key-distribution scheme which requires n bits of memory for legitimate players, an adversary with memory $\Omega(n^2)$ can obtain the key except for an arbitrarily small amount of remaining entropy[8]. It follows that password-based identification schemes in the classical bounded-storage model can only be secure against adversaries with memory at most $O(n^2)$. This holds even for identification schemes with only passive security and without security against man-in-the-middle attacks. Roughly, the reduction works as follows. Alice and Bob agree on a public set of two keys $\{w_0, w_1\}$. Alice picks $a \in_R \{0,1\}$, Bob picks $b \in_R \{0,1\}$, and they run the identification scheme with keys w_a and w_b respectively. The outcome of the identification is then made public from which Bob determines a. We argue that if the identification fails, i.e. $a \neq b$, then a is a secure bit. Thus, on average, a has entropy (close to) $\frac{1}{2}$ from an eavesdropper's point of view. Consider $w' \notin \{w_0, w_1\}$. By the security property of the identification scheme, Alice and thus also a passive eavesdropper Eve cannot distinguish between Bob having used w_b or w'. Similarly, we can then switch Alice's key w_a to w_{1-a} and Bob's switched key w' to w_{1-b} without changing Eve's view. Thus, Eve cannot distinguish an execution with $a = 0$ from one with $a = 1$ if $a \neq b$.

This limitation of the classical bounded-storage model is in sharp contrast with what we achieve in this paper, the honest players need no quantum memory at all while our identification scheme remains secure against adversaries with quantum memory linear in the total number of qubits sent. The same separation between the two models was shown for OT and bit commitment[4, 3].

Finally, if one settles for the bounded-quantum-storage model, then in principle one could take a generic construction for general two-party secure-function-

[1] In fact, we believe that the proof from [18] can be extended to cover secure computation of equality of strings, which is equivalent to password-based identification. This would mean that we could prove the impossibility result directly, without the detour via a secure AND computation. Details are omitted due to the space limitation.

evaluation (SFE) based on OT together with the OT scheme from [4, 3] in order to implement a SFE for string equality and thus password-based identification. However, this approach leads to a highly impractical solution, as the generic construction requires many executions of OT, whereas our solution is comparable with *one* execution of the OT scheme from [4, 3]. Furthermore, SFE does not automatically take care of a man-in-the-middle attack, thus additional work would need to be done using this approach.

2 Preliminaries

2.1 Notation and Terminology

QUANTUM STATES. The state of a *qubit* can be described by a vector in the 2-dimensional Hilbert space $\mathbb{C}^2$ in case of a *pure* state, and by a density matrix/operator on $\mathbb{C}^2$ in the general case of a *mixed* state. Similarly, an *n-qubit state* is characterized by a vector in the n-fold tensor product $(\mathbb{C}^2)^{\otimes n}$ in case of a *pure* n-qubit state, and by a density matrix/operator on $(\mathbb{C}^2)^{\otimes n}$ in case of a *mixed* n-qubit state. The pair $\{|0\rangle, |1\rangle\}$ denotes the standard basis, also known as computational or rectilinear or "+"-basis, for $\mathbb{C}^2$. When the context requires, we also write $|0\rangle_+$ and $|1\rangle_+$ instead of $|0\rangle$ respectively $|1\rangle$. The diagonal or "$\times$"-basis is defined as $\{|0\rangle_\times, |1\rangle_\times\}$ where $|0\rangle_\times = (|0\rangle + |1\rangle)/\sqrt{2}$ and $|1\rangle_\times = (|0\rangle - |1\rangle)/\sqrt{2}$. Measuring a qubit in the $+$-basis (resp. $\times$-basis) means applying the measurement described by projectors $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ (resp. projectors $|0\rangle_\times\langle 0|_\times$ and $|1\rangle_\times\langle 1|_\times$). The notation generalizes to n-qubit states: For $x = (x_1, \ldots, x_n) \in \{0,1\}^n$ and $\theta = (\theta_1, \ldots, \theta_n) \in \{+, \times\}^n$, we let $|x\rangle_\theta$ be the n-qubit state $|x\rangle_\theta = |x_1\rangle_{\theta_1} \cdots |x_n\rangle_{\theta_n}$; and measuring a n-qubit state in basis $\theta \in \{+, \times\}^n$ means applying the measurement described by projections $|x\rangle_\theta\langle x|_\theta$ with $x \in \{0,1\}^n$.

The behavior of a (mixed) quantum state in a register E is fully described by its density matrix ρ_{E}. In order to simplify language, we tend to be a bit sloppy and use E as well as ρ_{E} as "naming" for the quantum state. We often consider cases where a quantum state E may depend on some classical random variable X in that the state is described by the density matrix ρ_{E}^x if and only if $X = x$. For an observer who has only access to the state E but not to X, the behavior of the state is determined by the density matrix $\rho_{\mathsf{E}} := \sum_x P_X(x)\rho_{\mathsf{E}}^x$, whereas the joint state, consisting of the classical X and the quantum state E, is described by the density matrix $\rho_{X\mathsf{E}} := \sum_x P_X(x)|x\rangle\langle x| \otimes \rho_{\mathsf{E}}^x$, where we understand $\{|x\rangle\}_{x\in\mathcal{X}}$ to be the standard (orthonormal) basis of $\mathbb{C}^{|\mathcal{X}|}$. More general, for any event $\mathcal{E}$ (defined by $P_{\mathcal{E}|X}(x) = P[\mathcal{E}|X\!=\!x]$ for all x), we write

$$\rho_{X\mathsf{E}|\mathcal{E}} := \sum_x P_{X|\mathcal{E}}(x)|x\rangle\langle x| \otimes \rho_{\mathsf{E}}^x \quad \text{and} \quad \rho_{\mathsf{E}|\mathcal{E}} := \mathrm{tr}_X(\rho_{X\mathsf{E}|\mathcal{E}}) = \sum_x P_{X|\mathcal{E}}(x)\rho_{\mathsf{E}}^x \ .$$

We also write $\rho_X := \sum_x P_X(x)|x\rangle\langle x|$ for the quantum representation of the classical random variable X (and similarly for $\rho_{X|\mathcal{E}}$). This notation extends naturally to quantum states that depend on several classical random variables.

Given X and E as above, by saying that there exists a random variable Y such that ρ_{XYE} satisfies some condition, we mean that ρ_{XE} can be understood as $\rho_{XE} = \mathrm{tr}_Y(\rho_{XYE})$ for some ρ_{XYE} (with classical Y) and that ρ_{XYE} satisfies the required condition.

X is independent of E (in that ρ_E^x does not depend on x) if and only if $\rho_{XE} = \rho_X \otimes \rho_E$, which in particular implies that no information on X can be learned by observing only E. Similarly, X is random and independent of E if and only if $\rho_{XE} = \frac{1}{|\mathcal{X}|}\mathbb{I} \otimes \rho_E$, where $\frac{1}{|\mathcal{X}|}\mathbb{I}$ is the density matrix of the fully mixed state of suitable dimension. Finally, if two states like ρ_{XE} and $\rho_X \otimes \rho_E$ are ε-close in terms of their trace distance $\delta(\rho,\sigma) = \frac{1}{2}\,\mathrm{tr}(|\rho - \sigma|)$, which we write as $\rho_{XE} \approx_\varepsilon \rho_X \otimes \rho_E$, then the real system ρ_{XE} "behaves" as the ideal system $\rho_X \otimes \rho_E$ except with probability ε in that for any evolution of the system no observer can distinguish the real from the ideal one with advantage greater than ε [20].

We also need to express that a random variable X is (close to) independent of a quantum state E *when given a random variable* Y. This means that when given Y, the state E gives no (or little) additional information on X. Formally, this is expressed by requiring that ρ_{XYE} is of the form (or close to)

$$\rho_{XYE} = \sum_{x,y} P_{XY}(x,y)|x\rangle\langle x| \otimes |y\rangle\langle y| \otimes \rho_E^y \ ,$$

where $\rho_E^y = \sum_x P_{X|Y=y}(x)\rho_E^{x,y}$ for all y. As shorthand for the right-hand side above, we define $\rho_{X\leftrightarrow Y\leftrightarrow E} := \sum_{x,y} P_{XY}(x,y)|x\rangle\langle x| \otimes |y\rangle\langle y| \otimes \rho_E^y$.[2] To further illustrate its meaning, notice that if the Y-register is measured and value y is obtained, the state $\rho_{X\leftrightarrow Y\leftrightarrow E}$ collapses to $(\sum_x P_{X|Y=y}(x)|x\rangle\langle x|) \otimes \rho_E^y$, so that indeed no further information on x can be obtained from the E-register. This notation naturally extends to $\rho_{X\leftrightarrow Y\leftrightarrow E|\mathcal{E}}$ simply by considering $\rho_{XYE|\mathcal{E}}$.

(Conditional) Smooth Min-Entropy. We briefly recall the notion of (conditional) *smooth* min-entropy[19, 21]. For more details, we refer to the aforementioned literature. Let X be a random variable over alphabet $\mathcal{X}$ with distribution P_X. The notion of min-entropy is given by $H_\infty(X) = -\log\big(\max_x P_X(x)\big)$. More general, for any event $\mathcal{E}$, $H_\infty(X\mathcal{E})$ may be defined similarly simply by replacing P_X by $P_{X\mathcal{E}}$. Note that the "distribution" $P_{X\mathcal{E}}$ is not normalized; $H_\infty(X\mathcal{E})$ is still well defined, though. For an arbitrary $\varepsilon \geq 0$, the smooth version $H_\infty^\varepsilon(X)$ is defined as follows. $H_\infty^\varepsilon(X)$ is the *maximum* of the standard min-entropy $H_\infty(X\mathcal{E})$, where the maximum is taken over all events $\mathcal{E}$ with $\Pr(\mathcal{E}) \geq 1 - \varepsilon$. As ε can be interpreted as an error probability, we typically require ε to be negligible in the security parameter n, denoted as $\varepsilon = negl(n)$.

For a pair of random variables X and Y, the *conditional* smooth min-entropy $H_\infty^\varepsilon(X|Y)$ is defined as $H_\infty^\varepsilon(X|Y) = \max_\mathcal{E} \min_y H_\infty(X\mathcal{E}|Y=y)$, where the quantification over $\mathcal{E}$ is over all events $\mathcal{E}$ (defined by $P_{\mathcal{E}|XY}$) with $\Pr(\mathcal{E}) \geq 1-\varepsilon$. The following lemma shows that for a small ε, smooth min-entropy is essentially as good as ordinary min-entropy; the proof is given in the full version[5].

[2] The notation is inspired by the classical setting where the corresponding independence of X and Z given Y can be expressed by saying that $X \leftrightarrow Y \leftrightarrow Z$ forms a Markov chain.

Lemma 2.1. *If $H_\infty^\varepsilon(X|Y) = r$ then there exists an event $\mathcal{E}'$ such that $P[\mathcal{E}'] \geq 1 - 2\varepsilon$ and $H_\infty(X|\mathcal{E}', Y{=}y) \geq r - 1$ for every y with $P_{Y\mathcal{E}'}(y) > 0$.*

2.2 Tools

A New Min-Entropy-Splitting Lemma. A technical tool, which will come in handy, is the following new entropy-splitting lemma, which may be of independent interest. Informally, it says that if for a list of random variables, every pair has high (smooth) min-entropy, then all of the random variables except one must have high (smooth) min-entropy. The version given here follows immediately from the version given and proven in the full version[5].

Lemma 2.2 (Entropy-Splitting Lemma). *Let $\varepsilon > 0$. Let $X_1, \ldots, X_m$ be a sequence of random variables over $\mathcal{X}_1, \ldots, \mathcal{X}_m$ such that $H_\infty^\varepsilon(X_i X_j) \geq \alpha$ for all $i \neq j$. Then there exists a random variable V over $\{1, \ldots, m\}$ such that for any independent random variable W over $\{1, \ldots, m\}$*

$$H_\infty^{2\sqrt{\varepsilon}}(X_W|VW, V{\neq}W) \geq \alpha/2 - \log(m) - \log(1/\varepsilon) \quad .$$

Quantum Uncertainty Relation. At the very core of our security proofs lies (a special case of) the quantum uncertainty relation from [3], that lower bounds the (smooth) min-entropy of the outcome when measuring an arbitrary n-qubit state in a random basis $\theta \in \{0,1\}^n$.

Theorem 2.3 (Uncertainty Relation[3]). *Let E be an arbitrary fixed n-qubit state. Let Θ be uniformly distributed over $\{+, \times\}^n$ (independent of E), and let $X \in \{0,1\}^n$ be the random variable for the outcome of measuring E in basis Θ. Then, for any $\lambda > 0$, the conditional smooth min-entropy is lower bounded by $H_\infty^\varepsilon(X|\Theta) \geq \left(\frac{1}{2} - \lambda\right)n$ with $\varepsilon = negl(n)$.*

Thus, ignoring negligibly small "error probabilities" and linear fractions that can be chosen arbitrarily small, the outcome of measuring any n-qubit state in a random basis has $n/2$ bits of min-entropy, given the basis.

Privacy Amplification. Finally, we recall the quantum-privacy-amplification theorem of Renner and König[20]. We give the simplified version as used in [4]. Recall that a class $\mathcal{F}$ of hash functions from $\mathcal{X}$ to $\mathcal{Y}$ is called (strongly) universal-2 if for any $x \neq x' \in \mathcal{X}$, and for F uniformly distributed over $\mathcal{F}$, the collision probability $P[F(x) = F(x')]$ is upper bounded by $1/|\mathcal{Y}|$, respectively, for the strong notion, the random variables $F(x)$ and $F(x')$ are uniformly and independently distributed over $\mathcal{Y}$.

Theorem 2.4 (Privacy Amplification[20, 4]). *Let X be a random variable distributed over $\{0,1\}^n$, and let E be a q-qubit state that may depend on X. Let F be the random and independent choice of a member of a universal-2 class of hash functions $\mathcal{F}$ from $\{0,1\}^n$ into $\{0,1\}^\ell$. Then*

$$\delta\big(\rho_{F(X)F\mathsf{E}}, \tfrac{1}{2^\ell}\mathbb{I} \otimes \rho_{F\mathsf{E}}\big) \leq \frac{1}{2}\, 2^{-\frac{1}{2}\left(H_\infty(X) - q - \ell\right)} \; .$$

3 The Identification Scheme

3.1 The Setting

We assume the honest user U and the honest server S to share some key $w \in \mathcal{W}$. We do not require $\mathcal{W}$ to be very large (i.e. $|\mathcal{W}|$ may not be lower bounded by the security parameter in any way), and w does not necessarily have to be uniformly distributed in $\mathcal{W}$. So, we may think of w as a human-memorizable password or PIN code. The goal of this section is to construct an identification scheme that allows U to "prove" to S that he knows w. The scheme should have the following security properties: a dishonest server S^* learns essentially no information on w beyond that he can come up with a guess w' for w and learns whether $w' = w$ or not, and similarly a dishonest user succeeds in convincing the verifier essentially only if he guesses w correctly, and if his guess is incorrect then the only thing he learns is that his guess is incorrect. This in particular implies that as long as there is enough entropy in w, the identification scheme may be safely repeated.

3.2 The Intuition

The scheme we propose is related to the (randomized) 1-2 OT scheme of [3]. In that scheme, Alice sends $|x\rangle_\theta$ to Bob, for random $x \in \{0,1\}^n$ and $\theta \in \{+,\times\}^n$. Bob then measures everything in basis $+$ or $\times$, depending on his choice bit c, so that he essentially knows half of x (where Alice used the same basis as Bob) and has no information on the other half (where Alice used the other basis), though, at this point, he does not know yet which bits he knows and which ones he does not. Then, Alice sends θ and two hash functions to Bob, and outputs the hash values s_0 and s_1 of the two parts of x, whereas Bob outputs the hash value s_c that he is able to compute from the part of x he knows. It is proven in [3] that no dishonest Alice can learn c, and for any quantum-memory-bounded dishonest Bob, at least one of the two strings s_0 and s_1 is random for Bob.

This scheme can be extended by giving Bob more options for measuring the quantum state. Instead of measuring all qubits in the $+$ or the $\times$ basis, he may measure using m different strings of bases, where any two possible basis-strings have large Hamming distance. Then Alice computes and outputs m hash values, one for each possible basis-string that Bob might have used. She reveals θ and the hash functions to Bob, so he can compute the hash value corresponding to the basis that he has used, and no other hash value. Intuitively, such an extended scheme leads to a randomized 1-m OT.

The scheme can now be transformed into a secure identification scheme as follows, where we assume (wlog) that $\mathcal{W} = \{1,\dots,m\}$. The user U, acting as Alice, and the server S, acting as Bob, execute the randomized 1-m OT scheme where S "asks" for the string indexed by his key w, such that U obtains random strings $s_1,\dots,s_m$ and S obtains s_w. Then, to do the actual identification, U sends s_w to S, who accepts if and only if it coincides with his string s_w. Intuitively, such a construction is secure against a dishonest server since unless he asks for the right string (by guessing w correctly) the string U sends him is random and thus gives no information on w. On the other hand, a dishonest user does not know which of the m strings S asked for and wants to see from him. We realize this intuitive idea in the next section. In the actual protocol, U does not have to explicitly compute all the s_i's, and also we only need a single hash function (to compute s_w). We also take care of some subtleties, for instance that the s_i are not necessarily random if Alice (i.e. the user) is dishonest.

3.3 The Basic Scheme

Let $\mathfrak{c} : \mathcal{W} \to \{+, \times\}^n$ be the encoding function of a binary code of length n with $m = |\mathcal{W}|$ codewords and minimal distance d. $\mathfrak{c}$ can be chosen such that n is linear in $\log(m)$ or larger, and d is linear in n. Furthermore, let $\mathcal{F}$ and $\mathcal{G}$ be strongly universal-2 classes of hash functions[3] from $\{0,1\}^n$ to $\{0,1\}^\ell$ and from $\mathcal{W}$ to $\{0,1\}^\ell$, respectively, for some parameter ℓ. For $x \in \{0,1\}^n$ and $I \subseteq \{1,\dots,n\}$, we define $x|_I \in \{0,1\}^n$ to be the restriction of x to the coordinates x_i with $i \in I$. If $|I| < n$ then applying $f \in \mathcal{F}$ to $x|_I$ is to be understood as applying f to $x|_I$ padded with sufficiently many 0's.

Q-ID:
1. U picks $x \in_R \{0,1\}^n$ and $\theta \in_R \{+,\times\}^n$, and sends state $|x\rangle_\theta$ to S.
2. S measures $|x\rangle_\theta$ in basis $c = \mathfrak{c}(w)$. Let x' be the outcome.
3. U picks $f \in_R \mathcal{F}$ and sends θ and f to S. Both compute $I_w := \{i : \theta_i = \mathfrak{c}(w)_i\}$.
4. S picks $g \in_R \mathcal{G}$ and sends g to U.
5. U computes and sends $z := f(x|_{I_w}) \oplus g(w)$ to S.
6. S accepts if and only if $z = z'$ where $z' := f(x'|_{I_w}) \oplus g(w)$.

Proposition 3.1 (User security). *Let the initial state of a dishonest server S^*, whose quantum memory at step 3 is bounded by q qubits, be independent of the honest user's key W. Then, the joint state $\rho_{W E_{\mathsf{S}^*}}$ after the execution of Q-ID is such that there exists a random variable W' that is independent of W and such that*

$$\rho_{W W' E_{\mathsf{S}^*} | W' \neq W} \approx_\varepsilon \rho_{W \leftrightarrow W' \leftrightarrow E_{\mathsf{S}^*} | W' \neq W} \quad ,$$

where $\varepsilon = negl(d - 4\log(m) - 4q - 4\ell)$.

The proposition guarantees that whatever a dishonest S^* does is essentially as good as trying to guess W by some arbitrary (but independent) W' and learning whether the guess was correct or not, but nothing beyond that. Such a property is obviously the best one can hope for, since S^* may always honestly execute the protocol with a guess for W and observe whether he accepts U.

We would like to point out that the security definition used in Proposition 3.1, and in fact any security claim in this paper, guarantees *sequential composability*, as the output state is guaranteed to have the same independency property as is required from the input state (except if the attacker guesses w).

[3] Actually, we only need $\mathcal{G}$ to be *strongly* universal-2.

Proof. For readability, we do not keep track of negligibly small error probabilities and of linear fractions that can be chosen arbitrarily small, but (sometimes) merely give some indication of a small error by using the word "essentially". It is straightforward but rather tedious to keep rigorous track of these errors.

We consider and analyze a *purified* version of *Q-ID* where in step 1, instead of sending $|x\rangle_\theta$ to S^* for a random x, U prepares a fully entangled state $2^{-n/2}\sum_x |x\rangle|x\rangle$ and sends the second register to S^* while keeping the first. Then, in step 3 when the memory bound has applied, he measures his register in the random basis $\theta \in_R \{+,\times\}^n$ in order to obtain x. Standard arguments imply that this purified version produces exactly the same common state, consisting of the classical information on U's side and S^*'s quantum state.

Recall that before step 3 is executed, the memory bound applies to S^*, which means that S^* has to measure all but q of the qubits he holds, which consists of his initial state and his part of the EPR pairs. Before doing the measurement, he may append an ancilla register and apply an arbitrary unitary transform. As a result of S^*'s measurement, S^* gets some outcome y, and the common state collapses to a $(n+q)$-qubit state (which depends on y), where the first n qubits are with U and the remaining q with S^*. The following analysis is for a fixed y, and works no matter what y is.

We use upper case letters W, X, Θ, F, G and Z for the random variables that describe the respective values w, x, θ etc. in an execution of the purified version of *Q-ID*. We write $X_j = X|_{I_j}$ for any j, and we let E'_{S^*} be S^*'s q-qubit state at step 3, after the memory bound has applied. Note that W is independent of X, Θ, F, G and E'_{S^*}.

For $1 \leq i \neq j \leq m$, fix the value of X, and correspondingly of X_i and X_j, at the positions where $\mathfrak{c}(i)$ and $\mathfrak{c}(j)$ coincide, and focus on the remaining (at least) d positions. The uncertainty relation (Theorem 2.3) implies that the restriction of X to these positions has essentially $d/2$ bits of min-entropy given Θ. Since every bit in the restricted X appears in one of X_i and X_j, the pair X_i, X_j also has essentially $d/2$ bits of min-entropy given Θ. Lemma 2.2 implies that there exists W' (called V in Lemma 2.2) such that if $W \neq W'$ then X_W has essentially $d/4 - \log(m)$ bits of min-entropy, given W and W' (and Θ). Privacy amplification then guarantees that $F(X_W)$ is ε'-close to random and independent of F, W, W', Θ and E'_{S^*}, conditioned on $W \neq W'$, where $\varepsilon' = \frac{1}{2} \cdot 2^{-\frac{1}{2}(d/4-\log(m)-q-\ell)}$. It follows that $Z = F(X_W) \oplus G(W)$ is ε'-close to random and independent of F, G, W, W', Θ and E'_{S^*}, conditioned on $W \neq W'$. Formally, we want to upper bound

$$\delta\big(\rho_{WW'E_{\mathsf{S}^*}|W'\neq W},\ \rho_{W\leftrightarrow W'\leftrightarrow E_{\mathsf{S}^*}|W'\neq W}\big)\ .$$

Since the output state E_{S^*} is, without loss of generality, obtained by applying some unitary transform to the set of registers $(Z, F, G, W', \Theta, E'_{\mathsf{S}^*})$, the distance above is equal to $\delta(\rho_{WW'(Z,F,G,\Theta,E'_{\mathsf{S}^*})|W'\neq W},\ \rho_{W\leftrightarrow W'\leftrightarrow(Z,F,G,\Theta,E'_{\mathsf{S}^*})|W'\neq W})$. We then get:

$$\begin{aligned}\rho_{WW'(Z,F,G,\Theta,E'_{\mathsf{S}^*})|W'\neq W} &\approx_{\varepsilon'} \tfrac{1}{2^\ell}\mathbb{I}\otimes\rho_{WW'(F,G,\Theta,E'_{\mathsf{S}^*})|W'\neq W}\\ &= \tfrac{1}{2^\ell}\mathbb{I}\otimes\rho_{W\leftrightarrow W'\leftrightarrow(F,G,\Theta,E'_{\mathsf{S}^*})|W'\neq W} \approx_{\varepsilon'} \rho_{W\leftrightarrow W'\leftrightarrow(Z,F,G,\Theta,E'_{\mathsf{S}^*})|W'\neq W}\ ,\end{aligned}$$

where approximations follow from privacy amplification and the exact equality comes from the independency of W, which, when conditioned on $W' \neq W$, translates to independency given W'. The claim follows, with $\varepsilon = 2\varepsilon'$. □

Proposition 3.2 (Server security). *Let the initial state of an (unbounded) dishonest user* U^* *be independent of the honest server's key* W*, and let* $H_\infty(W) \geq 1$. *Then, there exists* W'*, independent of* W*, such that if* $W \neq W'$ *then* S *accepts with probability at most* $m^2/2^\ell$*, and the common state* $\rho_{W\mathsf{U}^*}$ *after the execution of* Q-ID *satisfies*

$$\rho_{WW'E_{\mathsf{U}^*}|W'\neq W} \approx_{m^2/2^{\ell-1}} \rho_{W\leftrightarrow W'\leftrightarrow \mathsf{U}^*|W'\neq W} \; .$$

The formal proof is given in the full version[5]. The idea is the following. We let U^* execute *Q-ID* with a server that is *unbounded* in quantum memory. Such a server can obviously obtain x and thus compute $s_j = f(x|_{I_j}) \oplus g(j)$ for all j. Note that s_w is the message z that U^* is required to send in the last step. Now, if the s_j's are all distinct, then z uniquely defines w' such that $z = s_{w'}$, and thus S accepts if and only if $w' = w$, and U^* does not learn anything beyond. The strong universal-2 property of g guarantees that the s_j's are all distinct except with probability $m^2/2^\ell$.

We call an identification scheme *ε-secure against impersonation attacks* if user and sender security are satisfied as in Propositions 3.1 and 3.2. The following holds.

Theorem 3.3. *If* $H_\infty(W) \geq 1$*, then the identification scheme* Q-ID *(with suitable choice of parameters) is ε-secure against impersonation attacks for any unbounded user and for any server with quantum memory bound* q*, where* $\varepsilon = negl\big(n - 33\log(m) - 11q\big)$.

Proof. We choose $\ell = \frac{1}{8}(d + 4\log(m) - 4q)$. Then user security holds except with an "error" negligible in $d - 4\log(m) - 4q - 4\ell = d/2 - 6\log(m) - 2q$, and thus negligible in $d - 12\log(m) - 4q$. And server security holds except with an "error" negligible in $\ell - 1 - 2\log(m) = \frac{1}{8}(d - 12\log(m) - 4q) - 1$, and thus negligible in $d - 12\log(m) - 4q$. Using a code $\mathfrak{c}$ which asymptotically meets the Gilbert-Varshamov bound[22], d may be chosen arbitrarily close to $n \cdot h^{-1}\big(1 - \log(m)/n\big)$, where h^{-1} is the inverse function of the binary entropy function $h : p \mapsto -\big(p \cdot \log(p) + (1-p)\cdot\log(1-p)\big)$ restricted to $0 < p \leq \frac{1}{2}$. For this d to be larger than $12\log(m)$, clearly n needs to be larger than $24\log(m)$, so that $h^{-1}\big(1 - \log(m)/n\big) > h^{-1}\big(1 - \frac{1}{24}\big)$ which turns out to be larger than $\frac{4}{11}$. The claim follows by normalizing $\frac{4}{11}n - 12\log(m) - 4q$ for n. □

3.4 An Error-tolerant Scheme

We now consider an imperfect quantum channel with "error rate" ϕ. The scheme *Q-ID* is sensitive to such errors in that they cause $x|_{I_w}$ and $x'|_{I_w}$ to be different and thus an honest server S is likely to reject an honest user U. This problem can be overcome by means of error-correcting techniques: U chooses a linear error-correcting code that allows to correct a ϕ-fraction of errors, and then in step 2, in addition to θ and f, U sends a description of the code and the syndrome s of $x|_{I_w}$ to S; this additional information allows S to recover $x|_{I_w}$ from its noisy version $x'|_{I_w}$ by standard techniques. However, this technique introduces a new problem: the syndrome s of $x|_{I_w}$ may give information on w to a dishonest server. Hence, to circumvent this problem, the code chosen by U must have the additional property that for a dishonest user, who has high min-entropy on $x|_{I_w}$, the syndrome s is (close to) independent of w.

This problem has recently been addressed and solved in the classical setting by Dodis and Smith[7]. They present a family of efficiently decodable linear codes allowing to correct a constant fraction of errors, and where the syndrome of a string is close to uniform if the string has enough min-entropy and the code is chosen at random from the family.[4] It remains to verify that their analysis can be translated to our setting where the adversary may have "quantum information".

Lemma 5 of [7] guarantees that for every $0 < \lambda < 1$ and for an infinite number of n''s there exists a δ-*biased* (as defined in [7]) family $\mathcal{C} = \{C_j\}_{j\in\mathcal{J}}$ of $[n', k', d']_2$-codes with $\delta < 2^{-\lambda n'/2}$, and which allows to efficiently correct a constant fraction of errors. Theorem 3.2 of [10] (which generalizes Lemma 4 in [7] to the quantum setting) guarantees that if a string Y has t bits of min-entropy then for a randomly chosen code $C_j \in \mathcal{C}$, the syndrome of Y is close to random and independent of j and any q-qubit state that may depend on Y, where the closeness is given by $\delta \cdot 2^{(n'+q-t)/2}$. In our application, $Y = X_W$, $n' \approx n/2$ and $t \approx d/4 - \log(m) - \ell$, where the additional loss of ℓ bits of entropy comes from learning the ℓ-bit string z. Choosing $\lambda = 1 - \frac{t}{2n'}$ gives an ensemble of code families that allow to correct a linear number of errors and the syndrome is ε-close to uniform given the quantum state, where $\varepsilon \leq 2^{-n'/2+t/4} \cdot 2^{(n'+q-t)/2} = 2^{-(t-2q)/4}$, which is exponentially small provided that there is a linear gap between t and $2q$. Thus, the syndrome gives essentially no additional information. The error rate ϕ that can be tolerated this way depends in a rather complicated way on λ, but choosing λ larger, for instance $\lambda = 1 - \frac{t+\nu q}{2n'}$ for a constant $\nu > 0$, allows to tolerate a higher error rate but requires q to be a smaller (but still constant) fraction of t.

Another imperfection has to be taken into account in current implementations of the quantum channel: imperfect sources. An imperfect source transmits more than one qubit in the same state with probability η independently each time a new transmission takes place. To deal with imperfect sources, we freely give away (x_i, θ_i) to the adversary when a multi-qubit transmission occurs in position i. It is not difficult to see that parameter ε in Proposition 3.1 then becomes essentially $\varepsilon = negl((1-\eta)d - 4\log(m) - 4q - 4\ell)$ in this case.

[4] As a matter of fact, the error correction in [7] is done by sending the string XOR'ed with a random code word, rather than sending the syndrome, but obviously the latter is equivalent to the first.

It follows that a quantum channel with error-rate ϕ and multi-pulse rate η, called the (ϕ, η)-weak quantum model in [4], can be tolerated for some small enough (but constant) ϕ and η.

4 Defeating Man-in-the-Middle Attacks

4.1 The Approach

In the previous section, we "only" proved security against impersonation attacks, but we did not consider a man-in-the-middle attack, where the attacker sits between an honest user and an honest server and controls their (quantum and classical) communication. And indeed, *Q-ID* is highly insecure against such an attack: the attacker may measure the first qubit in, say, basis $+$, and then forward the collapsed qubit (together with the remaining untouched ones) and observe if S accepts the session. If not, then the attacker knows that he introduced an error and hence that the first qubit must have been encoded and measured using the $\times$-basis, which gives him one bit of information on the key w. The error-tolerant scheme seems to prevent this particular attack, but it is by no means clear that it is secure against *any* man-in-the-middle attack.

To defeat a man-in-the-middle attack that tampers with the quantum communication, we perform a check of correctness on a random subset. The check allows to detect if the attacker tampers too much with the quantum communication, and the scheme can be aborted before sensitive information is leaked to the attacker. In order to protect the classical communication, one might use a standard information-theoretic authentication code. However, the key for such a code can only be securely used a limited number of times. A similar problem occurs in QKD: even though a successful QKD execution produces fresh key material that can be used in the next execution, the attacker can have the parties run out of authentication keys by repeatedly enforcing the executions to fail. In order to overcome this problem, we will use some special authentication scheme allowing to re-use the key under certain circumstances, as discussed in Sect. 4.3.

4.2 The Setting

Similar to before, we assume that the user U and the server S share a not necessarily uniform, low-entropy key w. In order to handle the stronger security requirements of this section, we have to assume that U and S in addition share a uniform high-entropy key k. We require that a man-in-the-middle attacker needs to guess w correctly in order to break the scheme, and if his guess is incorrect then he learns no more information on w besides that his guess is wrong, and he essentially learns no information on k. Furthermore, we require security against impersonation attacks, as defined in the previous section, *even if the dishonest party knows* k. It follows that k can for instance be stored on a smartcard, and security is still guaranteed even if the smartcard gets stolen, assuming that the theft is noticed and the corresponding party does/can not execute the scheme anymore. We would also like to stress that by our security notion, not only w but also k may be safely reused, even if the scheme was under attack.

4.3 An Additional Tool: Extractor MACs

An important tool used in this section is an authentication scheme, i.e., a Message Authentication Code (MAC), that also acts as an extractor, meaning that if there is high min-entropy in the message, then the key-tag pair cannot be distinguished from the key and a random tag. Such a MAC, introduced in [6], is called an extractor MAC, EXTR-MAC for short. For instance $MAC^*_{\alpha,\beta}(x) = [\alpha x] + \beta$, where $\alpha, x \in GF(2^n)$, $\beta \in GF(2^\ell)$ and $[\,.\,]$, denotes truncation to the ℓ first bits, is an EXTR-MAC: impersonation and substitution probability are $1/2^\ell$, and, for an arbitrary message X, a random key $K = (A, B)$ and the corresponding tag $T = [A \cdot X] + B$, the tag-key pair (T, K) is $2^{-(H_2(X)-\ell)/2}$-close to (U, K), where U is the uniform distribution, respectively, ρ_{TKE} is $2^{(H_2(X)-\ell-q)/2}$-close to $\frac{1}{2^\ell}\mathbb{I} \otimes \rho_{KE} = \frac{1}{2^\ell}\mathbb{I} \otimes \rho_K \otimes \rho_E$ if we allow a q-qubit state E that may depend only on X. A useful feature of an EXTR-MAC is that if an adversary gets to see the tag of a message on which he has high min-entropy, then the key for the MAC can be safely re-used (sequentially). Indeed, closeness of the real state, ρ_{TKE}, to the ideal state, $\frac{1}{2^\ell}\mathbb{I} \otimes \rho_{KE} = \frac{1}{2^\ell}\mathbb{I} \otimes \rho_K \otimes \rho_E$, means that no matter how the state evolves, the real state behaves like the ideal one (except with small probability), but of course in the ideal state, K is still "fresh" and can be reused.

4.4 The Scheme

As for *Q-ID*, let $\mathfrak{c} : \mathcal{W} \to \{+, \times\}^n$ be the encoding function of a binary code of length n with $m = |\mathcal{W}|$ codewords and minimal distance d. For some parameter ℓ, let $\mathcal{F}$, $\mathcal{G}$ and $\mathcal{H}$ be strongly universal-2 classes of hash functions from $\{0,1\}^n$ to $\{0,1\}^\ell$, $\mathcal{W}$ to $\{0,1\}^\ell$, and $\{0,1\}^n$ to $\{0,1\}^{2\ell}$, respectively. Also, let $MAC : \{0,1\}^{2\ell} \times \{0,1\}^* \to \{0,1\}^\ell$ be a standard MAC for a message of arbitrary length L, with an 2ℓ-bit key and an error probability at most $\lceil L/\ell \rceil \cdot 2^{-\ell}$, and let $MAC^* : \mathcal{K} \times \mathcal{M} \to \{0,1\}^\ell$ be an EXTR-MAC with an arbitrary key space $\mathcal{K}$, a (finite) message space $\mathcal{M}$ that will become clear later, and an error probability $2^{-\ell}$. Furthermore, let $\{syn_j\}_{j\in\mathcal{J}}$ be the family of syndrome functions[5] corresponding to a family $\mathcal{C} = \{C_j\}_{j\in\mathcal{J}}$ of linear error correcting codes of size $n' = n/2$, as discussed in Section 3.4: any C_j allows to efficiently correct a δ-fraction of errors for some constant $\delta > 0$, and for a random $j \in \mathcal{J}$, the syndrome of a string with $t = d/4 - \log(m) - 5\ell$ bits of min-entropy is $2^{-(t-2q)/4}$-close to uniform (given j and any q-qubit state). Finally, we let $\ell^* \leq \ell$ be a parameter linear in $n - \ell$, whose exact value will be specified in the proof.

Recall, by the set-up assumption, the user U and the server S share a password $w \in \mathcal{W}$ as well as a uniform high-entropy key, which we define to be a random authentication key $k \in \mathcal{K}$. The scheme is given in the box below.

Proposition 4.1 (Security against man-in-the-middle). *Let the initial state of a man-in-the-middle attacker with quantum memory q be independent of the*

[5] We agree on the following convention: for a bit string y of arbitrary length, $syn_j(y)$ is to be understood as $syn_j(y0\cdots0)$ with enough padded zeros if its bit length is smaller than n', and as $\big(syn_j(y'), y''\big)$, where y' consist of the first n' and y'' of the remaining bits of y, if its bit length is bigger than n'.

Q-ID$^+$:

1. U picks $x \in_R \{0,1\}^n$ and $\theta \in_R \{+,\times\}^n$, and sends the n-qubit state $|x\rangle_\theta$ to S. Write $I_w := \{i : \theta_i = \mathfrak{c}(w)_i\}$.
2. S picks a random subset $T \subset \{1,\ldots,n\}$ of size ℓ^*, it computes $c = \mathfrak{c}(w)$, replaces every c_i with $i \in T$ by $c_i \in_R \{+,\times\}$ and measures $|x\rangle_\theta$ in basis c. Let x' be the outcome, and let $test' := x'|_T$.
3. U sends θ, $j \in_R \mathcal{J}$, $s := syn_j(x|_{I_w})$, $f \in_R \mathcal{F}$, $h \in_R \mathcal{H}$ and $tag^* := MAC^*_k(\theta, j, s, f, h, x|_{I_w})$ to S.
4. S picks $g \in \mathcal{G}$, and sends T and g to U.
5. U sends $test := x|_T$, $z := f(x|_{I_w}) \oplus g(w)$ and $tag := MAC_{h(x|_{I_w})}(g, T, test, z)$ to S.
6. S recovers $x|_{I_w}$ from $x'|_{I_w}$ with the help of $test$ and s, and it accepts if and only if (1) both MAC's verify correctly, (2) $test$ coincides with $test'$ wherever the bases coincide, and (3) $z = f(x|_{I_w}) \oplus g(w)$.

keys W and K. Then, there exists W', independent of W, such that the common state ρ_{KWE} after the execution of Q-ID$^+$ satisfies

$$\rho_{KWW'E|W'\neq W} \approx_\varepsilon \rho_K \otimes \rho_{W\leftrightarrow W'\leftrightarrow E|W'\neq W} \quad ,$$

where $\varepsilon = negl(d - 4\log(m) - 8q - 20\ell)$.

Proof. We use capital letters (W, Θ, etc.) for the values (w, θ, etc.) occurring in the scheme whenever we view them as random variables, and we write X_W and X'_W for the random variables taking values $x|_{I_w}$ and $x'|_{I_w}$, respectively. To simplify the argument, we neglect error probabilities that are of order ε, as well as linear fractions that can be chosen arbitrarily small. We merely give indication of a small error by (sometimes) using the word "essentially".

First note that due to the security of the MAC and its key, if the attacker substitutes θ, j, s, f or h in step 3, or if S recovers an incorrect string as $x|_{I_w}$, then S will reject at the end of the protocol. We can define W' (independent of W) as in the proof of Proposition 3.1 such that if $W \neq W'$ then X_W has essentially $d/4 - \log(m)$ bits of min-entropy, given W, W' and Θ. Furthermore, given $TAG^*, F(X_W), H(X_W), TEST$ (as well as K, F, H, T, W, W' and Θ), X_W has still essentially $t = d/4 - \log(m) - 5\ell$ bits of min-entropy, if $W \neq W'$. By the property of the code family $\mathcal{C}$, it follows that if $t > 2q$ with a linear gap then the syndrome $S = syn_J(X_W)$ is essentially random and independent of $J, TAG^*, F(X_W), H(X_W), TEST, K, F, H, T, W, W', \Theta$ and E, conditioned on $W \neq W'$. Furthermore, it follows from the privacy-amplifying property of MAC^* and of f and h that if $d/4 - \log(m) - 5\ell > q$ with a linear gap, then the set of values $(TAG^*, F(X_W), H(X_W))$ is essentially random and independent of $K, F, H, TEST, T, W, W', \Theta$ and E, conditioned on $W \neq W'$. Finally, K is independent of the rest, and E is independent of $K, F, H, TEST, T, W, \Theta$. It follows that $\rho_{KWW'E|W'\neq W} \approx \rho_K \otimes \rho_{W\leftrightarrow W'\leftrightarrow E|W'\neq W}$, *before* he learns S's decision to accept or reject.

It remains to argue that S's decision does not give any additional information on W. We will make a case distinction, which does not depend on w, and we will show for both cases that S's decision to accept or reject is independent of w, which proves the claim. But first, we need the following observation. Recall that outside of the test set T, S measured in the bases dictated by w, but within T in random bases. Let I'_w be the subset of positions $i \in I_w$ with $c_i = \mathfrak{c}(w)_i$ (and thus also $= \theta_i$), and let $T' = T \cap I'_w$. In other words, we remove the positions where S measured in the "wrong" basis. The size of T' is essentially $\ell^*/4$, and given its size, it is a random subset of I'_w of size $|T'|$. It follows from the theory of random sampling, specifically from Lemma 4 of [16], that $\nu\big(x|_{I'_w}, x'|_{I'_w}\big)$ essentially equals $\nu\big(x|_{T'}, x'|_{T'}\big)$ (except with probability negligible in the size of T'), where $\nu(\cdot,\cdot)$ denotes the fraction of errors between the two input strings. Due to some technical reason, for the sampling technique to work it is required that $|T'|$ is upper bounded by $\alpha \cdot |I'_w|$, where the constant $\alpha > 0$ depends on the allowed tolerance in estimating the error fraction, and as such on δ, the fraction of errors the code C_j is able to correct. We refer to [16] for more details. Important for us is that ℓ^* can be chosen linear in $n-\ell$. Furthermore, since the set $V = \{i \in T : \theta_i = c_i\}$ of positions where U and S compare x and x' is a superset of T' of essentially twice the size, $\nu\big(x|_V, x'|_V\big)$ is essentially lower bounded by $\frac{1}{2}\,\nu\big(x|_{T'}, x'|_{T'}\big)$. Putting things together, we get that $\nu\big(x|_{I'_w}, x'|_{I'_w}\big)$ is essentially upper bounded by $2\,\nu\big(x|_V, x'|_V\big)$. Also note that $\nu\big(x|_V, x'|_V\big)$ does not depend on w. We can now do the case distinction: *Case 1:* If $\nu\big(x|_V, x'|_V\big) \le \frac{\delta}{2}$ (minus an arbitrarily small value), then $x|_{I'_w}$ and $x'|_{I'_w}$ differ in at most a δ-fraction of their positions, and thus S correctly recovers $x|_{I_w}$ (using $test = x|_T$ to get $x|_{I_w \setminus I'_w}$ and using s to correct the rest), no matter what w is, and it follows that S's decision only depends on the attacker's behavior, but not on w. *Case 2:* Otherwise, either S cannot correctly recover $x|_{I_w}$ and thus rejects, or it can correctly recover $x|_{I_w}$ and hence can verify tag with the correct key $h(x|_{I_w})$. S is therefore guaranteed to get the correct $test = x|_T$ (or else rejects) and thus rejects as $test$ and $test'$, restricted to V, differ in more than a $\frac{\delta}{2}$-fraction of their positions. Hence, S always rejects in case 2. □

For a dishonest user or server who knows k (but not w), breaking *Q-ID*$^+$ is equivalent to breaking *Q-ID*, up to a change in the parameters. Doing the maths on the parameters, it thus follows:

Theorem 4.2. *If $H_\infty(W) \ge 1$, then the identification scheme Q-ID$^+$ is ε-secure against a man-in-the-middle attacker with quantum memory bound q, and, even with a leaked k, Q-ID$^+$ is ε-secure against impersonation attacks for any unbounded user and for any server with quantum memory bound q, where $\varepsilon = negl(n - 100\log(m) - 19q)$.*

It is easy to see that *Q-ID*$^+$ can tolerate a noisy quantum communication up to any error rate $\phi < \delta$. Similar to the discussion in Section 3.4, tolerating a higher error rate requires the bound on the adversary's quantum memory to be smaller but still linear in the number of qubits transmitted. Imperfect sources can also be addressed in a similar way as for *Q-ID*. It follows that *Q-ID*$^+$ can also be shown secure in the (ϕ,η)-weak quantum model provided ϕ and η are small enough constants.

5 Application to QKD

As already pointed out in Section 4.1, current QKD schemes have the shortcoming that if there is no classical channel available that is authenticated by physical means, and thus messages need to be authenticated by an information-theoretic authentication scheme, an attacker can force the parties to run out of authentication keys simply by making an execution (or several executions if the parties share more key material) fail. Even worse, even if there is no attacker, but some execution(s) of the QKD scheme fails due to a technical problem, parties could run out of authentication keys. This shortcoming could make the technology impractical in situations where denial of service attacks or technical interruptions often occur.

The identification scheme $\textsf{\textit{Q-ID}}^+$ from the previous section immediately gives a QKD scheme *in the bounded-quantum-storage model* that allows to re-use the authentications key(s). Actually, we can inherit the key-setting from $\textsf{\textit{Q-ID}}^+$, where there are two keys, a human-memorizable password and a uniform, high-entropy key, where security is still guaranteed even if the latter gets stolen and the theft is noticed. In order to agree on a secret key sk, the two parties execute $\textsf{\textit{Q-ID}}^+$, and extract sk from $x|_{I_w}$ by applying yet another strongly universal-2 function, for instance chosen by U in step 3, where n needs to be increased accordingly to have the additional necessary amount of entropy in $x|_{I_w}$. The analysis of $\textsf{\textit{Q-ID}}^+$ immediately implies that if honest S accepts, then he is convinced to share sk with the legitimate U which knows w. In order to convince U, S can then use part of sk to one-time-pad encrypt w, and send it to U. The rest of sk is then a secure secret key, shared between U and S. In order to have a better "key rate", instead of using sk (minus the part used for the one-time-pad encryption) as secret key, one can also run a standard QKD scheme on top of $\textsf{\textit{Q-ID}}^+$ and use sk as a one-time authentication key.

References

1. Y. Aumann, Y. Z. Ding, and M. O. Rabin. Everlasting security in the bounded storage model. *IEEE Transactions on Information Theory*, 48(6):1668–1680, June 2002.
2. C. Cachin, C. Crépeau, and J. Marcil. Oblivious transfer with a memory-bounded receiver. In *39th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 493–502, 1998.
3. I. B. Damgård, S. Fehr, R. Renner, L. Salvail, and C. Schaffner. A tight high-order entropic quantum uncertainty relation with applications. 2007. In these CRYPTO'07 proceedings.
4. I. B. Damgård, S. Fehr, L. Salvail, and C. Schaffner. Cryptography in the bounded quantum-storage model. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 449–458, 2005.
5. I. B. Damgård, S. Fehr, L. Salvail, and C. Schaffner. Secure identification and QKD in the bounded-quantum-storage model. 2007. Available at `http://eprint.iacr.org/2007/`.

6. Y. Dodis, J. Katz, L. Reyzin, and A. Smith. Robust fuzzy extractors and authenticated key agreement from close secrets. In *Advances in Cryptology—CRYPTO '06*, volume 4117 of *Lecture Notes in Computer Science*, pages 232–250. Springer, 2006.
7. Y. Dodis and A. Smith. Correcting errors without leaking partial information. In *37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 654–663, 2005.
8. S. Dziembowski and U. M. Maurer. On generating the initial key in the bounded-storage model. In *Advances in Cryptology—EUROCRYPT '04*, volume 3027 of *Lecture Notes in Computer Science*, pages 126–137. Springer, 2004.
9. C. Elliott, D. Pearson, and G. Troxel. Quantum cryptography in practice. In *SIGCOMM '03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 227–238, 2003.
10. S. Fehr and C. Schaffner. Randomness extraction via delta-biased masking in the presence of a quantum attacker. Available at `http://eprint.iacr.org/2007/`, 2007.
11. U. Feige, A. Fiat, and A. Shamir. Zero knowledge proofs of identity. In *19th Annual ACM Symposium on Theory of Computing (STOC)*, pages 210–217, 1987.
12. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Advances in Cryptology—CRYPTO '86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, 1986.
13. R. Gennaro and Y. Lindell. A framework for password-based authenticated key exchange. In *Advances in Cryptology—EUROCRYPT '03*, volume 2656 of *Lecture Notes in Computer Science*, pages 524–543. Springer, 2003.
14. J. Katz, R. Ostrovsky, and M. Yung. Efficient password-authenticated key exchange using human-memorable passwords. In *Advances in Cryptology—EUROCRYPT '01*, volume 2045 of *Lecture Notes in Computer Science*, pages 473–492. Springer, 2001.
15. H.-K. Lo. Insecurity of quantum secure computations. *Physical Review A*, 56(2):1154–1162, 1997.
16. H.-K. Lo, H. F. Chau, and M. Ardehali. Efficient quantum key distribution scheme and a proof of its unconditional security. *Journal of Cryptology*, 18(2):133–165, March 2005.
17. U. M. Maurer. A provably-secure strongly-randomized cipher. In *Advances in Cryptology—EUROCRYPT '90*, volume 473 of *Lecture Notes in Computer Science*, pages 361–373. Springer, 1990.
18. J. B. Nielsen, T. B. Pedersen, and L. Salvail. Secure two-party quantum computation against semi-honest adversaries. In preparation, 2007.
19. R. Renner. *Security of Quantum Key Distribution.* PhD thesis, ETH Zürich, 2005. `http://arxiv.org/abs/quant-ph/0512258`.
20. R. Renner and R. König. Universally composable privacy amplification against quantum adversaries. In *Theory of Cryptography Conference (TCC)*, volume 3378 of *Lecture Notes in Computer Science*, pages 407–425. Springer, 2005.
21. R. Renner and S. Wolf. Simple and tight bounds for information reconciliation and privacy amplification. In *Advances in Cryptology—ASIACRYPT 2005*, Lecture Notes in Computer Science, pages 199–216. Springer, 2005.
22. C. Thommesen. The existence of binary linear concatenated codes with reed - solomon outer codes which asymptotically meet the gilbert- varshamov bound. *IEEE Transactions on Information Theory*, 29(6):850–853, 1983.