Cryptography from quantum uncertainty in the presence of quantum side information

Bouman, N.J.

5

Hybrid Security of Password-Based Identification

The content of this chapter is based on joint work with Serge Fehr, Carlos González-Guillén and Christian Schaffner [BFGS12].

Chapter Contents

5.1 Introduction

In this chapter, we propose a new entropic uncertainty relation. Furthermore, we present a modified version of the quantum identification protocol QID introduced in Section 2.11, that we will refer to as NEWQID. We will show how the new uncertainty relation can be used to prove NEWQID secure in the bounded-quantum-storage model (BQSM). Moreover, we will introduce another security model in this chapter, which we call the *single-qubit-operations model* (SQOM), and show that NEWQID is also secure in this model.

5.1.1 A New Uncertainty Relation

Uncertainty relations are quantitative characterizations of the uncertainty principle of quantum mechanics, which expresses that for certain pairs of measurements, there exists no state for which the measurement outcome is determined for *both* measurements: at least one of the outcomes must be somewhat uncertain. *Entropic* uncertainty relations express this uncertainty in at least one of the measurement outcomes by means of an entropy measure, usually the Shannon entropy. Our new entropic uncertainty relation distinguishes itself from previously known uncertainty relations by the following collection of features:

1. It uses the *min-entropy* as entropy measure, rather than the Shannon entropy. Such an uncertainty relation is sometimes also called a *high-order* entropic uncertainty relation.[1] Since privacy amplification needs a lower bound on the min-entropy, high-order entropic uncertainty relations are useful tools in quantum cryptography.

2. It lower bounds the uncertainty in the measurement outcome for *all but one* measurements, chosen from an *arbitrary* (and arbitrarily large) family of possible measurements. This is clearly *stronger* than typical entropic uncertainty relations that lower bound the uncertainty on *average* (over the choice of the measurement).

3. The measurements can be chosen to be qubit-wise measurements, in the computational or Hadamard basis, and thus the uncertainty relation is applicable to settings that can be implemented using current technology.

[1] This is because the min-entropy coincides with the Rényi entropy H_α of high(est) order $\alpha = \infty$ (see Section 2.3.1).

To the best of our knowledge, no previous entropic uncertainty relation satisfies (1) and (2) simultaneously, let alone in combination with (3). Indeed, as pointed out in a recent overview article by Wehner and Winter [WW10], little is known about entropic uncertainty relations for more than two measurement outcomes, and even less when additionally considering min-entropy.

Explanation by means of a Simpler Entropic Uncertainty Relation

To explain our new uncertainty relation, we find it helpful to first discuss a simpler variant, which does not satisfy (1), and which follows trivially from known results. Fix an arbitrary family $\{\mathcal{B}_1, \ldots, \mathcal{B}_m\}$ of bases for a given quantum system, and let us denote the state space of this given system by $\mathcal{H}$. The *maximum overlap* of such a family is defined as the real number

$$c := \max\{|\langle\phi|\psi\rangle| : |\phi\rangle \in \mathcal{B}_j, |\psi\rangle \in \mathcal{B}_k, 1 \leq j < k \leq m\}.$$

Let $d := -\log(c^2)$. Furthermore, let $\rho \in \mathcal{D}(\mathcal{H})$ be an arbitrary quantum state, and let X denote the measurement outcome when ρ is measured in one of the bases. We model the choice of the basis by a random variable J, so that $H(X|J=j)$ denotes the Shannon entropy of the measurement outcome when ρ is measured in basis $\mathcal{B}_j$. It follows immediately from Maassen and Uffink's uncertainty relation [MU88] that

$$H(X|J=j) + H(X|J=k) \geq -\log(c^2) = d \quad \forall j \neq k.$$

As a direct consequence, there exists a choice j' for the measurement so that $H(X|J=j) \geq \frac{d}{2}$ for all $j \in \{1, \ldots, m\}$ with $j \neq j'$. In other words, for any state ρ there exists j' so that unless the choice for the measurement coincides with j', which happens with probability at most $\max_j P_J(j)$, there is at least $d/2$ bits of entropy in the outcome X.

Our new high-order entropic uncertainty relation shows that this very statement essentially still holds when we replace Shannon by min-entropy, except that j' becomes randomized: for any ρ, there exists a *random variable* J', independent of J, such that[2]

$$H_{\min}(X|J=j, J'=j') \gtrsim \frac{d}{2} \quad \forall\, j, j' \in [m] \text{ such that } j \neq j'$$

no matter what the distribution of J is. Thus, unless the measurement J coincides with J', there is roughly $d/2$ bits of min-entropy in the outcome X. Furthermore,

[2] The rigorous version of the approximate inequality $\gtrsim$ is stated in Theorem 5.3.

since J' is *independent* of J, the probability that J coincides with J' is at most $\max_j P_J(j)$, as is the case for a fixed J'.

Note that we have no control over (the distribution of) J'. We can merely guarantee that it exists and is independent of J. It may be insightful to interpret J' as a *virtual guess* for J, guessed by the party that prepares ρ, and whose goal is to have little uncertainty in the measurement outcome X. The reader may think of the following specific way of preparing ρ: sample j' according to some arbitrary distribution J', and then prepare the state as the, say, first basis vector of $\mathcal{B}_{j'}$. If the resulting mixture ρ is then measured in some basis $\mathcal{B}_j$, sampled according to an arbitrary (independent) distribution J, then unless $j = j'$ (i.e., our guess for j was correct), there is obviously lower bounded uncertainty in the measurement outcome X (assuming a non-trivial maximum overlap). Our uncertainty relation can be understood as saying that for *any* state ρ, no matter how it is prepared, there exists such a (virtual) guess J', which exhibits this very behavior: if it differs from the actual choice for the measurement then there is lower bounded uncertainty in the measurement outcome X. As an immediate consequence, we can for instance say that X has min-entropy at least $d/2$, except with a probability that is given by the probability of guessing J, e.g., except with probability $1/m$ if the measurement is chosen uniformly at random from the family. This is clearly the best we can hope for.

We stress that because the min-entropy is more conservative than the Shannon entropy, our high-order entropic uncertainty relation does not follow from its simpler Shannon-entropy version. Neither can it be deduced in an analogue way; the main reason being that for fixed pairs $j \neq k$, there is no strong lower bound on $H_{\min}(X|J=j) + H_{\min}(X|J=k)$, in contrast to the case of Shannon entropy. More precisely and more generally, the *average* uncertainty $\frac{1}{|J|}\sum_j H_{\min}(X|J=j)$ does not allow a lower bound higher than $\log|J|$. To see this, consider the following example for $|J| = 2$ (the example can easily be extended to arbitrary $|J|$). Suppose that ρ is the uniform mixture of two pure states, one giving no uncertainty when measured in basis j, and the other giving no uncertainty when measured in basis k. Then, $\frac{1}{2}H_{\min}(X|J=j) + \frac{1}{2}H_{\min}(X|J=k) = 1$. Because of a similar reason, we cannot hope to get a good bound for all but a *fixed* choice of j'; the probabilistic nature of J' is necessary (in general). Hence, compared to bounding the average uncertainty, the all-but-one form of our uncertainty relation not only makes our uncertainty relation stronger in that uncertainty for all-but-one implies uncertainty on average (yet not vice versa), but it also allows for *more* uncertainty.

By using asymptotically good error-correcting codes, one can construct families

$\{\mathcal{B}_1, \ldots, \mathcal{B}_m\}$ of bases that have a large value of d, and thus for which our uncertainty relation guarantees a large amount of min-entropy (we discuss this in more detail in Section 5.2.1). These families consist of qubit-wise measurements in the computational or the Hadamard basis, and thus are implementable with current technology.

The proof of our new uncertainty relation comprises a rather involved probability reasoning to prove the existence of the random variable J' and builds on earlier work presented in [Sch07].

Quantum Identification with Hybrid Security

As an application of our entropic uncertainty relation, we propose a new *quantum identification protocol* (for an introduction into quantum identification, see Section 2.11). Our uncertainty relation gives us the right tool to prove security of the new quantum identification protocol in the BQSM. The distinguishing feature of our new protocol is that it also offers some security in case the assumption underlying the BQSM fails to hold. Indeed, we additionally prove security of our new protocol against a dishonest server that has unbounded quantum storage capabilities and can reliably store all the qubits communicated during an execution of the protocol, but is restricted to non-adaptive single-qubit operations and measurements. [3] This is in sharp contrast to protocol `QID` by Damgård *et al.* (Section 2.11.1), which completely breaks down against a dishonest server that can store all the communicated qubits in a quantum memory and postpone the measurements until the user announces the correct measurement bases. On the downside, our protocol only offers security in case of a perfect single-qubit (e.g., single-photon) source, because multi-qubit emissions reveal information about w. Hence, given the immature state of single-qubit-source technology (as of 2012), our protocol is currently mainly of theoretical interest.

We want to stress that proving security of our protocol in this *single-qubit-operations model* (SQOM) is non-trivial. Indeed, as we will see, standard tools like privacy amplification are not applicable. Our proof involves certain properties of random linear codes and makes use of Diaconis and Shahshahani's XOR inequality (Theorem 2.8, see also [Dia88]).

[3] Because secure identification belongs to the class of secure 2PC functionalities, it is well known that *some* restriction is necessary (for references, see Section 1.3.2).

5.1.2 Related Work

The study of *entropic* uncertainty relations, whose origin dates back to 1957 with the work of Hirschman [Hir57], has received a lot of attention over the last decade due to their various applications in quantum information processing. We refer the reader to [WW10] for a recent overview on entropic uncertainty relations. Most of the known entropic uncertainty relations are of the form

$$\frac{1}{|J|}\sum_j H_\alpha(X|J=j) \geq h\,,$$

where H_α is the Rényi entropy.[4] I.e., most uncertainty relations only give a lower bound on the entropy of the measurement outcome X *on average* over the (random) choice of the measurement. As argued in Section 5.1.1, the bound h on the *min*-entropy can be at most $\log|J|$, no matter the range of X. Furthermore, an uncertainty relation of this form only guarantees that there is uncertainty in X for *some* measurement(s), but does not specify precisely for how many, and certainly it does not guarantee uncertainty for *all but one* measurements. The same holds for the high-order entropic uncertainty relation from [DFR$^+$07], which considers an exponential number of measurement settings and guarantees that except with negligible probability over the (random) choice of the measurement, there is lower-bounded min-entropy in the outcome. On the other hand, the high-order entropic uncertainty relation from [DFSS05] only considers *two* measurement settings and guarantees lower-bounded min-entropy with probability (close to) $\frac{1}{2}$.

The uncertainty relation we know of that comes closest to ours is Lemma 2.13 in [FHS11]. Using our notation, it shows that X is ϵ-close to having roughly $d/2$ bits of min-entropy (i.e., the same bound we get), but only for all but an ϵ-fraction of all the m possible choices for the measurement j, where ϵ is about $\sqrt{2/m}$.

With respect to our application, backing up the security of the identification protocol by Damgård *et al.* [DFSS07] against an adversary that can overcome the quantum-memory bound assumed by the BQSM was also the goal of [DFL$^+$09]. However, the solution proposed there relies on an unproven computational hardness assumption, and as such, strictly speaking, can be broken by an adversary in the SQOM, i.e., by storing qubits and measuring them later qubit-wise and performing (possibly infeasible) classical computations. On the other hand, by *assuming* a lower bound on the hardness of the underlying computational problem against quantum machines, the security of the protocol in [DFL$^+$09] holds against an adversary with much

[4] See Section 2.3.1 for the definition of H_α. Nevertheless, for most known uncertainty relations $\alpha = 1$, i.e., the Shannon entropy.

more quantum computing power than our protocol in the SQOM, which restricts the adversary to single-qubit operations.

We hope that with future research on this topic, new quantum identification (or other cryptographic) protocols will be developed with security in the same spirit as our protocol, but with a more relaxed restriction on the adversary's quantum computation capabilities, for instance that he can only perform a limited number of quantum computation steps, and in every step he can only act on a limited number of qubits coherently.

5.2 An All-But-One Entropic Uncertainty Relation

Throughout this section, $\{\mathcal{B}_1, \ldots, \mathcal{B}_m\}$ is an arbitrary but fixed family of bases for the state space $\mathcal{H}$ of a quantum system. For simplicity, we restrict our attention to an n-qubit system, such that $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ for $n \in \mathbb{N}$, but our results immediately generalize to arbitrary quantum systems. We write the 2^n basis vectors of the j-th basis $\mathcal{B}_j$ as $\mathcal{B}_j = \{|x\rangle_j : x \in \{0,1\}^n\}$. Let c be the maximum overlap of $\{\mathcal{B}_1, \ldots, \mathcal{B}_m\}$ as defined in Section 5.1.1.

In order to obtain our entropic uncertainty relation that lower bounds the min-entropy of the measurement outcome for all but one measurement, we first state an uncertainty relation that expresses uncertainty by means of the probability measure of given sets.

Theorem 5.1 (cf. Thm. 4.18 in [Sch07]) *Let ρ be an arbitrary state of n qubits. For $j \in [m]$, let $Q^j(\cdot)$ be the distribution of the outcome when ρ is measured in the $\mathcal{B}_j$-basis, i.e., $Q^j(x) := \langle x|_j \, \rho \, |x\rangle_j$ for any $x \in \{0,1\}^n$. And for all subsets $\mathcal{X} \subset \{0,1\}^n$, let $Q^j(\mathcal{X}) := \sum_{x \in \mathcal{X}} Q^j(x)$. Then, for any family $\{\mathcal{L}^j\}_{j \in [m]}$ of subsets $\mathcal{L}^j \subset \{0,1\}^n$, it holds that*

$$\sum_{j \in [m]} Q^j(\mathcal{L}^j) \leq 1 + c\,(m-1) \cdot \max_{\substack{j,k \in [m] \\ j \neq k}} \sqrt{|\mathcal{L}^j||\mathcal{L}^k|}.$$

A special case of Theorem 5.1, obtained by restricting the family of bases to the specific choice $\{\mathcal{B}_+, \mathcal{B}_\times\}$ with $\mathcal{B}_+ = \{|x\rangle : x \in \{0,1\}^n\}$ and $\mathcal{B}_\times = \{H^{\otimes n}|x\rangle : x \in \{0,1\}^n\}$ (i.e., either the computational or Hadamard basis for all qubits), is an uncertainty relation that was proven and used in the original paper about the BQSM [DFSS05]. The proof of Theorem 5.1 goes along similar lines as the proof in the journal version of [DFSS05] for the special case outlined above. The proof of Theorem 5.1 can be found in [Sch07], as well as in [BFGS12].

In the same spirit as Corollary 4.17 in [Sch07] (see also the full version of [DFSS05]), we reformulate above uncertainty relation in terms of a "good event" $\mathcal{E}$, which occurs with reasonable probability, and if it occurs, then the measurement outcomes have high min-entropy. The statement is obtained by choosing the sets $\mathcal{L}^j$ in Theorem 5.1 appropriately.

Because we now switch to entropy notation, it will be convenient to work with a measure of overlap between bases that is logarithmic in nature and *relative* to the number n of qubits. Hence, we define

$$\delta := -\frac{1}{n} \log c^2 \,.$$

We will later see that for "good" choices of bases, δ stays constant for growing n.

Corollary 5.2 *Let ρ be an arbitrary n-qubit state, let J be a random variable over $[m]$ (with arbitrary distribution P_J), and let X be the outcome when measuring ρ in basis $\mathcal{B}_J$.*[5] *Then, for any $\epsilon \in \mathbb{R}$ such that $0 < \epsilon < \delta/4$, there exists an event $\mathcal{E}$ such that*

$$\sum_{j\in[m]} \Pr[\mathcal{E}|J=j] \geq (m-1) - (2m-1)\cdot 2^{-\epsilon n}$$

and

$$H_{\min}(X|J=j,\mathcal{E}) \geq \Big(\frac{\delta}{2} - 2\epsilon\Big)n$$

for $j \in [m]$ with $P_{J|\mathcal{E}}(j) > 0$.

Proof. For $j \in [m]$ define

$$\mathcal{S}^j := \{x \in \{0,1\}^n : Q^j(x) \leq 2^{-(\delta/2-\epsilon)n}\}$$

to be the sets of strings with small probabilities and denote by $\mathcal{L}^j := \overline{\mathcal{S}}^j$ their complements.[6] Note that for all $x \in \mathcal{L}^j$, we have that $Q^j(x) > 2^{-(\delta/2-\epsilon)n}$ and therefore $|\mathcal{L}^j| < 2^{(\delta/2-\epsilon)n}$. It follows from Theorem 5.1 that

$$\begin{aligned}\sum_{j\in[m]} Q^j(\mathcal{S}^j) &= \sum_{j\in[m]} (1 - Q^j(\mathcal{L}^j)) \geq m - (1 + (m-1)\cdot 2^{-\epsilon n}) \\ &= (m-1) - (m-1)2^{-\epsilon n}.\end{aligned}$$

We define $\mathcal{E} := \{X \in \mathcal{S}^J \wedge Q^J(\mathcal{S}^J) \geq 2^{-\epsilon n}\}$ to be the event that $X \in \mathcal{S}^J$ and at the same time the probability that this happens is not too small. Then $\Pr[\mathcal{E}|J=$

[5] I.e., $P_{X|J}(x|j) = Q^j(x)$, using the notation from Theorem 5.1.

[6] Here's the mnemonic: $\mathcal{S}$ for the strings with *s*mall probabilities, $\mathcal{L}$ for *l*arge.

$j] = \Pr[X \in \mathcal{S}^j \wedge Q^j(\mathcal{S}^j) \geq 2^{-\epsilon n} | J = j]$ either vanishes (if $Q^j(\mathcal{S}^j) < 2^{-\epsilon n}$) or else equals $Q^j(\mathcal{S}^j)$. In either case, $\Pr[\mathcal{E}|J=j] \geq Q^j(\mathcal{S}^j) - 2^{-\epsilon n}$ holds and thus the first claim follows by summing over $j \in [m]$ and using the derivation above. Furthermore, let $p = \max_j P_J(j)$, then $\Pr[\bar{\mathcal{E}}] = \sum_{j\in[m]} P_J(j) \Pr[\bar{\mathcal{E}}|J=j] \leq p \sum_{j\in[m]} \Pr[\bar{\mathcal{E}}|J=j] \leq p(m - (\sum_{j\in[m]} Q^j(\mathcal{S}^j) - 2^{-\epsilon n})) \leq p(1 + (2m-1)\cdot 2^{-\epsilon n})$, and $\Pr[\mathcal{E}] \geq (1-p) - p(2m-1)\cdot 2^{-\epsilon n}$

Regarding the second claim, in case $J = j$, we have

$$\begin{aligned} H_{\min}(X|J=j,\mathcal{E}) &= -\log\left(\max_{x\in\mathcal{S}^j} \frac{Q^j(x)}{Q^j(\mathcal{S}^j)}\right) \geq -\log\left(\frac{2^{-(\delta/2-\epsilon)n}}{Q^j(\mathcal{S}^j)}\right) \\ &= (\delta/2 - \epsilon)n + \log(Q^j(\mathcal{S}^j)). \end{aligned}$$

As $Q^j(\mathcal{S}^j) \geq 2^{-\epsilon n}$ by definition of $\mathcal{E}$, we have $H_{\min}(X|J=j,\mathcal{E}) \geq (\delta/2 - 2\epsilon)n$. □

We are now ready to state and prove our new all-but-one entropic uncertainty relation.

Theorem 5.3 *Let ρ be an arbitrary n-qubit state, let J be a random variable over $[m]$ (with arbitrary distribution P_J), and let X be the outcome when measuring ρ in basis $\mathcal{B}_J$. Then, for any $\epsilon \in \mathbb{R}$ such that $0 < \epsilon < \delta/4$, there exists a random variable J' with joint distribution $P_{JJ'X}$ such that (1) J and J' are independent and (2) there exists an event Ω with $\Pr[\Omega] \geq 1 - 2\cdot 2^{-\epsilon n}$ such that*[7]

$$H_{\min}(X|J=j, J'=j', \Omega) \geq \Big(\frac{\delta}{2} - 2\epsilon\Big)n - 1$$

for all $j, j' \in [m]$ with $j \neq j'$ and $P_{JJ'|\Omega}(j,j') > 0$.

Note that, as phrased, Theorem 5.3 requires that J is fixed and known, and only then the existence of J' can be guaranteed. This is actually not necessary. By looking at the proof, we see that J' can be defined simultaneously in all m probability spaces $P_{X|J=j}$ with $j \in [m]$, without having assigned a probability distribution to J yet, so that the resulting random variable J' we obtain by assigning an *arbitrary* probability distribution P_J to J, satisfies the claimed properties. This in particular implies that the (marginal) distribution of J' is fully determined by ρ.

The idea of the proof of Theorem 5.3 is to (try to) define the random variable J' in such a way that the event $J \neq J'$ coincides with the "good event" $\mathcal{E}$ from

[7] Instead of introducing such an event Ω, we could also express the min-entropy bound by means of the *smooth* min-entropy of X given $J = j$ and $J' = j'$.

Corollary 5.2. It then follows immediately from Corollary 5.2 that $H_{\min}(X|J = j, J' \neq J) \geq (\delta/2 - 2\epsilon)n$, which is already close to the actual min-entropy bound we need to prove. This approach dictates that if the event $\mathcal{E}$ does not occur, then J' needs to *coincide* with J. Vice versa, if $\mathcal{E}$ does occur, then J' needs to be *different* to J. However, it is a priori unclear *how* to choose J' different from J in case $\mathcal{E}$ occurs. There is only one way to set J' to be equal to J, but there are many ways to set J' to be different from J (unless $m = 2$). It needs to be done in such a way that without conditioning on $\mathcal{E}$ or its complement, J and J' are independent.

Somewhat surprisingly, it turns out that the following does the job. To simplify this informal discussion, we assume that the sum of the m probabilities $\Pr[\mathcal{E}|J = j]$ from Corollary 5.2 equals $m - 1$ exactly. It then follows that the corresponding complementary probabilities, $\Pr[\bar{\mathcal{E}}|J = j]$ for the m different choices of $j \in [m]$, add up to 1 and thus form a probability distribution. J' is now chosen, in the above spirit depending on the event $\mathcal{E}$, so that its marginal distribution $P_{J'}$ coincides with this probability distribution: $P_{J'}(j') = \Pr[\bar{\mathcal{E}}|J = j']$ for all $j' \in [m]$. Thus, in case the event $\mathcal{E}$ occurs, J' is chosen according to this distribution but conditioned on being different from the value j, taken on by J. The technical details, and how to massage the argument in case the sum of the $\Pr[\mathcal{E}|J = j]$'s is not exactly $m - 1$, are worked out in the proof below.

Proof of Theorem 5.3. From Corollary 5.2 we know that for any $0 < \epsilon < \delta/4$, there exists an event $\mathcal{E}$ such that $\sum_{j\in[m]} \Pr[\mathcal{E}|J = j] = m - 1 - \alpha$, and thus $\sum_{j\in[m]} \Pr[\bar{\mathcal{E}}|J = j] = 1 + \alpha$, for $\alpha \in \mathbb{R}$ such that $-1 \leq \alpha \leq (2m - 1)2^{-\epsilon n}$. We make the case distinction between $\alpha = 0$, $\alpha > 0$ and $\alpha < 0$. We start with case $\alpha = 0$, we subsequently prove the other two cases by reducing them to the case $\alpha = 0$ by "inflating" and "deflating" the event $\mathcal{E}$ appropriately. The approach for the case $\alpha = 0$ is to define J' in such way that $\mathcal{E} \iff J \neq J'$, i.e., the event $J \neq J'$ coincides with the event $\mathcal{E}$. The min-entropy bound from Corollary 5.2 then immediately translates to $H_{\min}(X|J = j, J' \neq J) \geq (\delta/2 - 2\epsilon)n$, and to $H_{\min}(X|J = j, J' = j') \geq (\delta/2 - 2\epsilon)n$ for $j' \neq j$ with $P_{JJ'}(j, j') > 0$, as we will show. What is not obvious about the approach is how to define J' when it is supposed to be different from J, i.e., when the event $\mathcal{E}$ occurs, so that in the end J and J' are independent.

Formally, we define J' by means of the following conditional probability distributions:

$$P_{J'|JX\bar{\mathcal{E}}}(j'|j, x) := \begin{cases} 1 & \text{if } j = j' \\ 0 & \text{if } j \neq j' \end{cases}$$

and

$$P_{J'|JX\mathcal{E}}(j'|j,x) := \begin{cases} 0 & \text{if } j = j' \\ \dfrac{\Pr[\bar{\mathcal{E}}|J=j']}{\Pr[\mathcal{E}|J=j]} & \text{if } j \neq j' \end{cases}$$

We assume for the moment that the denominator in the latter expression does not vanish for any j; we take care of the case where it does later. Trivially, $P_{J'|JX\bar{\mathcal{E}}}$ is a proper distribution, with non-negative probabilities that add up to 1, and the same holds for $P_{J'|JX\mathcal{E}}$:

$$\sum_{j'\in[m]} P_{J'|JX\bar{\mathcal{E}}} = \sum_{j'\in[m]\setminus\{j\}} P_{J'|JX\bar{\mathcal{E}}} = \sum_{j'\in[m]\setminus\{j\}} \frac{\Pr[\bar{\mathcal{E}}|J=j']}{\Pr[\mathcal{E}|J=j]} = 1$$

where we used that $\sum_{j\in[m]} \Pr[\bar{\mathcal{E}}|J=j] = 1$ (because $\alpha = 0$) in the last equality. Furthermore, it follows immediately from the definition of J' that $\bar{\mathcal{E}} \implies J = J'$ and $\mathcal{E} \implies J \neq J'$. Hence, $\mathcal{E} \iff J \neq J'$, and thus the bound from Corollary 5.2 translates to $H_{\min}(X|J=j, J' \neq J) \geq (\delta/2 - 2\epsilon)n$. It remains to argue that J' is independent of J, and that the bound also holds for $H_{\min}(X|J = j, J' = j')$ whenever $j \neq j'$.

The latter follows immediately from the fact that conditioned on $J \neq J'$ (which is equivalent to $\mathcal{E}$), X, J and J' form a Markov chain $X \leftrightarrow J \leftrightarrow J'$, and thus, given $J = j$, additionally conditioning on $J' = j'$ does not change the distribution of X. For the independence of J and J', consider the joint probability distribution of J and J', given by

$$\begin{aligned} P_{JJ'}(j,j') &= P_{J'J\mathcal{E}}(j',j) + P_{J'J\bar{\mathcal{E}}}(j',j) \\ &= P_J(j)\Pr[\mathcal{E}|J=j]P_{J'|J\mathcal{E}}(j'|j) + P_J(j)\Pr[\bar{\mathcal{E}}|J=j]P_{J'|J\bar{\mathcal{E}}}(j'|j) \\ &= P_J(j)\Pr[\bar{\mathcal{E}}|J=j'], \end{aligned}$$

where the last equality follows by separately analyzing the cases $j = j'$ and $j \neq j'$. It follows immediately that the marginal distribution of J' is

$$P_{J'}(j') = \sum_j P_{JJ'}(j,j') = \Pr[\bar{\mathcal{E}}|J=j'],$$

and thus $P_{JJ'} = P_J \cdot P_{J'}$.

What is left to do for the case $\alpha = 0$ is to deal with the case where there exists j^* with $\Pr[\mathcal{E}|J=j^*] = 0$. Since $\sum_{j\in[m]} \Pr[\bar{\mathcal{E}}|J=j] = 1$, it holds that $\Pr[\bar{\mathcal{E}}|J=j] = 0$ for $j \neq j^*$. This motivates to define J' as $J' := j^*$ with probability 1. Note that

this definition directly implies that J' is independent from J. Furthermore, by the above observations: $\mathcal{E} \iff J \neq J'$. This concludes the case $\alpha = 0$.

Next, we consider the case $\alpha > 0$. The idea is to "inflate" the event $\mathcal{E}$ so that α becomes 0, i.e., to define an event $\mathcal{E}'$ that contains $\mathcal{E}$ (meaning that $\mathcal{E} \implies \mathcal{E}'$) so that $\sum_{j\in[m]} \Pr[\mathcal{E}'|J = j] = m - 1$, and to define J' as in the case $\alpha = 0$ (but now using $\mathcal{E}'$). Formally, we define $\mathcal{E}'$ as the disjoint union $\mathcal{E}' = \mathcal{E} \vee \mathcal{E}_\circ$ of $\mathcal{E}$ and an event $\mathcal{E}_\circ$. The event $\mathcal{E}_\circ$ is defined by means of $\Pr[\mathcal{E}_\circ|\mathcal{E}, J = j, X = x] = 0$, so that $\mathcal{E}$ and $\mathcal{E}_\circ$ are indeed disjoint, and $\Pr[\mathcal{E}_\circ|J = j, X = x] = \alpha/m$, so that indeed

$$\begin{aligned}\sum_{j\in[m]} \Pr[\mathcal{E}'|J = j] &= \sum_{j\in[m]} (\Pr[\mathcal{E}|J = j] + \Pr[\mathcal{E}_\circ|J = j]) \\ &= (m - 1 - \alpha) + \alpha = m - 1\,.\end{aligned}$$

We can now apply the analysis of the case $\alpha = 0$ to conclude the existence of J', independent of J, such that $J \neq J' \iff \mathcal{E}'$ and thus $(J \neq J') \wedge \bar{\mathcal{E}}_\circ \iff \mathcal{E}' \wedge \bar{\mathcal{E}}_\circ \iff \mathcal{E}$. Setting $\Omega := \bar{\mathcal{E}}_\circ$, it follows that

$$H_{\min}(X|J = j, J \neq J', \Omega) = H_{\min}(X|J = j, \mathcal{E}) \geq (\delta/2 - 2\epsilon)n\,,$$

where $\Pr[\Omega] = 1 - \Pr[\mathcal{E}_\circ] = 1 - \alpha/m \geq 1 - (2m - 1)2^{-\epsilon n}/m \geq 1 - 2 \cdot 2^{-\epsilon n}$. Finally, using similar reasoning as in the case $\alpha = 0$, it follows that the same bound holds for $H_{\min}(X|J = j, J' = j', \Omega)$ whenever $j \neq j'$. This concludes the case $\alpha > 0$.

Finally, we consider the case $\alpha < 0$. The approach is the same as above, but now $\mathcal{E}'$ is obtained by "deflating" $\mathcal{E}$. Specifically, we define $\mathcal{E}'$ by means of $\Pr[\mathcal{E}'|\bar{\mathcal{E}}, J = j, X = x] = \Pr[\mathcal{E}'|\bar{\mathcal{E}}] = 0$, so that $\mathcal{E}'$ is contained in $\mathcal{E}$, and $\Pr[\mathcal{E}'|\mathcal{E}, J = j, X = x] = \Pr[\mathcal{E}'|\mathcal{E}] = \frac{m-1}{m-1-\alpha}$, so that

$$\sum_{j\in[m]} \Pr[\mathcal{E}'|J = j] = \sum_{j\in[m]} \Pr[\mathcal{E}'|\mathcal{E}] \cdot \Pr[\mathcal{E}|J = j] = m - 1\,.$$

Again, from the $\alpha = 0$ case we obtain J', independent of J, such that the event $J \neq J'$ is equivalent to the event $\mathcal{E}'$.

It follows that

$$\begin{aligned}H_{\min}(X|J = j, J \neq J') &= H_{\min}(X|J = j, \mathcal{E}') = H_{\min}(X|J = j, \mathcal{E}', \mathcal{E}) \\ &\geq H_{\min}(X|J = j, \mathcal{E}) - \log(P[\mathcal{E}'|\mathcal{E}, J = j]) \geq (\delta/2 - 2\epsilon)n - 1\,,\end{aligned}$$

where the second equality holds because $\mathcal{E}' \implies \mathcal{E}$, the first inequality holds because additionally conditioning on $\mathcal{E}'$ increases the probabilities of X conditioned

on $J = j$ and $\mathcal{E}$ by at most a factor $1/P[\mathcal{E}'|\mathcal{E}, J = j])$, and the last inequality holds by Corollary 5.2) and because $P[\mathcal{E}'|\mathcal{E}, J = j]) = \frac{m-1}{m-1-\alpha} \geq \frac{1}{2}$, where the latter holds since $\alpha \geq -1$. Finally, using similar reasoning as in the previous cases, it follows that the same bound holds for $H_{\min}(X|J = j, J' = j')$ whenever $j \neq j'$. This concludes the proof. □

5.2.1 Constructing Good Families of Bases

Here, we discuss some interesting choices for the family $\{\mathcal{B}_1, \ldots, \mathcal{B}_m\}$ of bases. We say that such a family is "good" if $\delta = -\frac{1}{n}\log(c^2)$ converges to a strictly positive constant as n tends to infinity. There are various ways to construct such families. For example, a family obtained through sampling according to the Haar measure will be good with overwhelming probability (a precise statement, in which "good" means $\delta = 0.9$, can be found at the very end of the proof of Theorem 2.5 of [FHS11]). The best possible constant $\delta = 1$ is achieved for a family of *mutually unbiased bases*. However, for arbitrary quantum systems (i.e., not necessarily multi-qubit systems) it is not well understood how large such a family may be, beyond that its size cannot exceed the dimension plus 1.

In the upcoming section, we will use the following simple and well-known construction. For an arbitrary binary code $\mathcal{C} \subset \mathbb{F}_2^n$ of size m, minimum distance d and encoding function $\mathfrak{c} : [m] \to \mathcal{C}$, we can construct a family $\{\mathcal{B}_1, \ldots, \mathcal{B}_m\}$ of bases as follows. We identify the jth codeword, i.e., $\mathfrak{c}(j) = (c_1, \ldots, c_n)$ for $j \in [m]$, with the basis $\mathcal{B}_j = \{H^{\mathfrak{c}(j)}|x\rangle : x \in \mathbb{F}_2^n\} = \{(H^{c_1} \otimes \cdots \otimes H^{c_n})|x\rangle : x \in \mathbb{F}_2^n\}$. In other words, $\mathcal{B}_j$ measures qubit-wise in the computational or the Hadamard basis, depending on the corresponding coordinate of $\mathfrak{c}(j)$. It is easy to see that the maximum overlap c of the family obtained this way is directly related to the minimum distance of $\mathcal{C}$, namely $\delta = -\frac{1}{n}\log(c^2)$ coincides with the relative minimal distance d/n of $\mathcal{C}$. Hence, choosing an asymptotically good code immediately yields a good family of bases.

5.3 A New Quantum Identification Protocol

Our main application of the new uncertainty relation is in proving security of a new password-based identification protocol in the quantum setting. Recall that in password-based identification, a user U wants to convince a server S that he (U) knows a password w, in such a way that only a negligible amount of information is leaked about w in case U is interacting with a dishonest server S*. *Vice versa*, a

dishonest user U^* (who does not know w) should not be able to gain information about w by interacting with S.

It is known that *without* any restriction on (one of) the dishonest participants, secure identification is impossible (even in the quantum setting). Indeed, if a quantum protocol is unconditionally secure against a dishonest user, then unavoidably it can be broken by a dishonest server with unbounded quantum storage and unbounded quantum computing power; this follows essentially from [Lo97] (see also [DFSS07]). Thus, the best one can hope for (for a protocol that is unconditionally secure against a dishonest user) is that in order to break it, unbounded quantum storage *and* unbounded quantum computing power are *necessary* for the dishonest server. Note that this is not the case for the existing quantum identification protocol `QID`, which we reviewed in Section 2.11.1: a dishonest server who can postpone the measurements of (most of) the qubits until the user announces the bases—by temporarily storing the qubits in a quantum memory—completely breaks the protocol. Thus, no quantum computing power at all is necessary to break `QID`, only sufficient quantum storage.

In this section, we propose a new identification protocol, `NEWQID`, which can be regarded as a first step towards closing the above gap. Like `QID`, our new protocol is secure against an unbounded dishonest user and against a dishonest server with limited quantum storage capabilities. Furthermore, and in contrast to `QID`, a minimal amount of quantum computation power is *necessary* to break the protocol, beyond sufficient quantum storage. Indeed, in addition to the security against a dishonest server with bounded quantum storage, we also prove security against a dishonest server that can store all the communicated qubits, but is restricted to measure them qubit-wise (in arbitrary qubit bases) at the end of the protocol execution. Thus, beyond sufficient quantum storage, quantum computation that involves *pairs* of qubits is necessary (and in fact sufficient) to break the new protocol.

Restricting the dishonest server to qubit-wise measurements may look restrictive; however, we stress that in order to break the protocol, the dishonest server needs to store many qubits *and* perform quantum operations on them that go beyond single-qubit operations; this may indeed be considerably more challenging than storing many qubits and measuring them qubit-wise. Furthermore, it turns out that proving security against such a dishonest server that is restricted to qubit-wise measurements is already challenging; indeed, standard techniques (e.g., privacy amplification) do not seem applicable here. Therefore, handling a dishonest server that can, say, act on *blocks* of qubits, must be left to future research.

The security properties that we want to achieve are given in Section 2.11. Similar to

QID, the new protocol will be shown to be unconditionally secure against dishonest users. The new uncertainty relation is the main ingredient for proving security against a dishonest server with bounded quantum storage. Our security proof against a dishonest server (having unbounded quantum storage) that is restricted to non-adaptive qubit-wise measurements uses very different techniques.

5.3.1 Description of Our New Protocol

Let $\mathcal{C} \subset \mathbb{F}_2^n$ be a binary code with minimum distance d, and let $\mathfrak{c} : \mathcal{W} \to \mathcal{C}$ be its encoding function. Let $m := |\mathcal{W}|$, and typically, $m < 2^n$. Let $\mathcal{F}$ be the class of all linear functions from $\{0,1\}^n$ to $\mathbb{F}_2^\ell$, where $\ell < n$, represented as $\ell \times n$ matrices over $\mathbb{F}_2$. Note that $\mathcal{F}$ is two-universal and coincides with the family $\mathcal{G}_1$ defined—and proved to be two-universal—in Section 2.4.1. Furthermore, let $\mathcal{G}$ be a strongly two-universal class of hash functions from $\mathcal{W}$ to $\mathbb{F}_2^\ell$. Protocol NEWQID is shown below.

---

1. U picks $x \xleftarrow{\mathrm{r}} \{0,1\}^n$ and sends $H^{\mathfrak{c}(w)}|x\rangle$ to S.
2. S measures in basis $\mathfrak{c}(w)$. Let x' be the outcome.
3. U picks $f \xleftarrow{\mathrm{r}} \mathcal{F}$ and sends it to S
4. S picks $g \xleftarrow{\mathrm{r}} \mathcal{G}$ and sends it to U
5. U computes and sends $z := f(x) \oplus g(w)$ to S
6. S accepts if and only if $z = z'$ where $z' := f(x') \oplus g(w)$

Protocol 5.1: Our new quantum password-based-identification protocol NEWQID. The difference between this new protocol and the existing protocol QID by Damgård *et al.* (see Protocol 2.1) is the way how the user prepares the state in step (1): in the new protocol the basis is chosen as a function of the password w, whereas in QID it is chosen at random and communicated in a later step in the protocol.

Note that our protocol is quite similar to QID (Section 2.11.1). The difference is that in our protocol, *both* parties, i.e., U and S, use $\mathfrak{c}(w)$ as basis for preparing/measuring the qubits in step (1) and (2), whereas in QID only S uses $\mathfrak{c}(w)$ and U uses a *random* basis $\theta \in \{0,1\}^n$ instead, and then U communicates θ to S and all the positions where θ and $\mathfrak{c}(w)$ differ are dismissed. Thus, in some sense, our new protocol is more natural since why should U use a random basis when he knows the right basis (i.e., the one that S uses)? In [DFSS07], using a random basis (for U) was crucial for their proof technique, which is based on an entropic uncertainty relation of a certain form, which asks for a random basis. However, using a random basis, which then

needs to be announced, renders the protocol insecure against a dishonest server S^* that is capable of storing all the communicated qubits and then measure them in the right basis once it has been announced. Our new uncertainty relation applies to the case where an n-qubit state is measured in a basis that is sampled from a code $\mathcal{C}$, and thus is applicable to the new protocol where U uses basis $\mathfrak{c}(w) \in \mathcal{C}$. Since this basis is common knowledge (to the honest participants), it does not have to be communicated, and as such a straightforward store-and-then-measure attack as above does not apply.

A downside of our protocol is that security only holds in case of a perfect quantum source, which emits exactly one qubit when triggered. Indeed, a multi-photon emission enables a dishonest server S^* to learn information on the basis used, and thus gives away information on the password w in our protocol. As such, our protocol is currently mainly of theoretical interest.

It is straightforward to verify that (in the ideal setting with perfect sources, no noise, etc.) `NEWQID` satisfies the correctness property (Definition 2.66) perfectly. In the upcoming sections, we give proofs for server and user security.

5.4 (Unconditional) Server Security

First, we argue security of `NEWQID` against an arbitrary dishonest user U^* (that is merely restricted by the laws of quantum mechanics).

Theorem 5.4 *`NEWQID` is ε-secure for the server with $\varepsilon = \binom{m}{2} 2^{-\ell}$.*

Proof. Clearly, from the steps (1) to (5) in the protocol `NEWQID`, U^* learns no information on W at all. The only information he may learn is by observing whether S accepts or not in step (6). Therefore, in order to prove server security, it suffices to show the existence of a random variable W', independent of W, with the property that S rejects whenever $W' \neq W$ (except with probability $\frac{1}{2}m(m-1)2^{-\ell}$) and that S accepts whenever $W' = W$.

We may assume that $\mathcal{W} = [m]$. Let $\rho_{WX'FGZE}$ be the state describing the password W, the variables X', F, G and Z occurring in the protocol from the server's point of view, and U^*'s quantum state E *before* observing S's decision to accept or reject. For any $w \in \mathcal{W}$, consider the state $\rho^w_{X'FGZE} := \rho_{X'FGZE|W=w}$. Note that the reduced state ρ^w_{FGZE} is the same for any $w \in \mathcal{W}$; this follows from the assumption that U^*'s initial state is independent of W and because F, G and Z are produced independently of W. We may thus write $\rho^w_{X'FGZE}$ as $\rho_{X'_wFGZE}$, and we can "glue together" the states $\rho_{X'_wFGZE}$ for all choices of w. This means, there exists a state

$\rho_{X'_1\cdots X'_m FGZE_1\cdots E_m}$ that correctly reduces to $\rho_{X'_w FGZE_w} = \rho_{X'_w FGZE}$ for any $w \in \mathcal{W}$, and conditioned on FGZ, we have that $X'_i E_i$ is independent of $X'_j E_j$ for any $i \neq j \in \mathcal{W}$. It is easy to see that for any $i \neq j \in \mathcal{W}$, G is independent of X'_i, X'_j and F. Therefore, by the strong two-universality of G, for any $i \neq j$ it holds that $Z'_i \neq Z'_j$ except with probability $2^{-\ell}$, where $Z'_w = F(X'_w) + G(w)$ for any w. Therefore, by the union bound, $Z'_1, \ldots, Z'_m$ are pairwise distinct and thus Z can coincide with at most one of the Z'_w's, except with probability $\varepsilon = \frac{1}{2}m(m-1)2^{-\ell}$. Let W' be defined such that $Z = Z'_{W'}$; if there is no such Z'_w then we let $W' = \perp$, and if there are more than one then we let it be the first. Recall, the latter can happen with probability at most ε. We now extend the state $\rho_{X'_1\cdots X'_m FGZW'E_1\cdots E_m}$ by W, chosen independently according to P_W. Clearly W' is independent of W. Furthermore, except with probability at most ε, if $W \neq W'$ then $Z \neq Z'_W$. Also note that $\rho_{X'_W FGZW'WE_W}$ is such that

$$\begin{aligned}\rho_{X'_W FGZWE_W} &= \sum_w P_W(w)\rho_{X'_w FGZE_w} \otimes |w\rangle\langle w| \\ &= \sum_w P_W(w)\rho^w_{X'FGZE} \otimes |w\rangle\langle w| = \rho_{X'FGZWE}.\end{aligned}$$

Thus, also with respect to the state $\rho_{X'FGZWE}$ there exist W', independent of W, such that if $W' \neq W$ then $Z \neq Z'$ except with probability at most ε. Finally, whenever $W = W'$ it follows by construction that $Z = Z'$ and S will always accept in this case. This was to be shown. □

5.5 User Security in the BQSM

Next, we consider a dishonest server S*, and first prove security of NEWQID in the *bounded-quantum-storage model*. In this model, as introduced in [DFSS05], it is assumed that the adversary (here S*) cannot store more than a fixed number of qubits, say q. The security proof of NEWQID in the bounded quantum storage model is very similar to the corresponding proof in [DFSS07] for their protocol, except that we use the new uncertainty relation from Section 5.2. Furthermore, since our uncertainty relation (Theorem 5.3) already guarantees the existence of the random variable W' as required by the security property, no *entropy-splitting* as in [DFSS07] is needed.

In the following, let $\delta := d/n$, i.e., the relative minimum distance of $\mathcal{C}$.

Theorem 5.5 *Let* S* *be a dishonest server whose quantum memory is at most q qubits at step 3 of* NEWQID*. Then, for any $0 < \kappa < \delta/4$,* NEWQID *is ε-secure for the*

user with

$$\varepsilon = 2^{-\frac{1}{2}((\delta/2-2\kappa)n-1-q-\ell)} + 4 \cdot 2^{-\kappa n}.$$

Proof. We consider and analyze a purified version of NEWQID where in step (1) instead of sending $H^{\mathfrak{c}(W)}|X\rangle$ to S^* for a uniformly distributed 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, U measures his register in the basis $\mathfrak{c}(W)$ in order to obtain X. Note that this procedure produces exactly the same common state as in the original (non-purified) version of NEWQID. Thus, we may just as well analyze this purified version.

The state of S^* consists of his initial state and his part of the EPR pairs, and may include an additional ancilla register. Before the memory bound applies, S^* may perform any unitary transformation on his composite system. When the memory bound is applied (just before step (3) is executed in NEWQID), S^* has to measure all but q qubits of his system. Let the classical outcome of this measurement be denoted by y, and let E' be the remaining quantum state of at most q qubits. The common state has collapsed to a $(n+q)$-qubit state and depends on y; the analysis below holds for any y. Next, U measures his n-qubit part of the common state in basis $\mathfrak{c}(W)$; let X denote the classical outcome of this measurement. By our new uncertainty relation (Theorem 5.3) and subsequently applying the min-entropy chain rule that is given in Proposition 2.62 (to take the q stored qubits into account) it follows that there exists W', independent of W, and an event Ω that occurs at least with probability $1 - 2 \cdot 2^{-\kappa n}$, such that

$$H_{\min}(X|E', W = w, W' = w', \Omega) \geq (\delta/2 - 2\kappa)n - 1 - q.$$

for any w, w' such that $w \neq w'$. Because U chooses F independently at random from a 2-universal family, privacy amplification guarantees that

$$d_{\mathsf{unif}}(F(X)|E'F, W = w, W' = w') \leq \varepsilon' := \frac{1}{2} \cdot 2^{-\frac{1}{2}((\delta/2-2\kappa)n-1-q-\ell)} + 2 \cdot 2^{-\kappa n},$$

for any w, w' such that $w \neq w'$. Recall that $Z = F(X) \oplus G(W)$. By security of the one-time pad it follows that

$$d_{\mathsf{unif}}(Z|E'FG, W = w, W' = w') \leq \varepsilon', \tag{5.1}$$

for any w, w' such that $w \neq w'$. To prove the claim, we need to bound,

$$\begin{aligned}
&\delta(\rho_{WW'E|W\neq W'}, \rho_{W\leftrightarrow W'\leftrightarrow E|W\neq W'}) \\
&\quad = \tfrac{1}{2}\|\rho_{WW'E'FGZ|W\neq W'} - \rho_{W\leftrightarrow W'\leftrightarrow E'FGZ|W\neq W'}\|_1 \\
&\quad \leq \tfrac{1}{2}\|\rho_{WW'E'FGZ|W\neq W'} - \rho_{WW'E'FG|W\neq W'} \otimes 2^{-\ell}\mathbb{I}\|_1 \\
&\qquad + \tfrac{1}{2}\|\rho_{WW'E'FG|W\neq W'} \otimes 2^{-\ell}\mathbb{I} - \rho_{W\leftrightarrow W'\leftrightarrow E'FGZ|W\neq W'}\|_1 \qquad (5.2)
\end{aligned}$$

where the equality follows by definition of trace distance (Definition 2.48) and the fact that the output state E is obtained by applying a unitary transformation to the set of registers (E', F, G, W', Z). The inequality is the triangle inequality; in the remainder of the proof, we will show that both terms in (5.2) are upper bounded by ε'.

$$\begin{aligned}
&\tfrac{1}{2}\|\rho_{WW'E'FGZ|W\neq W'} - \rho_{WW'E'FG|W\neq W'} \otimes 2^{-\ell}\mathbb{I}\|_1 \\
&\quad = \sum_{w\neq w'} P_{WW'|W\neq W'}(w,w')\, d_{\mathsf{unif}}(Z|E'FG, W=w, W'=w') \leq \varepsilon',
\end{aligned}$$

where the latter inequality follows from (5.1).For the other term, we reason as follows:

$$\begin{aligned}
&\tfrac{1}{2}\|\rho_{WW'E'FG|W\neq W'} \otimes 2^{-\ell}\mathbb{I} - \rho_{W\leftrightarrow W'\leftrightarrow E'FGZ|W\neq W'}\|_1 \\
&\quad = \tfrac{1}{2}\sum_{w\neq w'} P_{WW'|W\neq W'}(w,w')\, \|\rho^{w,w'}_{E'FG|W\neq W'} \otimes 2^{-\ell}\mathbb{I} - \rho^{w'}_{E'FGZ|W\neq W'}\|_1 \\
&\quad = \tfrac{1}{2}\sum_{w\neq w'} P_{WW'|W\neq W'}(w,w')\, \|\rho^{w,w'}_{E'FG|W\neq W'} \otimes 2^{-\ell}\mathbb{I} \\
&\qquad - \sum_{\substack{w'' \\ \text{s.t. } w''\neq w'}} P_{W|W',W\neq W'}(w''|w')\rho^{w'',w'}_{E'FGZ|W\neq W'}\|_1 \\
&\quad = \tfrac{1}{2}\sum_{w'} P_{W'|W\neq W'}(w')\, \|\sum_{\substack{w \\ \text{s.t. } w\neq w'}} P_{W|W',W\neq W'}(w|w')\rho^{w,w'}_{E'FG|W\neq W'} \otimes 2^{-\ell}\mathbb{I} \\
&\qquad - \sum_{\substack{w'' \\ \text{s.t. } w''\neq w'}} P_{W|W',W\neq W'}(w''|w')\rho^{w'',w'}_{E'FGZ|W\neq W'} \sum_{\substack{w \\ \text{s.t. } w\neq w'}} P_{W|W',W\neq W'}(w|w')\|_1 \\
&\quad = \tfrac{1}{2}\sum_{w\neq w'} P_{WW'|W\neq W'}(w,w')\, \|\rho^{w,w'}_{E'FG|W\neq W'} \otimes 2^{-\ell}\mathbb{I} - \rho^{w,w'}_{E'FGZ|W\neq W'}\|_1 \\
&\quad = \sum_{w\neq w'} P_{WW'|W\neq W'}(w,w')\, d_{\mathsf{unif}}(Z|E'FG, W=w, W'=w') \leq \varepsilon',
\end{aligned}$$

where the first equality follows by definition of conditional independence (the quantum version, see (2.10) on page 81) and by a basic property of the trace distance; the third and fourth equality follow by linearity of the trace distance. The inequality on the last line follows from (5.1). This proves the claim. □

5.6 User Security in the Single-Qubit-Operations Model

We now consider a dishonest server S^* that can store an unbounded number of qubits. Clearly, against such a S^*, Theorem 5.5 provides no security guarantee anymore. We show here that there is still *some* level of security left. Specifically, we show that NEWQID is still secure against a dishonest server S^* that can reliably store all the communicated qubits and measure them qubit-wise and non-adaptively at the end of the protocol. This feature distinguishes our identification protocol from the protocol from [DFSS07], which completely breaks down against such an attack.

5.6.1 The Model

Formally, a dishonest server S^* in the SQOM is modeled as follows.

1. S^* may reliably store the n-qubit state $H^{\mathfrak{c}(w)}|x\rangle = H^{\mathfrak{c}(w)_1}|x_1\rangle \otimes \cdots \otimes H^{\mathfrak{c}(w)_n}|x_n\rangle$ received in step (1) of NEWQID.

2. At the end of the protocol, in step (5), S^* chooses an arbitrary sequence $\theta = (\theta_1, \ldots, \theta_n)$, where each θ_i describes an arbitrary orthonormal basis of $\mathbb{C}^2$, and measures each qubit $H^{\mathfrak{c}(w)_i}|x_i\rangle$ in basis θ_i to observe $Y_i \in \mathbb{F}_2$. Hence, we assume that S^* *measures all qubits at the end of the protocol.*

3. The choice of θ may depend on all the classical information gathered during the execution of the protocol, but we assume a *non-adaptive* setting where θ_i does not depend on Y_j for $i \neq j$, i.e., S^* has to choose θ entirely before performing any measurement.

Considering complete projective measurements acting on individual qubits, rather than general single-qubit POVMs, may be considered a restriction of our model. Nonetheless, general POVM measurements can always be described by projective measurements on a bigger system. In this sense, restricting to projective measurements is consistent with the requirement of single-qubit operations. It seems non-trivial to extend our security proof to general single-qubit POVMs.

The restriction to non-adaptive measurements (item 3) is rather strong, even though the protocol from [DFSS07] already breaks down in this non-adaptive setting. The restriction was introduced as a stepping stone towards proving the adaptive case. Up to now, we have unfortunately not yet succeeded in doing so, hence we leave the adaptive case for future research.

We also leave for future research the case of a less restricted dishonest server S^* that can do measurements on blocks that are less stringently bounded in size. Whereas the adaptive versus non-adaptive issue appears to be a proof-technical problem (`NEWQID` looks secure also against an adaptive S^*), allowing measurements on larger blocks will require a new protocol, since `NEWQID` becomes insecure when S^* can do measurements on blocks of size 2, as we show in Section 5.6.5.

5.6.2 No Privacy Amplification

One might expect that proving security of `NEWQID` in the SQOM, i.e., against a dishonest server S^* that is restricted to single-qubit operations should be straightforward, but actually the opposite is true, for the following reason. Even though it is not hard to show that after his measurements, S^* has lower bounded uncertainty in x (except if he was able to guess w), it is not clear how to conclude that $f(x)$ is close to random so that z does not reveal a significant amount of information about w. The reason is that standard privacy amplification fails to apply here. Indeed, the model allows S^* to postpone the measurement of all qubits to step (5) of the protocol. The hash function f, however, is chosen and sent already in step (3). This means that S^* can choose his measurements in step (5) depending on f. As a consequence, the distribution of x from the point of view of S^* may depend on the choice of the hash function f, in which case the privacy-amplification theorem does not give any guarantees.

5.6.3 Single-Qubit Measurements

Consider an arbitrary sequence $\theta = (\theta_1, \ldots, \theta_n)$ where each θ_i describes an orthonormal basis of $\mathbb{C}^2$. Let $|\psi\rangle$ be an n-qubit system of the form

$$|\psi\rangle = H^{b_1}|x_1\rangle \otimes \cdots \otimes H^{b_n}|x_n\rangle,$$

where x and b are arbitrary in $\mathbb{F}_2^n$. Measuring $|\psi\rangle$ qubit-wise in basis θ results in a measurement outcome $Y = (Y_1, \ldots, Y_n) \in \mathbb{F}_2^n$. Suppose that x, b and θ are in fact realizations of the random variables X, B and Θ respectively. It follows

immediately from the product structure of the state $|\psi\rangle$ that

$$P_{Y|XB\Theta}(y|x,b,\theta) = \prod_{i=0}^{n} P_{Y_i|X_iB_i\Theta_i}(y_i|x_i,b_i,\theta_i),$$

i.e., the random variables Y_i are statistically independent conditioned on arbitrary fixed values for X_i, B_i and Θ_i but such that $P_{X_iB_i\Theta_i}(x_i,b_i,\theta_i) > 0$.

Lemma 5.6 *The distribution $P_{Y_i|X_iB_i\Theta_i}(y_i|x_i,b_i,\theta_i)$ exhibits the following symmetries:*

$$P_{Y_i|X_iB_i\Theta_i}(0|0,b_i,\theta_i) = P_{Y_i|X_iB_i\Theta_i}(1|1,b_i,\theta_i)$$

and

$$P_{Y_i|X_iB_i\Theta_i}(0|1,b_i,\theta_i) = P_{Y_i|X_iB_i\Theta_i}(1|0,b_i,\theta_i)$$

for all $i \in [n]$, for all b_i and θ_i with $P_{X_iB_i\Theta_i}(\xi,b_i,\theta_i) > 0$ for all $\xi \in \mathbb{F}_2$.

Proof. Let $\alpha, \beta \in \mathbb{C}$ be such that $\theta_i := \{\bar{\alpha}|0\rangle + \bar{\beta}|1\rangle, \bar{\beta}|0\rangle - \bar{\alpha}|1\rangle\}$. (We can always find such α and β.) Writing out the measurement explicitly gives

$$\begin{aligned} P_{Y_i|X_iB_i\Theta_i}(0|x_i,b_i,\theta_i) &= |(\alpha\langle 0| + \beta\langle 1|)H^{b_i}|x_i\rangle|^2 \qquad \text{and} \\ P_{Y_i|X_iB_i\Theta_i}(1|x_i,b_i,\theta_i) &= |(\beta\langle 0| - \alpha\langle 1|)H^{b_i}|x_i\rangle|^2. \end{aligned}$$

Hence, it suffices to prove that

$$|(\alpha\langle 0| + \beta\langle 1|)H^{b_i}|x_i\rangle|^2 = |(\beta\langle 0| - \alpha\langle 1|)H^{b_i}|x_i \oplus 1\rangle|^2 \tag{5.3}$$

for every $x_i, b_i \in \mathbb{F}_2$.

We first show (5.3) for $b_i = 0$. Let σ_1 be the first Pauli matrix defined by $\sigma_1|a\rangle = |a \oplus 1\rangle$ for every $a \in \mathbb{F}_2$. It follows immediately from the definition that σ_1 is a unitary matrix and it is easy to see that σ_1 is Hermitian. Then,

$$\begin{aligned} |(\alpha\langle 0| + \beta\langle 1|)|x_i\rangle|^2 &= |(\alpha\langle 0| + \beta\langle 1|)\sigma_1\sigma_1|x_i\rangle|^2 = |(\alpha\langle 1| + \beta\langle 0|)|x_i \oplus 1\rangle|^2 \\ &= |(\beta\langle 0| - \alpha\langle 1|)|x_i \oplus 1\rangle|^2 \end{aligned}$$

The last equation follows because the expression equals either $|\alpha|^2$ or $|\beta|^2$ (depending on $x_i \in \mathbb{F}_2$), hence we may freely change the sign of α. For $b_i = 1$, we have

$$|(\alpha\langle 0| + \beta\langle 1|)H|x_i\rangle|^2 = |(\alpha\langle 0| + \beta\langle 1|)(|0\rangle + (-1)^{x_i}|1\rangle)|^2 = |\alpha + (-1)^{x_i}\beta|^2$$

and

$$|(\beta\langle 0| - \alpha\langle 1|)H|x_i \oplus 1\rangle|^2 = |(\beta\langle 0| - \alpha\langle 1|)(|0\rangle - (-1)^{x_i}|1\rangle)|^2 = |\beta + (-1)^{x_i}\alpha|^2.$$

We see that those expressions are equal for every $x_i \in \mathbb{F}_2$. □

The symmetry characterized in Lemma 5.6 coincides with that of the *binary symmetric channel*, i.e., we can view Y as a "noisy version" of X, where this noise—produced by the measurement—is independent of X.

Formally, we can write Y as

$$Y = X \oplus \Delta, \tag{5.4}$$

where the random variable $\Delta = (\Delta_1, \ldots, \Delta_n) \in \mathbb{F}_2^n$ thus represents the error between the random variable $X \in \mathbb{F}_2^n$ that is "encoded" in the quantum state and the measurement outcome $Y \in \mathbb{F}_2^n$. By substituting (5.4) in Lemma 5.6, we get the following corollary.

Corollary 5.7 (Independence Between Δ and X) *For every $i \in [n]$ it holds that*

$$P_{\Delta_i|X_iB_i\Theta_i}(\delta_i|x_i, b_i, \theta_i) = P_{\Delta_i|B_i\Theta_i}(\delta_i|b_i, \theta_i)$$

for all $\delta_i \in \{0, 1\}$ and for all x_i, b_i and θ_i such that $P_{X_iB_i\Theta_i}(x_i, b_i, \theta_i) > 0$.

Furthermore, since the random variables Y_i are statistically independent conditioned on fixed values for X_i, B_i and Θ_i, it follows that the Δ_i are statistically independent conditioned on fixed values for B_i and Θ_i.

Definition 5.8 (Quantized Basis) For any orthonormal basis $\theta_i = \{|v_1\rangle, |v_2\rangle\}$ on $\mathbb{C}^2$, we define the *quantized basis* of θ_i as

$$\hat{\theta}_i := j^* \in \mathbb{F}_2, \quad \text{where } j^* \in \arg\max_{j \in \mathbb{F}_2} \max_{k \in \{1,2\}} |\langle v_k | H^j | 0 \rangle|.$$

If both $j \in \mathbb{F}_2$ attain the maximum, then j^* is chosen arbitrarily from $\mathbb{F}_2$. The quantized basis of the sequence $\theta = (\theta_1, \ldots, \theta_n)$ is naturally defined as the element-wise application of the above, resulting in $\hat{\theta} \in \mathbb{F}_2^n$.

We will use the bias (see Section 2.2.3) as a measure for the predictability of Δ_i.

Theorem 5.9 *When measuring the qubit $H^{b_i}|x_i\rangle$ for any $x_i, b_i \in \mathbb{F}_2$ in any orthonormal basis θ_i on $\mathbb{C}^2$ for which the quantized basis $\hat{\theta}_i$ is the complement of b_i, i.e., $\hat{\theta}_i = b_i \oplus 1$, then the bias of $\Delta_i \in \mathbb{F}_2$, where $\Delta_i = Y_i \oplus x_i$ and $Y_i \in \mathbb{F}_2$ is the measurement outcome, is upper bounded by*

$$\text{bias}(\Delta_i) \leq \frac{1}{\sqrt{2}}.$$

Since the theorem holds for any $x_i \in \mathbb{F}_2$ and since Corollary 5.7 guarantees that Δ_i is independent from an arbitrary random variable X_i, the theorem also applies when we replace x_i by the random variable X_i.

In order to prove Theorem 5.9, we need the following lemma.

Lemma 5.10 *If, for any orthonormal basis θ_i on $\mathbb{C}^2$, there exists a bit $b_i \in \mathbb{F}_2$ so that when measuring the qubit $H^{b_i}|x_i\rangle$ for any $x_i \in \mathbb{F}_2$ in the basis θ_i to obtain $Z_i \in \mathbb{F}_2$ it holds that*

$$\mathrm{bias}(Z_i) \geq 1/\sqrt{2},$$

then it holds that when measuring the qubit $H^{b_i \oplus 1}|x_i\rangle$ in the basis θ_i to obtain $Y_i \in \mathbb{F}_2$,

$$\mathrm{bias}(Y_i) \leq 1/\sqrt{2}.$$

Proof. First note that for any $x_i, b_i \in \mathbb{F}_2$ and any orthonormal basis θ_i on $\mathbb{C}^2$, measuring a state $H^{b_i}|x_i\rangle$ in $\theta_i = \{|v\rangle, |w\rangle\}$ where $|v\rangle = \alpha|0\rangle + \beta|1\rangle$ and $|w\rangle = \beta|0\rangle - \alpha|1\rangle$ gives the same outcome distribution (up to permutations) as when measuring one of the basis states of θ_i (when viewed as a quantum state), say $|w\rangle$, using the basis $\{H^{b_i}|x_i\rangle, H^{b_i}|x_i \oplus 1\rangle\}$. To see why this holds, note that it follows immediately that $|\langle w|H^{b_i}|x_i\rangle|^2 = |\langle x_i|H^{b_i}|w\rangle|^2$. Furthermore, we have already shown in the proof of Lemma 5.6 that

$$|\langle v|H^{b_i}|x_i\rangle|^2 = |\langle w|H^{b_i}|x_i \oplus 1\rangle|^2$$

holds.

Hence, we can apply Theorem 5.1 with $\rho = |w\rangle\langle w|$ (this implies that $n = 1$), $m = 2$ and $\mathcal{B}_0$ and $\mathcal{B}_1$ are the computational and Hadamard basis respectively. The maximum overlap between those bases is $c = 1/\sqrt{2}$. Theorem 5.1 gives us that

$$p_{\max}^{\{|0\rangle,|1\rangle\}} + p_{\max}^{\{|+\rangle,|-\rangle\}} \leq 1 + \frac{1}{\sqrt{2}},$$

where $p_{\max}^{\{|0\rangle,|1\rangle\}}$ and $p_{\max}^{\{|+\rangle,|-\rangle\}}$ respectively denote the maximum probability in the distribution obtained by measuring in the computational and Hadamard basis. By simple manipulations we can write this as a bound on the sum of the biases:

$$\begin{aligned} \frac{2}{\sqrt{2}} &\geq (2p_{\max}^{\{|0\rangle,|1\rangle\}} - 1) + (2p_{\max}^{\{|+\rangle,|-\rangle\}} - 1) \\ &= \mathrm{bias}(Y_i) + \mathrm{bias}(Z_i). \end{aligned} \tag{5.5}$$

From this relation, the claim follows immediately. □

Following [Scho7], we want to remark that both biases in (5.5) are equal to $1/\sqrt{2}$ when θ_i is the *Breidbart basis*, which is the basis that is precisely "in between" the

computational and the Hadamard basis:[8]

$$|v\rangle = \cos(\tfrac{\pi}{8})|0\rangle + \sin(\tfrac{\pi}{8})|1\rangle \qquad \text{and} \qquad |w\rangle = \sin(\tfrac{\pi}{8})|0\rangle - \cos(\tfrac{\pi}{8})|1\rangle.$$

Proof of Theorem 5.9 . Let $\theta_i = \{|v_0\rangle, |v_1\rangle\}$. We will make a case distinction based on the value of

$$\mu := \max_{k\in\mathbb{F}_2} |\langle v_k|H^{\hat{\theta}_i}|0\rangle|. \tag{5.6}$$

If $\mu \leq \cos(\pi/8)$, then we also have that $\max_{k\in\mathbb{F}_2} |\langle v_k|H^{b_i}|x_i\rangle| \leq \cos(\pi/8)$ where $b_i = \hat{\theta}_i \oplus 1$, this holds by definition of the quantized basis (Definition 5.8). Then, the probability of obtaining outcome $Y_i = k^*$, where $k^* \in \mathbb{F}_2$ achieves the maximum in (5.6), is bounded by

$$P_{Y_i}(k^*) = |\langle v_{k^*}|H^{b_i}|x_i\rangle|^2 \leq \cos^2(\pi/8) = \tfrac{1}{2} + \tfrac{1}{2\sqrt{2}}.$$

Hence,

$$\mathrm{bias}(\Delta_i) = \mathrm{bias}(Y_i) = |P_{Y_i}(k^*) - (1 - P_{Y_i}(k^*))| = |2P_{Y_i}(k^*) - 1| \leq \tfrac{1}{\sqrt{2}}.$$

If $\mu > \cos(\pi/8)$, then when measuring the state $H^{\hat{\theta}_i}|x_i\rangle$ in θ_i to obtain $Z_i \in \mathbb{F}_2$, we have that $\mathrm{bias}(Z_i) > 1/\sqrt{2}$ (this follows from similar computations as performed above). We now invoke Lemma 5.10 to conclude that when measuring the state $H^{b_i}|x_i\rangle$ in θ_i to obtain Y_i, $\mathrm{bias}(\Delta_i) = \mathrm{bias}(Y_i) < \frac{1}{\sqrt{2}}$. □

5.6.4 User Security of NEWQID

We are now ready to state and prove the security of NEWQID against a dishonest user in the SQOM.

Theorem 5.11 (User Security) *Let S^* be a dishonest server with unbounded quantum storage that is restricted to non-adaptive single-qubit operations, as specified in Section 5.6.1. Then, for any $\beta \in \mathbb{R}$ such that $0 < \beta < \frac{1}{4}$, user security (as defined in Definition 2.67) holds with*

$$\varepsilon \leq \tfrac{1}{2} 2^{\frac{1}{2}\ell - \frac{1}{4}(\frac{1}{4}-\beta)d} + \tbinom{m}{2} 2^{2\ell} \exp(-2d\beta^2)$$

Note that d is typically linear in n whereas ℓ is chosen independently of n, hence the expression above is negligible in d.

[8] In [Sch07], the corresponding state is called the "Hadamard-invariant state."

To prove Theorem 5.11 we need the following technical lemma and corollary. Recall that $\mathcal{F}$ denotes the class of all linear functions from $\mathbb{F}_2^n$ to $\mathbb{F}_2^\ell$, where $\ell < n$, represented as $\ell \times n$ matrices over $\mathbb{F}_2$. When $F \in \mathcal{F}$ acts on an n-bit vector $x \in \mathbb{F}_2^n$, we prefer the notation $F(x)$ over matrix-product notation Fx.[9] Furthermore, we write $\mathrm{span}(F)$ for the *row* span of F: the set of vectors obtained by making all possible $\mathbb{F}_2$ linear combinations of the rows of F, i.e., the set $\{sF : \forall s \in \mathbb{F}_2^\ell\}$, where s should be interpreted as a row vector and sF denotes a vector-matrix product. For two vectors $v, w \in \mathbb{F}_2^n$, the *Schur product* is defined as the element-wise product $v \odot w := (v_1 w_1, v_2 w_2, \ldots, v_n w_n) \in \mathbb{F}_2^n$, and the *inner product* between v and w is given by $v \bullet w := v_1 w_1 \oplus \cdots \oplus v_n w_n \in \mathbb{F}_2$. For an n-bit vector vector $v = (v_1, \ldots, v_n)$ in $\mathbb{F}_2^n$, we write $|v|$ for its Hamming weight (as defined in Section 3.1.3), and, for any subset $\mathcal{I} \subseteq [n]$, we write $v_\mathcal{I}$ for the restricted vector $(v_i)_{i \in \mathcal{I}} \in \mathbb{F}_2^{|\mathcal{I}|}$.

Lemma 5.12 *Let n, k and ℓ be arbitrary positive integers, let $0 < \beta < \frac{1}{4}$ and let $\mathcal{I} \subset [n]$ such that $|\mathcal{I}| \geq k$, and let F be uniform over $\mathcal{F} = \mathbb{F}_2^{\ell \times n}$. Then, it holds except with probability $2^{2\ell} \exp(-2k\beta^2)$ (the probability is over the random matrix F) that*

$$\left|(f \odot g)_\mathcal{I}\right| > (\tfrac{1}{4} - \beta)k \qquad \forall f, g \in \mathrm{span}(F) \setminus \{\mathbf{0}\}$$

Proof. Without loss of generality, we will assume that $|\mathcal{I}| = k$. Now take arbitrary but non-zero vectors $r, s \in \mathbb{F}_2^\ell$ and let $V := rF$ and $W := sF$. We will analyze the case $r \neq s$; the case $r = s$ is similar but simpler. Because each element of F is an independent random bit, and r and s are non-zero and $r \neq s$, V and W are independent and uniformly distributed n-bit vectors with expected relative Hamming weight $1/2$. Hence, on average $|(V \odot W)_\mathcal{I}|$ equals $k/4$. Furthermore, using Hoeffding's inequality (Theorem 2.11), we may conclude that

$$\Pr\left[\frac{k}{4} - |(V \odot W)_\mathcal{I}| > \beta k\right] = \Pr\left[|(V \odot W)_\mathcal{I}| < (\tfrac{1}{4} - \beta)k\right] \leq \exp(-2k\beta^2)\,.$$

Finally, the claim follows by applying the union bound over the choice of r and s (each 2^ℓ possibilities). □

Recall that $\mathcal{C} \subset \mathbb{F}_2^n$ is a binary code with minimum distance d, $\mathfrak{c}(\cdot)$ its encoding function, and that $m := |\mathcal{W}|$.

[9] When using matrix-product notation ambiguities could arise, e.g., in subscripts of probability distributions like P_{FX}: then it is not clear whether this means the joint distribution of F and X or the distribution of F acting on X?

Corollary 5.13 *Let $0 < \beta < \frac{1}{4}$, and let F be uniformly distributed over $\mathcal{F}$. Then, F has the following property except with probability $\binom{m}{2}2^{2\ell}\exp(-2d\beta^2)$: for any string $s \in \mathbb{F}_2^n$ (possibly depending on the choice of F), there exists at most one $\widetilde{c} \in \mathcal{C}$ such that for any code word $c \in \mathcal{C}$ different from $\widetilde{c}$, it holds that*

$$|f \odot (c \oplus s)| \geq \tfrac{1}{2}(\tfrac{1}{4} - \beta)d \qquad \forall f \in \operatorname{span}(F) \setminus \{\mathbf{0}\}$$

We prove the statement by arguing for two $\widetilde{c}$'s and showing that they must be identical. In the proof, we will make use of elementary properties of the Schur product and the Hamming weight:

1. $|a| \geq |a \odot b|$ for all $a, b \in \mathbb{F}_2^n$. (Follows immediately.)
2. $|a \odot b| + |a \odot c| \geq |a \odot (b \oplus c)|$ for all $a, b, c \in \mathbb{F}_2^n$.
   Proof. $|a \odot (b \oplus c)| = |a \odot b \oplus a \odot c| \leq |a \odot b| + |a \odot c|$, where the equality is the distributivity of the Schur product, and the inequality is the triangle inequality for the Hamming weight. □

Proof. By Lemma 5.12 with $\mathcal{I} := \{i \in [n] : c_i \neq c'_i\}$ for $c, c' \in \mathcal{C}$, and by applying the union bound over all possible pairs (c, c'), we obtain that except with probability $\binom{m}{2}2^{2\ell}\exp(-2d\beta^2)$ (over the choice of F), it holds that

$$|f \odot g \odot (c \oplus c')| > (\tfrac{1}{4} - \beta)d \tag{5.7}$$

for all $f, g \in \operatorname{span}(F) \setminus \{\mathbf{0}\}$ and all $c, c' \in \mathcal{C}$ with $c \neq c'$.

Now, for such an F, and for every choice of $s \in \mathbb{F}_2^n$, consider $\widetilde{c}_1, \widetilde{c}_2 \in \mathcal{C}$ and $f_1, f_2 \in \operatorname{span}(F) \setminus \{\mathbf{0}\}$ such that

$$|f_1 \odot (\widetilde{c}_1 \oplus s)| < \tfrac{1}{2}(\tfrac{1}{4} - \beta)d \quad \text{and} \quad |f_2 \odot (\widetilde{c}_2 \oplus s)| < \tfrac{1}{2}(\tfrac{1}{4} - \beta)d.$$

We will show that this implies $\widetilde{c}_1 = \widetilde{c}_2$, which proves the claim. Indeed, we can write

$$\begin{aligned}(\tfrac{1}{4} - \beta)d &> |f_1 \odot (\widetilde{c}_1 \oplus s)| + |f_2 \odot (\widetilde{c}_2 \oplus s)| \\ &\geq |f_1 \odot f_2 \odot (\widetilde{c}_1 \oplus s)| + |f_1 \odot f_2 \odot (\widetilde{c}_2 \oplus s)| \geq |f_1 \odot f_2 \odot (\widetilde{c}_1 \oplus \widetilde{c}_2)|\end{aligned}$$

where the second inequality is property (1) from above applied twice and the third inequality is property (2). This contradicts (5.7) unless $\widetilde{c}_1 = \widetilde{c}_2$. □

Proof of Theorem 5.11. Consider an execution of NEWQID, with a dishonest server S^* as described in Section 5.6.1. We let W, X and Z be the random variables that describe the values w, x and z occurring in the protocol.

From NEWQID's description, we see that F is uniform over $\mathcal{F}$. Hence, by Corollary 5.13 it will be "good" (in the sense that the bound from Corollary 5.13 holds) except with probability $\binom{m}{2} 2^{2\ell} \exp(-2d\beta^2)$. From here, we consider a fixed choice for F and condition on the event that it is "good," we will take the probability that F is "bad" into account at the end of the analysis. Although we have fixed F, we will keep using capital notation for it, to emphasize that F is a matrix. We also fix $G = g$ for an arbitrary g; the analysis below holds for any such choice.

Let Θ describe the qubit-wise measurement performed by S^* at the end of the execution, and Y the corresponding measurement outcome. By the non-adaptivity restriction and by the requirement in Definition 2.67 that S^* is initially independent of W, we may conclude that, once G and F are fixed, Θ is a function of Z. (Recall that $Z = F(X) \oplus g(W)$.)

We will define W' with the help of Corollary 5.13. Let $\hat{\Theta}$ be the quantized basis of Θ, as defined in Definition 5.8. Given a fixed value θ for Θ, and thus a fixed value $\hat{\theta}$ for $\hat{\Theta}$, we set s, which is a variable that occurs in Corollary 5.13, to $s = \hat{\theta}$. Corollary 5.13 now guarantees that there exists *at most one* $\widetilde{c}$. If $\widetilde{c}$ indeed exists, then we choose w' such that $\mathfrak{c}(w') = \widetilde{c}$. Otherwise, we pick $w' \in \mathcal{W}$ arbitrarily (any choice will do). Note that this defines the random variable W', and furthermore note that $Z \to \Theta \to \hat{\Theta} \to W'$ forms a Markov chain. Moreover, by the choice of w' it immediately follows from Corollary 5.13 that for all $w \neq w'$ and for all $f \in \operatorname{span}(F) \setminus \{\mathbf{0}\}$ it holds that

$$\left| f \odot (\mathfrak{c}(w) \oplus \hat{\theta}) \right| \geq \tfrac{1}{2}(\tfrac{1}{4} - \beta)d. \tag{5.8}$$

We will make use of this bound later in the proof.

Since the model (Section 5.6.1) enforces the dishonest server to measure all qubits at the end of the protocol, the system $E = (Y, Z, \Theta)$ is classical and hence the trace-distance-based user-security definition (Definition 2.67) simplifies to a bound on the statistical distance between distributions. I.e., it is sufficient to prove that

$$\mathrm{SD}(P_{EW|W'=w',W'\neq W}, P_{W|W'=w',W\neq W'} P_{E|W'=w',W\neq W'}) \leq \varepsilon$$

holds for any w'. Consider the distribution that appears above as the first argument to the statistical distance, i.e., $P_{EW|W'=w',W'\neq W}$. By substituting $E = (Y, Z, \Theta)$,

it factors as follows[10]

$$\begin{aligned} P_{YZ\Theta W|W',W\neq W'} &= P_{W|W',W\neq W'}\, P_{Z\Theta|WW',W\neq W'}\, P_{Y|Z\Theta WW',W\neq W'} \\ &= P_{W|W',W\neq W'}\, P_{Z\Theta|W',W\neq W'}\, P_{Y|F(X)\Theta WW',W\neq W'}, \end{aligned} \tag{5.9}$$

where the equality $P_{Z\Theta|WW',W\neq W'} = P_{Z\Theta|W',W\neq W'}$ holds by the following argument: Z is independent of W (since $F(X)$ acts as one-time pad) and $Z \to \Theta \to W'$ is a Markov chain, and S^* (who computes Θ from Z) is initially independent of W by Definition 2.67, hence W is independent of Z, Θ and W', which implies the above equality. The equality $P_{Y|Z\Theta WW',W\neq W'} = P_{Y|F(X)\Theta WW',W\neq W'}$ holds by the observation that given W, Z is uniquely determined by $F(X)$ and vice versa.

In the remainder of this proof we will show that

$$d_{\mathsf{unif}}(Y|F(X) = u, \Theta = v, W = w, W' = w') \leq \tfrac{1}{2} 2^{\frac{\ell}{2} - \frac{1}{4}(\frac{1}{4} - \beta)d},$$

for all u, v, w such that $w \neq w'$, where w' is determined by v. This then implies that the rightmost factor in (5.9) is essentially independent of W, and concludes the proof.

To simplify notation, we define $\mathcal{E}$ to be the event

$$\mathcal{E} := \{F(X) = u, \Theta = v, W = w, W' = w'\}$$

for fixed but arbitrary choices u, v and w such that $w \neq w'$, where w' is determined by v. We show closeness to the uniform distribution by using the XOR inequality from Diaconis *et al.* (Theorem 2.8), i.e., we use the inequality

$$d_{\mathsf{unif}}(Y|\mathcal{E}) \leq \tfrac{1}{2} \Big[\sum_{\alpha} \mathrm{bias}(\alpha \bullet Y|\mathcal{E})^2 \Big]^{\frac{1}{2}},$$

where the sum is over all α in $\mathbb{F}_2^n \setminus \{\mathbf{0}\}$. We split this sum into two parts, one for $\alpha \in \mathrm{span}(F)$ and one for α not in $\mathrm{span}(F)$, and analyze the two parts separately.

Since X is uniformly distributed, it follows that for any $\alpha \notin \mathrm{span}(F)$, it holds that $P_{\alpha \bullet X|F(X)}(\cdot|u) = \frac{1}{2}$ (for any u). We conclude that

$$\begin{aligned} \tfrac{1}{2} &= P_{\alpha \bullet X|F(X)} = P_{\alpha \bullet X|F(X)W} = P_{\alpha \bullet X|F(X)\Theta WW'} \\ &= P_{\alpha \bullet Y|F(X)\Theta WW'} = P_{\alpha \bullet Y|\mathcal{E}} \quad \forall \alpha \notin \mathrm{span}(F). \end{aligned}$$

[10] Note that Convention 2.2 applies here.

The second equality follows since W is independent of X. The third equality holds by the fact that Θ is computed from $F(X) \oplus g(W)$ and W' is determined by Θ. The fourth equality follows by the security of the one-time pad, i.e., recall that $Y = X \oplus \Delta$, where by Corollary 5.7 it holds that $\Delta \in \mathbb{F}_2^n$ is independent of X when conditioned on fixed values for $B = \mathfrak{c}(W)$ and Θ. Hence, it follows that $\mathrm{bias}(\alpha \bullet Y|\mathcal{E}) = 0$ for $\alpha \notin \mathrm{span}(F)$.

For any non-zero $\alpha \in \mathrm{span}(F)$, we can write

$$\begin{aligned}
\mathrm{bias}(\alpha \bullet Y|\mathcal{E}) &= \mathrm{bias}(\alpha \bullet (X \oplus \Delta)|\mathcal{E}) \\
&= \mathrm{bias}(\alpha \bullet X \oplus \alpha \bullet \Delta|\mathcal{E}) && \text{(distributivity of dot product)} \\
&= \mathrm{bias}(\alpha \bullet X|\mathcal{E})\mathrm{bias}(\alpha \bullet \Delta|\mathcal{E}) && \text{(Corollary 5.7)} \\
&\leq \mathrm{bias}(\alpha \bullet \Delta|\mathcal{E}) && (\mathrm{bias}(\alpha \bullet X) \leq 1) \\
&= \prod_{i \in [n]} \mathrm{bias}(\alpha_i \cdot \Delta_i|\mathcal{E}) && (\Delta_i \text{ independent}) \\
&= \prod_{i \in [n]:\alpha_i = 1} \mathrm{bias}(\Delta_i|\mathcal{E}) \\
&\leq \prod_{\substack{i \in [n]:\alpha_i = 1 \\ \hat{\theta}_i = \mathfrak{c}(w)_i \oplus 1}} 2^{-\frac{1}{2}} && \text{(Theorem 5.9)} \\
&= 2^{-\frac{1}{2}|\alpha \odot (\mathfrak{c}(w) \oplus \hat{\theta})|} \leq 2^{-\frac{1}{4}(\frac{1}{4} - \beta)d} && \text{(by (5.8))}
\end{aligned}$$

Combining the two parts, we get

$$\begin{aligned}
d_{\mathsf{unif}}(Y|\mathcal{E}) &\leq \tfrac{1}{2}\Big[\sum_{\alpha} \mathrm{bias}(\alpha \bullet Y|\mathcal{E})^2\Big]^{\frac{1}{2}} \\
&= \tfrac{1}{2}\Big[\sum_{\alpha \in \mathrm{span}(F) \setminus \{\mathbf{0}\}} \mathrm{bias}(\alpha \bullet Y|\mathcal{E})^2 + 0\Big]^{\frac{1}{2}} \leq \tfrac{1}{2} 2^{\frac{\ell}{2} - \frac{1}{4}(\frac{1}{4} - \beta)d}.
\end{aligned}$$

Incorporating the error probability of having a "bad" F completes the proof. □

5.6.5 Attack against NEWQID using Operations on Pairs of Qubits

We present an attack with which the dishonest server S^* can discard two passwords in one execution of NEWQID using coherent operations on pairs of qubits.

Before discussing this attack, we first explain a straightforward strategy by which S^* can discard one password per execution: S^* chooses a candidate password $\hat{w} \in \mathcal{W}$

and measures the state $H^{\mathfrak{c}(W)}|X\rangle$ qubit-wise in the basis $H^{\mathfrak{c}(\hat{w})}$ to obtain $Y \in \mathbb{F}_2^n$. S^* then computes $F(Y) \oplus g(\hat{w})$ and compares this to $Z = F(X) \oplus g(W)$, which he received from the user. If indeed $Z = F(Y) \oplus g(\hat{w})$, then it is very likely that $W = \hat{w}$, i.e., that S^* guessed the password correctly.

Let us now explain the attack, which is obtained by modifying the above strategy. The attack is based on the following observation [DFSS05]: if S^* can perform Bell measurements on qubit pairs $H^a|x_1\rangle \otimes H^a|x_2\rangle$, for $a, x_1, x_2 \in \mathbb{F}_2$, then he can learn the parity of $x_1 \oplus x_2$ for both choices of a simultaneously. This strategy can also be adapted to determine both parities of a pair in which the first qubit is encoded in a basis that is opposite to that of the second qubit, i.e., by appropriately applying a Hadamard gate prior to applying the Bell measurement.

Let the first bit of Z be equal to $f \bullet X \oplus g(W)_1$,[11] where $f \in \operatorname{span}(F) \setminus \{\mathbf{0}\}$. Let $\hat{w}_1, \hat{w}_2 \in \mathcal{W}$ be two candidate passwords. With the trick from above, S^* can measure the positions in the set

$$\mathcal{P} := \{i \in [n] : f_i = 1, \mathfrak{c}(\hat{w}_1)_i = 1 \oplus \mathfrak{c}(\hat{w}_2)_i\}$$

pairwise (assuming $|\mathcal{P}|$ to be even) using Bell measurements, while measuring the positions where $\mathfrak{c}(\hat{w}_1)$ and $\mathfrak{c}(\hat{w}_2)$ coincide using ordinary single-qubit measurements. This allows him to compute both "check bits" corresponding to both passwords *simultaneously*, i.e., those check bits coincide with $f \bullet Y_1 \oplus g(\hat{w}_1)_1$ and $f \bullet Y_2 \oplus g(\hat{w}_2)_1$, where $Y_1 \in \mathbb{F}_2^n$ and $Y_2 \in \mathbb{F}_2^n$ are the outcomes that S^* would have obtained if he had measured all qubits qubit-wise in either $\mathfrak{c}(\hat{w}_1)$ or $\mathfrak{c}(\hat{w}_2)$, respectively. If both these check bits are different from the bit Z_1, then S^* can discard both w_1 and w_2.

We have seen that in the *worst case*, the attack is capable of discarding two passwords in one execution, and hence clearly violates the security definition. On *average*, however, the attack seems to discard just one password per execution, i.e., a candidate password cannot be discarded if its check bit is consistent with Z_1, which essentially happens with probability $1/2$. This raises the question whether the security definition is unnecessarily strong, because it seems that not being able to discard more than one password on average would be sufficient. Apart from this, it might be possible to improve the attack, e.g., by selecting the positions where to measure pairwise in a more clever way, as to obtain multiple check bits (corresponding to multiple fs in the span of F) per candidate password, thereby increasing the probability of discarding a wrong candidate password.

[11] By $g(W)_1$ we mean the first bit of $g(W)$.

5.7 Conclusion

We view our work related to NEWQID as a first step in a promising line of research, aimed at achieving security in multiple models simultaneously. The main open problem in the context of the SQOM is to reprove our results in a more general model in which the dishonest server S^* can choose his basis adaptively. Also, it would be interesting to see whether similar results can be obtained in a model where the adversary is restricted to performing quantum operations on blocks of several qubits.