

Improving the Security of Quantum Protocols via Commit-and-Open

Ivan Damgård¹, Serge Fehr^{2,*}, Carolin Lunemann^{1,**},
Louis Salvail^{3,***}, and Christian Schaffner^{2,†}

¹ DAIMI, Aarhus University, Denmark
{ivan,carolin}@cs.au.dk

² Centrum Wiskunde & Informatica (CWI) Amsterdam, The Netherlands
{s.fehr,c.schaffner}@cwi.nl

³ Université de Montréal (DIRO), QC, Canada
salvail@iro.umontreal.ca

Abstract. We consider two-party quantum protocols starting with a transmission of some random BB84 qubits followed by classical messages. We show a general “compiler” improving the security of such protocols: if the original protocol is secure against an “almost honest” adversary, then the compiled protocol is secure against an arbitrary computation-bounded (quantum) adversary. The compilation preserves the number of qubits sent and the number of rounds up to a constant factor. The compiler also preserves security in the bounded-quantum-storage model (BQSM), so if the original protocol was BQSM-secure, the compiled protocol can only be broken by an adversary who has large quantum memory *and* large computing power. This is in contrast to known BQSM-secure protocols, where security breaks down completely if the adversary has larger quantum memory than expected. We show how our technique can be applied to quantum identification and oblivious transfer protocols.

1 Introduction

We consider two-party quantum protocols for mutually distrusting players Alice and Bob. Such protocols typically start by Alice sending n random BB84 qubits to Bob who is supposed to measure them. Then some classical exchange of messages follows. Several protocols following this pattern have been proposed, implementing Oblivious Transfer (OT), Commitment, and Password-Based Identification [1,4,7,8].

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

** Supported by the MOBISEQ project funded by NABIIT, Denmark.

*** Supported by the QUSEP project of the Danish Natural Science Research Council, and by the QuantumWorks Network.

† Supported by EU fifth framework project QAP IST 015848 and the NWO VICI project 2004-2009.

In more details, the first step of the protocol consists of Alice choosing random binary strings $x = x_1, \dots, x_n$ and $\theta = \theta_1, \dots, \theta_n$. She then prepares n particles where x_i is encoded in the state of the i 'th particle using basis θ_i . Bob chooses a basis string $\hat{\theta} = \hat{\theta}_1, \dots, \hat{\theta}_n$ and measures the i 'th particle in basis $\hat{\theta}_i$. If Bob plays honestly, he learns x_i whenever $\hat{\theta}_i = \theta_i$ and else gets a random independent result.

Protocols of the form we consider here are typically unconditionally secure against cheating by Alice, but can (in their basic form) be broken easily by Bob, if he does not measure the qubits immediately. This is because the protocol typically asks Alice to reveal θ at a later stage, and Bob can then measure the qubits with $\hat{\theta} = \theta$ and learn more information than he was supposed to.

In this paper, we show a general “compiler” that can be used to improve security against such an attack. We assume that the original protocol implements some two-party functionality \mathcal{F} with statistical security against Bob if he is *benign*, meaning that he treats the qubits “almost honestly”, a notion we make more precise below. Then we show that the compiled protocol also implements \mathcal{F} , but now with security against *any* computationally bounded (quantum) Bob (note that we cannot in general obtain unconditional security against both Alice and Bob, not even using quantum communication [13]). The compiled protocol preserves unconditional security against Alice and has the same number of transmitted qubits and rounds as the original one up to a constant factor.

By benign behavior of Bob, we mean that after having received the qubits, two conditions are satisfied: First, Bob’s quantum storage is essentially of size zero (note that it would be exactly zero if he had measured the qubits). Second, there exists a basis string $\hat{\theta}$ such that the uncertainty about x is essentially as it would be if Bob had really measured in bases $\hat{\theta}$, namely 1 bit for every position where $\hat{\theta}$ differs from θ .

Thus, with our compiler, one can build a protocol for any two-party functionality by designing a protocol that only has to be secure if Bob is benign. We note that proofs for known protocols typically go through under this assumption. For instance, our compiler can easily be applied to the quantum identification protocols of [7] and the OT protocol of [1].

The compiler is based on a computational assumption; namely we assume the existence of a classical commitment scheme with some special properties, similar to the commitment schemes used in [5] but with an additional extraction property, secure against a quantum adversary. A good candidate is the cryptosystem of Regev [16]. For efficiency, we use a common reference string which allows us to use Regev’s scheme in a simple way and, since it is relatively efficient, we get a protocol that is potentially practical. It is possible to generate the reference string from scratch, but this requires a more complicated non-constant round protocol [9].

The reader may ask whether it is really interesting to improve the security of quantum protocols for classical tasks such as identification or OT using a computational assumption. Perhaps it would be a more practical approach to use the same assumption to build *classical* protocols for the same tasks, secure against

quantum attacks? To answer this, it is important to point out that our compiler also preserves security in the bounded-quantum-storage model (BQSM) [6], and this feature allows us to get security properties that classical protocols cannot achieve. In the BQSM, one assumes that Bob can only keep in his quantum memory a limited number of qubits received from Alice. With current state of the art, it is much easier to transmit and measure qubits than it is to store them for a non-negligible time, suggesting that the BQSM and the subsequently proposed noisy-quantum-storage model [20] are reasonable. On the other hand, if the assumption fails and the adversary can perfectly store all qubits sent, the known protocols can be easily broken. In contrast, by applying our compiler, one obtains new protocols where the adversary must have large quantum storage *and* large computing power to break the protocol.¹

The basic technique we use to construct the compiler was already suggested in connection with the first quantum OT protocol from [1]: we try to force Bob to measure by asking him to commit (using a classical scheme) to all his basis choices and measurement results, and open some of them later. While classical intuition suggests that the commitments should force Bob to measure (almost) all the qubits, it has proved very tricky to show that the approach really works against a quantum adversary. In fact, it was previously very unclear what exactly the commit-and-open approach forces Bob to do. Although some partial results for OT have been shown [21,2], the original OT protocol from [1] has never been proved secure for a concrete unconditionally hiding commitment scheme – which is needed to maintain unconditional security against Alice. In this paper, we develop new quantum information-theoretic tools (that may be of independent interest) to characterize what commit-and-open achieves in general, namely it forces Bob to be benign. This property allows us to apply the compiler to any two-party functionality and in particular to show that the OT from [1] is indeed secure when using an appropriate commitment scheme.

2 Preliminaries

We assume the reader to be familiar with the basic notation and concepts of quantum information processing [14]. In this paper, the computational or $+$ -basis is defined by the pair $\{|0\rangle, |1\rangle\}$ (also written as $\{|0\rangle_+, |1\rangle_+\}$). The pair $\{|0\rangle_\times, |1\rangle_\times\}$ denotes the diagonal or \times -basis, where $|0\rangle_\times = (|0\rangle + |1\rangle)/\sqrt{2}$ and $|1\rangle_\times = (|0\rangle - |1\rangle)/\sqrt{2}$. We write $|x\rangle_\theta = |x_1\rangle_{\theta_1} \otimes \cdots \otimes |x_n\rangle_{\theta_n}$ for the n -qubit state where string $x = (x_1, \dots, x_n) \in \{0, 1\}^n$ is encoded in bases $\theta = (\theta_1, \dots, \theta_n) \in \{+, \times\}^n$. For $S \subseteq \{1, \dots, n\}$ of size s , we denote by $\bar{S} := \{1, \dots, n\} \setminus S$ the complement of S and define $x|_S \in \{0, 1\}^s$ and $\theta|_S \in \{+, \times\}^s$ to be the restrictions $(x_i)_{i \in S}$ and $(\theta_i)_{i \in S}$, respectively. For two strings $x, y \in \{0, 1\}^n$, we define the *Hamming distance* between x and y as $d_H(x, y) := |\{i : x_i \neq y_i\}|$.

¹ For the case of identification[7], the compiled protocol is not only secure against adversaries trying to impersonate Alice or Bob, but can also be made secure against man-in-the-middle attacks, where again the adversary must have large quantum storage and large computing power to break the protocol.

We use upper case letters for the random variables in the proofs that describe the respective values in the protocol. Given a bipartite quantum state ρ_{XE} , we say that X is *classical* if ρ_{XE} is of the form $\rho_{XE} = \sum_{x \in \mathcal{X}} P_X(x) |x\rangle\langle x| \otimes \rho_E^x$ for a probability distribution P_X over a finite set \mathcal{X} , i.e. the state of the quantum register E depends on the classical random variable X in the sense that E is in state ρ_E^x exactly if $X = x$. This naturally extends to states with two or more classical registers.

For a state ρ_{XE} as above, X is *independent* of register E if $\rho_{XE} = \rho_X \otimes \rho_E$, where $\rho_X = \sum_x P_X(x) |x\rangle\langle x|$ and $\rho_E = \sum_x P_X(x) \rho_E^x$. We also need to express that a random variable X is independent of a quantum state E when given a random variable Y . Independence means that when given Y , the state E gives no additional information on X . Formally, adopting the notion introduced in [7], we require that ρ_{XYE} equals $\rho_{X \leftrightarrow Y \leftrightarrow E}$, where the latter is defined as

$$\rho_{X \leftrightarrow Y \leftrightarrow E} := \sum_{x,y} P_{XY}(x,y) |x\rangle\langle x| \otimes |y\rangle\langle y| \otimes \rho_E^y,$$

where $\rho_E^y := \sum_x P_{X|Y}(x|y) \rho_E^{x,y}$. In other words, $\rho_{XYE} = \rho_{X \leftrightarrow Y \leftrightarrow E}$ precisely if $\rho_E^{x,y} = \rho_E^y$ for all x and y .

Full (conditional) independence is often too strong a requirement, and it usually suffices to be “close” to such a situation. Closeness of two states ρ and σ is measured in terms of their trace distance $\delta(\rho, \sigma) = \frac{1}{2} \text{tr}(|\rho - \sigma|)$, where for any operator A , $|A|$ is defined as $|A| := \sqrt{AA^\dagger}$.

A quantum algorithm consists of a family $\{C_n\}_{n \in \mathbb{N}}$ of quantum circuits and is said to run in *polynomial time*, if the number of gates of C_n is polynomial in n . Two families of quantum states $\{\rho_n\}_{n \in \mathbb{N}}$ and $\{\sigma_n\}_{n \in \mathbb{N}}$ are called *quantum-computationally indistinguishable*, denoted $\rho \stackrel{q}{\approx} \sigma$, if any polynomial-time quantum algorithm has negligible advantage (in n) of distinguishing ρ_n from σ_n . Analogously, we call them *statistically indistinguishable*, $\rho \stackrel{s}{\approx} \sigma$, if their trace distance $\delta(\rho_n, \sigma_n)$ is negligible in n .

Definition 2.1 (Min-Entropy). *The min-entropy of a random variable X with probability distribution P_X is defined as $H_\infty(X) := -\log(\max_x P_X(x))$.*

Definition 2.2 (Max-Entropy). *The max-entropy of a density matrix ρ is defined as $H_0(\rho) := \log(\text{rank}(\rho))$.*

We will make use of the following properties of a pure state that can be written as a “small superposition” of basis vectors; the proof is given in the full version [3].

Lemma 2.3. *Let $|\varphi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$ be of the form $|\varphi_{AE}\rangle = \sum_{i \in J} \alpha_i |i\rangle |\varphi_E^i\rangle$, where $\{|i\rangle\}_{i \in I}$ is a basis of \mathcal{H}_A and $J \subseteq I$. Then, the following holds.*

1. *Let $\tilde{\rho}_{AE} = \sum_{i \in J} |\alpha_i|^2 |i\rangle\langle i| \otimes |\varphi_E^i\rangle\langle \varphi_E^i|$, and let W and \tilde{W} be the outcome of measuring A of $|\varphi_{AE}\rangle$ respectively of $\tilde{\rho}_{AE}$ in some basis $\{|w\rangle\}_{w \in \mathcal{W}}$. Then,*

$$H_\infty(W) \geq H_\infty(\tilde{W}) - \log |J|.$$

2. *The reduced density matrix $\rho_E = \text{tr}_A(|\varphi_{AE}\rangle\langle \varphi_{AE}|)$ has max-entropy*

$$H_0(\rho_E) \leq \log |J|.$$

3 Definition of Security

In order to define security of our two-party protocols, we follow the framework put forward by Fehr and Schaffner in [10]. We are interested in quantum protocols that implement *classical functionalities* such as oblivious transfer. Such primitives are often used as building blocks in more complicated classical (multi-party) protocols which implement advanced tasks. Therefore, it is natural to restrict our focus on quantum protocols that run in a classical environment and have classical in- and outputs. A two-party quantum protocol $\Pi = (\mathbf{A}_m, \mathbf{B}_m)$ consists of an infinite family of interactive quantum circuits for players Alice and Bob indexed by the security parameter m (in our case, m will also be the number of qubits transmitted). To ease notation, we often leave the dependence on m implicit. A classical non-reactive two-party *ideal functionality* \mathcal{F} is given by a conditional probability distribution $P_{\mathcal{F}(U,V)|UV}$, inducing a pair of random variables $(X, Y) = \mathcal{F}(U, V)$ for every joint distribution of U and V . The definition of correctness of a protocol is straightforward.

Definition 3.1 (Correctness). *A protocol $\Pi = (\mathbf{A}, \mathbf{B})$ correctly implements an ideal classical functionality \mathcal{F} , if for every distribution of the input values U and V , the resulting common output satisfies*

$$(U, V, (X, Y)) \stackrel{\circ}{\approx} (U, V, \mathcal{F}(U, V)).$$

Let us denote by $out_{\hat{\mathbf{A}}, \hat{\mathbf{B}}}^{\mathcal{F}}$ the joint output² of the “ideal-life” protocol, where Alice and Bob forward their inputs to \mathcal{F} and output whatever they obtain from \mathcal{F} . And we write $out_{\hat{\mathbf{A}}, \hat{\mathbf{B}}'}^{\mathcal{F}}$ for the joint output of the execution of this protocol with a dishonest Bob with strategy $\hat{\mathbf{B}}'$ (and similarly for a dishonest Alice). Note that Bob’s possibilities in the ideal world are very limited: he can produce some classical input V for \mathcal{F} from his input quantum state V' , and then he can prepare and output a quantum state Y' which might depend on \mathcal{F} ’s classical reply Y .

3.1 Information-Theoretic Security

We define information-theoretic security using the real/ideal-world paradigm, which requires that by attacking a protocol in the real world the dishonest party cannot achieve (significantly) more than when attacking the corresponding functionality in the ideal world. To be consistent with the framework used in [10], we restrict the joint input state, consisting of a classical input to the honest party and a possibly quantum input to the dishonest party, to a special form: in case of a dishonest Bob (and correspondingly for a dishonest Alice), we require that Bob’s input consists of a classical part Z and a quantum part V' , such that the joint state $\rho_{UZV'}$ satisfies $\rho_{UZV'} = \rho_{U \leftrightarrow Z \leftrightarrow V'}$, i.e., that V' is correlated with Alice’s input only via the classical Z . We call a joint input state of that form (respectively of the form $\rho_{U'ZV} = \rho_{U' \leftrightarrow Z \leftrightarrow V}$ in case of dishonest Alice) a

² We use a slightly different notation here than in [10]. Our notation $out_{\hat{\mathbf{A}}, \hat{\mathbf{B}}}^{\mathcal{F}}$ does not mention the name of the input registers and corresponds to $(\mathcal{F}_{\hat{\mathbf{A}}, \hat{\mathbf{B}}})\rho_{UV}$ in [10].

legitimate input state. As shown in [10], this restriction on the input state leads to a meaningful security definition with a composition theorem that guarantees sequential composition within *classical* outer protocols. Furthermore, the results of Section 4 also hold when quantifying over all (possibly non-legitimate) joint input states.

Definition 3.2 (Unconditional security against dishonest Alice). *A protocol $\Pi = (A, B)$ implements an ideal classical functionality \mathcal{F} unconditionally securely against dishonest Alice, if for any real-world adversary A' there exists an ideal-world adversary \hat{A}' such that for any legitimate input state, it holds that the outputs in the real and ideal world are statistically indistinguishable, i.e.*

$$\text{out}_{A',B}^{\Pi} \stackrel{s}{\approx} \text{out}_{\hat{A},\hat{B}}^{\mathcal{F}}.$$

Definition 3.3 (Unconditional security against dishonest Bob). *A protocol $\Pi = (A, B)$ implements an ideal classical functionality \mathcal{F} unconditionally securely against dishonest Bob, if for any real-world adversary B' there exists an ideal-world adversary \hat{B}' such that for any legitimate input state, it holds that the outputs in the real and ideal world are statistically indistinguishable, i.e.*

$$\text{out}_{A,B'}^{\Pi} \stackrel{s}{\approx} \text{out}_{\hat{A},\hat{B}'}^{\mathcal{F}}.$$

It has been shown in Theorem 5.1 in [10] that protocols fulfilling the above definitions compose sequentially as follows. For a classical real-life protocol Σ which makes at most k oracle calls to functionalities $\mathcal{F}_1, \dots, \mathcal{F}_k$, it is guaranteed that whatever output Σ produces, the output produced when the oracle calls are replaced by ε -secure protocols is at distance at most $O(k\varepsilon)$.

3.2 Computational Security in the CRS Model

One can define security against a computationally bounded dishonest Bob analogously to information-theoretic security with the two differences that the input given to the parties has to be sampled by an efficient quantum algorithm and that the output states should be computationally indistinguishable.

In the common-reference-string (CRS) model, all participants in the real-life protocol $\Pi_{A,B}$ have access to a classical public string ω which is chosen before any interaction starts according to a distribution only depending on the security parameter. On the other hand, the participants in the “ideal-life” protocol $\mathcal{F}_{\hat{A},\hat{B}}$ interacting only with the ideal functionality do not make use of the string ω . Hence, an ideal-world adversary \hat{B}' , that operates by simulating the real world to the adversary B' , is free to choose ω in any way he wishes.

In order to define computational security against a dishonest Bob in the CRS model, we consider a polynomial-size quantum circuit, called *input sampler*, which takes as input the security parameter m and the CRS ω (chosen according to its distribution) and produces the input state $\rho_{UZV'}$; U is Alice’s classical input to the protocol, and Z and V' denote the respective classical and quantum information given to dishonest Bob. We call the input sampler *legitimate* if $\rho_{UZV'} = \rho_{U \leftrightarrow Z \leftrightarrow V'}$.

In the following and throughout the article, we let $\mathfrak{B}_{\text{poly}}$ be the family of all *polynomial-time* quantum strategies for dishonest Bob B' .

Definition 3.4 (Computational security against dishonest Bob). *A protocol $\Pi = (A, B)$ implements an ideal classical functionality \mathcal{F} computationally securely against dishonest Bob, if for any real-world adversary $B' \in \mathfrak{B}_{\text{poly}}$ who has access to the common reference string ω , there exists an ideal-world adversary $\hat{B}' \in \mathfrak{B}_{\text{poly}}$ not using ω such that for any efficient legitimate input sampler, it holds that the outputs in the real and ideal world are q -indistinguishable, i.e.*

$$\text{out}_{A, B'}^{\Pi} \stackrel{q}{\approx} \text{out}_{\hat{A}, \hat{B}'}^{\mathcal{F}}.$$

In the full version [3], we show that also the computational security definition, as given here, allows for (sequential) composition of quantum protocols into classical outer protocols.

4 Improving the Security via Commit-and-Open

4.1 Security Against Benign Bob

In this paper, we consider quantum two-party protocols that follow a particular but very typical construction design. These protocols consist of two phases, called *preparation* and *post-processing* phase, and are as specified in Figure 1. We call a protocol that follows this construction design a *BB84-type* protocol.

PROTOCOL Π

Preparation: A chooses $x \in_R \{0, 1\}^n$ and $\theta \in_R \{+, \times\}^n$ and sends $|x\rangle_{\theta}$ to B, and B chooses $\hat{\theta} \in_R \{+, \times\}^n$ and obtains $\hat{x} \in \{0, 1\}^n$ by measuring $|x\rangle_{\theta}$ in basis $\hat{\theta}$.

Post-processing: Arbitrary classical communication and local computations.

Fig. 1. Generic BB84-type quantum protocol Π

The following definition captures information-theoretic security against a somewhat mildly dishonest Bob who we call a *benign* (dishonest) Bob. Such a dishonest Bob is benign in that, in the preparation phase, he does not deviate too much from what he is supposed to do; in the post-processing phase though, he may be arbitrarily dishonest.

To make this description formal, we fix an arbitrary choice of θ and an arbitrary value for the classical information, z , which B' may obtain as a result of the preparation phase (i.e. $z = (\hat{\theta}, \hat{x})$ in case B' is actually honest). Let X denote the random variable describing the bit-string x , where we understand the distribution P_X of X to be conditioned on the fixed values of θ and z . Furthermore,

let ρ_E be the state of B' 's quantum register E after the preparation phase. Note that, still with fixed θ and z , ρ_E is of the form $\rho_E = \sum_x P_X(x)\rho_E^x$, where ρ_E^x is the state of B' 's quantum register in case X takes on the value x . In general, the ρ_E^x 's may be mixed, but we can think of them as being reduced pure states: $\rho_E^x = \text{tr}_R(|\psi_{ER}^x\rangle\langle\psi_{ER}^x|)$ for a suitable register R and pure states $|\psi_{ER}^x\rangle$; we then call the state $\rho_{ER} = \sum_x P_X(x)|\psi_{ER}^x\rangle\langle\psi_{ER}^x|$ a *pointwise purification* (with respect to X) of ρ_E .

Obviously, in case B' is honest, X_i is fully random whenever $\theta_i \neq \hat{\theta}_i$, so that $H_\infty(X|_I | X|_{\bar{I}} = x|_{\bar{I}}) = d_H(\theta|_I, \hat{\theta}|_I)$ for every $I \subseteq \{1, \dots, n\}$ and every $x|_{\bar{I}}$, and B' does not store any non-trivial quantum state so that R is “empty” and $H_0(\rho_{ER}) = H_0(\rho_E) = 0$. A benign Bob B' is now specified to behave close-to-honestly in the preparation phase: he produces an auxiliary output $\hat{\theta}$ after the preparation phase, and given this output, we are in a certain sense close to the ideal situation where Bob really measured in basis $\hat{\theta}$ as far as the values of $H_\infty(X|_I | X|_{\bar{I}} = x|_{\bar{I}})$ and $H_0(\rho_{ER})$ are concerned.³ We now make this precise:

Definition 4.1 (Unconditional security against benign Bob). *A BB84-type quantum protocol Π securely implements \mathcal{F} against a β -benign Bob for some parameter $\beta \geq 0$, if it securely implements \mathcal{F} according to Definition 3.3, with the following two modifications:*

1. *The quantification is over all B' with the following property: after the preparation phase B' either aborts, or else produces an auxiliary output $\hat{\theta} \in \{+, \times\}^n$. Moreover, the joint state of A, B' (after $\hat{\theta}$ has been output) is statistically indistinguishable from a state for which it holds that for any fixed values of θ , $\hat{\theta}$ and z , for any subset $I \subseteq \{1, \dots, n\}$, and for any $x|_{\bar{I}}$*

$$H_\infty(X|_I | X|_{\bar{I}} = x|_{\bar{I}}) \geq d_H(\theta|_I, \hat{\theta}|_I) - \beta n \quad \text{and} \quad H_0(\rho_{ER}) \leq \beta n \quad (1)$$

where ρ_{ER} is the pointwise purification of ρ_E with respect to X .

2. *B' 's running-time is polynomial in the running-time of B' .*

4.2 From Benign to Computational Security

We show a generic compiler which transforms any BB84-type protocol into a new quantum protocol for the same task. The compiler achieves that if the original protocol is unconditionally secure against dishonest Alice and unconditionally secure against *benign* Bob, then the compiled protocol is still unconditionally secure against dishonest Alice and it is *computationally secure* against *arbitrary* dishonest Bob.

The idea behind the construction of the compiler is to incorporate a commit-scheme and force Bob to behave benignly by means of a commit-and-open

³ The reason why we consider the *pointwise purification* of ρ_E is to prevent Bob from artificially blowing up $H_0(\rho_{ER})$ by locally generating a large mixture or storing an unrelated mixed input state.

procedure. Figure 2 shows the compilation of an arbitrary BB84-type protocol Π . The quantum communication is increased from n to $m = n/(1 - \alpha)$ qubits, where $0 < \alpha < 1$ is some additional parameter that can be arbitrarily chosen. The compiled protocol also requires 3 more rounds of interaction.

PROTOCOL $\mathcal{C}^\alpha(\Pi)$

Preparation: A chooses $x \in_R \{0, 1\}^m$ and $\theta \in_R \{+, \times\}^m$ and sends $|x\rangle_\theta$ to B. Then, B chooses $\hat{\theta} \in_R \{+, \times\}^m$ and obtains $\hat{x} \in \{0, 1\}^m$ by measuring $|x\rangle_\theta$ in basis $\hat{\theta}$.

Verification:

1. B commits to $\hat{\theta}$ and \hat{x} position-wise: $c_i := \text{Commit}((\hat{\theta}_i, \hat{x}_i), r_i)$ with randomness r_i for $i = 1, \dots, m$. He sends the commitments to A.
2. A sends a random test subset $T \subset \{1, \dots, m\}$ of size αm . B opens c_i for all $i \in T$, and A checks that the openings were correct and that $x_i = \hat{x}_i$ whenever $\theta_i = \hat{\theta}_i$. If all tests are passed, A accepts, otherwise, she rejects and aborts.
3. The tested positions are discarded by both parties: A and B restrict x and θ , respectively $\hat{\theta}$ and \hat{x} , to $i \in \bar{T}$.

Post-processing: As in Π (with x, θ, \hat{x} and $\hat{\theta}$ restricted to the positions $i \in \bar{T}$).

Fig. 2. Compiled protocol $\mathcal{C}^\alpha(\Pi)$

We need to specify what kind of commitment scheme to use. In order to preserve unconditional security against dishonest Alice, the commitment scheme needs to be unconditionally hiding, and so can at best be computationally binding. However, for a plain computationally binding commitment scheme, we do not know how to reduce the computational security of $\mathcal{C}^\alpha(\Pi)$ against dishonest Bob to the computational binding property of the commitment scheme.⁴ Therefore, we use a commitment scheme with additional properties: we require a *keyed* commitment scheme $\text{Commit}_{\text{pk}}$, where the corresponding public key pk is generated by one of two possible key-generation algorithms: \mathcal{G}_H or \mathcal{G}_B . For a key pk_H generated by \mathcal{G}_H , the commitment scheme $\text{Commit}_{\text{pk}_H}$ is unconditionally hiding, whereas the other generator, \mathcal{G}_B , actually produces a key *pair* (pk_B, sk) , so that the secret key sk allows to efficiently extract m from $\text{Commit}_{\text{pk}_B}(m, r)$, and as such $\text{Commit}_{\text{pk}_B}$ is unconditionally binding. Furthermore, we require pk_H and pk_B to be computationally indistinguishable, even against quantum attacks. We call such a commitment scheme a *dual-mode* commitment scheme.⁵ As a candidate for implementing such a system, we propose the public-key encryption scheme of Regev [16], which is based on a worst-case lattice assumption and is not known

⁴ Classically, this would be done by a rewinding argument, but this fails to work for a quantum Bob.

⁵ The notions of dual-mode *cryptosystems* and of meaningful/meaningless encryptions, as introduced in [15] and [12], are similar in spirit but differ slightly technically.

to be breakable even by (efficient) quantum algorithms. Regev does not explicitly state that the scheme has the property we need, but this is implicit in his proof that the underlying computational assumption implies semantic security. For simplicity and efficiency, we consider the common-reference-string model, and we assume the key pkB for the commitment scheme, generated according to \mathcal{G}_B , to be contained in the CRS. We sketch in Section 6 how to avoid the CRS model, at the cost of a non constant-round construction where the parties generate the CRS jointly by means of a coin-tossing protocol (see [9] for details).

We sometimes write $C_{\text{pkH}}^\alpha(\Pi)$ for the compiled protocol $C^\alpha(\Pi)$ to stress that a key pkH produced by \mathcal{G}_H is used for the dual-mode commitment scheme, and we write $C_{\text{pkB}}^\alpha(\Pi)$ when a key pkB produced by \mathcal{G}_B is used instead.

Theorem 4.2. *Let Π be a BB84-type protocol, unconditionally secure against dishonest Alice and against β -benign Bob for some constant $\beta > 0$. Consider the compiled protocol $C^\alpha(\Pi)$ for an arbitrary $\alpha > 0$, where the commitment scheme is instantiated by a dual-mode commitment scheme as described above. Then, $C^\alpha(\Pi)$ is unconditionally secure against dishonest Alice and computationally secure against dishonest Bob in the CRS model.*

We now prove this theorem, which assumes noise-free quantum communication; we explain in Section 4.4 how to generalize it for a noisy quantum channel. Correctness and unconditional security against dishonest Alice are obvious; the latter is due to the unconditional hiding property of the commitment scheme. As for computational security against dishonest Bob, according to Definition 3.4, we need to prove that for every real-world adversary $B' \in \mathfrak{B}_{\text{poly}}$ attacking $C^\alpha(\Pi)$, there exists a suitable ideal-world adversary $\hat{B}' \in \mathfrak{B}_{\text{poly}}$ attacking \mathcal{F} such that

$$\text{out}_{A,B'}^{C^\alpha(\Pi)} \stackrel{q}{\approx} \text{out}_{\hat{A},\hat{B}'}^{\mathcal{F}} .$$

First, note that by the computational indistinguishability of pkH and pkB ,

$$\text{out}_{A,B'}^{C^\alpha(\Pi)} = \text{out}_{A,B'}^{C_{\text{pkH}}^\alpha(\Pi)} \stackrel{q}{\approx} \text{out}_{A,B'}^{C_{\text{pkB}}^\alpha(\Pi)} . \tag{2}$$

Then, we construct an adversary $B'_o \in \mathfrak{B}_{\text{poly}}$ who attacks the unconditional security against benign Bob of protocol Π , and which satisfies

$$\text{out}_{A,B'}^{C_{\text{pkB}}^\alpha(\Pi)} = \text{out}_{A_o,B'_o}^\Pi , \tag{3}$$

where A_o honestly executes Π . We define B'_o in the following way. Consider the execution of $C^\alpha(\Pi)$ between A and B' . We split A into two players A_o and \tilde{A} , where we think of \tilde{A} as being placed in between A_o and B' , see Figure 3. A_o plays honest Alice’s part of Π while \tilde{A} acts as follows: It receives n qubits from A_o , produces $\alpha n / (1 - \alpha)$ random BB84 qubits of its own and interleaves them randomly with those received and sends the resulting $m = n / (1 - \alpha)$ qubits to B' . It then does the verification step of $C^\alpha(\Pi)$ with B' , asking to have commitments corresponding to its own qubits opened. If this results in accept, it lets A_o finish the protocol with B' . Note that the pair (A_o, \tilde{A}) does exactly the

same as A ; however, we can also move the actions of \tilde{A} to Bob’s side, and define B'_o as follows. B'_o samples (pk_B, sk) according to \mathcal{G}_B and executes Π with A by locally running A and B' , using pk_B as CRS. If \tilde{A} accepts the verification then B'_o outputs $\hat{\theta} \in \{0, 1\}^n$ (as required from a *benign* Bob), obtained by decrypting the unopened commitments with the help of sk ; else, B'_o aborts at this point. It is now clear that Equation (3) holds: exactly the same computation takes place in both “experiments”, the only difference being that they are executed partly by different entities. The last step is to show that

$$out_{A_o, B'_o}^{\Pi} \stackrel{s}{\approx} out_{\tilde{A}, \hat{B}'}^{\mathcal{F}}, \tag{4}$$

for some \hat{B}' . It is clear that the theorem follows from (2) - (4) together.




Fig. 3. Constructing an attacker B'_o against Π from an attacker B' against $C^\alpha(\Pi)$

Now (4) actually claims that \hat{A}, \hat{B}' successfully simulate A_o and B'_o executing Π , and this claim follows by assumption of benign security of Π if we show that B'_o is β -benign according to Definition 4.1 for any $\beta > 0$. We show this in the following subsection, i.e., the joint state of A_o, B'_o after the preparation phase is statistically indistinguishable from a state ρ_{Ideal} which satisfies the bounds (1) from Definition 4.1.

4.3 Completing the Proof: Bounding Entropy and Memory Size

First recall that A_o executing Π with B'_o can equivalently be thought of as A executing $C_{pk_B}^\alpha(\Pi)$ with B' . Furthermore, a joint state of A, B' is clearly also a joint state of A_o, B'_o .

To show the existence of ρ_{Ideal} as promised above, it therefore suffices to show such a state for A, B' . In other words, we need to show that the execution of $C_{pk_B}^\alpha(\Pi)$ with honest Alice A and arbitrarily dishonest Bob B' will, after verification, be close to a state where (1) holds. To show this closeness, we consider an equivalent EPR-pair version, where Alice creates m EPR pairs $(|00\rangle + |11\rangle)/\sqrt{2}$, sends one qubit in each pair to Bob and keeps the others in register A . Alice measures her qubits only when needed: she measures the qubits within T in Step 2 of the verification phase, and the remaining qubits at the end of the verification phase. With respect to the information Alice and Bob obtain, this EPR version is *identical* to the original protocol $C_{pk_B}^\alpha(\Pi)$: the only difference is the point in time when Alice obtains certain information. Furthermore, we can also do the

following modification without affecting (1). Instead of measuring her qubits in T in *her* basis $\theta|_T$, she measures them in *Bob's* basis $\hat{\theta}|_T$; however, she still verifies only whether $x_i = \hat{x}_i$ for those $i \in T$ with $\theta_i = \hat{\theta}_i$. Because the positions $i \in T$ with $\theta_i \neq \hat{\theta}_i$ are not used in the protocol at all, this change has no effect. As the commitment scheme is unconditionally binding if key pkB is used, Bob's basis $\hat{\theta}$ is well defined by his commitments (although hard to compute), even if Bob is dishonest. The resulting scheme is given in Figure 4.

PROTOCOL EPR- $C_{\text{pkB}}^\alpha(\Pi)$

Preparation: A prepares m EPR pairs and sends the second qubit in each pair to Bob while keeping the others in register $A = A_1 \cdots A_m$. B chooses $\hat{\theta} \in_R \{+, \times\}^m$ and obtains $\hat{x} \in \{0, 1\}^m$ by measuring the received qubits in basis $\hat{\theta}$.

Verification:

1. B commits to $\hat{\theta}$ and \hat{x} position-wise: $c_i := \text{Commit}((\hat{\theta}_i, \hat{x}_i), r_i)$ with randomness r_i for $i = 1, \dots, m$. He sends the commitments to A.
2. A sends a random test subset $T \subset \{1, \dots, m\}$ of size αm . B opens c_i for all $i \in T$. A chooses $\theta \in_R \{+, \times\}^m$, measures registers A_i with $i \in T$ in basis $\hat{\theta}_i$ to obtain x_i , and she checks that the openings were correct and that $x_i = \hat{x}_i$ whenever $\theta_i = \hat{\theta}_i$ for $i \in T$. If all tests are passed, A accepts, otherwise, she rejects and aborts the protocol.
3. A measures the remaining registers in basis $\theta|_{\bar{T}}$ to obtain $x|_{\bar{T}}$. The tested positions are discarded by both parties: A and B restrict x and θ , respectively $\hat{\theta}$ and \hat{x} , to the positions $i \in \bar{T}$.

Post-processing: As in Π (with x, θ, \hat{x} and $\hat{\theta}$ restricted to the positions $i \in \bar{T}$).

Fig. 4. EPR version of $C_{\text{pkB}}^\alpha(\Pi)$

We consider an execution of the scheme from Figure 4 with an honest Alice A and a dishonest Bob B', and we fix $\hat{\theta}$ and \hat{x} , determined by Bob's commitments. Let $|\varphi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$ be the state of the joint system right before Step 2 of the verification phase. Since in the end, we are anyway interested in the pointwise purification of Bob's state, we may indeed assume this state to be pure; if it is not, then we purify it and carry the purifying register R along with E . Clearly, if B' had honestly done his measurements then $|\varphi_{AE}\rangle = |\hat{x}\rangle_{\hat{\theta}} \otimes |\varphi_E\rangle$ for some $|\varphi_E\rangle \in \mathcal{H}_E$. In this case, the quantum memory E would be empty: $H_0(|\varphi_E\rangle\langle\varphi_E|) = 0$. Moreover, X , obtained by measuring $A|_{\bar{T}}$ in basis $\theta|_{\bar{T}}$, would contain $d_H(\theta|_{\bar{T}}, \hat{\theta}|_{\bar{T}})$ random bits. We show that the verification phase enforces these properties, at least approximately in the sense of (1), for an arbitrary dishonest Bob B'.

In the following, $r_H(\cdot, \cdot)$ denotes the relative Hamming distance between two strings, i.e., the Hamming distance divided by their length. Recall that $T \subset \{1, \dots, m\}$ is random subject to $|T| = \alpha m$. Furthermore, for a fixed $\hat{\theta}$ but a randomly chosen θ , the subset $T' = \{i \in T : \theta_i = \hat{\theta}_i\}$ is a random subset

(of arbitrary size) of T . Let the random variable $Test$ describe the choice of $test = (T, T')$ as specified above, and consider the state

$$\rho_{TestAE} = \rho_{Test} \otimes |\varphi_{AE}\rangle\langle\varphi_{AE}| = \sum_{test} P_{Test}(test) |test\rangle\langle test| \otimes |\varphi_{AE}\rangle\langle\varphi_{AE}|$$

consisting of the classical $Test$ and the quantum state $|\varphi_{AE}\rangle$.

Lemma 4.3. *For any $\varepsilon > 0$, $\hat{x} \in \{0, 1\}^m$ and $\hat{\theta} \in \{+, \times\}^m$, the state ρ_{TestAE} is negligibly close (in m) to a state*

$$\tilde{\rho}_{TestAE} = \sum_{test} P_{Test}(test) |test\rangle\langle test| \otimes |\tilde{\varphi}_{AE}^{test}\rangle\langle\tilde{\varphi}_{AE}^{test}|$$

where for any $test = (T, T')$:

$$|\tilde{\varphi}_{AE}^{test}\rangle = \sum_{x \in B_{test}} \alpha_x^{test} |x\rangle_{\hat{\theta}} |\psi_E^x\rangle$$

for $B_{test} = \{x \in \{0, 1\}^m \mid r_H(x|_{\bar{T}}, \hat{x}|_{\bar{T}}) \leq r_H(x|_{T'}, \hat{x}|_{T'}) + \varepsilon\}$ and arbitrary coefficients $\alpha_x^{test} \in \mathbb{C}$.

In other words, we are close to a situation where for *any* choice of T and T' and for *any* outcome $x|_T$ when measuring $A|_T$ in basis $\hat{\theta}|_T$, the relative error $r_H(x|_{T'}, \hat{x}|_{T'})$ gives an upper bound (which holds with probability 1) on the relative error $r_H(x|_{\bar{T}}, \hat{x}|_{\bar{T}})$ one would obtain by measuring the remaining subsystems A_i with $i \in \bar{T}$ in basis $\hat{\theta}_i$.

Proof. For any $test$ we let $|\tilde{\varphi}_{AE}^{test}\rangle$ be the renormalized projection of $|\varphi_{AE}\rangle$ into the subspace $\text{span}\{|x\rangle_{\hat{\theta}} \mid x \in B_{test}\} \otimes \mathcal{H}_E$ and let $|\tilde{\varphi}_{AE}^{test\perp}\rangle$ be the renormalized projection of $|\varphi_{AE}\rangle$ into the orthogonal complement, such that $|\varphi_{AE}\rangle = \varepsilon_{test} |\tilde{\varphi}_{AE}^{test}\rangle + \varepsilon_{test}^\perp |\tilde{\varphi}_{AE}^{test\perp}\rangle$ with $\varepsilon_{test} = \langle\tilde{\varphi}_{AE}^{test}|\varphi_{AE}\rangle$ and $\varepsilon_{test}^\perp = \langle\tilde{\varphi}_{AE}^{test\perp}|\varphi_{AE}\rangle$. By construction, $|\tilde{\varphi}_{AE}^{test}\rangle$ is of the form required in the statement of the lemma. A basic property of the trace norm of pure states gives

$$\delta(|\varphi_{AE}\rangle\langle\varphi_{AE}|, |\tilde{\varphi}_{AE}^{test}\rangle\langle\tilde{\varphi}_{AE}^{test}|) = \sqrt{1 - |\langle\tilde{\varphi}_{AE}^{test}|\varphi_{AE}\rangle|^2} = |\varepsilon_{test}^\perp|.$$

This last term corresponds to the square root of the probability, when given $test$, to observe a string $x \notin B_{test}$ when measuring subsystem A of $|\varphi_{AE}\rangle$ in basis $\hat{\theta}$. Furthermore, using elementary properties of the trace norm and Jensen's inequality gives

$$\begin{aligned} \delta(\rho_{TestAE}, \tilde{\rho}_{TestAE})^2 &= \left(\sum_{test} P_{Test}(test) \delta(|\varphi_{AE}\rangle\langle\varphi_{AE}|, |\tilde{\varphi}_{AE}^{test}\rangle\langle\tilde{\varphi}_{AE}^{test}|) \right)^2 \\ &= \left(\sum_{test} P_{Test}(test) |\varepsilon_{test}^\perp| \right)^2 \leq \sum_{test} P_{Test}(test) |\varepsilon_{test}^\perp|^2, \end{aligned}$$

where the last term is the probability to observe a string $x \notin B_{test}$ when choosing $test$ according to P_{Test} and measuring subsystem A of $|\varphi_{AE}\rangle$ in basis $\hat{\theta}$. This situation, though, is a classical sampling problem, for which it is well known that for any measurement outcome x , the probability (over the choice of $test$) that $x \notin B_{test}$ is negligible in m (see e.g. [11]). \square

In combination with Lemma 2.3 on “small superpositions of product states”, and writing h for the binary entropy function $h(\mu) = -(\mu \log(\mu) + (1 - \mu) \log(1 - \mu))$ as well as using that $|\{y \in \{0, 1\}^n \mid d_H(y, \hat{y}) \leq \mu n\}| \leq 2^{h(\mu)n}$ for any $\hat{y} \in \{0, 1\}^n$ and $0 \leq \mu \leq \frac{1}{2}$, we can conclude the following (the proof is given in [3]).

Corollary 4.4. *Let $\tilde{\rho}_{TestAE}$ be of the form as in Lemma 4.3 (for given ε , \hat{x} and $\hat{\theta}$). For any fixed $test = (T, T')$ and for any fixed $x|_T \in \{0, 1\}^{\alpha m}$ with $err := r_H(x|_{T'}, \hat{x}|_{T'}) \leq \frac{1}{2}$, let $|\psi_{AE}\rangle$ be the state to which $|\tilde{\varphi}_{AE}^{test}\rangle$ collapses when for every $i \in T$ subsystem A_i is measured in basis $\hat{\theta}_i$ and x_i is observed, where we understand A in $|\psi_{AE}\rangle$ to be restricted to the registers A_i with $i \in \bar{T}$. Finally, let $\sigma_E = \text{tr}_A(|\psi_{AE}\rangle\langle\psi_{AE}|)$ and let the random variable X describe the outcome when measuring the remaining $n = (1 - \alpha)m$ subsystems of A in basis $\theta|_{\bar{T}} \in \{+, \times\}^n$. Then, for any subset $I \subseteq \{1, \dots, n\}$ and any $x|_I$,⁶*

$$H_\infty(X|_I \mid X|_{\bar{I}} = x|_{\bar{I}}) \geq d_H(\theta|_I, \hat{\theta}|_I) - h(err + \varepsilon)n \quad \text{and} \quad H_0(\sigma_E) \leq h(err + \varepsilon)n.$$

Thus, the number of errors between the measured $x|_{T'}$ and the given $\hat{x}|_{T'}$ gives us a bound on the min-entropy of the outcome when measuring the remaining subsystems of A , and on the max-entropy of the state of subsystem E .

The claim to be shown now follows by combining Lemma 4.3 and Corollary 4.4. Indeed, the ideal state ρ_{Ideal} we promised is produced by putting A and B' in the state $\tilde{\rho}_{TestAE}$ defined in Lemma 4.3, and running Steps 2 and 3 of the verification phase. This state is negligibly close to the real state since by Lemma 4.3 we were negligibly close to the real state before these operations. Corollary 4.4 guarantees that (1) is satisfied.

4.4 In the Presence of Noise

In the description of the compiler \mathcal{C}^α and in its analysis, we assumed the quantum communication to be noise-free. Indeed, if the quantum communication is noisy honest Alice is likely to reject an execution with honest Bob. It is straightforward to generalize the result to noisy quantum communication: In Step 2 in the verification phase of $\mathcal{C}^\alpha(II)$, Alice rejects and aborts if the relative number of errors between x_i and \hat{x}_i for $i \in T$ with $\theta_i = \hat{\theta}_i$ exceeds the error probability ϕ induced by the noise in the quantum communication by some small $\varepsilon' > 0$. By Hoeffding’s inequality [11], this guarantees that honest Alice does not reject honest Bob except with exponentially small probability. Furthermore, proving the security of this “noise-resistant” compiler goes along the exact same lines

⁶ Below, $\theta|_I$ (and similarly $\hat{\theta}|_I$) should be understood as first restricting the m -bit vector θ to \bar{T} , and then restricting the resulting n -bit vector $\theta|_{\bar{T}}$ to I : $\theta|_I := (\theta|_{\bar{T}})|_I$.

as for the original compiler. The only difference is that when applying Corollary 4.4, the parameter err has to be chosen as $err = \phi + \varepsilon'$, so that (1) holds for $\beta = h(err + \varepsilon) = h(\phi + \varepsilon' + \varepsilon)$ and thus the claim of Theorem 4.2 hold for any $\beta > h(\phi)$ (by choosing $\varepsilon, \varepsilon' > 0$ small enough). This allows us to generalize the results from the Section 5 to the setting of noisy quantum communication.

4.5 Bounded-Quantum-Storage Security

In this section we show that our compiler preserves security in the bounded-quantum-storage model (BQSM). In this model, one of the players (Bob in our case) is assumed be able to store only a limited number of qubits beyond a certain point in the protocol. BQSM-secure OT and identification protocols are known [4,7], but they can be efficiently broken if the memory bound does not hold. Therefore, by the theorem below, applying the compiler produces protocols with better security, namely the adversary needs large quantum storage *and* large computing power to succeed.

Consider a BB84-type protocol Π , and for a constant $0 < \gamma < 1$, let $\mathfrak{B}_{\text{BQSM}}^\gamma(\Pi)$ be the set of dishonest players B' that store only γn qubits after a certain point in Π , where n is the number of qubits sent initially. Protocol Π is said to be unconditionally secure against γ -BQSM Bob, if it satisfies Definition 3.3 with the restriction that the quantification is over all dishonest $B' \in \mathfrak{B}_{\text{BQSM}}^\gamma(\Pi)$.

Theorem 4.5. *If Π is unconditionally secure against γ -BQSM Bob, then $\mathcal{C}^\alpha(\Pi)$ (for an $0 < \alpha < 1$) is unconditionally secure against $\gamma(1-\alpha)$ -BQSM Bob.*

Proof. Exactly as in the proof of Theorem 4.2, given dishonest Bob B' attacking $\mathcal{C}^\alpha(\Pi)$, we construct dishonest Bob B'_0 attacking the original protocol Π . The only difference here is that we let B'_0 generate the CRS “correctly” as pkH sampled according to \mathcal{G}_H . It follows by construction of B'_0 that $\text{out}_{A,B'}^{\mathcal{C}^\alpha(\Pi)} = \text{out}_{A_0,B'_0}^\Pi$. Also, it follows by construction of B'_0 that if $B' \in \mathfrak{B}_{\text{BQSM}}^{\gamma(1-\alpha)}(\mathcal{C}^\alpha(\Pi))$ then $B'_0 \in \mathfrak{B}_{\text{BQSM}}^\gamma(\Pi)$, since B'_0 requires the same amount of quantum storage as B' but communicates an α -fraction fewer qubits. It thus follows that there exists \hat{B}' such that $\text{out}_{A_0,B'_0}^\Pi \stackrel{s}{\approx} \text{out}_{\hat{A},\hat{B}'}^{\mathcal{F}}$. This proves the claim. \square

5 Applications

5.1 Oblivious Transfer

We discuss a protocol that securely implements one-out-of-two oblivious transfer of strings of length ℓ (i.e. 1-2 OT $^\ell$). In 1-2 OT $^\ell$, the sender A sends two l -bit strings s_0 and s_1 to the receiver B . B can choose which string to receive (s_k) but does not learn anything about the other one (s_{1-k}). On the other hand, A does not learn B 's choice bit k . The protocol is almost identical to the 1-2 OT 1 introduced in [1], but uses hash functions instead of parity values to mask the inputs s_0 and s_1 . The resulting scheme, called 1-2 QOT $^\ell$, is presented in Figure 5, where \mathcal{F} denotes a suitable family of universal hash functions with range $\{0, 1\}^\ell$ (as specified in [4]). We assume that $\ell = \lfloor \lambda n \rfloor$ for some constant $\lambda > 0$.

PROTOCOL 1-2 QOT^ℓ :

Preparation: A chooses $x \in_R \{0, 1\}^n$ and $\theta \in_R \{+, \times\}^n$ and sends $|x\rangle_\theta$ to B, and B chooses $\hat{\theta} \in_R \{0, 1\}^n$ and obtains $\hat{x} \in \{0, 1\}^n$ by measuring $|x\rangle_\theta$ in basis $\hat{\theta}$.

Post-processing: 1. A sends θ to B.

2. B partitions all positions $1 \leq i \leq n$ in two subsets according to his choice bit $k \in \{0, 1\}$: the “good” subset $I_k := \{i : \theta_i = \hat{\theta}_i\}$ and the “bad” subset $I_{1-k} := \{i : \theta_i \neq \hat{\theta}_i\}$. B sends (I_0, I_1) to A.
3. A sends descriptions of $f_0, f_1 \in_R \mathcal{F}$ together with $m_0 := s_0 \oplus f_0(x|_{I_0})$ and $m_1 := s_1 \oplus f_1(x|_{I_1})$.
4. B computes $s_k = m_k \oplus f_k(\hat{x}|_{I_k})$.

Fig. 5. Protocol for String OT

Theorem 5.1. *Protocol 1-2 QOT^ℓ is unconditionally secure against β -benign Bob for any $\beta < \frac{1}{8} - \frac{\lambda}{2}$.*

The proof and precise definition of OT security is given in [3]. By combining Theorem 5.1 with Theorem 4.2, and the results of [4] (realizing that the same analysis also applies to 1-2 QOT^ℓ) with Theorem 4.5, we obtain the following hybrid-security result.

Corollary 5.2. *Let $0 < \alpha < 1$ and $\lambda < \frac{1}{8}$. Then protocol $C^\alpha(1-2 \text{QOT}^\ell)$ is computationally secure against dishonest Bob and unconditionally secure against $\gamma(1-\alpha)$ -BQSM Bob with $\gamma < \frac{1}{4} - 2\lambda$.*

5.2 Password-Based Identification

We want to apply our compiler to the quantum password-based identification scheme from [7]. Such an identification scheme allows a user A to identify herself to server B by means of a common (possibly non-uniform and low-entropy) password $w \in \mathcal{W}$, such that dishonest A' cannot delude honest server B with probability better than trying to guess the password, and dishonest B' learns no information on A's password beyond trying to guessing it and learn whether the guess is correct or not.

In [7], using quantum-information-theoretic security definitions, the proposed identification scheme was proven to be unconditionally secure against arbitrary dishonest Alice and against quantum-memory-bounded dishonest Bob. In [10] it was then shown that these security definitions imply simulation-based security as considered here, with respect to the functionality \mathcal{F}_{ID} given in Figure 6.⁷

⁷ Actually, the definition and proof from [7] guarantees security only for a slightly weaker functionality, which gives some unfair advantage to dishonest A' in case she guesses the password correctly; however, as discussed in [10], the protocol from [7] does implement functionality \mathcal{F}_{ID} .

Functionality \mathcal{F}_{ID} : Upon receiving $w_A, w_B \in \mathcal{W}$ from user Alice and from server Bob, respectively, \mathcal{F}_{ID} outputs the bit $y := (w_A \stackrel{?}{=} w_B)$ to Bob. In case Alice is dishonest, she may choose $w_A = \perp$ (where $\perp \notin \mathcal{W}$). For any choice of w_A the bit y is also output to dishonest Alice.

Fig. 6. The Ideal Password-Based Identification Functionality

We cannot directly apply our compiler to the identification scheme as given in [7], since it is *not* a BB84-type protocol. The protocol does start with a preparation phase in which Alice sends BB84 qubits to Bob, but Bob does not measure them in a random basis but in a basis determined by his password $w_B \in \mathcal{W}$; specifically, Bob uses as basis the encoding $\mathbf{c}(w_B)$ of w_B with respect to a code $\mathbf{c} : \mathcal{W} \rightarrow \{+, \times\}^n$ with “large” minimal distance. However, it is easy to transform the original protocol from [7] into a BB84-type protocol without affecting security: We simply let Bob apply a *random shift* κ to the code, which Bob only announces to Alice in the post-processing phase, and then Alice and Bob complete the protocol with the shifted code. The resulting protocol QID is described in Figure 7, where \mathcal{F} and \mathcal{G} are suitable families of (strongly) universal hash functions (we refer to [7] for the exact specifications). It is not hard to see that this modification does not affect security as proven in [7] (and [10]).

PROTOCOL QID :

- Preparation:* A chooses $x \in_R \{0, 1\}^n$ and $\theta \in_R \{+, \times\}^n$ and sends $|x\rangle_\theta$ to B, and B chooses $\hat{\theta} \in_R \{0, 1\}^n$ and obtains $\hat{x} \in \{0, 1\}^n$ by measuring $|x\rangle_\theta$ in basis $\hat{\theta}$.
- Post-processing:*
1. B computes a string $\kappa \in \{+, \times\}^n$ such that $\hat{\theta} = \mathbf{c}(w) \oplus \kappa$ (we think of $+$ as 0 and \times as 1 so that \oplus makes sense). He sends κ to A and we define $\mathbf{c}'(w) := \mathbf{c}(w) \oplus \kappa$.
 2. A sends θ and $f \in_R \mathcal{F}$ to B. Both compute $I_w := \{i : \theta_i = \mathbf{c}'(w)_i\}$.
 3. B sends $g \in_R \mathcal{G}$.
 4. A sends $z := f(x|_{I_w}) \oplus g(w)$ to B.
 5. B accepts if and only if $z = f(\hat{x}|_{I_w}) \oplus g(w)$.

Fig. 7. Protocol for Secure Password-Based Identification

Theorem 5.3. *If the code $\mathbf{c} : \mathcal{W} \rightarrow \{+, \times\}^n$ can correct at least δn errors in polynomial-time for a constant δ , then protocol QID is unconditionally secure against β -benign Bob for any $\beta < \frac{\delta}{4}$.*

Proof. For any given benign Bob B' , we construct \hat{B}' as follows. \hat{B}' runs locally a copy of B' and simulates Alice’s actions by running A faithfully except for the following modifications. After the preparation phase, \hat{B}' gets $\hat{\theta}$ and κ from B' and

attempts to decode $\hat{\theta} \oplus \kappa$. If this succeeds, it computes w' such that $c(w')$ is the decoded codeword. Otherwise an arbitrary w' is chosen. Then, \hat{B}' submits w' as Bob's input w_B to \mathcal{F}_{ID} and receives output $y \in \{0, 1\}$. If $y = 1$ then \hat{B}' faithfully completes A's simulation using w' as w ; else, \hat{B}' completes the simulation by using a random z' instead of z . In the end, \hat{B}' outputs whatever B' outputs.

We need to show that the state output by \hat{B}' (respectively B') above is statistically close to the state output by B' when executing QID with real A. Note that if $w' = w_A$, then the simulation of A is perfect and thus the two states are equal. If $w' \neq w_A$ then the simulation is not perfect: the real A would use $z = f(x|_{I_{w_A}}) \oplus g(w_A)$ instead of random z' . It thus suffices to argue that $f(x|_{I_w})$ is statistically close to random and independent of the view of B' for any fixed $w \neq w'$. Note that this is also what had to be proven in [7], but under a different assumption, namely that B' has bounded quantum memory, rather than that he is benign; nevertheless, we can recycle part of the proof.

Recall from the definition of a benign Bob that the common state after the preparation phase is statistically close to a state for which it is guaranteed that $H_\infty(X|_I) \geq d_H(\theta|_I, \hat{\theta}|_I) - \beta n$ for any $I \subseteq \{1, \dots, n\}$, and $H_0(\rho_{ER}) \leq \beta n$. By the closeness of these two states, switching from the real state to the "ideal" state (which satisfies these bounds) has only a negligible effect on the state output by \hat{B}' ; thus, we may assume these bounds to hold.

Now, if decoding of $\hat{\theta} \oplus \kappa$ succeeded, it is at Hamming distance at most δn from $c(w')$. Since the distance from here to the (distinct) codeword $c(w)$ is greater than $2\delta n$, we see that $\hat{\theta} \oplus \kappa$ is at least δn away from $c(w)$. The same is true if decoding failed, since then $\hat{\theta} \oplus \kappa$ is at least δn away from any codeword. It follows that $c'(w) = c(w) \oplus \kappa$ has Hamming distance at least δn from $\hat{\theta}$. Furthermore, for arbitrary $\varepsilon > 0$ and except with negligible probability, the Hamming distance between $\theta|_{I_w} = c'(w)|_{I_w}$ and $\hat{\theta}|_{I_w}$ is at least essentially $(\delta/2 - \varepsilon)n$. Therefore, we can conclude that $H_\infty(X|_{I_w}) \geq (\delta/2 - \varepsilon - \beta)n$ and $H_0(\rho_{ER}) \leq \beta n$. But now, if such bounds hold such that $H_\infty(X|_{I_w}) - H_0(\rho_{ER})$ is positive and linear in n , which is the case here by the choice of parameters, then we can step into the proof from [7] and conclude by privacy amplification [17] that z is close to random and independent of E . This finishes the proof. \square

By combining Theorem 5.3 with Theorem 4.2, and the results of [7] with Theorem 4.5, we obtain the following hybrid-security result.

Corollary 5.4. *Let $0 < \alpha < 1$ and $|\mathcal{W}| \leq 2^{\nu n}$. If the code $c : \mathcal{W} \rightarrow \{+, \times\}^n$ can correct δn errors for a constant $\delta > 0$ in polynomial-time, then protocol $\mathcal{C}^\alpha(\text{QID})$ is computationally secure against dishonest Bob and unconditionally secure against $\gamma(1-\alpha)$ -BQSM Bob with $\gamma < \frac{\delta}{2} - \nu$.*

Families of codes as required in these results, correcting a constant fraction of errors efficiently and with constant information rate are indeed known, see [18].

In the full version [3], we discuss how to obtain hybrid security against *man-in-the-middle* attacks by means of incorporating the techniques used in [7] to obtain security in the BQSM against such attacks.

6 Doing without a Common Reference String

We can get rid of the CRS assumption by instead generating a reference string from scratch using a coin-flip protocol. In [9], such a coin-flip protocol is described and proved secure against quantum adversaries using Watrous' quantum rewinding method [19]. Note that for our compiler, we want the CRS to be an unconditionally hiding public key, and when using Regev's cryptosystem, a uniformly random string (as output by the coin-flip) does indeed determine such a key, except with negligible probability.

References

1. Bennett, C.H., Brassard, G., Crépeau, C., Skubiszewska, M.-H.: Practical quantum oblivious transfer. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 351–366. Springer, Heidelberg (1991)
2. Crépeau, C., Dumais, P., Mayers, D., Salvail, L.: Computational collapse of quantum state with application to oblivious transfer. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 374–393. Springer, Heidelberg (2004)
3. Damgård, I.B., Fehr, S., Lunemann, C., Salvail, L., Schaffner, C.: Improving the security of quantum protocols (2009), <http://arxiv.org/abs/0902.3918>
4. Damgård, I.B., Fehr, S., Renner, R., Salvail, L., Schaffner, C.: A tight high-order entropic quantum uncertainty relation with applications. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 360–378. Springer, Heidelberg (2007)
5. Damgård, I.B., Fehr, S., Salvail, L.: Zero-knowledge proofs and string commitments withstanding quantum attacks. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 254–272. Springer, Heidelberg (2004)
6. Damgård, I.B., Fehr, S., Salvail, L., Schaffner, C.: Cryptography in the bounded quantum-storage model. In: 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 449–458 (2005), <http://arxiv.org/abs/quant-ph/0508222v2>
7. Damgård, I.B., Fehr, S., Salvail, L., Schaffner, C.: Secure identification and QKD in the bounded-quantum-storage model. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 342–359. Springer, Heidelberg (2007)
8. Damgård, I.B., Fehr, S., Salvail, L., Schaffner, C.: Cryptography in the bounded-quantum-storage model. *SIAM Journal on Computing* 37(6), 1865–1890 (2008)
9. Damgård, I.B., Lunemann, C.: Quantum-secure coin-flipping and applications (2009), <http://arxiv.org/abs/0903.3118>
10. Fehr, S., Schaffner, C.: Composing quantum protocols in a classical environment. In: Theory of Cryptography Conference (TCC). LNCS, vol. 5444, pp. 350–367. Springer, Heidelberg (2009)
11. Hoeffding, W.: Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301), 13–30 (1963)
12. Kol, G., Naor, M.: Games for exchanging information. In: TCC 2008. LNCS, vol. 4948, pp. 423–432. Springer, Heidelberg (2008)
13. Lo, H.-K.: Insecurity of quantum secure computations. *Physical Review A* 56(2), 1154–1162 (1997)
14. Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge (2000)

15. Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 554–571. Springer, Heidelberg (2008)
16. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: 37th Annual ACM Symposium on Theory of Computing (STOC), pp. 84–93 (2005)
17. Renner, R., König, R.: Universally composable privacy amplification against quantum adversaries. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 407–425. Springer, Heidelberg (2005)
18. Sipser, M., Spielman, D.A.: Expander codes. IEEE Transactions on Information Theory 42(6), 1710–1722 (1996)
19. Watrous, J.: Zero-knowledge against quantum attacks. In: 38th Annual ACM Symposium on Theory of Computing (STOC), pp. 296–305 (2006), <http://www.cs.uwaterloo.ca/~watrous/papers.html>
20. Wehner, S., Schaffner, C., Terhal, B.M.: Cryptography from noisy storage. Physical Review Letters 100(22), 220–502 (2008)
21. Yao, A.C.-C.: Security of quantum protocols against coherent measurements. In: 27th Annual ACM Symposium on the Theory of Computing (STOC), pp. 67–75 (1995)