Cryptography In the Bounded Quantum-Storage Model[*]

Ivan B. Damgård[†‡] Serge Fehr[§] Louis Salvail[†‡¶]
Christian Schaffner[†||]

August 29, 2005

Abstract

We initiate the study of two-party cryptographic primitives with unconditional security, assuming that the adversary's *quantum* memory is of bounded size. We show that oblivious transfer and bit commitment can be implemented in this model using protocols where honest parties need no quantum memory, whereas an adversarial player needs quantum memory of size at least $n/2$ in order to break the protocol, where n is the number of qubits transmitted. This is in sharp contrast to the classical bounded-memory model, where we can only tolerate adversaries with memory of size quadratic in honest players' memory size. Our protocols are efficient, non-interactive and can be implemented using today's technology. On the technical side, a new entropic uncertainty relation involving min-entropy is established.

1 Introduction

It is well known that non-trivial 2-party cryptographic primitives cannot be securely implemented if only error-free communication is available and there is no limitation assumed on the computing power and memory of the players. Fundamental examples of such primitives are bit commitment (BC) and oblivious transfer (OT). In BC, a committer C commits himself to a choice of a bit b by exchanging information with a verifier V. We want that V does not learn b (we say the commitment is hiding), yet C can later chose to reveal b in a convincing way, i.e., only the value fixed at commitment time will be accepted by V (we say the commitment is binding). In (Rabin) OT, a sender S sends a bit b to a receiver R by executing some protocol in such a way that R receives b with probability $\frac{1}{2}$ and nothing with probability $\frac{1}{2}$, yet S does not learn what was received.

[*] This is the full version of a paper published at FOCS 2005 [8].

[†] Basic Research in Computer Science (BRICS), funded by the Danish National Research Foundation, Department of Computer Science, University of Århus, {ivan|salvail|chris}@brics.dk.

[‡] FICS, Foundations in Cryptography and Security, funded by the Danish Natural Sciences Research Council.

[§] CWI, Amsterdam, Netherlands, fehr@cwi.nl

[¶] Supported in part by the European project PROSECCO.

[||] Supported by the European project SECOQC.

Informally, BC is not possible with unconditional security since hiding means that when 0 is committed, exactly the same information exchange could have happened when committing to a 1. Hence, even if 0 was actually committed to, C could always compute a complete view of the protocol consistent with having committed to 1, and pretend that this was what he had in mind originally. A similar type of argument shows that OT is also impossible in this setting.

One might hope that allowing the protocol to make use of quantum communication would make a difference. Here, information is stored in qubits, i.e., in the state of two-level quantum mechanical systems, such as the polarization state of a single photon. It is well known that quantum information behaves in a way that is fundamentally different from classical information, enabling, for instance, unconditionally secure key exchange between two honest players. However, in the case of two mutually distrusting parties, we are not so fortunate: even with quantum communication, unconditionally secure BC and OT remain impossible [15, 17].

There are, however, several scenarios where these impossibility results do not apply, namely:

- if the computing power of players is bounded,
- if the communication is noisy,
- if the adversary is under some physical limitation, e.g., the size of the available memory is bounded.

The first scenario is the basis of many well known solutions based on plausible but unproven complexity assumptions, such as hardness of factoring or discrete logarithms. The second scenario has been used to construct both BC and OT protocols in various models for the noise [5, 6, 9]. The third scenario is our focus here. In this model, OT and BC can be done using classical communication assuming, however, quite restrictive bounds on the adversary's memory size [2, 10], namely it can be at most quadratic in the memory size of honest players. Such an assumption is on the edge of being realistic, it would clearly be more satisfactory to have a larger separation between the memory size of honest players and that of the adversary. However, this was shown to be impossible [13].

In this paper, we study for the first time what happens if instead we consider protocols where quantum communication is used and we place a bound on the adversary's *quantum* memory size. There are two reasons why this may be a good idea: first, if we do not bound the classical memory size, we avoid the impossibility result of [13]. Second, the adversary's goal typically is to obtain a certain piece of classical information, however, converting quantum information to classical by measuring may irreversibly destroy information, and we may be able to arrange it such that the adversary cannot afford to loose information this way, while honest players can.

It turns out that this is indeed possible: we present protocols for both BC and OT in which n qubits are transmitted, where honest players need *no quantum memory*, but where the adversary must store at least $n/2$ qubits to break the protocol. We emphasize that no bounds are assumed on the adversary's computing power, nor on his classical memory. This is clearly much more promising than the classical case, not only from a theoretical point of view, but also in practice: while sending qubits and measuring them immediately as they arrive is well within reach of current technology, storing even a single qubit for more than fraction of a second is a formidable technological challenge. Furthermore, we show that our protocols also work in a non-ideal setting where we allow the quantum source to be imperfect and the quantum communication to be noisy.

Our protocols are non-interactive, only one party sends information when doing OT, commitment or opening. Furthermore, the commitment protocol has the interesting property that the only message is sent to the committer, i.e., it is possible to commit while only *receiving* information. Such a scheme clearly does not exist without a bound on the committer's memory, even under computational assumptions and using quantum communication: a corrupt committer could always store (possibly quantumly) all the information sent, until opening time, and only then follow the honest committer's algorithm to figure out what should be sent to convincingly open a 0 or a 1. Note that in the classical bounded-storage model, it is known how to do time-stamping that is non-interactive in our sense: a player can time-stamp a document while only receiving information [18]. However, no reasonable BC or protocol that time-stamps a bit exist in this model. It is straightforward to see that any such protocol can be broken by an adversary with classical memory of size twice that of an honest player, while our protocol requires no memory for the honest players and remains secure against any adversary not able to store more than half the size of the quantum transmission.

We also note that it has been shown earlier that BC is possible using quantum communication, assuming a different type of physical limitation, namely a bound on the size of coherent measurement that can be implemented [20]. This limitation is incomparable to ours: it does not limit the total size of the memory, instead it limits the number of bits that can be simultaneously operated on to produce a classical result. Our adversary has a limit on the total memory size, but can measure all of it coherently. The protocol from [20] is interactive, and requires a bound on the maximal measurement size that is sublinear in n.

On the technical side, we use the quantum privacy amplification result by Renner and König [19] together with a proof technique by Shor and Preskill [21] where we purify the actions of honest players. This makes no difference from the adversary's point of view, but makes proofs go through more easily. We combine this with a new technical result that may be seen as a new type of uncertainty relation involving min-entropy (Theorem 3.7 and Corollary 3.8).

2 Preliminaries

2.1 Notation and Quantum Stuff

For a set $I = \{i_1, i_2, \ldots, i_\ell\} \subseteq \{1, \ldots, n\}$ and a n-bit string $x \in \{0,1\}^n$, we define $x|_I := x_{i_1} x_{i_2} \cdots x_{i_\ell}$. For $x \in \{0,1\}^n$, we write $B^{\delta n}(x)$ for the set of all n-bit strings at Hamming distance at most δn from x. Note that the number of elements in $B^{\delta n}(x)$ is the same for all x, we denote it by $B^{\delta n} := |B^{\delta n}(x)|$. For $x, y \in \{0,1\}^n$, $x \cdot y \in \{0,1\}$ denotes the (standard) in-product of x and y. For a probability distribution Q over n-bit strings and a set $L \subseteq \{0,1\}^n$, we abbreviate the (overall) probability of L with $Q(L) := \sum_{x \in L} Q(x)$. All logarithms in this paper are to base two. We denote by $h(p)$ the binary entropy function $h(p) := -\big(p \cdot \log p + (1-p) \cdot \log(1-p)\big)$. We denote by $negl(n)$ any function of n smaller than any polynomial provided n is sufficiently large.

The pair $\{|0\rangle, |1\rangle\}$ denotes the computational or rectilinear or "+" basis for the 2-dimensional complex Hilbert space $\mathbb{C}^2$. The diagonal or "×" basis is defined as $\{|0\rangle_\times, |1\rangle_\times\}$ where $|0\rangle_\times = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $|1\rangle_\times = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. Measuring a qubit in the $+$-basis (resp. $\times$-basis) means applying the measurement described by projectors $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ (resp. projectors $|0\rangle_\times\langle 0|_\times$ and $|1\rangle_\times\langle 1|_\times$). When the context requires it, we write $|0\rangle_+$ and $|1\rangle_+$ instead of $|0\rangle$ respectively $|1\rangle$; and for any $x \in \{0,1\}^n$ and $r \in \{+, \times\}$, we write $|x\rangle_r = \bigotimes_{i=1}^n |x_i\rangle_r$. If we want to choose the $+$ or $\times$-basis according to the bit $b \in \{0,1\}$, we write $\{+, \times\}_{[b]}$.

2.2 Quantum Probability Theory

As basis for the security definitions and proofs of our protocols, we are using the formalism introduced in [19], which we briefly summarize here. A *random state* $\boldsymbol{\rho}$ is a random variable, with distribution $P_{\boldsymbol{\rho}}$, whose range is the set of density operators of a fixed Hilbert space. The view of an observer (which is ignorant of the value of $\boldsymbol{\rho}$) is given by the quantum system described by the density operator $[\boldsymbol{\rho}] := \sum_\rho P_{\boldsymbol{\rho}}(\rho)\rho$. In general, for any event $\mathcal{E}$, we define $[\boldsymbol{\rho}|\mathcal{E}] := \sum_\rho P_{\boldsymbol{\rho}|\mathcal{E}}(\rho)\rho$. If $\boldsymbol{\rho}$ is dependent on some classical random variable X, with joint distribution $P_{X\boldsymbol{\rho}}$, we also write ρ_x instead of $[\boldsymbol{\rho}|X = x]$. Note that ρ_x is a density operator (for any fixed x) whereas ρ_X is again a random state. The overall quantum system is then given by $[\{X\} \otimes \boldsymbol{\rho}] = \sum_x P_X(x)\,\{x\} \otimes \rho_x$, where $\{x\} := |x\rangle\langle x|$ is the *state representation* of x and $\{X\}$ the corresponding random state. Obviously, $[\{X\} \otimes \boldsymbol{\rho}] = [\{X\}] \otimes [\boldsymbol{\rho}]$ if and only if ρ_X is independent of X, where the latter in particular implies that no information on X can be learned by observing only $\boldsymbol{\rho}$. Furthermore, if $[\{X\} \otimes \boldsymbol{\rho}]$ and $[\{X\}] \otimes [\boldsymbol{\rho}]$ are ε-close in terms of their trace distance $\delta(\rho, \sigma) = \frac{1}{2}\operatorname{tr}(|\rho - \sigma|)$, then the real system $[\{X\} \otimes \boldsymbol{\rho}]$ "behaves" as the ideal system $[\{X\}] \otimes [\boldsymbol{\rho}]$ except with probability ε [19] in that for any evolution of the system no observer can distinguish the real from the ideal one with advantage greater than $\varepsilon/2$ (or ε, depending on the exact definition of advantage). By slight abuse of notation, we usually simply write X instead of $\{X\}$. Henceforth, we use UNIF to denote a random variable with range $\{0,1\}$, uniformly distributed and independent of anything else.

When reviewing the privacy amplification theorem from [19], we briefly address the generalization of the classical *Rényi entropy* $H_\alpha(X)$ of order α of a random variable X to the Rényi entropy $S_\alpha(\rho)$ of order α of a density operator ρ. Otherwise, though, we are only using the classical Rényi entropy of order ∞, commonly known as the *min-entropy* $H_\infty(X) = -\log \max_x P_X(x)$.

2.3 Privacy Amplification

In this paper, we only use privacy amplification with one-bit output. A class H_n of hashing functions from $\{0,1\}^n$ to $\{0,1\}$ is called *two-universal* if for any pair $x, y \in \{0,1\}^n$ with $x \neq y$

$$\big|\{f \in \mathrm{H}_n : f(x) = f(y)\}\big| \leq \frac{|\mathrm{H}_n|}{2}.$$

Several two-universal classes of hashing functions are such that evaluating and picking a function uniformly and at random in H_n can be done efficiently [3, 22].

Theorem 2.1 **([19]).** *Let X be distributed over $\{0,1\}^n$, and let $\boldsymbol{\rho}$ be a random state of q qubits*[1]*. Let F be the random variable corresponding to the random choice (with uniform distribution and independent from X and $\boldsymbol{\rho}$) of a member of a two-universal class of hashing functions H_n. Then*

$$\begin{aligned}\delta([F(X) \otimes F \otimes \boldsymbol{\rho}], [\text{UNIF}] \otimes [F \otimes \boldsymbol{\rho}]) &\leq \frac{1}{2} 2^{-\frac{1}{2}(S_2([\{X\}\otimes\boldsymbol{\rho}])-S_0([\boldsymbol{\rho}])-1)} \\ &\leq \frac{1}{2} 2^{-\frac{1}{2}(H_\infty(X)-q-1)}. \qquad (1)\end{aligned}$$

The first inequality is the original theorem from [19], and (1) follows by observing that $S_2([\{X\} \otimes \boldsymbol{\rho}]) \geq H_2(X) \geq H_\infty(X)$. In this paper, we only use this weaker version of the theorem.

Note that if the rightmost term of (1) is negligible, i.e. say smaller than $2^{-\varepsilon n}$, then this situation is $2^{-\varepsilon n}$-close to the ideal situation where $F(X)$ is perfectly uniform and independent of $\boldsymbol{\rho}$ and F. In particular, the situations $F(X) = 0$ and $F(X) = 1$ are statistically indistinguishable given $\boldsymbol{\rho}$ and F [14].

The following lemma is a direct consequence of Theorem 2.1. In Section 4, this lemma will be useful for proving the binding condition of our commitment scheme. Recall that for $X \in \{0,1\}^n$, $B^{\delta n}(X)$ denotes the set of all n-bit strings at Hamming distance at most δn from X and $B^{\delta n} := |B^{\delta n}(X)|$ is the number of such strings.

Lemma 2.2. *Let X be distributed over $\{0,1\}^n$, let $\boldsymbol{\rho}$ be a random state of q qubits and let $\hat{X}$ be a guess for X given $\boldsymbol{\rho}$. Then, for all $\delta < \frac{1}{2}$ it holds that*

$$\Pr\left[\hat{X} \in B^{\delta n}(X)\right] \leq 2^{-\frac{1}{2}(H_\infty(X)-q-1)+\log(B^{\delta n})}.$$

In other words, given a quantum memory of q qubits arbitrarily correlated with a classical random variable X, the probability to find $\hat{X}$ at Hamming distance at most δn from X where $nh(\delta) < \frac{1}{2}(H_\infty(X) - q)$ is negligible.

Proof: Here is a strategy to try to bias $F(X)$ when given $\hat{X}$ and $F \in_R \mathrm{H}_n$: Sample $X' \in_R B^{\delta n}(\hat{X})$ and output $F(X')$. Note that, using p_{succ} as a short hand for the probability $\Pr\left[\hat{X} \in B^{\delta n}(X)\right]$ to be bounded,

$$\begin{aligned}\Pr\left[F(X') = F(X)\right] &= \frac{p_{\text{succ}}}{B^{\delta n}} + \left(1 - \frac{p_{\text{succ}}}{B^{\delta n}}\right)\frac{1}{2} \\ &= \frac{1}{2} + \frac{p_{\text{succ}}}{2 \cdot B^{\delta n}},\end{aligned}$$

where the first equality follows from the fact that if $X' \neq X$ then, as H_n is two-universal, $\Pr\left[F(X) = F(X')\right] = \frac{1}{2}$. Since the probability of correctly guessing a binary $F(X)$ given F and $\boldsymbol{\rho}$ is always upper bounded by $\frac{1}{2}+\delta([F(X)\otimes F\otimes\boldsymbol{\rho}], [\text{UNIF}]\otimes[F\otimes\boldsymbol{\rho}])$, in combination with Theorem 2.1 the above results in

$$\frac{1}{2} + \frac{p_{\text{succ}}}{2 \cdot B^{\delta n}} \leq \frac{1}{2} + \frac{1}{2} 2^{-\frac{1}{2}(H_\infty(X)-q-1)}$$

and the claim follows immediately. □

[1] Remember that $\boldsymbol{\rho}$ can be correlated with X in an arbitrary way. In particular, we can think of $\boldsymbol{\rho}$ as an attempt to store the n-bit string X in q qubits.

3 Rabin Oblivious Transfer

3.1 The Definition

A protocol for Rabin Oblivious Transfer (ROT) between sender Alice and receiver Bob allows for Alice to send a bit b through an erasure channel to Bob. Each transmission delivers b or an erasure with probability $\frac{1}{2}$. Intuitively, a protocol for ROT is secure if

- sender Alice gets no information on whether b was received or not, no matter what she does, and
- receiver Bob gets no information about b with probability at least $\frac{1}{2}$, no matter what he does.

In this paper, we are considering quantum protocols for ROT. This means that while in- and outputs of the honest senders are classical, described by random variables, the protocol may contain quantum computation and quantum communication, and the view of a dishonest player is quantum, and is thus described by a random state.

Any such (two-party) protocol is specified by a family $\{(\mathsf{S}_n, \mathsf{R}_n)\}_{n>0}$ of pairs of interactive quantum circuits (i.e. interacting through a quantum channel). Each pair is indexed by a security parameter $n > 0$, where S_n and R_n denote the circuits for sender Alice and receiver Bob, respectively. In order to simplify the notation, we often omit the index n, leaving the dependency on it implicit.

For the formal definition of the security requirements of a ROT protocol, let us fix the following notation. Let B denote the binary random variable describing S's input bit b, and let A and B' denote the binary random variables describing R's two output bits, where the meaning is that A indicates whether the bit was received or not. Furthermore, for a dishonest sender $\tilde{\mathsf{S}}$ (respecively $\tilde{\mathsf{R}}$) let $\boldsymbol{\rho}_{\tilde{\mathsf{S}}}$ ($\boldsymbol{\rho}_{\tilde{\mathsf{R}}}$) denote the random state describing $\tilde{\mathsf{S}}$'s ($\tilde{\mathsf{R}}$'s) view of the protocol. Note that for a fixed candidate protocol for ROT, and for a fixed input distribution P_B, depending on whether we consider two honest S and R, a dishonest $\tilde{\mathsf{S}}$ and an honest R, or an honest S and a dishonest $\tilde{\mathsf{R}}$, the corresponding joint distribution $P_{BAB'}$, $P_{\boldsymbol{\rho}_{\tilde{\mathsf{S}}}AB'}$ respectively $P_{B\boldsymbol{\rho}_{\tilde{\mathsf{R}}}}$ is uniquely determined.

Definition 3.1. *A two-party (quantum) protocol* (S, R) *is a* **(statistically) secure ROT** *if the following holds.*

Correctness: *For honest* S *and* R

$$\Pr\left[B = B' | A = 1\right] \geq 1 - negl(n)\,.$$

Privacy: *For any* $\tilde{\mathsf{S}}$

$$\delta([A \otimes \boldsymbol{\rho}_{\tilde{\mathsf{S}}}], [\text{UNIF}] \otimes [\boldsymbol{\rho}_{\tilde{\mathsf{S}}}]) \leq negl(n)\,.$$

Obliviousness: *For any* $\tilde{\mathsf{R}}$ *there exists an event* $\mathcal{E}$ *with* $P[\mathcal{E}] \geq \frac{1}{2} - negl(n)$ *such that*

$$\delta([B \otimes \boldsymbol{\rho}_{\tilde{\mathsf{R}}} | \mathcal{E}], [B] \otimes [\boldsymbol{\rho}_{\tilde{\mathsf{R}}} | \mathcal{E}]) \leq negl(n)\,.$$

If any of the above trace distances equals 0, then the corresponding property is said to hold **perfectly**. *If one of the properties only holds with respect to a restricted class $\mathfrak{S}$ of $\tilde{\mathsf{S}}$'s respectively $\mathfrak{R}$ of $\tilde{\mathsf{R}}$'s, then this property is said to hold and the protocol is said to be secure* **against** $\mathfrak{S}$ *respectively* $\mathfrak{R}$.

Privacy requires that the joint quantum state is essentially the same as when A is uniformly distributed and independent of the senders's view, and obliviousness requires that there exists some event which occurs with probability at least $\frac{1}{2}$ (the event that the receiver does not receive the bit) and under which the joint quantum state is essentially the same as when B is distributed (according to P_B) independently of the receiver's view.

3.2 The Protocol

We introduce a quantum protocol for ROT that will be shown perfectly private (against any sender) and statistically oblivious against any quantum memory-bounded receiver.

The protocol is very simple (see Figure 1): S picks $x \in_R \{0,1\}^n$ and sends to R n qubits in state either $|x\rangle_+$ or $|x\rangle_\times$ each chosen with probability $\frac{1}{2}$. R then measures all received qubits either in the rectilinear or in the diagonal basis. With probability $\frac{1}{2}$, R picked the right basis and gets x, while any $\tilde{\mathsf{R}}$ that is forced to measure part of the state (due to a memory bound) can only have full information on x in case the $+$-basis was used *or* in case the $\times$-basis was used (but not in both cases). Privacy amplification using any two-universal class of hashing functions H_n allows to obtain a proper ROT. (In order to avoid aborting, we specify that if a dishonest $\tilde{\mathsf{S}}$ refuses to participate, or sends data in incorrect format, then R samples its output bits a and b' both at random in $\{0,1\}$.)

QOT(b):

1. S picks $x \in_R \{0,1\}^n$, and $r \in_R \{+,\times\}$.
2. S sends $|\psi\rangle := |x\rangle_r$ in basis r to R.
3. R picks $r' \in_R \{+,\times\}$ and measures all qubits of $|\psi\rangle$ in basis r'. Let $x' \in \{0,1\}^n$ be the result.
4. S announces r, $f \in_R \mathrm{H}_n$, and $s := b \oplus f(x)$.
5. R outputs $a := 1$ and $b' := s \oplus f(x')$ if $r' = r$ and else $a := 0$ and $b' := 0$.

Figure 1. Protocol for Rabin QOT

As we shall see in Section 3.5, the security of the QOT protocol against receivers with bounded-size quantum memory holds as long as the bound applies before Step 4 is reached. An equivalent protocol is obtained by purifying the sender's actions. Although QOT is easy to implement, the purified or EPR-based version depicted in Figure 2 is easier to prove secure. A similar approach was taken in the Shor-Preskill proof of security for the BB84 quantum key distribution scheme [21].

EPR-QOT(b):

1. S prepares n EPR pairs each in state $|\Omega\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$.
2. S sends one half of each pair to R and keeps the other halves.
3. R picks $r' \in_R \{+, \times\}$ and measures all received qubits in basis r'. Let $x' \in \{0,1\}^n$ be the result.
4. S picks $r \in_R \{+, \times\}$, and measures all kept qubits in basis r. Let $x \in \{0,1\}^n$ be the outcome. S announces r, $f \in_R \mathrm{H}_n$, and $s := b \oplus f(x)$.
5. R outputs $a := 1$ and $b' := s \oplus f(x')$ if $r' = r$ and else $a := 0$ and $b' := 0$.

Figure 2. Protocol for EPR-based Rabin QOT

Notice that while QOT requires no quantum memory for honest players, quantum memory for S seems to be required in EPR-QOT. The following Lemma shows the strict equivalence between QOT and EPR-QOT.

Lemma 3.2. QOT *is secure if and only if* EPR-QOT *is secure.*

The proof follows easily after observing that S's choices of r and f, together with the measurements all commute with R's actions. Therefore, they can be performed right after Step 1 with no change for R's view. Modifying EPR-QOT that way results in QOT.

Lemma 3.3. EPR-QOT *is perfectly private.*

Proof: It is straightforward to verify that no information about whether R has received the bit is leaked to any sender $\tilde{\mathsf{S}}$, since R does not send anything, i.e. EPR-QOT is non-interactive! □

3.3 Modeling Dishonest Receivers

We model dishonest receivers in EPR-QOT under the assumption that the maximum size of their quantum storage is bounded. These adversaries are only required to have bounded quantum storage when they reach Step 4 in EPR-QOT. Before that, the adversary can store and carry out quantum computations involving any number of qubits. Apart from the restriction on the size of the quantum memory available to the adversary, no other assumption is made. In particular, the adversary is not assumed to be computationally bounded and the size of its classical memory is not restricted.

Definition 3.4. *The set $\mathfrak{R}_\gamma$ denotes all possible quantum dishonest receivers $\{\tilde{\mathsf{R}}_n\}_{n>0}$ in* QOT *or* EPR-QOT *where for each $n > 0$, $\tilde{\mathsf{R}}_n$ has quantum memory of size at most γn when Step 4 is reached.*

In general, the adversary $\tilde{\mathsf{R}}$ is allowed to perform any quantum computation compressing the n qubits received from S into a quantum register M of size at most γn

when Step 4 is reached. More precisely, the compression function is implemented by some unitary transform C acting upon the quantum state received and an ancilla of arbitrary size. The compression is performed by a measurement that we assume in the computational basis without loss of generality. Before starting Step 4, the adversary first applies a unitary transform C:

$$2^{-n/2} \sum_{x\in\{0,1\}^n} |x\rangle \otimes C|x\rangle|0\rangle \mapsto 2^{-n/2} \sum_{x\in\{0,1\}^n} |x\rangle \otimes \sum_y \alpha_{x,y}|\varphi_{x,y}\rangle^M |y\rangle^Y,$$

where for all x, $\sum_y |\alpha_{x,y}|^2 = 1$. Then, a measurement in the computational basis is applied to register Y providing classical outcome y. The result is a quantum state in register M of size γn qubits. Ignoring the value of y to ease the notation, the re-normalized state of the system is now in its most general form when Step 4 is reached:

$$|\psi\rangle = \sum_{x\in\{0,1\}^n} \alpha_x|x\rangle \otimes |\varphi_x\rangle^M,$$

where $\sum_x |\alpha_x|^2 = 1$.

3.4 Uncertainty Relation

We first prove a general uncertainty result and derive from that a corollary that plays the crucial role in the security proof of EPR-QOT. The uncertainty result concerns the situation where the sender holds an arbitrary quantum register of n qubits. He may measure them in either the +- or the ×-basis. We are interested in the distribution of both these measurement results, and we want to claim that they cannot *both* be "very far from uniform". One way to express this is to say that a distribution is very non-uniform if one can identify a subset of outcomes that has much higher probability than for a uniform choice. Intuitively, the theorem below says that such sets cannot be found for both of the sender's measurements.

Theorem 3.5. *Let the density matrix ρ^A describe the state of a n-qubit register A. Let $Q^+(\cdot)$ and $Q^\times(\cdot)$ be the respective distributions of the outcome when register A is measured in the $+$-basis respectively the $\times$-basis. Then, for any two sets $L^+ \subset \{0,1\}^n$ and $L^\times \subset \{0,1\}^n$ it holds that*

$$Q^+(L^+) + Q^\times(L^\times) \leq \left(1 + \sqrt{2^{-n}|L^+||L^\times|}\right)^2.$$

Proof: We can purify register A by adding a register B, such that the state of the composite system is pure. It can then be written as $|\psi\rangle^{AB} = \sum_{x\in\{0,1\}^n} \alpha_x|x\rangle^A|\varphi_x\rangle^B$ for some complex amplitudes α_x and normalised state vectors $|\varphi_x\rangle$.

Clearly, $Q^+(x) = |\alpha_x|^2$. To give a more explicit form of the distribution $Q^\times$, we apply the Hadamard transformation to register A:

$$(H^{\otimes n} \otimes \mathbb{1}^B)|\psi\rangle = \sum_{z\in\{0,1\}^n} |z\rangle \otimes \sum_{x\in\{0,1\}^n} 2^{-\frac{n}{2}}(-1)^{x\cdot z}\alpha_x|\varphi_x\rangle$$

and obtain

$$Q^\times(z) = \left| \sum_{x\in\{0,1\}^n} 2^{-\frac{n}{2}}(-1)^{x\cdot z}\alpha_x|\varphi_x\rangle \right|^2.$$

Let $\overline{L}^+$ denote the complement of L^+ and p its probability $Q^+(\overline{L}^+)$. We can now split the sum in $Q^\times(z)$ in the following way:

$$\begin{aligned}
Q^\times(z) &= \left| \sum_{x\in\{0,1\}^n} 2^{-\frac{n}{2}}(-1)^{x\cdot z}\alpha_x|\varphi_x\rangle \right|^2 \\
&= \left| \sqrt{p}\sum_{x\in\overline{L}^+} 2^{-\frac{n}{2}}(-1)^{x\cdot z}\frac{\alpha_x}{\sqrt{p}}|\varphi_x\rangle + \sum_{x\in L^+} 2^{-\frac{n}{2}}(-1)^{x\cdot z}\alpha_x|\varphi_x\rangle \right|^2 \\
&= \left| \sqrt{p}\cdot\zeta_z|v_z\rangle + \sum_{x\in L^+} 2^{-\frac{n}{2}}(-1)^{x\cdot z}\alpha_x|\varphi_x\rangle \right|^2
\end{aligned}$$

where $|v_z\rangle$ is defined as follows: For the normalised state $|v\rangle := \sum_{x\in\overline{L}^+}\frac{\alpha_x}{\sqrt{p}}|x\rangle|\varphi_x\rangle$, $\zeta_z|v_z\rangle$ is the z-component of the state $H^{\otimes n}|v\rangle = \sum_z \zeta_z|z\rangle\otimes|v_z\rangle$. It therefore holds that $\sum_z|\zeta_z|^2 = 1$.

To upperbound the amplitudes provided by the sum over L^+, we notice that the amplitude is maximized when all unit vectors $|\varphi_x\rangle$ point in the same direction and when $(-1)^{x\cdot z}\alpha_x = |\alpha_x|$. More formally,

$$\begin{aligned}
\left| \sum_{x\in L^+} 2^{-\frac{n}{2}}(-1)^{x\cdot z}\alpha_x|\varphi_x\rangle \right| &\leq 2^{-\frac{n}{2}}\sum_{x\in L^+}|\alpha_x| \\
&\leq 2^{-\frac{n}{2}}\sqrt{|L^+|}\sqrt{\sum_{x\in L^+}|\alpha_x|^2} \qquad (2)\\
&\leq 2^{-\frac{n}{2}}\sqrt{|L^+|},
\end{aligned}$$

where (2) is obtained from the Cauchy-Schwarz inequality. Using ℓ^+ and $\ell^\times$ as shorthands for $|L^+|$ respectively $|L^\times|$, we conclude that

$$\begin{aligned}
Q^\times(L^\times) &= \sum_{z\in L^\times} Q^\times(z) \\
&\leq \sum_{z\in L^\times}\left(|\sqrt{p}\cdot\zeta_z|v_z\rangle| + 2^{-\frac{n}{2}}\sqrt{\ell^+}\right)^2 \\
&\leq p\sum_{z\in L^\times}|\zeta_z|^2 + 2\cdot 2^{-\frac{n}{2}}\sqrt{\ell^+}\sum_{z\in L^\times}|\zeta_z| + \ell^\times\cdot 2^{-n}\ell^+ \\
&\leq p + 2\cdot 2^{-\frac{n}{2}}\sqrt{\ell^+}\sqrt{\ell^\times\sum_{z\in L^\times}|\zeta_z|^2} + 2^{-n}\ell^+\ell^\times \qquad (3)\\
&\leq p + 2\sqrt{2^{-n}\ell^+\ell^\times} + 2^{-n}\ell^+\ell^\times \\
&= 1 - Q^+(L^+) + 2\sqrt{2^{-n}\ell^+\ell^\times} + 2^{-n}\ell^+\ell^\times. \qquad (4)
\end{aligned}$$

Inequality (3) follows again from Cauchy-Schwarz while in (4), we use the definition of p. The claim of the proposition follows after re-arranging the terms. □

This theorem yields a meaningful bound as long as $|L^+| \cdot |L^\times| < (\sqrt{2}-1)^2 \cdot 2^n$, e.g. if L^+ and $L^\times$ both contain less than $2^{n/2}$ elements. If for $r \in \{+, \times\}$, L^r contains only the n-bit string with the maximal probability of Q^r, we obtain as a corollary a slightly weaker version of a known relation (see (9) in [16]).

Corollary 3.6. *Let q_∞^+ and $q_\infty^\times$ be the maximal probabilities of the distributions Q^+ and $Q^\times$ from above. It then holds that $q_\infty^+ \cdot q_\infty^\times \leq \frac{1}{4}(1+c)^4$ where $c = 2^{-n/2}$.*

Theorem 3.5 can be generalised to more than two mutually unbiased bases. We call different sets $\mathcal{B}^0, \mathcal{B}^1, \ldots, \mathcal{B}^N$ of bases of the complex Hilbert space $\mathbb{C}^{2^n}$ *mutually unbiased*, if for all $i \neq j \in \{0, \ldots, N\}$, it holds that

$$\forall |\varphi\rangle \in \mathcal{B}^i \ \forall |\psi\rangle \in \mathcal{B}^j : |\langle \varphi | \psi \rangle|^2 = 2^{-n}.$$

Theorem 3.7. *Let the density matrix ρ^A describe the state of a n-qubit register A and let $\mathcal{B}^0, \mathcal{B}^1, \ldots, \mathcal{B}^N$ be mutually unbiased bases of register A. Let $Q^0(\cdot), Q^1(\cdot), \ldots, Q^N(\cdot)$ be the distributions of the outcome when register A is measured in bases $\mathcal{B}^0, \mathcal{B}^1, \ldots, \mathcal{B}^N$, respectively. Then, for any sets $L^0, L^1, \ldots, L^N \subset \{0,1\}^n$, it holds that*

$$\sum_{i=0}^{N} Q^i(L^i) \leq 1 - \binom{N+1}{2} + \sum_{0 \leq j < k \leq N} \left(1 + \sqrt{2^{-n}|L^j||L^k|}\right)^2.$$

Proof: Like in the proof of Theorem 3.5, we can purify register A by adding a register B. The composite state can then be written as $|\psi\rangle^{AB} = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle^A |\varphi_x\rangle^B$ for some complex amplitudes α_x and normalised state vectors $|\varphi_x\rangle$.

We prove the statement by induction over N: For $N = 1$, by applying an appropriate unitary transform to the whole system, we can assume without loss of generality that $\mathcal{B}^0$ is the standard $+$-basis.

Let us denote by T the matrix of the basis change from $\mathcal{B}^0$ to $\mathcal{B}^1$. As the inner product between states $|\phi\rangle \in \mathcal{B}^0$ and $|\phi'\rangle \in \mathcal{B}^1$ is always $|\langle \phi | \phi' \rangle| = 2^{-n/2}$, it follows that all entries of T are complex numbers of the form $2^{-n/2} \cdot e^{i\lambda}$ for real $\lambda \in \mathbb{R}$.

It is easy to verify that the same proof as for Theorem 3.5 applies after replacing the Hadamard transform $H^{\otimes n}$ on the sender's part by T and using the above observation about the entries of T.

For the induction step from N to $N{+}1$, we define $p := Q^0(\overline{L}^0)$, $|v\rangle := \sum_{x \in \overline{L}^0} \frac{\alpha_x}{\sqrt{p}} |x\rangle |\varphi_x\rangle$, and let $\zeta_z^j |v_z^j\rangle$ be the z-component of the state $|v\rangle$ transformed into basis $\mathcal{B}^j$. As in the proof of Theorem 3.5, using ℓ_i as a short hand for $\left|L^i\right|$, it follows:

$$\begin{aligned}
\sum_{i=1}^{N} Q^i(L^i) &= \sum_{i=1}^{N} \sum_{z \in L^i} Q^i(z) \\
&\leq \sum_{i=1}^{N} \sum_{z \in L^i} \left(\sqrt{p}\,\left|\zeta_z^i |v_z^i\rangle\right| + 2^{-n/2}\sqrt{\ell_0}\right)^2 \\
&\leq p \cdot \sum_{i=1}^{N} \sum_{z \in L^i} |\zeta_z^i|^2 + \sum_{i=1}^{N} \left(2 \cdot \sqrt{2^{-n}\ell_0\ell_i} + 2^{-n}\ell_0\ell_i\right)
\end{aligned}$$

$$\leq p \cdot \sum_{i=1}^{N} P^i(L^i) + \sum_{i=1}^{N} \left(1 - \sqrt{2^{-n}\ell_0\ell_i}\right)^2 - N$$

where the distributions P^i are obtained by measuring register A of the normalised state $|v\rangle$ in the mutually unbiased bases $\mathcal{B}^1, \mathcal{B}^2, \ldots, \mathcal{B}^N$. We apply the induction hypothesis to the sum of $P^i(L^i)$:

$$\begin{aligned}
\sum_{i=1}^{N} Q^i(L^i) &\leq p \cdot \sum_{i=1}^{N} P^i(L^i) + \sum_{i=1}^{N} \left(1 + \sqrt{2^{-n}\ell_0\ell_i}\right)^2 - N \\
&\leq \left[1 - Q^0(L^0)\right] \left[\sum_{1 \leq j < k \leq N} \left(1 + \sqrt{2^{-n}\ell_j\ell_k}\right)^2 + 1 - \binom{N}{2} \right] \\
&\quad + \sum_{i=1}^{N} \left(1 - \sqrt{2^{-n}\ell_0\ell_i}\right)^2 - N \\
&\leq -Q^0(L^0) + 1 - \binom{N+1}{2} + \sum_{0 \leq j < k \leq N} \left(1 + \sqrt{2^{-n}\ell_j\ell_k}\right)^2
\end{aligned}$$

where the last inequality follows by observing that the term in the right bracket is at least 1 and rearranging the terms. This completes the induction step and the proof of the proposition. □

Analogous to Corollary 3.6, we derive an uncertainty relation about the sum of the min-entropies of up to $2^{\frac{n}{4}}$ distributions.

Corollary 3.8. *For an $\varepsilon > 0$, let $0 < N < 2^{(\frac{1}{4}-\varepsilon)n}$. For $i = 0, \ldots, N$, let H_∞^i be the min-entropies of the distributions Q^i from the theorem above. Then,*

$$\sum_{i=0}^{N} H_\infty^i \geq (N+1)\big(\log(N+1) - \mathit{negl}(n)\big).$$

Proof: For $i = 0, \ldots, N$, we denote by q_∞^i the maximal probability of Q^i and let L^i be the set containing only the n-bit string x with this maximal probability q_∞^i. Theorem 3.7 together with the assumption about N assures $\sum_{i=0}^{N} q_\infty^i \leq 1 + \mathit{negl}(n)$. By the inequality of the geometric and arithmetic mean follows:

$$\begin{aligned}
\sum_{i=0}^{N} H_\infty^i &= -\log \prod_{i=0}^{N} q_\infty^i \geq -\log \left(\frac{1 + \mathit{negl}(n)}{N+1}\right)^{N+1} \\
&= (N+1)\big(\log(N+1) - \mathit{negl}(n)\big).
\end{aligned}$$

□

3.5 Security Against Dishonest Receivers

In this section, we show that EPR-QOT is secure against any dishonest receiver having access to a quantum storage device of size strictly smaller than half the number of qubits received at Step 2.

In our setting, we use Theorem 3.5 to lowerbound the overall probability of strings with small probabilities in the following sense. For $0 \leq \gamma + \kappa \leq 1$, define

$$S^+ := \left\{x \in \{0,1\}^n : Q^+(x) \leq 2^{-(\gamma+\kappa)n}\right\} \text{ and}$$
$$S^\times := \left\{z \in \{0,1\}^n : Q^\times(z) \leq 2^{-(\gamma+\kappa)n}\right\}$$

to be the sets of strings with small probabilities and denote by $L^+ := \overline{S}^+$ and $L^\times := \overline{S}^\times$ their complements. (Here's the mnemonic: S for the strings with Small probabilities, L for Large.) Note that for all $x \in L^+$, we have that $Q^+(x) > 2^{-(\gamma+\kappa)n}$ and therefore $|L^+| < 2^{(\gamma+\kappa)n}$. Analogously, we have $|L^\times| < 2^{(\gamma+\kappa)n}$. For the ease of notation, we abbreviate the probabilities that strings with small probabilities occur as follows: $q^+ := Q^+(S^+)$ and $q^\times := Q^\times(S^\times)$. The next corollary now immediately follows from Theorem 3.5.

Corollary 3.9. *Let $\gamma + \kappa < \frac{1}{2}$. For the probability distributions Q^+, $Q^\times$ and the sets S^+, $S^\times$ defined above, we have*

$$q^+ + q^\times := Q^+(S^+) + Q^\times(S^\times) \geq 1 - negl(n).$$

Theorem 3.10. *For all $\gamma < \frac{1}{2}$,* QOT *is secure against $\mathfrak{R}_\gamma$.*

Proof: After Lemmata 3.2 and 3.3, it remains to show that EPR-QOT is oblivious against $\mathfrak{R}_\gamma$. Since $\gamma < \frac{1}{2}$, we can find $\kappa > 0$ with $\gamma + \kappa < \frac{1}{2}$. Consider a dishonest receiver in EPR-QOT $\tilde{\mathsf{R}}$ with quantum memory of size γn.

Using the notation from Section 3.1, we show that there exists an event $\mathcal{E}$ such that $P[\mathcal{E}] \geq \frac{1}{2} - negl(n)$ as well as $\delta([B \otimes \boldsymbol{\rho}_{\tilde{\mathsf{R}}}|\mathcal{E}], [B] \otimes [\boldsymbol{\rho}_{\tilde{\mathsf{R}}}|\mathcal{E}]) \leq negl(n)$, as required by the obliviousness condition of Definition 3.1. Let X denote the random variable describing the outcome x of S's measurement (in basis r) in Step 4 of EPR-QOT. We implicitely understand the distribution of X to be conditioned on the classical outcome y of the measurement $\tilde{\mathsf{R}}$ performs, as described in Section 3.3. We define $\mathcal{E}$ to be the event $X \in S^r$. Note that $\mathcal{E}$ is independent of B and thus $[B|\mathcal{E}] = [B]$. Furthermore, due to the uniform choice of r, and using Corollary 3.9, $P[\mathcal{E}] = \frac{1}{2}(q^+ + q^\times) \geq \frac{1}{2} - negl(n)$.

In order to show the second condition, we have to show that whenever $\mathcal{E}$ occurs, the dishonest receiver cannot distinguish the situation where $B = 0$ is sent from the one where $B = 1$ is sent. As the bit B is masked by the output of the hash function $F(X)$ in Step 4 of EPR-QOT (where the random variable F represents the random choice for f), this is equivalent to distinguish between $F(X) = 0$ and $F(X) = 1$. This situation is exactly suited for applying Theorem 2.1, which says that $F(X) = 0$ is indistinguishable from $F(X) = 1$ whenever the right-hand side of (1) is negligible.

In the case $r = +$, we have

$$\begin{aligned} H_\infty(X|X \in S^+) &= -\log\left(\max_{x \in S^+} \frac{Q^+(x)}{q^+}\right) \\ &\geq -\log\left(\frac{2^{-(\gamma+\kappa)n}}{q^+}\right) = \gamma n + \kappa n + \log(q^+). \end{aligned} \tag{5}$$

If $q^+ \geq 2^{-\frac{\kappa}{2}n}$ then $H_\infty(X|X \in S^+) \geq \gamma n + \frac{\kappa}{2}n$ and indeed the right-hand side of (1) decreases exponentially when conditioning on $X \in S^+$. The corresponding holds for the case $r = \times$.

Finally, if $q^+ < 2^{-\frac{\kappa}{2}n}$ (or similarly $q^\times < 2^{-\frac{\kappa}{2}n}$) then instead of as above we define $\mathcal{E}$ as the *empty event* if $r = +$ and as the event $X \in S^\times$ if $r = \times$. It follows that $P[\mathcal{E}] = \frac{1}{2} \cdot q^\times \geq \frac{1}{2} - negl(n)$ as well as $H_\infty(X|\mathcal{E}) = H_\infty(X|X \in S^\times) \geq \gamma n + \kappa n + \log(q^\times) \geq \gamma n + \frac{\kappa}{2}n$ (for n large enough), both by Corollary 3.9 and the bound on q^+. □

3.6 Weakening The Assumptions

Observe that QOT requires error-free quantum communication, in that a transmitted bit b, that is encoded by the sender and measured by the receiver using the same basis, is always received as b. And it requires a perfect quantum source which on request produces *one* qubit in the right state, e.g. *one* photon with the right polarization. Indeed, in case of noisy quantum communication, an honest receiver in QOT is likely to receive an incorrect bit, and the obliviousness of QOT is vulnerable to imperfect sources that once in while transmit more than one qubit in the same state: a malicious receiver $\tilde{\mathsf{R}}$ can easily determine the basis $r \in \{+, \times\}$ and measure all the following qubits in the right basis. However, current technology only allows to approximate the behavior of single-photon sources and of noise-free quantum communication. It would be preferable to find a variant of QOT that allows to weaken the technological requirements put upon the honest participants.

In this section, we present such a protocol based on BB84 states [1], BB84-QOT (see Figure 3). The security proof follows essentially by adapting the security analysis of QOT in a rather straightforward way, as will be discussed later.

Let us consider a quantum channel with an error probability $\phi < \frac{1}{2}$, i.e., ϕ denotes the probability that a transmitted bit b, that is encoded by the sender and measured by the receiver using the same basis, is received as $1-b$. In order not to have the security rely on any level of noise, we assume the error probability to be zero when considering a *dishonest* receiver. Also, let us consider a quantum source which produces two or more qubits (in the same state), rather than just one, with probability $\eta < 1 - \phi$. We call this the (ϕ, η)-weak quantum model.

In order to deal with noisy quantum communication, we need to do error-correction without giving the adversary too much information. For this, we use *secure sketches*, as introduced in [11]. A (ℓ, m, ϕ)-secure sketch[2] is a randomized function $S : \{0,1\}^\ell \to \{0,1\}^*$ such that (1) for any $w \in \{0,1\}^\ell$ and for w' received from w by flipping each bit (independently) with probability ϕ, the string w can be recovered from w' and $S(w)$ except with negligible probability (in ℓ), and (2) for all random variables W over $\{0,1\}^\ell$, the "average min-entropy" of W given $S(W)$ is at least $H_\infty(W) - m$. We would like to point out that the notion of average min-entropy used in [11] and here differs slightly from the standard notion $H_\infty(W|S(W))$, but it implies that for any $\Delta > 0$, the probability that $S(W)$ takes on a value y such that $H_\infty(W|S(W) = y) \geq H_\infty(W) - m - \Delta$ is at least $1 - 2^{-\Delta}$ (which is sufficient for our purpose).

Consider the protocol BB84-QOT in the (ϕ, η)-weak quantum model shown in Figure 3. For simplicity, we assume n to be even. The protocol uses a $(\frac{n}{2}, \alpha\frac{n}{2}, \phi)$-secure sketch S. We will argue later that α can be chosen arbitrarily close to (but greater than) $h(\phi)$. Like before, the memory bound in BB84-QOT applies before Step 4.

By the properties of a secure sketch, it is obvious that R receives the correct bit b if $r' = r$, except with negligible probability. Also, since there is no communication from

[2] Note that our definition of a secure sketch differs slightly from the one given in [11].

BB84-QOT(b):

1. S picks $x \in_R \{0,1\}^n$ and a random index set $I_+ \subset_R \{1,\ldots,n\}$ of size $\frac{n}{2}$ and sets $I_\times := \{1,\ldots,n\} \setminus I_+$.
2. For $i = 1,2,\ldots,n$: If $i \in I_+$, S sends $|x_i\rangle_+$ to R. If otherwise $i \in I_\times$, S sends $|x_i\rangle_\times$.
3. R picks $r' \in_R \{+,\times\}$ and measures all qubits in basis r'. Let $x' \in \{0,1\}^n$ be the result.
4. S picks $r \in_R \{+,\times\}$ and announces r, I_r, $y := S(x|_{I_r})$, $f \in_R \mathrm{H}_{n/2}$, and $s := b \oplus f(x|_{I_r})$.
5. R can recover $x|_{I_r}$ from $x'|_{I_r}$ and y, and outputs $a := 1$ and $b' := s \oplus f(x|_{I_r})$ if $r' = r$ and else $a := 0$ and $b' := 0$.

Figure 3. Protocol for the BB84 version of Rabin QOT

R to S, BB84-QOT is clearly private. Similar as for protocol QOT, in order to argue about obliviousness we compare BB84-QOT with a purified version shown in Figure 4. BB84-EPR-QOT runs in the $(\phi, 0)$-weak quantum model, and the imperfectness of the quantum source assumed in BB84-QOT is simulated by S in BB84-EPR-QOT so that there is no difference from R's point of view. We would like to point out that the way S chooses the set I_r is more complicated than necessary; this is for proof-technical reasons, as will be clear later.

The security equivalence between BB84-QOT (in the (ϕ, η)-weak quantum model) and BB84-EPR-QOT (in the $(\phi, 0)$-weak quantum model) is omitted here as it follows essentially along the same lines as in Section 3.2. The main difference here is that additionally one has to argue that the distribution of the "imperfectly generated qubits" (within the sets I_+ and $I_\times$) is the same as in BB84-QOT. As a matter of fact, it is not perfectly the same, but it is obviously the same conditioned on the event that the number of "imperfectly generated qubits" with basis $+$ and the number of those with basis $\times$ are both at most $(\eta + \varepsilon)n/2$ (in which case S does not abort in BB84-EPR-QOT). This event, though, happens with overwhelming probability by Bernstein's law of large numbers. This is good enough.

Theorem 3.11. *In the (ϕ, η)-weak quantum model,* BB84-QOT *is secure against $\mathfrak{R}_\gamma$ for any $\gamma < \frac{1-\eta}{4} - \frac{h(\phi)}{2}$ (if parameter α is appropriately chosen).*

Proof Sketch: It remains to show that BB84-EPR-QOT is oblivious against $\mathfrak{B}_\gamma$ (in the $(\phi, 0)$-weak quantum model). The reasoning goes exactly along the lines of the proof of Theorem 3.10, except that we restrict our attention to those i's which are in J. Write $n' = |J| = (1 - \eta - \varepsilon)n/2$, and let γ' be such that $\gamma n = \gamma' n'$, i.e., $\gamma' = 2\gamma/(1 - \eta - \varepsilon)$. It then follows as in the proof of Theorem 3.10 that

$$\begin{aligned} H_\infty\big(X|_J \big| X|_J \in S^+\big) &\geq \gamma' n' + \kappa n' + \log(q^+) \\ &= \gamma n + \kappa(1 - \eta - \varepsilon)n/2 + \log(q^+) \end{aligned}$$

BB84-EPR-QOT(b):

1. S prepares n EPR pairs each in state $|\Omega\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. Additionally, S samples $\theta \in \{+,\times\}^n$ such that $\theta_i = +$ for exactly $\frac{n}{2}$ indices i, and S initializes $I'_+ := \emptyset$ and $I'_\times := \emptyset$.
2. For every $i \in \{1,\dots,n\}$, S does the following. With probability $1-\eta$ S sends one half of the i-th pair to R and keeps the other half. While with probability η S replaces I'_{θ_i} by $I'_{\theta_i} \cup \{i\}$ and sends two or more qubits in the same state $|x_i\rangle_{\theta_i}$ to R where $x_i \in_R \{0,1\}$.
3. R picks $r' \in_R \{+,\times\}$ and measures all received qubits in basis r'. Let $x' \in \{0,1\}^n$ be the result.
4. S picks a random index set $J \subset_R \{1,\dots,n\} \setminus (I'_+ \cup I'_\times)$ of size $(1-\eta-\varepsilon)n/2$ (where $\varepsilon > 0$ is sufficiently small). Then, S picks $r \in_R \{+,\times\}$, chooses a random index set $I_r \subset \{1,\dots,n\}$ of size $\frac{n}{2}$ subject to $J \cup I'_r \subseteq I_r$ (respectively aborts if that is not possible) and for each $i \in I_r \setminus I'_r$ measures the corresponding qubit in basis r. Let x_i be the corresponding outcome, and let $x|_{I_r}$ be the collection of all x_i's with $i \in I_r$. S announces r, I_r, $y = S(x|_{I_r})$, $f \in_R \mathrm{H}_{n/2}$, and $s = b \oplus f(x|_{I_r})$.
5. R can recover $x|_{I_r}$ from $x'|_{I_r}$ and y, and outputs $a := 1$ and $b' := s \oplus f(x|_{I_r})$, if $r' = r$ and else $a := 0$ and $b' := 0$.

Figure 4. Protocol for EPR-based Rabin QOT, BB84 version

Property (2) of a secure sketch then implies that, except with negligible probability, y is such that

$$\begin{aligned} &H_\infty\big(X|_{I_r} \big| X|_J \in S^+, S(X|_{I_r}) = y\big) \\ &\qquad \geq \gamma n + \kappa(1-\eta-\varepsilon)n/2 + \log(q^+) - \alpha n/2 - \varepsilon n \end{aligned}$$

Similar as in the proof of Theorem 3.10, one can consider the cases $q^+ \geq 2^{-\varepsilon n}$ and $q^+ < 2^{-\varepsilon n}$, and in both cases argue that the min-entropy in question is larger than $\gamma n + \varepsilon n$ (which then completes the proof by referring to Theorem 2.1) if $\kappa(1-\eta-\varepsilon) > \alpha + 4\varepsilon$, where $\varepsilon > 0$ may be arbitrarily small and κ has to satisfy $\kappa < \frac{1}{2} - \gamma' = \frac{1}{2} - 2\gamma/(1-\eta-\varepsilon)$. This can be achieved (by choosing ε appropriately) if $\alpha < \kappa(1-\eta) < (1-\eta)/2 - 2\gamma$, which can be achieved (by choosing κ appropriately) if

$$\gamma < \frac{1-\eta}{4} - \frac{\alpha}{2}\,.$$

By the assumed restriction on γ, this inequality can be satisfied if α is chosen arbitrarily close to $h(\phi)$. But this follows in a straightforward way from a result in [11], where it is shown that every (efficiently decodable) error correcting code induces an (efficient) secure sketch (with related parameters), combined with the fact that for every $\alpha > h(\phi)$ there exists an efficiently decodable code of large enough length ℓ, with rate $R = 1 - \alpha$ and which (except with negligible probability) corrects errors introduced with probability ϕ (see [4] and the reference therein). □

4 Quantum Commitment Scheme

In this section, we present a BC scheme from a committer C with bounded quantum memory to an unbounded receiver V. The scheme is peculiar since in order to commit to a bit, the committer does not send anything. During the committing stage information only goes from V to C. The security analysis of the scheme uses similar techniques as the analysis of EPR-QOT.

4.1 The Protocol

The objective of this section is to present a bounded quantum-memory BC scheme COMM (see Figure 5). Intuitively, a commitment to a bit b is made by measuring random BB84-states in basis $\{+, \times\}_{[b]}$.

COMM(b):

1. V picks $x \in_R \{0,1\}^n$ and $r \in_R \{+, \times\}^n$.
2. V sends x_i in the corresponding bases $|x_1\rangle_{r_1}, |x_2\rangle_{r_2}, \ldots, |x_n\rangle_{r_n}$ to C.
3. C commits to the bit b by measuring all qubits in basis $\{+, \times\}_{[b]}$. Let $x' \in \{0,1\}^n$ be the result.
4. To open the commitment, C sends b and x' to V.
5. V verifies that $x_i = x'_i$ for those i where $r_i = \{+, \times\}_{[b]}$. V accepts if and only if this is the case.

Figure 5. Protocol for quantum commitment

As for the OT-protocol of Section 3.2, we present an equivalent EPR-version of the protocol that is easier to analyze (see Figure 6).

EPR-COMM(b):

1. V prepares n EPR pairs each in state $|\Omega\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$.
2. V sends one half of each pair to C and keeps the other halves.
3. C commits to the bit b by measuring all received qubits in basis $\{+, \times\}_{[b]}$. Let $x' \in \{0,1\}^n$ be the result.
4. To open the commitment, C sends b and x' to V.
5. V measures all his qubits in basis $\{+, \times\}_{[b]}$ and obtains $x \in \{0,1\}^n$. He chooses a random subset $I \subseteq \{1, \ldots, n\}$. V verifies that $x_i = x'_i$ for all $i \in I$ and accepts if and only if this is the case.

Figure 6. Protocol for EPR-based quantum commitment

Lemma 4.1. COMM *is secure if and only if* EPR-COMM *is secure.*

Proof: The proof uses similar reasoning as the one for Lemma 3.2. First, it clearly makes no difference, if we change Step 5 to the following:

5'. V chooses the subset I, measures all qubits with index in I in basis $\{+,\times\}_{[b]}$ and all qubits not in I in basis $\{+,\times\}_{[1-b]}$. V verifies that $x_i = x'_i$ for all $i \in I$ and accepts if and only if this is the case.

Finally, we can observe that the view of C does not change if V would have done his choice of I and his measurement already in Step 1. Doing the measurements at this point means that the qubits to be sent to C collapse to a state that is distributed identically to the state prepared in the original scheme. The EPR-version is therefore equivalent to the original commitment scheme from C's point of view. □

It is clear that EPR-COMM is hiding, i.e., that the commit phase reveals no information on the committed bit, since no information is transmitted to V at all. Hence we have

Lemma 4.2. EPR-COMM *is perfectly hiding.*

4.2 Modeling Dishonest Committers

A dishonest committer $\tilde{\mathsf{C}}$ with bounded memory of at most γn qubits in EPR-COMM can be modeled very similarly to the dishonest OT-receiver $\tilde{\mathsf{R}}$ from Section 3.3: $\tilde{\mathsf{C}}$ consists first of a circuit acting on all n qubits received, then of a measurement of all but at most γn qubits, and finally of a circuit that takes the following input: a bit b that $\tilde{\mathsf{C}}$ will attempt to open, the γn qubits in memory, and some ancilla in a fixed state. The output is a string $x' \in \{0,1\}^n$ to be sent to V at the opening stage.

Definition 4.3. *We define $\mathfrak{C}_\gamma$ to be the class of all committers $\{\tilde{\mathsf{C}}_n\}_{n>0}$ in* EPR-COMM *that, at the start of the opening phase (i.e. at Step 4), have a quantum memory of size at most γn qubits.*

We adopt the binding condition for quantum BC from [12]:

Definition 4.4. *A (quantum) BC scheme is* **(statistically) binding** *against $\mathfrak{C}$ if for all $\{\tilde{\mathsf{C}}_n\}_{n>0} \in \mathfrak{C}$, the probability $p_b(n)$ that $\tilde{\mathsf{C}}_n$ opens $b \in \{0,1\}$ with success satisfies*

$$p_0(n) + p_1(n) \leq 1 + negl(n).$$

In the next section, we show that EPR-COMM is binding against $\mathfrak{C}_\gamma$ for any $\gamma < \frac{1}{2}$.

Note that the binding condition given here in Definition 4.4 is weaker than the classical one, where one would require that a bit b exists such that $p_b(n)$ is negligible. But it is the best that can be achieved for a general quantum adversary who can always commit to 0 and 1 in superposition. However, an adversary with bounded quantum storage cannot necessarily maintain a commitment in superposition since the memory compression may force a collapse. Indeed, in upcoming work, we show that commitment schemes exist satisfying the stronger binding condition in the bounded quantum-storage model [7]. While the weaker condition is sufficient for many applications, the stronger one seems to be necessary in some cases (see the conclusion).

4.3 Security Proof Of The Commitment Scheme

Note that the first three steps of EPR-QOT and EPR-COMM (i.e. before the memory bound applies) are exactly the same! This allows us to reuse Corollary 3.9 and the analysis of Section 3.5 to prove the binding property of EPR-COMM.

Theorem 4.5. *For any $\gamma < \frac{1}{2}$,* COMM *is perfectly hiding and statistically binding against $\mathfrak{C}_\gamma$.*

The proof is given below. It boils down to showing that essentially $p_0(n) \leq 1 - q^+$ and $p_1(n) \leq 1 - q^\times$. The binding property then follows immediately from Corollary 3.9. The intuition behind $p_0(n) \leq 1 - q^+ := 1 - Q^+(S^+)$ is that a committer has only a fair chance in opening to 0 if x measured in $+$-basis has a large probability, i.e., $x \notin S^+$. The following proof makes this intuition precise by choosing the ε and δ's correctly.

Proof: It remains to show that EPR-COMM is binding against $\mathfrak{C}_\gamma$. Let $\kappa > 0$ be such that $\gamma + \kappa < \frac{1}{2}$. For the parameters κ and γ considered here, define Q^+, S^+ and q^+ as well as $Q^\times$, $S^\times$ and $q^\times$ as in Section 3.5. Furthermore, let $0 < \delta < \frac{1}{2}$ be such that $h(\delta) < \kappa/2$, where h is the binary entropy function, and choose $\varepsilon > 0$ small enough such that $h(\delta) < (\kappa - \varepsilon)/2$. This guarantees that $B^{\delta n} \leq 2^{(\kappa-\varepsilon)n/2}$ for all (sufficiently large) n. For every n we distinguish between the following two cases. If $q^+ \geq 2^{-\varepsilon n/2}$ then

$$H_\infty(X|X \in S^+) \geq \gamma n + \kappa n + \log(q^+) \geq \gamma n + \Big(\kappa - \frac{\varepsilon}{2}\Big)n$$

where the first inequality is argued as in (5). Applying Lemma 2.2, it follows that any guess $\hat{X}$ for X satisfies

$$\Pr\big[\hat{X} \in B^{\delta n}(X) \,|\, X \in S^+\big] \leq 2^{-\frac{1}{2}(H_\infty(X|X\in S^+)-\gamma n-1)+\log(B^{\delta n})} \leq 2^{-\frac{\varepsilon}{4}n+\frac{1}{2}}.$$

However, if $\hat{X} \notin B^{\delta n}(X)$ then sampling a random subset of the positions will detect an error except with probability not bigger than $2^{-\delta n}$. Hence,

$$\begin{aligned} p_0(n) &= (1-q^+) \cdot p_{0|X \notin S^+} + q^+ \cdot p_{0|X\in S^+} \\ &\leq 1 - q^+ + q^+ \cdot \big(2^{-\delta n}(1 - 2^{-\frac{\varepsilon}{4}n+\frac{1}{2}}) + 2^{-\frac{\varepsilon}{4}n+\frac{1}{2}}\big). \end{aligned}$$

If on the other hand $q^+ < 2^{-\varepsilon n/2}$ then trivially

$$p_0(n) \leq 1 = 1 - q^+ + q^+ < 1 - q^+ + 2^{-\varepsilon n/2}.$$

In any case we have $p_0(n) \leq 1 - q^+ + \mathit{negl}(n)$.

Analogously, we derive $p_1(n) \leq 1 - q^\times + \mathit{negl}(n)$ and conclude that

$$p_0(n) + p_1(n) \leq 2 - q^+ - q^\times + \mathit{negl}(n) \leq 1 + \mathit{negl}(n), \tag{6}$$

where (6) is obtained from Corollary 3.9. □

4.4 Weakening The Assumptions

As argued earlier, assuming that a party can produce single qubits (with probability 1) is not reasonable given current technology. Also the assumption that there is no noise on the quantum channel is impractical. It can be shown that a straightforward modification of COMM remains secure in the (ϕ, η)-weak quantum model as introduced in Section 3.6, with $\phi < \frac{1}{2}$ and $\eta < 1 - \phi$.

Let COMM' be the modification of COMM where in Step 5 V accepts if and only if $x_i = x_i'$ for all *but about a ϕ-fraction* of the i where $r_i = \{+, \times\}_{[b]}$. More precisely, for all but a $(\phi + \varepsilon)$-fraction, where $\varepsilon > 0$ is sufficiently small.

Theorem 4.6. *In the (ϕ, η)-weak quantum model,* COMM' *is perfectly hiding and it is binding against $\mathfrak{C}_\gamma$ for any γ satisfying $\gamma < \frac{1}{2}(1-\eta) - 2h(\phi)$.*

Proof Sketch: Using Bernstein's law of large numbers, one can argue that for *honest* C and V, the opening of a commitment is accepted except with negligible probability. The hiding property holds using the same reasoning as in Lemma 4.2. And the binding property can be argued essentially along the lines of Theorem 4.5, with the following modifications. Let J denote the set of indices i where V succeeds in sending a single qubit. We restrict the analysis to those i's which are in J. By Bernstein's law of large numbers, the cardinality of J is about $(1-\eta)n$ (meaning within $(1 - \eta \pm \varepsilon)n$), except with negligible probability. Thus, restricting to these i's has the same effect as replacing γ by $\gamma/(1-\eta)$ (neglecting the $\pm\varepsilon$ to simplify notation). Assuming that $\tilde{\mathsf{C}}$ knows every x_i for $i \notin J$, for all x_i's with $i \in J$ he has to be able to guess all but about a $\phi/(1-\eta)$-fraction correctly, in order to be successful in the opening. However, $\tilde{\mathsf{C}}$ succeeds with only negligible probability if

$$\phi/(1-\eta) < \delta\,.$$

Additionally, δ must be such that

$$h(\delta) < \frac{\kappa}{2} \qquad \text{with} \qquad \frac{\gamma}{1-\eta} + \kappa < \frac{1}{2}\,.$$

Both restrictions on δ hold (respectively can be achieved by choosing κ appropriately) if

$$2\,h\left(\frac{\phi}{1-\eta}\right) + \frac{\gamma}{1-\eta} < \frac{1}{2}\,.$$

Using the fact that $h(\nu p) \leq \nu h(p)$ for any $\nu \geq 1$ and $0 \leq p \leq \frac{1}{2}$ such that $\nu p \leq 1$, this is clearly satisfied if $2h(\phi) + \gamma < \frac{1}{2}(1-\eta)$. This proves the claim. □

5 Conclusion And Further Research

We have shown how to construct ROT and BC securely in the bounded quantum-storage model. Our protocols require no quantum memory for honest players and remain secure provided the adversary has only access to quantum memory of size bounded by a large fraction of all qubits transmitted. Such a gap between the amount of storage required for honest players and adversaries is not achievable by classical means. All our protocols are non-interactive and can be implemented using current technology.

In this paper, we only considered ROT of one bit per invocation. Our technique can easily be extended to deal with string ROT, essentially by using a class of two-universal functions with range $\{0,1\}^{\ell n}$ rather than $\{0,1\}$, for some ℓ with $\gamma + \ell < \frac{1}{2}$ (respectively $< \frac{1-\eta}{4} - \frac{h(\phi)}{2}$ for BB84-QOT).

Although other flavors of OTs can be constructed from ROT using standard reductions, a more direct approach would give a better ratio storage-bound/communication-complexity. Recent extensions of this work have shown that a 1-2 OT protocol built along the lines of BB84-QOT is secure against adversaries with bounded quantum memory [7]. Interestingly, the techniques used are quite different from the ones of this paper (which appear to fail in case of 1-2 OT), and they additionally allow to analyse and prove secure the BC COMM with respect to the stronger security definition, as discussed in section 4.2.

COMM can easily be transformed into a *string* commitment scheme simply by committing bitwise, at the cost of a corresponding blow-up of the communication complexity. In order to prove this string commitment secure, though, it is necessary that COMM is secure with respect to the stronger security definition.

How to construct and in particular prove secure a more efficient string commitment scheme is still an open problem. Furthermore, it is still unsolved how to construct and prove secure a 1-m OT protocol, more efficient than via the general reduction.

Acknowledgements. We would like to thank the anonymous referees for useful comments. The authors are also grateful to Renato Renner for enlightning discussions.

References

[1] C. H. Bennett and G. Brassard. *Quantum cryptography: Public key distribution and coin tossing.* In Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing, pages 175–179, 1984.

[2] C. Cachin, C. Crépeau, and J. Marcil. *Oblivious Transfer with a Memory-Bounded Receiver.* In Proceedings of the 39th IEEE Symposium on Foundations of Computer Science — FOCS 1998, pages 493–502, 1998.

[3] J. L. Carter and M. N. Wegman. *Universal classes of hash functions.* In Proceedings of the 9th ACM Symposium on Theory of Computing — STOC 1977, pages 106–112, 1977.

[4] C. Crépeau. *Efficient Cryptographic Protocols Based on Noisy Channels.* In Advances in Cryptology — EUROCRYPT 1997, pages 306–317, 1997.

[5] C. Crépeau and J. Kilian. *Achieving Oblivious Transfer Using Weakened Security Assumptions.* In Proceedings of the 29th IEEE Symposium on Foundations of Computer Science — FOCS 1988, pages 42–53, 1988.

[6] I. B. Damgård, S. Fehr, K. Morozov, and L. Salvail. *Unfair Noisy Channels and Oblivious Transfer.* In Theory of Cryptography — TCC 2004, pages 355–373, 2004.

[7] I. B. Damgård, S. Fehr, L. Salvail, and C. Schaffner. *1-2 OT in the Bounded Quantum-Storage Model with Applications.* In preparation, 2005.

[8] I. B. Damgård, S. Fehr, L. Salvail, and C. Schaffner. *Cryptography In the Bounded Quantum-Storage Model.* In Proceedings of the 46th IEEE Symposium on Foundations of Computer Science — FOCS 2005, October 2005.

[9] I. B. Damgård, J. Kilian, and L. Salvail. *On the (Im)possibility of Basing Oblivious Transfer and Bit Commitment on Weakened Security Assumptions.* In Advances in Cryptology — EUROCRYPT 1999, pages 56–73, 1999.

[10] Y. Z. Ding, D. Harnik, A. Rosen, and R. Shaltiel. *Constant-Round Oblivious Transfer in the Bounded Storage Model.* In Theory of Cryptography — TCC 2004, pages 446–472, 2004.

[11] Y. Dodis, L. Reyzin, and A. Smith. *Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data.* In Advances in Cryptology — EUROCRYPT 2004, pages 523–540, 2004.

[12] P. Dumais, D. Mayers, and L. Salvail. *Perfectly Concealing Quantum Bit Commitment from any Quantum One-Way Permutation.* In Advances in Cryptology — EUROCRYPT 2000, pages 300–315, 2000.

[13] S. Dziembowski and U. M. Maurer. *On Generating the Initial Key in the Bounded-Storage Model.* In Advances in Cryptology — EUROCRYPT 2004, pages 126–137, 2004.

[14] C. A. Fuchs and J. van de Graaf. *Cryptographic Distinguishability Measures for Quantum-Mechanical States.* IEEE Transactions on Information Theory, 45(4):1216–1227, 1999.

[15] H.-K. Lo and H. F. Chau. *Is quantum bit commitment really possible?* Physical Review Letters, 78:3410–3413, April 1997.

[16] H. Maassen and J. B. M. Uffink. *Generalized entropic uncertainty relations.* Physical Review Letters, 60(12):1103–1106, March 1988.

[17] D. Mayers. *Unconditionally secure quantum bit commitment is impossible.* Physical Review Letters, 78:3414–3417, April 1997.

[18] T. Moran, R. Shaltiel, and A. Ta-Shma. *Non-interactive Timestamping in the Bounded Storage Model.* In Advances in Cryptology — CRYPTO 2004, pages 460–476, 2004.

[19] R. Renner and R. König. *Universally Composable Privacy Amplification Against Quantum Adversaries.* In Theory of Cryptography — TCC 2005, pages 407–425, February 2005.

[20] L. Salvail. *Quantum Bit Commitment from a Physical Assumption.* In Advances in Cryptology — CRYPTO 1998, pages 338–353, 1998.

[21] P. W. Shor and J. Preskill. *Simple Proof of Security of the BB84 Quantum Key Distribution Protocol.* Physical Review Letters, 85:441–444, July 2000.

[22] M. N. Wegman and J. L. Carter. *New Classes and Applications of Hash Functions.* In Proceedings of the 20th IEEE Symposium on Foundations of Computer Science — FOCS 1979, pages 175–182, 1979.