Efficient Unconditional Oblivious Transfer from Almost any Noisy Channel

Claude Crépeau[1⋆], Kirill Morozov[2⋆⋆], and Stefan Wolf[3⋆⋆⋆]

[1] School of Computer Science, McGill University, Montreal, Canada
`Claude.Crepeau@McGill.ca`
[2] BRICS[†], Aarhus University, Denmark
`kirill@brics.dk`
[3] Département d'Informatique et recherche opérationnelle, Université de Montréal, Canada
`wolf@iro.umontreal.ca`

Abstract. *Oblivious transfer (OT)* is a cryptographic primitive of central importance, in particular in two- and multi-party computation. There exist various protocols for different variants of OT, but any such realization from scratch can be broken in principle by at least one of the two involved parties if she has sufficient computing power—and the same even holds when the parties are connected by a quantum channel. We show that, on the other hand, if *noise*—which is inherently present in any physical communication channel—is taken into account, then OT can be realized in an unconditionally secure way for both parties, i.e., even against dishonest players with unlimited computing power. We give the exact condition under which a general noisy channel allows for realizing OT and show that only "trivial" channels, for which OT is obviously impossible to achieve, have to be excluded. Moreover, our realization of OT is efficient: For a security parameter $\alpha > 0$—an upper bound on the probability that the protocol fails in any way—the required number of uses of the noisy channel is of order $O(\log(1/\alpha)^{2+\varepsilon})$ for any $\varepsilon > 0$.

1 Introduction and Motivation

Cryptographic security can either stem from the fact that an adversary's information about the data of interest is zero (or limited), or that these data are difficult to access from her information. The second type of security is based on the hardness of certain computational problems and depends on assumptions on the adversary such as her computing power. Unfortunately, there do not exist proven lower bounds on the complexity of solving particular problems

⋆ Supported by Canada's NSERC and Québec's FQRNT.
⋆⋆ Work carried out while author was with School of Computer Science, McGill University, Montreal, Canada.
⋆⋆⋆ Supported by Canada's NSERC.
† Basic Research in Computer Science (`www.brics.dk`), funded by the Danish National Research Foundation.

that are directly useful in a cryptographic context. The first type of security is called *information-theoretic* security. It can be realized under specific assumptions—such as bounds on an adversary's memory space [17], [12]—, or from no assumptions at all, in which case this type of security is also called *unconditional*. Clearly, this is the most desirable type of security—but it has its price and can generally not be generated from scratch; this is true for encryption, where its price is a secret key of a certain length [20], as well as for so-called *two-* or *multi-party computation*. Examples of two-party tasks that can be shown impossible from scratch in an unconditionally secure way with respect to both parties simultaneously are the computation of the *OR function* or *bit commitment*.

A primitive of particular importance in the context of secure two-party computation is *oblivious transfer (OT)* due to its universality: From information-theoretically secure OT, any two-party computation can be realized in an unconditionally secure way. OT or, more precisely, chosen one-out-of-two bit OT [13], is the following primitive involving a sender A and a receiver B: A sends two bits b_0 and b_1, B inputs a choice bit c and receives b_c, but remains ignorant about b_{1-c}. The sender A, on the other hand, does not learn c. This variant of OT was shown equivalent—given that a small but non-zero failure probability can be accepted—to the original, so-called *Rabin OT* [19] in [6]; Rabin OT in fact corresponds to a binary erasure channel with erasure probability $1/2$.

It is, therefore, a natural question whether OT can as well be realized from other noisy communication channels between the parties. In fact, also in different contexts, such as encryption, noisy channels have proven useful as simple *information-theoretic primitives* [18] allowing for achieving tasks such as (almost) perfectly secret message transmission [10], [24]. In [7], it was shown that any non-trivial *binary-symmetric channel (BSC)* allows for realizing OT in polynomial time, and in [8] a more efficient construction was given and later shown to work as well for *any* non-trivial BSC in [15], [21].

In the present paper, we generalize this result to *arbitrary discrete memoryless channels (DMCs)*—characterized by a conditional probability distribution $P_{Y|X}$. More precisely, we first define *triviality*: Intuitively speaking, a channel $P_{Y|X}$ is trivial if, after removal of all input symbols whose output distribution can be generated by combining other input symbols, the channel is a parallel composition of capacity-zero channels. The main result of our paper then states that any *non*-trivial channel, and only those channels, allow for—efficiently—realizing OT.

Main Result. *Let two players A and B be connected by a non-trivial channel $P_{Y|X}$. Then, for any $\alpha > 0$, there exists a protocol for unconditionally secure OT from A to B with failure probability at most α, where the number of uses of the channel is of order $O(\log(1/\alpha)^{2+\varepsilon})$ for any $\varepsilon > 0$. Trivial channels, on the other hand, do not allow for realizing OT in an unconditional way.*

In [22], the problem of realizing bit commitment from discrete channels was studied. They showed that string commitment with positive rate can be achieved, i.e., the length of the committed string divided by the number of channel uses is bounded from below by a strictly positive constant. They have extended their notion of *commitment capacity* to *OT capacity* [23]. In this context, they independently used the same notion of non-triviality of discrete channels.

The rest of this paper is organized as follows. In Section 2, we briefly review some notions, facts, and constructions from coding theory and information-theoretic cryptography. Section 3 introduces the notion of non-triviality of a discrete memoryless channel. In particular, we prove a property of non-trivial channels that is crucial for the construction of the OT protocol, which is presented in Section 4.

2 Preliminaries

2.1 Coding Theory

We briefly review some basic facts from coding theory. For a more detailed discussion, we refer to, for instance, [16].

A *binary error-correction code* with code-word length or size n, dimension k, and minimal distance d is a subset of cardinality 2^k of $\{0,1\}^n$—the code words—such that for any two elements v, w of this set $d_H(v,w) \geq d$ holds, where d_H denotes the *Hamming distance* between two bit strings. Of particular importance is the special case of *linear codes*, where the subset of code words is in fact a k-dimensional linear subspace of $\{0,1\}^n$. In this case, the code is called a $[n,k,d]$-code and can be represented by a $k \times n$ matrix G, the *generating matrix*, or, alternatively, by the $n \times (n-k)$ *parity-check matrix* H.

In our protocol presented in Section 4, we use a special class of linear codes, so-called *concatenated codes* [14]. Such codes allow for correcting an asymptotically optimal number of errors: For any $\varphi > 0$ there exists $\rho > 1$ such that for all[4] $R < 1 - h(\varphi)$—the latter expression is the capacity of a BSC with bit error probability φ—and sufficiently large N there exists a linear code with length N and dimension at least RN, failing to correct φN uniformly distributed errors only with probability at most $\rho^{(R-1+h(\varphi))N}$.

The idea of concatenated codes is as follows. A straight-forward Las Vegas construction algorithm combines a power-of-two ($N = 2^n$) size $[N, (1-\alpha)N, \alpha N - 1]$ *Extended Reed-Solomon (outer) code* over the field $\mathbf{F}_{2^n}$ to a rather good (inner) code of size n selected at random among all linear codes $[n, \kappa n, \delta n]$ of appropriate dimension κn. The resulting concatenated code has parameters $[Nn, (1-\alpha)\kappa Nn, \alpha\delta Nn]$ and is able to efficiently correct up to nearly δNn errors on average if they are uniformly distributed (because only very few errors will be uncorrected by the inner code). The error correction procedure uses a brute-force search for the nearest codeword on the inner code and the Berlekamp-Massey algorithm for the outer Extended Reed-Solomon code. Both of these algorithms run in polynomial-time with respect to the global code size Nn.

[4] Here, $h(x) = -(x \log x + (1-x)\log(1-x))$ is the *binary entropy function*. All logarithms are binary.

In our protocols, the information transmitted will not be a codeword but only a syndrome $syn(w) = H^T w$—the noisy versions of the information bits are already known to the receiver. From this syndrome, the decoding algorithm allows for recovering w, given its noisy version.

2.2 Privacy Amplification

Privacy amplification is a general technique for distribution uniformizing or—in a cryptographic context—concentrating an adversary's uncertainty. Privacy amplification was first proposed in the context of quantum key agreement for the special case of deterministic side information [2] and later in general [1]. On the other hand, the effect of additional side information, in our case the syndrome the receiver learns, was studied in [5]. Roughly speaking, the number of bits by which the resulting almost secret string will be shorter corresponds to the length of this side information.

For the following, we can restrict ourselves to the special case where one party knows a noisy version—independently bit by bit—of an original string. This case is simpler than the general case since one can deal with typical sequences and almost-uniform distributions.

Let V be a uniformly distributed n-bit string and let W be generated by independently sending each bit over a BSC with error probability φ. Let, furthermore, $syn : \{0,1\}^n \to \{0,1\}^t$ be a linear function and G be a random variable corresponding to the random choice, according to the uniform distribution, of a function from a 2-universal class of functions [9] $\{0,1\}^n \to \{0,1\}^s$ (for instance, G can be a random linear function mapping n bits to s bits). Then we have, except with exponentially (in n) small probability,

$$H(G(V) \mid syn(V) = syn(v), W, G) \geq s - 2^{-\Omega(h(\varphi)n-t-s)} .$$

3 Trivial Versus Non-Trivial Discrete Memoryless Channels

As a first step towards deriving our main result, we prove a property of non-trivial channels. Intuitively, we show the existence of two particular input symbols of such a channel to which the sender can restrict herself in the OT protocol. A crucial point hereby is that, roughly speaking, she can be forced to use only these two symbols—since her failure to do so will be detected by the receiver.

Definition 1. Let $P_{Y|X}$ be a DMC. We call an input symbol $x \in \mathcal{X}$ *redundant* if its output distribution $P_{Y|X=x}$ can be written as a linear combination of the other output distributions as follows:

$$P_{Y|X=x} = \sum_{x' \in \mathcal{X} \setminus \{x\}} \mu_{x'} P_{Y|X=x'}$$

with $\mu_{x'} \in [0,1]$.

Definition 2. We call a channel $P_{Y|X}$ *trivial* if there exist, after removal of all redundant input symbols, partitions of the (remaining) ranges $\mathcal{X}$ of X and $\mathcal{Y}$ of Y, $\mathcal{X} = \mathcal{X}_1 \cup \cdots \cup \mathcal{X}_n$, $\mathcal{Y} = \mathcal{Y}_1 \cup \cdots \cup \mathcal{Y}_n$, and channels $P_{Y_i|X_i}$, where the ranges of X_i and Y_i are $\mathcal{X}_i$ and $\mathcal{Y}_i$, respectively, such that

$$P_{Y|X=x}(y) = \begin{cases} P_{Y_i|X_i=x}(y) & \text{if } x \in \mathcal{X}_i,\ y \in \mathcal{Y}_i\ , \\ 0 & \text{if } x \in \mathcal{X}_i,\ y \in \mathcal{Y}_j,\ i \neq j \end{cases}$$

holds and such that the capacity of the channel $P_{Y_i|X_i}$ is 0 for all i.

The mentioned well-known result that unconditionally secure OT is impossible to realize by noiseless communication immediately carries over to trivial channels. In Section 4 we will show that, on the other hand, any *other* channel *does* allow for realizing OT. Non-triviality is, therefore, a necessary and sufficient condition for a channel to allow for achieving OT in an unconditionally secure way with respect to both parties. We first give an alternative characterization of non-triviality of a channel.

Theorem 1. *Let $P_{Y|X}$ be a non-trivial channel. Then there exist $x_1, x_2 \in \mathcal{X}$ with the following properties.*

1. $P_{Y|X=x_1} \neq P_{Y|X=x_2}$.
2. *There exists $y \in \mathcal{Y}$ such that $P_{Y|X=x_1}(y) > 0$ and $P_{Y|X=x_2}(y) > 0$.*
3. *Let, for $\lambda, \mu_i \in [0,1]$,*

$$\lambda P_{Y|X=x_1} + (1-\lambda) P_{Y|X=x_2} = \sum_i \mu_i P_{Y|X=x_i}\ .$$

Then $\mu_i > 0$ implies that $P_{Y|X=x_i} = \tau P_{Y|X=x_1} + (1-\tau) P_{Y|X=x_2}$ holds for some $\tau \in [0,1]$.

Remark 1. Intuitively speaking, Theorem 1 states that there are two particular input symbols $x_1, x_2 \in \mathcal{X}$ of the channel with the following properties. If a sender is supposed to use only these two symbols as channel inputs (with certain probabilities or frequencies, say p and $1-p$, respectively), then the receiver can—if the channel is used a large number N of times—detect whenever the sender fails to do so if the latter cheats $\Omega(\sqrt{N})$ times. The only exception is the use of input symbols $x \notin \{x_1, x_2\}$ whose output distribution over $\mathcal{Y}$ is a convex linear combination of the output distributions of x_1 and x_2, i.e., if

$$P_{Y|X=x} = \tau P_{Y|X=x_1} + (1-\tau) P_{Y|X=x_2}$$

holds for some $\tau \in [0,1]$. In our context—where the sender tries to maximize the information he has about the resulting output—, this is, however, not a problem because using x leaves him with less information than if he had used x_1 with probability β and x_2 with probability $1-\beta$, and then forgot what he sent.

Proof. Because of the non-triviality of the channel, there exist two non-redundant input symbols x_1 and x_2' and $y \in \mathcal{Y}$ such that $P_{Y|X=x_1} \neq P_{Y|X=x_2'}$, $P_{Y|X=x_1}(y) > 0$, and $P_{Y|X=x_2'}(y) > 0$ hold.

Let us now interpret $P_{Y|X=x}$, for any $x \in \mathcal{X}$, as a point in $\mathbf{R}^{|\mathcal{Y}|-1}$, where the different coordinates correspond to the probabilities $P_{Y|X=x}(y)$ (which sum up to 1). In the following, we will consider the *convex hull* of the set of points

$$\{P_{Y|X=x} \,|\, x \in \mathcal{X}\} \subseteq \mathbf{R}^{|\mathcal{Y}|-1} \ . \tag{1}$$

We call $P_{Y|X=x_0}$ a *spanning point* of the convex hull if the convex hull of $\{P_{Y|X=x} \,|\, x \in \mathcal{X} \setminus \{x_0\}\}$ is strictly smaller than the one of (1).

Since the spanning points of the hull correspond to non-redundant inputs, we can conclude that there exist two spanning points $P_{Y|X=x_1}$ and $P_{Y|X=x_2'}$ of the convex hull such that there exists $y \in \mathcal{Y}$ with $P_{Y|X=x_1}(y) > 0$ and $P_{Y|X=x_2'}(y) > 0$.

Let us now look at the connections between $P_{Y|X=x_1}$ and all other points $P_{Y|X=x}$, $x \in \mathcal{X}$, and let v_x be the unity vector parallel to the vector in $\mathbf{R}^{|\mathcal{Y}|-1}$ connecting $P_{Y|X=x_1}$ and $P_{Y|X=x}$. In a similar way as for points, we define convex linear combinations and the convex hull for these vectors. Let $\{v_x \,|\, x \in \mathcal{A}\}$, where $\mathcal{A} \subseteq \mathcal{X}$, be the set of spanning vectors.

We will first argue that there exists $x_2 \in \mathcal{A}$ with $P_{Y|X=x_2}(y) > 0$, and secondly, that the representation of any linear combination of x_1 and x_2 as a linear combination of *all points* $P_{Y|X=x}$ is unique—modulo points that are *themselves* linear combinations of x_1 and x_2.

Assume that for all $x \in \mathcal{A}$, we have $P_{Y|X=x}(y) = 0$. Then the same is true also for all distributions in the convex hull of these points. On the other hand, the connection between x_1 and x_2' has a non-empty intersection with this convex hull by definition of $\mathcal{A}$. Since every distribution in this intersection is a convex linear combination of $P_{Y|X=x_1}$ and $P_{Y|X=x_2'}$—both non-zero in y— there exists a point x_2 in $\mathcal{A}$ with $P_{Y|X=x_2}(y) > 0$.

By construction, x_1 and x_2 have now the following properties. First, they satisfy $P_{Y|X=x_1} > 0$ and $P_{Y|X=x_2} > 0$. Second, any convex linear combination of $P_1 = P_{Y|X=x_1}$ and $P_2 = P_{Y|X=x_2}$ cannot be represented as a convex linear combination involving points $P_{Y|X=x}$ *not* lying on the line connecting P_1 and P_2; this would contradict the fact that P_2 is a spanning point of the connections of P_1 to all other points P; indeed, the line from P_1 to P_2 could in this case be represented as a linear combination of the lines connecting P_1 with the external points occurring in the linear combination. This observation concludes the proof. □

4 A Protocol for Efficient Oblivious Transfer from any Non-Trivial Channel

In this section we describe a protocol for OT based on an arbitrary non-trivial DMC and give, hence, a proof of our main result stated above. Our protocol is an adaptation of the protocol from [8] for the general case (where, at the same time, we reduce the required number of channel uses from *cubic* to, roughly, *quadratic* order in $\log(1/\alpha)$). We develop the protocol in three steps. In Section 4.1, the original channel is used to obtain a binary-symmetric erasure channel with error; in Section 4.2, this is transfered into a weak form of OT vulnerable to active attacks by the sender A; in Section 4.3, finally, we derive the final protocol avoiding these attacks by statistical analysis by the receiver.

4.1 Binary-Symmetric Erasure Channel with Error from any Non-Trivial Channel

¿From a non-trivial channel $P_{Y|X}$, we first construct a *binary erasure channel with error*. We encode the bits to be transmitted over the DMC as pairs of two fixed distinct input symbols $x_1, x_2 \in \mathcal{X}$ chosen according to Theorem 1: "0" is encoded as x_1x_2 and "1" as x_2x_1. When this is repeated many times, B gets the sent bits with different error rates, depending on the actual output symbols received. We will have B make a decision on 0 or 1 only when he receives certain specific pairs—otherwise, he will decide on erasure Δ. More precisely, B will accept only the pairs which give him the best estimate of what has been sent; we will call these the *most informative pairs*. Note that there might even be output symbols y which allow for deciding *with certainty* whether x_1 or x_2 has been sent. Note, however, that the choice of x_1 and x_2 *guarantees* that there exist pairs which are *not* conclusive with certainty. The crucial point is that there are at least two different levels of conclusiveness, and it is the difference between the two that will be used in the protocol. In the following, we will call the pairs providing B with the best *a posteriori* probabilities *good pairs* and denote them by y_1y_2 and y_2y_1, respectively.

Let $\mathcal{Y}'$ be the set of y with $P_{Y|X=x_1}(y) > 0$ or $P_{Y|X=x_2}(y) > 0$. Formally, the *most informative pair* (y_1, y_2) is the pair $(y, \overline{y}) \in \mathcal{Y}' \times \mathcal{Y}'$, $y \neq \overline{y}$, that achieves the following minimum:

$$\varphi = \min_{(y,\overline{y}) \in \mathcal{Y}' \times \mathcal{Y}'} \frac{P_{Y|X=x_1}(\overline{y}) P_{Y|X=x_2}(y)}{P_{Y|X=x_1}(\overline{y}) P_{Y|X=x_2}(y) + P_{Y|X=x_1}(y) P_{Y|X=x_2}(\overline{y})} \ . \tag{2}$$

Note that (2) is symmetric with respect to x_1 and x_2. The resulting channel is, hence, a binary-symmetric erasure channel (BSEC), i.e., a binary-input channel with some erasure probability and a certain bit-error probability.

Protocol 1 $P_{Y|X} \rightarrow \mathrm{BSEC}(r)$

1. A sends x_1x_2 if $r = 0$ and x_2x_1 if $r = 1$.
2. B returns $\begin{cases} 0 & \text{if } y_1y_2 \text{ is received,} \\ 1 & \text{if } y_2y_1 \text{ is received,} \\ \Delta & \text{if any other pair is received.} \end{cases}$

4.2 Passively Secure OT

The BSEC obtained above is not a Rabin OT: B might get some information even when deciding on Δ, and there are bit errors. We now describe a protocol, based on the obtained BSEC, for realizing OT under the assumption that A behaves correctly.

In the "weak OT" protocol, A sends $2n$ random bits $r_1, r_2, \ldots, r_{2n}$ to B using BSEC. B should receive roughly $2p_g n$ of them as *good pairs* and $2(1-p_g)n$ "bad" ones, where p_g denotes the probability that B decides on either 0 or 1, but not Δ, in an execution of BSEC given that A is honest, i.e.,

$$p_g = (P_{Y|X=x_1}(y_1)P_{Y|X=x_2}(y_2) + P_{Y|X=x_2}(y_1)P_{Y|X=x_1}(y_2))/2 \ .$$

B then forms two sets I_0 and I_1 of size n if $p_g > 1/2$ and, otherwise, size $n' = (p_g - \gamma)n$, $\gamma > 0$. By the index sets I_0 and I_1, B defines two bit-strings r'_{I_0}, r'_{I_1} such that r'_{I_c} should contain only good pairs.

Let now φ be the bit error probability of the BSEC. The players now establish a code—according to the discussion in Section 2.1—which exactly allows for correcting (except with small probability) all errors of a set consisting only of good pairs. More precisely, the errors are corrected by having A send the syndromes of the two words $syn(r_{I_0})$ and $syn(r_{I_1})$. Using r'_{I_c} and $syn(r_{I_c})$, B can recover r_{I_c} except with small probability. On the other hand, this correction information is *not* sufficient to find out *both* words r_{I_c} and $r_{I_{1-c}}$ as long as the dimension of the code does not exceed $(1 - h(\varphi))n'$.

Finally, a linear privacy amplification function is used to extract one bit per string, such that one of the two bits may be recovered, but not both. This function is the scalar product (we denote it as "$\odot$") with a random n'-bit string m. (Note that *string* OT instead of *bit* OT could be obtained using hashing to a string as the privacy-amplification function.)

Protocol 2 BSEC $\rightarrow \widehat{\mathrm{OT}}(b_0, b_1)(c)$

1. A picks $2n$ random bits r_i, $i = 1, \ldots, 2n$, and sends them to B as BSEC(r_i); B receives r'_i.
2. B picks and sends two disjoint sets I_0, I_1, $|I_0| = |I_1| = n'$, such that $r'_i \neq \Delta$ holds for all $i \in I_c$.
3. A and B agree on a parity-check matrix H of a concatenated code C with parameters $[n', k = (1 - h(\varphi))n', d]$ correcting $\psi\varphi n'$ errors, $\psi > 1$.
4. (a) A computes and sends $s_0 = syn(r_{I_0})$ and $s_1 = syn(r_{I_1})$,
   (b) picks and sends a random n'-bit word m, and
   (c) computes and sends $\hat{b}_0 = b_0 \oplus (m \odot r_{I_0})$ and $\hat{b}_1 = b_1 \oplus (m \odot r_{I_1})$.
5. (a) B recovers r_{I_c} using r'_{I_c}, s_c and the decoding algorithm of C and
   (b) computes and returns $\hat{b}_c \oplus (m \odot r_{I_c})$.

Let us discuss why B is unable to cheat in the weak OT protocol. In fact, the chosen code is such that *complete* error correction is possible only if B collects all the good pairs into one of the two sets. Suppose first that $p_g > 1/2$ holds.

Then there exists a constant fraction of bad bits the error rate of which is at least $\varphi' > \varphi$, where φ' is the error rate of the second most informative pairs. Assume for simplicity that the fraction of the second most informative bits is $1 - p_g$, which is the worst case (from A's viewpoint). A dishonest B is not able to put more than $p_g n$ good bits in at least one of the sets I_0 and I_1. The bits of this set do not contain more than $((1 - h(\varphi))p_g + (1 - h(\varphi'))(1 - p_g))n$ bits of Shannon information about the original string with high probability. Therefore, at least $(h(\varphi)p_g + h(\varphi')(1 - p_g))n$ parity-check bits are needed to correct all the errors in each set with high probability; however, $syn(r_{I_0})$, $syn(r_{I_1})$ each contain $(h(\varphi) + \delta)n$ bits only. Thus, at least one of the two words r_{I_0}, r_{I_1} will be undetermined by at least $n(h(\varphi') - h(\varphi))(1 - p_g)$ bits. ¿From the results sketched in Section 2.2, one can conclude that after privacy amplification, B only has an exponentially small amount of information about the corresponding bit. The case of $p_g \leq 1/2$ can be treated in a similar way.

Unfortunately, the weak OT protocol is not secure against cheating by A with the objective of figuring out B's choice bit c. For instance, A can send *incorrect pairs:* x_1x_1 or x_2x_2 instead of x_1x_2 and x_2x_1, hereby increasing the probability that it is received as a bad pair (i.e., $r'_i = \Delta$) by B. Alternatively, A can use any other input symbols but x_1 and x_2 (we call them *forbidden* input symbols) whose support intersects with those of x_1 and x_2. Finally, she can send an incorrect syndrome at Step 4.

In the first and second active attacks, incorrect pairs are more likely to end up in the "bad" set, thus indicating to A which one of I_0 and I_1 is more likely to be the "good" and the "bad" set, respectively, and hence what B's choice is. In the third attack, if A renders only one of the syndromes incorrect, then B will abort or not, depending on which bit he is trying to get.

4.3 The Complete OT Protocol

The main idea is now, as in [8], to avoid cheating by A by repeating the weak OT protocol many times in such a way that A has to cheat in a substantial fraction of all executions of BSEC (namely, in more than the square root of the total number of executions) in order to gain useful information. This, however, can be detected by B when he analyzes his output statistically.

More precisely, Protocol $\widehat{\mathrm{OT}}$ is repeated $\lceil n^{1+\varepsilon} \rceil$ times, $0 < \varepsilon < 1$; thus, we apply BSEC $2\lceil n^{2+\varepsilon} \rceil$ times in total. In order to cheat, A will have to send at least $\lceil n^{1+\varepsilon} \rceil$ wrong pairs (i.e., she forms the pair incorrectly or uses forbidden symbols) in these executions. This will, however, lead to a detectable bias in the output distribution (with probability almost 1). If, on the other hand, A uses *less than* $\lceil n^{1+\varepsilon} \rceil$ incorrect pairs, she finds out nothing about c. Similarly, if A sends wrong syndromes in the protocol for $\widehat{\mathrm{OT}}$ she will, each time, be detected by B with probability 1/2. If she uses $n^{1+\varepsilon}$ such faulty syndromes it is, hence, only with exponentially small probability that B will not detect her cheating.

Let $n_\varepsilon = \left[n^{1+\varepsilon} \right]$, where $[\cdot]$ means rounding up to the next odd integer, and $n'_\varepsilon = n \cdot n_\varepsilon$. The instances are combined by requesting $b_{l,0} \oplus b_{l,1} = b_0 \oplus b_1$ for

$1 \leq l \leq n_\varepsilon$. Let

$$b_{0,0} = \bigoplus_{l=1}^{n_\varepsilon} b_{l,0} \quad \text{and} \quad b_{0,1} = \bigoplus_{l=1}^{n_\varepsilon} b_{l,1} \ .$$

Then we get

$$\bigoplus_{l=1}^{n_\varepsilon} b_{l,c_l} = b_{0,z} \quad \text{for} \quad z = \bigoplus_{l=1}^{n_\varepsilon} c_l \ .$$

Thus, in order to find out which of $b_{0,0}$ or $b_{0,1}$ B is trying to receive, A must find out all the c_l.

Let $\psi > 1$. An extra index l is added to each variable of the lth iteration of $\widehat{\mathrm{OT}}$. Let us denote by $q_{l,\bar{i}} \in \mathcal{Y}'$ the $\bar{i}$th output symbol ($1 \leq \bar{i} \leq 4n'_\varepsilon$) and as $r_{l,i} \in \{0,1\}$ the ith output bit ($1 \leq \bar{i} \leq 2n'_\varepsilon$) received by B in the lth iteration of $\widehat{\mathrm{OT}}$. Let

$$\delta = \min_{y \in \mathcal{Y},\, \overline{x} \in \mathcal{X} \setminus \{x_1, x_2\}} \left| \frac{P_{Y|X=x_1}(y) + P_{Y|X=x_2}(y)}{2} - P_{Y|X=\overline{x}}(y) \right| .$$

Roughly speaking, δ is the closest the sender can get to "the middle point" between the distributions $P_{Y|X=x_1}$ and $P_{Y|X=x_2}$ using forbidden symbols (except the symbols lying on the line between x_1 and x_2, as discussed in Section 3).

Protocol 3 $\widehat{\mathrm{OT}} \to \mathrm{OT}$

1. A picks n_ε random bits $b_{1,0}, b_{2,0}, \ldots b_{n_\varepsilon,0}$ and sets $b_{l,1} = b_0 \oplus b_1 \oplus b_{l,0}$ for $1 \leq l \leq n_\varepsilon$.
2. B picks n_ε random bits $c_1, c_2, \ldots, c_{n_\varepsilon}$.
3. **Repeat** for $l = 1, \ldots, n_\varepsilon$
   (a) A runs $\widehat{\mathrm{OT}}(b_{l,0}, b_{l,1})(c_l)$ with B who gets b'_l,
   (b) **if** $d_H(r_{l,I_{l,c_l}}, r'_{l,I_{l,c_l}}) > \psi\varphi n'$ **then** B aborts.
4. **if** for some j, $1 \leq j \leq |\mathcal{Y}'| - 1$:

$$\left| \#\{l, \bar{i} | q_{l,\bar{i}} = y_j\} - 2n'_\varepsilon \big(P_{Y|X=x_1}(y_j) + P_{Y|X=x_1}(y_j)\big) \right| > \frac{\delta}{2(|\mathcal{X}| - 2)} n_\varepsilon \ ,$$

   then B aborts **else if**

$$\#\{l, i \mid r_{l,i} = y_1 y_2 \text{ or } r_{l,i} = y_2 y_1\} < 2p_g n'_\varepsilon - \frac{2p_g - 1}{4} n_\varepsilon \ ,$$

   then B aborts **else** B computes and sends $c' = c \oplus \left(\bigoplus_{l=1}^{n_\varepsilon} c_l \right)$.
5. A computes and sends $\hat{b}_0 = b_0 \oplus \left(\bigoplus_{l=1}^{n_\varepsilon} b_{l,c'} \right)$ and $\hat{b}_1 = b_1 \oplus \left(\bigoplus_{l=1}^{n_\varepsilon} b_{l,1-c'} \right)$ to B.
6. B computes and returns $\hat{b}_c \oplus \left(\bigoplus_{l=1}^{n_\varepsilon} b'_l \right)$.

The test in Step 3 of the protocol is to decide whether the syndrome sent by A was valid: If the decoded word has Hamming distance larger than $\psi\varphi n'$ to the received string, then the syndrome was wrong.

We briefly argue that the tests of Step 4 achieve their goals. Let $y \in \mathcal{Y}$ and

$$z_{i,j}^{(y)} = \begin{cases} 0 & \text{if } q_{l,i} \neq y \\ 1 & \text{if } q_{l,i} = y. \end{cases}$$

When A sends only x_1 and x_2, we have for all y that

$$E\left[\sum_{i=1}^{n_\varepsilon}\sum_{j=1}^{4n} z_{i,j}^{(y)}\right] = 4n'_\varepsilon \frac{P_{Y=y|X=x_1} + P_{Y=y|X=x_2}}{2}$$

holds, i.e., B expects to see the "middle distribution" between $P_{Y|X=x_1}$ and $P_{Y|X=x_2}$ for all y. Because of Theorem 1 and the choice of x_1 and x_2, A cannot simulate this "middle point" using the forbidden symbols. Therefore, all she can do is send other symbols in order to get as close as possible to the target distribution, however, she cannot get closer than δ.

For the second test of Step 4 the idea is that the receiver calculates the overall number of accepted symbols y_1y_2 and y_2y_1:

$$w_{i,j} = \begin{cases} 1 & \text{if } r_{l,i} = y_1y_2 \text{ or } r_{l,i} = y_2y_1, \\ 0 & \text{otherwise.} \end{cases}$$

Then,

$$E\left[\sum_{i=1}^{n_\varepsilon}\sum_{j=1}^{2n} w_{i,j}\right] = 2p_g n'_\varepsilon$$

holds, where p_g is, as above, the probability to receive a good pair given that x_1x_2 or x_1x_2 was sent. If the actual number of good pairs received is too low, A must have used the incorrect pairs x_1x_1 or x_2x_2; hence, the receiver aborts.

Theorem 2 follows from Bernstein's law of large numbers. Note, hereby, that $n_\varepsilon = [n^{1+\varepsilon}] > \sqrt{n'_\varepsilon} = \sqrt{n \cdot n_\varepsilon}$.

Theorem 2. *There exist constants $\rho_1 < 1$, $\rho_2 < 1$ such that when A does not use the forbidden symbols then*

$$Prob\left[\sum_{i=1}^{n_\varepsilon}\sum_{j=1}^{4n}\left|z_{i,j}^{(y)} - 2n'_\varepsilon\big(P_{Y|X=x_1}(y_j) + P_{Y|X=x_1}(y_j)\big)\right| > \frac{\delta}{2(|\mathcal{X}|-2)}n_\varepsilon\right] < \rho_1^n$$

holds, whereas, when she cheats n_ε times,

$$Prob\left[\sum_{i=1}^{n_\varepsilon}\sum_{j=1}^{4n}\left|z_{i,j}^{(y)} - 2n'_\varepsilon\big(P_{Y|X=x_1}(y_j) + P_{Y|X=x_1}(y_j)\big)\right| < \frac{\delta}{2(|\mathcal{X}|-2)}n_\varepsilon\right] < \rho_1^n$$

holds; if A does not use incorrect pairs, then we have

$$Prob\left[\sum_{i=1}^{n_\varepsilon}\sum_{j=1}^{2n} w_{i,j} < 2p_g n'_\varepsilon - \frac{2p_g-1}{4} n_\varepsilon\right] < \rho_2^n \ ,$$

whereas, when she cheats n_ε times,

$$Prob\left[\sum_{i=1}^{n_\varepsilon}\sum_{j=1}^{2n} w_{i,j} > 2p_g n'_\varepsilon - \frac{2p_g-1}{4} n_\varepsilon\right] < \rho_2^n .$$

Finally, if A is honest, then the probability that more than $\psi\varphi n'$ transmission errors occur is exponentially small. Thus, an honest A is unlikely to fail the test of Step 3, while a dishonest A who deliberately sends a wrong syndrome will be detected with probability $1/2$ if B picks this syndrome.

This concludes the analysis of the protocol, and, hence, the proof of our main result.

5 Concluding Remarks

All computationally secure cryptography is based on assumptions on a possible adversary, and, hence, threatened by any progress in algorithm design and computer engineering. Functionalities of central importance such as encryption, authentication, or multi-party computation cannot, however, be realized in an unconditionally secure way without any given information-theoretic primitive to start from. In the case of oblivious transfer—as for encryption—, however, this initial primitive can be as simple as *noise*, which is an inherent property of any physical communication channel. More precisely, we have shown that OT can be realized in an unconditionally secure way from almost any discrete memoryless noisy channel. This result should be seen in the context of a number of recent results with the common objective to realize cryptographic functionalities in an unconditional way from simple primitives or weak assumptions.

A non-asymptotic analysis of the presented protocol — the concrete values of failure probability depending on the number of channel uses — is out of scope of this paper. Such analysis for the particular case of BSC can be found in [15].

We propose as open problems to realize string OT with non-zero rate in the sense of [22]. A useful result in this context might be a generic reduction of string OT to bit OT based on privacy amplification [4]. Another open problem is to realize OT from more general channels, such as channels *with memory*, or to give a complete characterization with respect to the use of *general unfair channels* [11].

References

1. Bennett, C. H., Brassard, G., Crépeau, C., and Maurer, U. M.: Generalized privacy amplification. In: IEEE Transactions on Information Theory, Vol. 41, Num. 6. IEEE (1995) 1915–1923.

2. Bennett, C. H., Brassard, G., and Robert, J.-M.: Privacy amplification by public discussion. In: SIAM Journal on Computing, vol. 17 (1988) 210–229.
3. Brassard, G., Chaum, D., Crépeau, C.: Minimum disclosure proofs of knowledge. J. of Computer and System Sciences, 37(2). Elsevier (1988) 156–189.
4. Brassard, G., Crépeau, C., and Wolf, S.: Oblivious transfers and privacy amplification. In: *Journal of Cryptology*, vol. 16, no. 4 (2003) 219–237.
5. Cachin, C.: Entropy measures and unconditional security in cryptography. Ph. D. Thesis, ETH Zürich, Hartung-Gorre Verlag, Konstanz, 1997.
6. Crépeau, C.: Equivalence between two flavours of oblivious transfer. In: Advances in Cryptology–CRYPTO '87. Lecture Notes in Computer Science, Vol. 293. Springer-Verlag (1988) 350–354.
7. Crépeau, C., Kilian, J.: Achieving oblivious transfer using weakened security assumptions. In: Proc. 29th Annual Symposium on the Foundations of Computer Science. IEEE (1988) 42–52.
8. Crépeau, C.: Efficient cryptographic primitives based on noisy channels. In: Advances in Cryptology–EUROCRYPT '97. Lecture Notes in Computer Science, Vol. 1233. Springer-Verlag (1997) 306–317.
9. Carter, J.L., Wegman, M.N.: Universal classes of hash functions. J. of Computer and System Sciences, 18. Elsevier (1979) 143–154.
10. Csiszár, I., and Körner, J.: Broadcast channels with confidential messages. In: *IEEE Trans. on Information Theory*, Vol. 24 (1978) 339–348.
11. Damgård, I., Kilian, J., Salvail, L.: On the (im)possibility of basing bit commitment and oblivious transfer on weakened security assumptions. In: Advances in Cryptology–EUROCRYPT '99. LNCS, vol. 1592. Springer-Verlag (1999) 56–73.
12. Dziembowski, S. and Maurer, U. M.: Tight security proofs for the bounded-storage model. In: *Proceedings of STOC 2002* (2002) 341–350.
13. Even, S., Goldreich, O., and Lempel, A.: A randomized protocol for signing contracts. In: Proceedings of CRYPTO '82. Plenum Press (1983) 205–210.
14. Forney, G.D.: Concatenated codes. MIT Press (1966).
15. Korjik, V., Morozov, K.: Generalized oblivious transfer protocols based on noisy channels. In: Proc. Workshop MMM ACNS 2001. LNCS, vol. 2052. Springer-Verlag (2001) 219-229.
16. MacWilliams, F.J., Sloane, N.J.A.: The theory of error-correcting codes. North-Holland (1977).
17. Maurer, U. M., Conditionally-perfect secrecy and a provably-secure randomized cipher. In: *Journal of Cryptology*, Vol. 5, No. 1 (1992) 53–66.
18. Maurer, U. M.: Information-theoretic cryptography. In: *Advances in Cryptology - CRYPTO '99*, LNCS, Vol. 1666. Springer-Verlag. (1999) 47–64.
19. Rabin, M.O.: How to exchange secrets by oblivious transfer. Technical Memo TR-81, Aiken Computation Laboratory, Harvard University (1981).
20. Shannon, C. E.: Communication theory of secrecy systems. In: *Bell System Technical Journal*, Vol. 28 (1949) 656–715.
21. Stebila, D., Wolf, S.: Efficient oblivious transfer from any non-trivial binary-symmetric channel. In: International Symposium on Information Theory (ISIT) (2002) 293.
22. Winter, A., Nascimento, A.C.A., Imai, H.: Commitment capacity of discrete memoryless channels. In: Cryptography and Coding. LNCS, vol. 2898. Springer-Verlag (2003) 35–51.
23. Winter, A., Nascimento, A.C.A.: Oblivious transfer from any genuine noise. Unpublished manuscript (2004).

24. Wyner, A. D.: The wire-tap channel. In: *Bell System Technical Journal*, Vol. 54, No. 8 (1975) 1355–1387.