One-Way Secret-Key Agreement and Applications to Circuit Polarization and Immunization of Public-Key Encryption*

Thomas Holenstein and Renato Renner

Department of Computer Science,
Swiss Federal Institute of Technology (ETH),
Zürich, Switzerland
{thomahol, renner}@inf.ethz.ch

Abstract. Secret-key agreement between two parties Alice and Bob, connected by an insecure channel, can be realized in an information-theoretic sense if the parties share many independent pairs of correlated and partially secure bits. We study the special case where only one-way communication from Alice to Bob is allowed and where, for each of the bit pairs, with a certain probability, the adversary has no information on Alice's bit. We give an expression which, for this situation, exactly characterizes the rate at which Alice and Bob can generate secret key bits.

This result can be used to analyze a slightly restricted variant of the problem of polarizing circuits, introduced by Sahai and Vadhan in the context of statistical zero-knowledge, which we show to be equivalent to secret-key agreement as described above. This provides us both with new constructions to polarize circuits, but also proves that the known constructions work for parameters which are tight.

As a further application of our results on secret-key agreement, we show how to immunize single-bit public-key encryption schemes from decryption errors and insecurities of the encryption, a question posed and partially answered by Dwork, Naor, and Reingold. Our construction works for stronger parameters than the known constructions.

1 Introduction

Consider two parties, Alice and Bob, connected by an authentic but otherwise fully insecure communication channel. It is well known that it is impossible for Alice and Bob to establish information-theoretically secure private communication (see [16,11]). In particular, they are unable to generate an unconditionally secure key. This changes dramatically if we additionally assume that Alice and Bob have access to some correlated randomness on which an adversary has only partial information.

The initial correlation shared by Alice and Bob can originate from various sources. For example, Wyner [20] and, subsequently, Csiszár and Körner [3] have

* Supported by the Swiss National Science Foundation, project no. 200020-103847/1.

studied a scenario where Alice and Bob are connected by a noisy channel on which an adversary has only limited access. Maurer [11] (cf. also [1]) proposed to consider a setting where a satellite broadcasts uniform random bits with low signal intensity, such that Alice, Bob, and also Eve cannot receive them perfectly. It has been shown that, in both settings, Alice and Bob can indeed generate an information-theoretically secure key and thus communicate secretly.

In this paper, we study one-way secret-key agreement, i.e., we assume that only one-way communication from Alice to Bob is allowed. We fully analyze the case where Alice and Bob hold many independent pairs of correlated bits, and where the only secrecy guarantee is that, for each of these pairs, with a certain probability, the adversary has no information about Alice's value. It turns out that this particular kind of information-theoretic secret-key agreement has interesting applications, even in the context of computational cryptography.

1.1 Secret-Key Agreement

Previous Work: Information-theoretically secure secret-key agreement from correlated information has first been proposed by Maurer in [11]. He considered a setting where Alice, Bob, and Eve hold many independent realizations of correlated random variables X , Y , and Z , respectively, with joint probability distribution P_{XYZ} . The (two-way) *secret-key rate* $S(X; Y|Z)$, i.e., the rate at which Alice and Bob can generate secret-key bits per realization of (X, Y, Z) , has further been studied in [1] and later in [12], where the *intrinsic information* $I(X; Y \downarrow Z)$ is defined and shown to be an upper bound on $S(X; Y|Z)$, which, however, is not tight [13].

For *one-way* communication, it is already implied by a result in [3] and has later been shown in [1] that the secret-key rate $S_{\rightarrow}(X; Y|Z)$ is given by the supremum of $H(U|ZV) - H(U|YV)$, taken over all possible random variables U and V obtained from X .¹ However, as this is a purely information-theoretic result, it does not directly imply that there exists an *efficient* key-agreement protocol.

Our Contributions: In Section 2, we show that $H(U|ZV) - H(U|YV)$ is the exact rate at which Alice and Bob can *efficiently* generate a secret key. The methods used to show this are not new, but as far as we know this result has not appeared anywhere else in the literature.

Furthermore, we study the class of distributions P_{XYZ} where X and Y are random variables over $\{0, 1\}$ with some bounded error $\Pr[X \neq Y]$, and where all that is known about Z is that, with a certain probability, it does not give any information on X .² This class will be important for our applications. Using novel techniques, we give an explicitly computable lower bound on the one-way

¹ This result is proven with respect to a slightly different definition of the secret-key rate than we use. For completeness, we thus provide a new proof for this.

² As the *exact* distribution of the initial randomness—especially the part held by Eve—is usually not known, it is natural to consider such classes.

secret-key rate as well as a tight characterization of the parameters for which one-way secret-key agreement is possible.

1.2 Circuit Polarization

Previous Work: In [17], Sahai and Vadhan introduced the promise problem *statistical difference*. This problem is defined for parameters α and β , $\alpha > \beta$ as follows: given two circuits which, on uniform random input, produce output distributed according to C_0 and C_1 with the promise that the statistical distance of the distributions is either bigger than α or smaller than β , decide which of the two is the case. If $\alpha^2 > \beta$, Sahai and Vadhan show (cf. also [18]) how to *polarize* such a pair of circuits, i.e., they give an efficient construction which takes a pair of circuits and outputs a pair of circuits such that, if the statistical distance of the initial pair was at least α to begin with, the statistical distance of the resulting distributions is very high (i.e., at least $1 - 2^{-k}$ for an arbitrary k), and if the statistical distance of the pair was at most β , then the resulting statistical distance is very small (i.e., at most 2^{-k}).

In order to achieve this only two operations are used, where one of them increases the statistical distance of the distributions at hand and the other reduces the distance. These operations share a certain similarity to operations used in secret-key agreement protocols (cf. [11] and [19]), and indeed, in [5], Dwork et al. note that their construction to immunize public-key encryption is inspired by [17].

Our Contributions: In this work, we make the connection anticipated in [5] explicit by showing that one-way secret-key agreement for the class of distributions given in Section 2.3 is equivalent to the task of circuit polarization, as long as one is restricted to black-box constructions (i.e., the description of the circuits given may not be used), only gives independent and uniform random inputs to the circuits, and directly outputs the samples of the circuits. These restrictions may seem quite strong at first, but the method given in [18] is of this form. Using our bounds for secret-key agreement, we show that such a polarization method *does only exist* if $\alpha^2 > \beta$, i.e., the bounds given in [18] are optimal for this class of constructions.

1.3 Immunization of Public-Key Encryption

Previous Work: Assume that a public-key encryption scheme for single bits is given, which has the property that the receiver may succeed in decrypting correctly only with probability $(1 + \alpha)/2$, and also that a potential eavesdropper Eve may have probability up to $(1 + \beta)/2$ to find the message, for some constants (or functions of a security parameter) α and β . In [5], the question was posed whether such a scheme can be used to get a public-key encryption scheme in the usual sense. Furthermore, the question was answered in the positive sense in two cases: if $\alpha^2 > c\beta$, for some absolute constant $c \gg 1$, a scheme is given. Also, for every constant $\beta < 1$ a construction which works for some constant $\alpha < 1$ is

given. However, this construction is not very strong: for example, for $\beta = 1/2$, the constant α is about $1 - 2^{-15}$. Note that Dwork et al. make no attempt to optimize these constants.

In [7] a similar question was asked for key agreement where Alice and Bob may communicate an arbitrary number of rounds.

Our Contributions: Using a lemma from [7], we improve the result of [5] and show that, for constants α and β , immunizing such an encryption scheme is possible if $\alpha^2 > \beta$. Furthermore we show that, in a setting which is sufficiently black-box, this is optimal.

1.4 Notation

Throughout the paper, we use calligraphic letters (e.g. $\mathcal{X}, \mathcal{Y}, \mathcal{U}$) to denote sets. Uppercase letters (X, Y, U) are used to denote random variables, and lowercase letters denote values of these random variables.

For distributions P_X and $P_{X'}$ over the same domain $\mathcal{X}$, we denote by $\|P_X - P_{X'}\| = \frac{1}{2} \sum_{x \in \mathcal{X}} |P_X(x) - P_{X'}(x)|$ the statistical distance between P_X and $P_{X'}$. If X and X' are the corresponding random variables we sometimes slightly abuse notation and write $\|X - X'\|$ instead.

The min-entropy (or Rényi entropy of order ∞) of a random variable X over $\mathcal{X}$ is defined as $H_\infty(X) := -\log(\max_{x \in \mathcal{X}} P_X(x))$, and the Rényi entropy of order zero is $H_0(X) := \log(|\{x \in \mathcal{X} | P_X(x) > 0\}|)$. More generally, the conditional Rényi entropies are

$$\begin{aligned} H_\infty(X|Y) &:= -\log\left(\max_{x \in \mathcal{X}, y \in \mathcal{Y}} P_{X|Y}(x|y)\right), \\ H_0(X|Y) &:= \log\left(\max_{y \in \mathcal{Y}} |\{x \in \mathcal{X} | P_{X|Y}(x|y) > 0\}|\right). \end{aligned}$$

Additionally, we use the following smoothed versions of these entropy measures [14], which are defined for any $\varepsilon \geq 0$:

$$\begin{aligned} H_\infty^\varepsilon(X) &:= \max_{P_{X'}: \|P_X - P_{X'}\| \leq \varepsilon} H_\infty(X'), \\ H_\infty^\varepsilon(X|Y) &:= \max_{P_{X'|Y'}: \|P_{XY} - P_{X'Y'}\| \leq \varepsilon} H_\infty(X'|Y'). \end{aligned}$$

For a random variable X , we write $U \leftarrow X$ if, for any other random variable Z , $U \leftrightarrow X \leftrightarrow Z$ is a Markov chain. In other words, one can think of U as being obtained from X by sending it through a channel without considering anything else.

2 One-Way Secret-Key Agreement

2.1 Notation and Definitions

A one-way secret-key agreement protocol has three important parameters, which are denoted by the same letters throughout the paper: the length m of the secret

key produced, a security parameter k , and the number n of instances of the initial random variables used. It will be convenient in applications to assume that, for given m and k , n can be computed by a function $n(k, m)$.

Definition 1 (Protocol). A one-way secret-key agreement (OW-SKA) protocol on $\mathcal{X} \times \mathcal{Y}$ consists of the function $n(k, m) : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$; a function family, called Alice, with parameters k and m , mapping n instances of X to a bit string $S_A \in \{0, 1\}^m$ (the secret key) and a bit string $\Gamma \in \{0, 1\}^*$ (the communication); and a function family, called Bob, with parameters k and m , mapping Γ and n instances of Y to a bit string $S_B \in \{0, 1\}^m$. The protocol is efficient if $n(k, m)$, Alice, and Bob can be computed by probabilistic Turing machines in time $\text{poly}(k, m)$. The rate of the protocol is $\lim_{k \rightarrow \infty} \lim_{m \rightarrow \infty} \frac{n(k, m)}{m}$.

The goal of secret-key agreement is to get a secure key (S_A, S_B) , i.e., two strings which are likely to be equal and look like a uniform random string to Eve. We can define this as follows:

Definition 2 (Secure Key). A pair (X, Y) over $\{0, 1\}^m \times \{0, 1\}^m$ of random variables is ε -secure with respect to Z if

$$\|P_{XYZ} - P_{UU} \times P_Z\| \leq \varepsilon,$$

where P_{UU} is the probability distribution over $\{0, 1\}^m \times \{0, 1\}^m$ given by

$$P_{UU}(x, y) = \begin{cases} 2^{-m} & \text{if } x = y \\ 0 & \text{otherwise.} \end{cases}$$

We say that a protocol is secure if it generates a 2^{-k} -secure key with respect to the information Eve has after the protocol execution, that is, the initial randomness $Z_1, \dots, Z_n$ and the communication Γ . In some cases it is desirable to have a protocol which works for a class of distributions rather than for a single distribution (since one may not know the exact distribution of the random variables).

Definition 3 (Secure protocol). A OW-SKA protocol on $\mathcal{X} \times \mathcal{Y}$ is secure on a probability distribution P_{XYZ} over $\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$ if, for any $k, m \in \mathbb{N}$, (S_A, S_B) is 2^{-k} -secure with respect to $(Z_1, \dots, Z_{n(k, m)}, \Gamma)$.

A protocol is secure on a set $\mathcal{P} = \{P_{XYZ}\}$ of tripartite probability distributions if it is secure for every distribution $P_{XYZ} \in \mathcal{P}$.

This way, we can study the secret-key rate of classes of distributions, and also of single distributions.

Definition 4 (One-way secret-key rate). The one-way secret key rate $S_{\rightarrow}(\mathcal{P})$ of a set $\mathcal{P} = \{P_{XYZ}\}$ of probability distributions is the supremum of the rate of any OW-SKA protocol which is secure on $\mathcal{P}$.

We also write $S_{\rightarrow}(X; Y|Z)$ to denote the one-way secret-key rate of a single distribution, i.e., $S_{\rightarrow}(X; Y|Z) := S_{\rightarrow}(\{P_{XYZ}\})$.

2.2 A General Expression for the One-Way Secret-Key Rate

In this section, we derive a simple expression for the one-way secret-key rate of a general tripartite probability distribution. As mentioned in the introduction, Theorem 1 has already been known to hold for general (not necessarily efficient) protocols [3,1].

Theorem 1. *Let P_{XYZ} be a probability distribution. Then*

$$S_{\rightarrow}(X;Y|Z) = \sup_{V \leftarrow U \leftarrow X} H(U|ZV) - H(U|YV).$$

Moreover, the same identity holds if only efficient secret-key agreement protocols are considered.

For the (two-way) secret-key rate no comparable expression is known. We prove Theorem 1 in two steps: We first give an efficient protocol for any rate which is below $\sup_{V \leftarrow U \leftarrow X} H(U|ZV) - H(U|YV)$ (Theorem 2) and then show that no protocol can achieve a higher rate (Theorem 3).

The protocol is based on the following proposition. A proof can be found in [6]; the idea is to concatenate a random linear code with a Reed-Solomon code such that the decoding can be done in polynomial time.

Proposition 1. *For any memoryless channel and any rate s below the capacity it is possible to design codes $\mathcal{C} : \mathcal{X}^\ell \rightarrow \mathcal{X}^n$ of growing length $\ell \rightarrow \infty$ with overall complexity (construction, encoding, and decoding) of order n^2 and decoding error probability $2^{-c_s \cdot n}$ where the constant c_s only depends on the channel and the rate s .*

Furthermore, we use the following from [15]:

Proposition 2. *Let P_{XYZ} be a probability distribution. For any $\varepsilon, \varepsilon' \geq 0$,*

$$\begin{aligned} H_{\infty}^{\varepsilon+\varepsilon'}(X|Y) &\geq H_{\infty}^{\varepsilon}(XY) - H_0(Y) - \log\left(\frac{1}{\varepsilon'}\right) \\ H_{\infty}^{\varepsilon+\varepsilon'}(XY) &\geq H_{\infty}^{\varepsilon}(X) + H_{\infty}^{\varepsilon'}(Y|X). \end{aligned}$$

More generally, the statement still holds if all entropies are conditioned on some additional random variable Z .

Also, we use the following from [8].³

Proposition 3. *Let $(X_1, Y_1), \dots, (X_n, Y_n)$ i.i.d. according to P_{XY} . Then,*

$$H_{\infty}^{\varepsilon}(X_1, \dots, X_n | Y_1, \dots, Y_n) \geq nH(X|Y) - 4\sqrt{n \log(1/\varepsilon)} \log(|\mathcal{X}|).$$

³ Note that a non-quantitative version of this statements follows directly from the asymptotic equipartition property (see, e.g., [4]). A slightly different quantitative version can be found in [9].

Also, we need the left-over hash-lemma, first given in [9] (see also [2]). The function Ext used is a two-universal hash-function.

Proposition 4 (Left-Over Hash-Lemma). *Let X be a random variable over $\{0, 1\}^n$. Let U^n and U^m be independent and uniform over $\{0, 1\}^n$ and $\{0, 1\}^m$, respectively. There exists an efficiently computable function $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ such that, if $H_\infty(X|Z) \geq m + 2 \log(1/\varepsilon)$, then $\|(\text{Ext}(X, U^n), U^n, Z) - (U^m, U^n, Z)\| \leq \varepsilon$.*

Lemma 1. *Let P_{XYZ} be an arbitrary tripartite probability distribution and let $r < H(X|Z) - H(X|Y)$. There exists a constant d_r (depending on P_{XYZ} and r) and an efficient OW-SKA protocol secure on P_{XYZ} such that $n \leq \max(m/r, k \cdot d_r)$.*

Proof. Let γ be such that $r + 3\gamma = H(X|Z) - H(X|Y)$. Let $\oplus$ be an arbitrary group operation over $\mathcal{X}$. For the channel which maps x to a pair $(X \oplus x, Y)$ (by choosing X and Y according to P_{XY} , this channel has capacity $H_0(X) - H(X|Y)$), we use Proposition 1 to get a code $\mathcal{C}$ with rate $s := H_0(X) - H(X|Y) - \gamma$.

Choose n such that $n \geq \frac{m}{r}$, $n \geq \frac{32k \log^2 |\mathcal{X}|}{\gamma^2}$, $n \geq \frac{2k}{c_s}$ and such that there exists a code of this length in the family guaranteed by Proposition 1. From the code of this length, Alice now chooses a random word $C = (C_1, \dots, C_n)$ and sends, for all i , $C_i \oplus X_i$ to Bob, who gets $(Y_i, C_i \oplus X_i)$. Using the property of the code, Bob can find the original codeword C with probability $1 - 2^{-c_s \cdot n} \geq 1 - 2^{-2k}$. Alice then sends a randomly chosen seed of a two-universal hash-function which maps the codeword to a string of length m . Both parties apply the hash-function and output S_A and S_B , respectively.

We show that Eve gets no information with probability 2^{-k} . For this, we set $\varepsilon := 2^{-2k}$. From Proposition 2 and using $H_\infty(C|Z^n) = ns$ (which follows from the fact that the codeword is chosen uniformly at random), we get

$$\begin{aligned} H_\infty^{2\varepsilon}(C|(X^n \oplus C)Z^n) &\geq H_\infty(C|Z^n) + H_\infty^\varepsilon(X^n \oplus C|CZ^n) - H_0(X^n \oplus C|Z^n) - \log\left(\frac{1}{\varepsilon}\right) \\ &= ns + H_\infty^\varepsilon(X^n|Z^n) - nH_0(X) - 2k. \end{aligned}$$

From Proposition 3 we get $H_\infty^\varepsilon(X^n|Z^n) \geq nH(X|Z) - 4 \log(|\mathcal{X}|) \sqrt{2nk}$. Together, we obtain

$$\begin{aligned} H_\infty^{2\varepsilon}(C|(X^n \oplus C)Z^n) &\geq \underbrace{n(H(X|Z) - H(X|Y) - \gamma)}_{=n(r+2\gamma)} - \underbrace{4 \log(|\mathcal{X}|) \sqrt{2nk}}_{=\sqrt{32nk \log^2 |\mathcal{X}|} \leq n\gamma} - 2k \\ &\geq nr + n\gamma - 2k. \end{aligned}$$

From Proposition 4 we see that it is possible to extract a secret key of length $nr + n\gamma - 6k > nr \geq m$ such that Eve gets no information except with probability $2\varepsilon + \varepsilon = 3 \cdot 2^{-2k} \leq 2^{-k}$. $\square$

Theorem 2. Let P_{XYZ} be an arbitrary probability distribution and let r be a constant satisfying $r < \sup_{V \leftarrow U \leftarrow X} H(U|ZV) - H(U|YV)$. There exists a constant d_r and an efficient OW-SKA protocol which is secure on P_{XYZ} and uses at most $\max(m/r, k \cdot d_r)$ instances of the initial random variables.

Proof. For any random variables U and V such that $V \leftarrow U \leftarrow X$, Alice can compute an instance of U and V locally from an instance of X , and then send V over the channel to Bob (and Eve). The result then follows from Lemma 1.

Theorem 3. Let P_{XYZ} be a probability distribution. Then

$$S_{\rightarrow}(X; Y|Z) \leq \sup_{V \leftarrow U \leftarrow X} H(U|ZV) - H(U|YV).$$

Proof (sketch). We show that $\sup_{V \leftarrow U \leftarrow X} H(U|ZV) - H(U|YV)$ does not increase by any step of a one-way key-agreement protocol. More precisely, it does not increase by local processing of either Alice or Bob, or sending a message from Alice to Bob. Furthermore, taking n copies of X , Y , and Z at most multiplies this quantity by n . Finally, if Alice and Bob share a secret key of length m , then this quantity is arbitrarily close to m (depending on k). Hence, the initial quantity is at least m . $\square$

Proof (Theorem 1). From Theorems 2 and 3. $\square$

2.3 The Secret Key Rate of a Class of Binary Distributions

In this section we study the one-way secret-key rate of a general class of distributions. Namely, for parameters α and β , we assume that Alice and Bob are given binary random variables X and Y which have the property that they are equal with probability at least $(1 + \alpha)/2$ (i.e., X and Y have correlation at least α). Furthermore, we assume that with probability $1 - \beta$, the random variable Z does not give any information about X . This class will also be of interest for Sections 3 and 4.

Definition 5. Let $\mathcal{D}(\alpha, \beta)$ be the set of probability distributions P_{XYZ} over $\{0, 1\} \times \{0, 1\} \times \mathcal{Z}$ satisfying

- $\Pr[X = 0] = \Pr[X = 1] = \frac{1}{2}$,
- $\Pr[X = Y] \geq \frac{1+\alpha}{2}$,
- there exists an event $\mathcal{E}$ such that $H(X|Z\mathcal{E}) = 1$ and $\Pr[\mathcal{E}] \geq 1 - \beta$.

It is not hard to see that we could similarly look at the distributions which satisfy $\|P_{Y|X=0} - P_{Y|X=1}\| \geq \alpha$ and $\|P_{Z|X=0} - P_{Z|X=1}\| \leq \beta$, where Y does not have to be binary. This condition implies that Bob can apply a function to Y such that a distribution from $\mathcal{D}(\alpha, \beta)$ results. Furthermore, all distributions in $\mathcal{D}(\alpha, \beta)$ satisfy this characterization.

Some distributions in $\mathcal{D}(\alpha, \beta)$ have a higher secret-key rate than others, of course. We will see that the following distribution has the lowest secret-key

rate of all distributions in $\mathcal{D}(\alpha, \beta)$. Intuitively, this distribution gives as much information to Eve as possible, and makes X and Y as independent as possible under the constraints of Definition 5. For a random variable X , let⁴ $\boxtimes_\lambda(X)$ be the random variable describing the output of a binary symmetric channel taking input X , i.e., $P_{\boxtimes_\lambda(X)|X=0}(0) = P_{\boxtimes_\lambda(X)|X=1}(1) = \frac{1+\lambda}{2}$.

Definition 6. For fixed α, β , the characteristic distribution P_{XYZ} of $\mathcal{D}(\alpha, \beta)$ is given by the following random process: we chose $X \in \{0, 1\}$ uniformly at random. Then, Y is given as $\boxtimes_\alpha(X)$, and Z over $\{0, 1, \perp\}$ is given as the output of an erasure channel with symmetric error probability $1 - \beta$ on input X , i.e., $\Pr[Z = X] = \beta$, independently of X , and $\Pr[Z = \perp] = 1 - \beta$.

We are now ready to formulate our main statement of this section, namely an easily computable expression for $S_{\rightarrow}(\mathcal{D}(\alpha, \beta))$:

Theorem 4. For any α, β , let P_{XYZ} be the characteristic distribution of $\mathcal{D}(\alpha, \beta)$. Then,

$$S_{\rightarrow}(\mathcal{D}(\alpha, \beta)) = \max_{\lambda} H(\boxtimes_\lambda(X)|Z) - H(\boxtimes_\lambda(X)|Y). \quad (1)$$

In particular, if $\alpha^2 > \beta$ then $S_{\rightarrow}(\mathcal{D}(\alpha, \beta)) \geq \frac{1}{7}(\alpha^2 - \beta)^2$ and if $\alpha^2 \leq \beta$ then $S_{\rightarrow}(\mathcal{D}(\alpha, \beta)) = 0$.

Since the term in the maximum of (1) only involves random variables whose distribution is explicitly known (cf. Definition 6) we can get the following form of it (where $h(x)$ is the binary entropy function):

$$\begin{aligned} g_{\alpha, \beta}(\lambda) &:= H(\boxtimes_\lambda(X)|Z) - H(\boxtimes_\lambda(X)|Y) \\ &= (1 - \beta) + \beta h\left(\frac{1 + \lambda}{2}\right) - h\left(\frac{1 + \alpha\lambda}{2}\right) \end{aligned} \quad (2)$$

In order to prove Theorem 4, we need a few properties of $g_{\alpha, \beta}$ (see also Fig. 1). As they can be obtained with standard tools from calculus, the (not very interesting) proof is omitted.

Lemma 2. Let the function $g_{\alpha, \beta} : [-1, 1] \rightarrow \mathbb{R}$ be as in (2). If $\alpha^2 \leq \beta$, then $g_{\alpha, \beta}(\lambda) \leq 0$ for all $\lambda \in [-1, 1]$ and $g_{\alpha, \beta}$ is concave. If $\alpha^2 > \beta$, then $g_{\alpha, \beta}$ has one local minimum at $\lambda = 0$ with $g_{\alpha, \beta}(0) = 0$ and two local maxima at $-\lambda^+$ and λ^+ , $\lambda^+ \in (0, 1]$ with $g_{\alpha, \beta}(-\lambda^+) = g_{\alpha, \beta}(\lambda^+) \geq \frac{1}{7}(\alpha^2 - \beta)^2$. Furthermore, $g_{\alpha, \beta}$ is concave in $[-1, -\lambda^+]$ and $[\lambda^+, 1]$.

We first give an upper bound on $S_{\rightarrow}(X; Y|Z)$ for the distribution from Definition 6.

Lemma 3. Let P_{XYZ} be the characteristic distribution of $\mathcal{D}(\alpha, \beta)$. Then, $S_{\rightarrow}(X; Y|Z) \leq \max_{\lambda} g_{\alpha, \beta}(\lambda)$, where $g_{\alpha, \beta}$ is defined by (2).

⁴ The symbol $\boxtimes$ is supposed to look like a binary symmetric channel, and can be pronounced as *noise*.

Fig. 1. Plot of $g_{\alpha, \beta}(\lambda)$ with $\alpha = 0.8$ and $\beta = 0.59$

Proof. We know that $S_{\rightarrow}(X; Y|Z) = \sup_{V \leftarrow U \leftarrow X} H(U|ZV) - H(U|YV)$ (Theorem 1). Let $P_{U|X}$ and $P_{V|U}$ be fixed channels. It is sufficient to show that $H(U|ZV) - H(U|YV) \leq \max_{\lambda} g_{\alpha, \beta}(\lambda)$.

We can rewrite $H(U|ZV) - H(U|YV)$ as

$$\begin{aligned} H(U|ZV) - H(U|YV) &= H(UZV) - H(UYV) - (H(ZV) - H(YV)) \\ &= H(Z|UV) - H(Y|UV) - (H(Z|V) - H(Y|V)). \end{aligned} \quad (3)$$

Consider now a fixed pair (u, v) . Setting $\frac{1+\lambda_{uv}}{2} := \Pr[X=0|U=u, V=v]$ and $\frac{1+\lambda_v}{2} := \Pr[X=0|V=v]$, a straightforward computation yields:

$$\begin{aligned} H(Z|U=u, V=v) - H(Y|U=u, V=v) &= h(\beta) + \beta h\left(\frac{1+\lambda_{uv}}{2}\right) - h\left(\frac{1+\alpha\lambda_{uv}}{2}\right) \\ H(Z|V=v) - H(Y|V=v) &= h(\beta) + \beta h\left(\frac{1+\lambda_v}{2}\right) - h\left(\frac{1+\alpha\lambda_v}{2}\right). \end{aligned}$$

Because $g_{\alpha, \beta}$ differs from these expressions only by a constant, together with (3) this gives

$$H(U|ZV) - H(U|YV) = \mathbf{E}_{uv}[g_{\alpha, \beta}(\lambda_{uv})] - \mathbf{E}_v[g_{\alpha, \beta}(\lambda_v)].$$

Using $\mathbf{E}_u[\lambda_{uv}] = \lambda_v$, where u is chosen according to the probability distribution $P_{U|V=v}$, we thus obtain

$$H(U|ZV) - H(U|YV) = \mathbf{E}_v \left[\mathbf{E}_u[g_{\alpha, \beta}(\lambda_{uv})] - g_{\alpha, \beta}(\mathbf{E}_u[\lambda_{uv}]) \right].$$

For every fixed v , we can use Lemma 2 to obtain the following upper bound on the term in the expectation:

$$\mathbf{E}_u[g_{\alpha, \beta}(\lambda_{uv})] - g_{\alpha, \beta}(\mathbf{E}_u[\lambda_{uv}]) \leq \max_{\lambda} g_{\alpha, \beta}(\lambda) - g_{\alpha, \beta}(0) = \max_{\lambda} g_{\alpha, \beta}(\lambda),$$

which can now be inserted in the above expression. $\square$

Next, we show that for every distribution in $\mathcal{D}(\alpha, \beta)$ we can achieve at least this rate by sending X over a fixed channel. As we want a protocol which works for *every* distribution in $\mathcal{D}(\alpha, \beta)$, it is important that this processing only depends on the parameters α and β .

Lemma 4. *Let α, β be fixed, $P_{XYZ} \in \mathcal{D}(\alpha, \beta)$, $g_{\alpha, \beta}$ as in (2), $\lambda \in [0, 1]$. Then $H(\mathbb{X}_\lambda(X)|Z) - H(\mathbb{X}_\lambda(X)|Y) \geq g_{\alpha, \beta}(\lambda)$.*

Proof. Using a simple calculation we see that $H(\mathbb{X}_\lambda(X)|Z) \geq (1 - \beta) + \beta h(\frac{1+\lambda}{2})$. To see that $H(\mathbb{X}_\lambda(X)|Y) \leq h(\frac{1+\alpha\lambda}{2})$, let B be a uniform random bit, which is independent of X and Y . Then we obtain $H(\mathbb{X}_\lambda(X)|Y) = H(\mathbb{X}_\lambda(X \oplus B)|Y \oplus B, B) \leq H(\mathbb{X}_\lambda(X \oplus B)|Y \oplus B) = h(\frac{1+\alpha\lambda}{2})$. $\square$

We are now ready to prove Theorem 4.

Proof (Theorem 4). From Theorem 1, Lemmata 2, 3, and 4. $\square$

Furthermore, together with the results of the previous section, we conclude that for any α, β with $\alpha^2 > \beta$ there exists an efficient one-way secret-key agreement protocol secure on $\mathcal{D}(\alpha, \beta)$.

Corollary 1. *Let α, β be constant with $\alpha^2 > \beta$. There exists an efficient one-way secret-key agreement protocol with rate $(\alpha^2 - \beta)^2/8$ which is secure on $\mathcal{D}(\alpha, \beta)$.*

Proof. From Theorems 1 and 4.⁵ $\square$

3 Circuit Polarization

3.1 Polarization and Oblivious Polarization

Circuit polarization was introduced by Sahai and Vadhan in [17] in the context of statistical zero knowledge. It can be described as follows: assume that two circuits are given, which on uniform random input yield output distributions C_0 and C_1 over $\{0, 1\}^\ell$, respectively. We look for an efficient method to *polarize* the circuits: if $\|C_0 - C_1\| \geq \alpha$, for some parameter α , the method should output circuits which are near disjoint, if $\|C_0 - C_1\| \leq \beta$, for some parameter β , then the method should output circuits which produce very close distributions.

In general, such a method uses a description of the circuits given. Here, we focus on methods which use the given circuits in a black-box manner, obviously and with random input only.

Definition 7. *An oblivious polarization method for parameters α and β is a randomized algorithm which, on input k and b , outputs “query bits” $Q_b^1, \dots, Q_b^n$ and a string R_b . For two distributions C_0 and C_1 it satisfies:*

$$\begin{aligned} \|C_0 - C_1\| \geq \alpha &\implies \|(C_{Q_0^1}, \dots, C_{Q_0^n}, R_0) - (C_{Q_1^1}, \dots, C_{Q_1^n}, R_1)\| \geq 1 - 2^{-k} \\ \|C_0 - C_1\| \leq \beta &\implies \|(C_{Q_0^1}, \dots, C_{Q_0^n}, R_0) - (C_{Q_1^1}, \dots, C_{Q_1^n}, R_1)\| \leq 2^{-k}. \end{aligned}$$

The method is efficient if the algorithm runs in time polynomial in k .

⁵ Technically speaking, Theorem 1 only guarantees that such a protocol exists for one single distribution, and in general the protocol *will* depend on the distribution at hand. Of course the protocol cannot depend on the distribution of $P_{Z|XY}$, but the distribution $P_{Y|X}$ can vary in $\mathcal{D}(\alpha, \beta)$, so we have to be careful. However, since the protocol just uses an error correcting code which is too strong for some distributions, it is easy to see that this is not a problem.

Note that the method given in [18] to polarize circuits is oblivious in this sense.⁶ The method given to *invert* the statistical distance is not oblivious (and cannot possibly be).

3.2 Equivalence of Polarization and Secret-Key Agreement

The goal of this section is to prove that an oblivious polarization method for parameters α and β is equivalent to a secret-key agreement protocol (for a one bit key) secure on $\mathcal{D}(\alpha, \beta)$, as defined in Section 2.3.

Theorem 5. *There exists an oblivious polarization method for parameters α and β if and only if there exists a one-way secret-key agreement protocol secure on $\mathcal{D}(\alpha, \beta)$. Moreover, there exists an efficient oblivious polarization method if and only if there exists a protocol with efficient encoding (i.e., Alice is efficient).*

We prove Theorem 5 in both directions separately, and start by showing that a polarization method implies the existence of a one-way secret-key agreement protocol:

Lemma 5. *Let an oblivious polarization method for parameters α, β be given. Then there exists a one-way secret-key agreement protocol which is secure on $\mathcal{D}(\alpha, \beta)$. Furthermore, if the polarization method is efficient, then Alice is efficient.*

Proof. It is sufficient to show how to get a one-way secret-key agreement protocol for $m := 1$ bit.

The number of random variables $n := n(k, 1)$ the protocol uses is set to the number of queries produced by the polarization method. Alice first simulates the polarization method with input k and a uniform random bit b which yields R_b and $Q_b^1, \dots, Q_b^n$. Subsequently, Alice sends R_b as well as $(X_1 \oplus Q_b^1, \dots, X_n \oplus Q_b^n)$ as communication to Bob, and outputs b as secret bit.

We show that Bob can find b with high probability from the communication and Y^n (this may not necessarily be efficient). Since $P_{XYZ} \in \mathcal{D}(\alpha, \beta)$ the random variables $C_0 := (X, Y)$ and $C_1 := (1 \oplus X, Y)$ satisfy $\|C_0 - C_1\| \geq \alpha$. Furthermore, $Y_1, \dots, Y_n$ and the communication gives Bob a sample of the distribution $(C_{Q_b^1}, \dots, C_{Q_b^n}, R_b)$. The definition of the polarization method now implies that a statistical test can find b except with probability exponentially small in k .

Also the protocol is secure against Eve: consider the random variable $D_0 := (Z, X)$ and the random variable $D_1 := (Z, X \oplus 1)$. Here, $P_{XYZ} \in \mathcal{D}(\alpha, \beta)$ implies $\|D_0 - D_1\| \leq \beta$, and Eve sees exactly a sample of $(D_{Q_b^1}, \dots, D_{Q_b^n}, R_b)$, which is independent of b except with probability exponentially small in k . $\square$

On the other hand, a one-way secret-key agreement protocol yields a polarization method:

⁶ In fact, R_0 and R_1 are empty in the method given.

Lemma 6. *Let a one-way secret-key agreement protocol secure on $\mathcal{D}(\alpha, \beta)$ be given. Then, there exists an oblivious polarization method for parameters α and β using $n(k, 1)$ copies of the given distribution. Furthermore, if Alice is efficient, then the polarization method is efficient.*

Proof. Throughout the proof we only need key agreement for one key bit and set $m := 1$. On input b and k , the polarization method first chooses random (uniform and independent) queries $Q_b^1, \dots, Q_b^n$. Then Alice is simulated with random variables $X_1 := Q_b^1, \dots, X_n := Q_b^n$, which yields communication C , and a secret bit S . The string R_b is then defined as $R_b := (C, S \oplus b)$.

We first show that $\|C_0 - C_1\| \geq \alpha$ implies that $\|(C_{Q_0^1}, \dots, C_{Q_0^n}, R_0) - (C_{Q_1^1}, \dots, C_{Q_1^n}, R_1)\|$ is exponentially close to 1. For this, it is enough to show how to find b from $(C_{Q_b^1}, \dots, C_{Q_b^n}, R_b)$ with probability almost 1. $\|C_0 - C_1\| \geq \alpha$ implies that there exists a function y (a statistical test) such that setting $Y_i := y(C_{Q_b^i})$ gives $\Pr[Y_i = Q_b^i] \geq \frac{1+\alpha}{2}$. Thus we can use the decoding algorithm Bob of the secret key agreement protocol to reconstruct S with very high probability. Since $S \oplus b$ is also given, we can find b .

Now assume that $\|C_0 - C_1\| \leq \beta$. Consider the tripartite probability distribution P_{XYZ} where $X = Y$ is a uniform random bit, and $Z = C_X$. It is not hard to see that $P_{XYZ} \in \mathcal{D}(\alpha, \beta)$. Thus, in the one-way secret-key agreement protocol (using this distribution) Eve will see exactly a sample of $(C_{Q_b^1}, \dots, C_{Q_b^n}, R_b)$ and the value $S \oplus b$. The properties of the protocol imply that this distribution is statistically independent (with high probability) of S . Furthermore, in the construction above only $S \oplus b$ depends on b , which implies the lemma. $\square$

Proof (Theorem 5). Follows from Lemmata 5 and 6.

Furthermore, since we know for which parameters α and β a protocol exists, we get:

Corollary 2. *There exists an (efficient) oblivious black-box polarization method for constant parameters α and β if and only if $\alpha^2 > \beta$.*

Proof. Using Theorem 4 and Corollary 1. Additionally, we observe that if $S_-(X; Y|Z) = 0$ then no one-way secret-key agreement protocol can exist, since one could use it to get a positive rate. $\square$

As mentioned before such a polarization method was already given in [18]. However, it was unknown that this is tight for oblivious methods.

3.3 Further Improvements

Note that instead of using the code as guaranteed in Proposition 1, we could have used a random linear code in this application (where the code is chosen by Alice and a description is sent as communication). In this case, the resulting polarization method is very efficient, as only $k \cdot \text{poly}((\alpha^2 - \beta)^{-1})$ copies of the circuits are needed. If this method is used in a statistical zero-knowledge proof

system however, the prover needs additional power since he needs to decode a random linear code.

Finally, a statistical zero knowledge proof for the promise problem statistical difference (with parameters α and β) can be realized as follows: the two given circuits are sampled obliviously and uniformly at random by the verifier, sending the samples to the prover. The information *which* circuit was sampled is used as random variables $X_1, \dots, X_n$ in a one-way secret-key agreement protocol, whose communication is also sent to the prover. Now, if the given instance produces distributions with statistical distance at least α , then the prover gets the same information as Bob does, and he can prove this to the verifier by sending back the secret key. If the circuits produce distributions with statistical distance at most β , the prover gets the same information as Eve does, and cannot find the secret key. Thus, it can be useful to use protocols which yield more than one secret bit, as this immediately reduces the error of the zero-knowledge proof.

4 Immunizing Bit Encryption Schemes

In this section we study the implications of our work on the task of immunizing bit encryption schemes. Thus, we assume that a public-key encryption scheme for bits is given, which has a certain probability of being correct, and a certain security.

Definition 8. A $(\alpha(k), \beta(k))$ -secure public-key bit encryption scheme is a triple (G, E, D) of probabilistic polynomial time algorithms such that

- Algorithm G , on input 1^k produces a pair $(\mathbf{pk}, \mathbf{sk})$.
- For a random bit $b \in \{0, 1\}$, $\Pr[D_{\mathbf{sk}}(E_{\mathbf{pk}}(b)) = b] > \frac{1+\alpha(k)}{2}$, where the probability is over the randomness of G (giving the pair $(\mathbf{pk}, \mathbf{sk})$), E , D , and the choice of b .
- For any polynomial time algorithm A , and a uniform random bit b : $\Pr[A(\mathbf{pk}, E_{\mathbf{pk}}(b)) = b] < \frac{1+\beta(k)}{2}$, where the probability is over the randomness of A , G , E , and the choice of b .

If such a scheme is (α, β) -secure for every function $1 - \alpha = \beta \in \frac{1}{\text{poly}(k)}$, we say that it is a secure public-key encryption scheme.

We can combine information-theoretic and computational protocols to obtain the following:

Lemma 7. Let $\alpha, \beta : \mathbb{N} \rightarrow [0, 1]$ be noticeable and computable in time $\text{poly}(k)$. Let (G, E, D) be a (α, β) -secure public-key bit encryption scheme. If there exists an efficient one-way secret-key agreement protocol secure on $\mathcal{D}(\alpha, \beta)$, then there exists a secure public-key encryption scheme (G', E', D') .

The proof of this is very similar to the corresponding Lemma in [7]. Due to space constraints it is only sketched here.

Proof (sketch). We only need one key bit, and therefore we will set $m = 1$ throughout the proof. On input 1^k , algorithm G' then does $n(k, 1)$ invocations of algorithm G with input 1^k . This gives a key pair $(\mathbf{pk}', \mathbf{sk}')$, such that both the public- and the secret-key are n -tuples $\mathbf{pk}' = (\mathbf{pk}'_1, \dots, \mathbf{pk}'_n)$ and $\mathbf{sk}' = (\mathbf{sk}'_1, \dots, \mathbf{sk}'_n)$.

To encrypt a bit b with public key $\mathbf{pk}'$, Alice first encrypts n random bits $X_1, \dots, X_n$ with the underlying scheme, i.e., X_i is encrypted with $E_{\mathbf{pk}'_i}$. It then uses the information-theoretic one-way secret-key agreement protocol, where the X_i are used as random variables. Let S_A be the resulting secret bit. The output of algorithm E' is then the encryption of $X_1, \dots, X_n$, the communication of the information-theoretic one-way secret key agreement protocol, and $S_A \oplus b$.

It is easy to see that the communication together with the secret key suffices to decode the encrypted bit. Furthermore, the security of the protocol can be shown using a standard hybrid argument together with the uniform hard-core lemma given in [7] (see also [10]). □

Lemma 7 together with Corollary 1 implies that a (α, β) -secure public-key cryptosystem can be used to get a secure public-key cryptosystem if $\alpha^2 > \beta$. For a limited class of reductions this is tight: a *strong black-box reduction* is a black-box reduction which allows Alice and Bob to use such a cryptosystem only in a way such that it can be modeled by an oracle where Alice and Bob obtain random bits X and Y , respectively, and an attacking algorithm obtains information Z .⁷

Theorem 6. *Let α and β be constants. There exists a strong black-box reduction from a (α, β) -secure public-key cryptosystem to a secure public-key cryptosystem if and only if $\alpha^2 > \beta$.*

Proof. If $\alpha^2 > \beta$, this is implied by Lemma 7 and Corollary 1.

Assume now that $\alpha^2 \leq \beta$, and assume that a reduction is given. It is easy to see that for suitably chosen random variables X , Y and Z an attacker can break every protocol in polynomial space from the information given. Consequently, by giving the attacker access to a PSPACE-complete oracle we can obtain a contradiction. □

References

1. Rudolph Ahlswede and Imre Csiszàr. Common randomness in information theory and cryptography—part I: Secret sharing. *IEEE Transactions on Information Theory*, 39(4):1121–1132, 1993.
2. Charles H. Bennett, Gilles Brassard, Claude Crépeau, and Ueli Maurer. Generalized privacy amplification. *IEEE Transaction on Information Theory*, 41(6):1915–1923, 1995.

⁷ As an example, this excludes the possibility of using the (α, β) -secure cryptosystem to obtain a one-way function.

3. Imre Csiszár and János Körner. Broadcast channels with confidential messages. *IEEE Transactions on Information Theory*, 22(6):644–654, 1978.
4. Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., first edition, 1991. ISBN 0-471-06259-6.
5. Cynthia Dwork, Moni Naor, and Omer Reingold. Immunizing encryption schemes from decryption errors. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 342–360, 2004.
6. Ilya I. Dumer. Concatenated codes and their multilevel generalizations. In V. S. Pless and W. C. Huffman, editors, *The Handbook of Coding Theory*, volume 2, chapter 23, pages 1191–1988. North-Holland, Elsevier, 1998.
7. Thomas Holenstein. Key agreement from weak bit agreement. In *Proceedings of the 37th STOC*, pages 664–673, 2005.
8. Thomas Holenstein and Renato Renner. On the smooth Rényi entropy of independently repeated random experiments. manuscript, 2005.
9. Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstract). In *Proceedings of the 21st STOC*, pages 12–24, 1989.
10. Russell Impagliazzo. Hard-core distributions for somewhat hard problems. In *36th FOCS*, pages 538–545, 1995.
11. Ueli Maurer. Secret key agreement by public discussion. *IEEE Transaction on Information Theory*, 39(3):733–742, 1993.
12. Ueli Maurer and Stefan Wolf. Unconditionally secure key agreement and the intrinsic conditional information. *IEEE Transaction on Information Theory*, 45(2):499–514, 1999.
13. Renato Renner and Stefan Wolf. New bounds in secret-key agreement: The gap between formation and secrecy extraction. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 562–577, 2003.
14. Renato Renner and Stefan Wolf. Smooth Rényi entropy and applications. In *Proceedings of 2004 IEEE International Symposium on Information Theory*, page 233. IEEE, 2004.
15. Renato Renner and Stefan Wolf. Simple and tight bounds for information reconciliation and privacy amplification. Manuscript, 2005.
16. Claude E. Shannon. Communication theory of secrecy systems. *Bell Systems Technical Journal*, 28:656–715, 1949.
17. Amit Sahai and Salil Vadhan. A complete promise problem for statistical zero-knowledge. In *The 38th FOCS*, pages 448–457, 1997.
18. Amit Sahai and Salil Vadhan. Manipulating statistical difference. In Panos Pardalos, Sanguthevar Rajasekaran, and José Rolim, editors, *DIMACS Series*, volume 43, pages 251–270, 1999.
19. Stefan Wolf. *Information-Theoretically and Computationally Secure Key Agreement in Cryptography*. PhD thesis, ETH Zürich, 1999.
20. Aaron D. Wyner. The wire-tap channel. *Bell Systems Technical Journal*, 54:1355–1387, 1975.