Trade-Offs in Information-Theoretic Multi-Party One-Way Key Agreement

Renato Renner[1] Stefan Wolf[2] Jürg Wullschleger[2]
[1] Centre for Quantum Computation, University of Cambridge, UK
E-mail: r.renner@damtp.cam.ac.uk
[2] Computer Science Department, ETH Zürich, Switzerland
E-mail: {wolf,wjuerg}@inf.ethz.ch

Abstract

We consider the following scenario involving three honest parties, Alice, Bob, and Carol, as well as an adversary, Eve. Each party has access to a single piece of information, jointly distributed according to some distribution P. Additionally, authentic public communication is possible from Alice to Carol and from Bob to Carol. Their goal is to establish two information-theoretically secret keys, one known to Alice and Carol, and one known to Bob and Carol. We derive joint bounds on the lengths of these keys. Our protocols combine distributed variants of Slepian-Wolf coding and the leftover hash lemma. The obtained bounds are expressed in terms of smooth Rényi entropies and show that these quantities are useful in this—single-serving—context as well.

1 Introduction

Consider the following scenario: Three parties, Alice, Bob, and Carol, as well as an adversary, Eve, each have access to a *single* realization of random variables X, Y, W, and Z, respectively, jointly distributed according to P_{XYWZ}. Furthermore, both Alice and Bob can send messages to Carol, but no other communication is possible between the parties (in particular, Alice and Bob cannot communicate). The goal of Alice, Bob, and Carol is to generate two *secret keys*, one of them known to Alice and Carol, and the other known to Bob and Carol. Secrecy means that Eve, who is assumed to have access to the entire communication between the parties, has almost no information about the two keys.

Our scenario is an extension of the two-party settings considered in [AC93, CK78, Mau93, RW05], and also partly fits into the general framework on information-theoretic key agreement with a helper proposed in [CN04]. An important distinction to the treatment in [AC93, CK78, Mau93, CN04], however, is that—similarly to [RW05]—we are concerned with the *single-serving* case, where only *single* realizations—in contrast to infinitely many independent and identically distributed (i.i.d.) realizations—of the random variables are available.

In a nutshell, our result shows that there is a direct trade-off between the lengths of the two resulting keys. This effect is a direct consequence of a novel single-serving version of the well-known *Slepian-Wolf coding* result [SW73], that acts as a distributed *information-reconciliation* protocol. The second ingredient of our key generation scheme is a multi-party version of *privacy amplification*, which essentially uses a distributed variant of the leftover hash lemma [ILL89] first proposed in [Wul07].

The quantitative results are expressed in terms of *(conditional) smooth Rényi entropies*. These entropy measures have been introduced as generalizations of the Shannon entropy in order to deal with single-serving scenarios [RW05] (see Section 2).[1] For the special case of i.i.d. distributions (i.e., n-fold product distributions $P^{\times n}$), smooth Rényi entropies asymptotically approach Shannon entropy. In particular, if the distribution P_{XYWZ} describing our scenario is of the form $P_{X'Y'W'Z'}^{\otimes n}$, for some large n, then the smooth Rényi entropies in our results can be replaced by the corresponding Shannon entropies (thus reproducing i.i.d. results as in [CN04]).

2 Smooth Min- and Max-Entropies

2.1 Motivation

Traditionally, operational quantities in information theory, i.e., quantities describing information-theoretic tasks such as channel coding, are defined asymptotically. More precisely, it is typically assumed that a certain functionality, e.g., a (memoryless) communication channel $P_{Y|X}$, can be invoked *many times independently*. The functionality is then characterized in terms of *asymptotic rates*. For example, the *capacity* $C^{\mathrm{asym}}(P_{Y|X})$ of a channel $P_{Y|X}$ is defined as the maximum rate at which bits can be transmitted per channel use such that the probability of a decoding error vanishes asymptotically as the number of channel uses approaches infinity. As shown in [Sha48], $C^{\mathrm{asym}}(P_{Y|X})$ can be expressed in terms of Shannon entropy,

$$C^{\mathrm{asym}}(P_{Y|X}) = \max_{P_X}(I(X;Y)) = \max_{P_X}(H(X) - H(X|Y)) \ . \tag{1}$$

[1] See also [Ren05] for a quantum information-theoretic version of smooth Rényi entropies.

Another example, situated in the area of cryptography, is *key agreement from correlated information* [AC93, CK78, Mau93]. Assume that two parties, Alice and Bob, as well as an adversary, Eve, have access to a source providing them with random variables X, Y, and Z, respectively. The goal of Alice and Bob is to generate a secret key, using only communication over an authentic, but otherwise fully insecure, communication channel. Under the assumption that the source emits *many independent* triples (X, Y, Z), the *key rate* $K^{\text{asym}}(P_{XYZ})$, i.e., the asymptotic rate at which key bits can be generated per invocation of the source, is bounded by an expression which only involves Shannon entropy,

$$K^{\text{asym}}(P_{XYZ}) \geq H(X|Z) - H(X|Y) \ . \tag{2}$$

In a realistic scenario, however, such an asymptotic viewpoint might not be fully satisfying. Firstly, any realistic device can only be accessed a finite number of times; this number might be smaller than the (usually unknown) minimum threshold which is needed for the asymptotic results to apply. Secondly, and even more importantly, the assumption of *independence* might not hold or, at least, be hard to justify. For instance, in cryptography, such an assumption typically translates to a condition on the behavior of the adversary.[2] Results depending on such assumptions are thus usually not sufficient for realistic applications. It is, therefore, natural to ask what happens *if the assumptions of independence and asymptoticity are dropped.* Ideally, one might want to completely drop the assumption that a resource is invoked many times. Smooth Rényi entropies are designed to deal with such general *single-serving* settings.

Recently, a variety of information-theoretic results have been generalized to the single-serving case. For instance, it has been shown in [RWW06] that the number $C^\varepsilon(P_{Y|X})$ of bits that can be sent by *one single use* of a communication channel $P_{Y|X}$ such that a decoding error occurs with probability at most ε is given in terms of the smooth Rényi entropy,

$$C^\varepsilon(P_{Y|X}) \approx \max_{P_X}\bigl(H^\varepsilon_{\min}(X) - H^\varepsilon_{\max}(X|Y)\bigr) \ . \tag{3}$$

Similarly, the number $K^\varepsilon(P_{XYZ})$ of ε-secure key bits that Alice and Bob can generate in the above described scenario is bounded by

$$K^\varepsilon(P_{XYZ}) \gtrapprox H^\varepsilon_{\min}(X|Z) - H^\varepsilon_{\max}(X|Y) \ . \tag{4}$$

[2] E.g., in quantum key distribution, perfect independence of the distributed data is only guaranteed if the adversary does not introduce any correlations during her attack on the quantum channel.

As indicated above, Shannon entropy can be seen as a special case of smooth Rényi entropy. In fact, any result involving smooth Rényi entropies can be specialised to a result for Shannon entropy by virtue of the relation

$$H(X|Y) = \lim_{\varepsilon \to 0} \lim_{n \to \infty} \frac{1}{n} H^{\varepsilon}_{\min}(X^n|Y^n) = \lim_{\varepsilon \to 0} \lim_{n \to \infty} \frac{1}{n} H^{\varepsilon}_{\max}(X^n|Y^n) , \qquad (5)$$

where (X^n, Y^n) denotes n independent pairs of random variables jointly distributed according to P_{XY}. For example, using this identity, it is easy to see that expressions (1) and (2) are special cases of (3) and (4), respectively.

2.2 Definition and Properties

Let X be a random variable with distribution P_X. The *max-entropy* of X is defined as the (binary) logarithm of the size of the support of P_X, i.e.,

$$H_{\max}(X) = \log \bigl|\{x \in \mathcal{X} : P_X(x) > 0\}\bigr| .$$

Similarly, the *min-entropy* of X is given by the negative logarithm of the maximum probability of P_X, i.e.,

$$H_{\min}(X) = -\log \max_x (P_X(x) .$$

Note that $H_{\min}(X) \leq H(X) \leq H_{\max}(X)$, i.e., the min- and max-entropies are lower and upper bounds for Shannon entropy (and also for any Rényi entropy H_α), respectively.

For random variables X and Y with joint distribution P_{XY}, the "conditional" versions of these entropic quantities are defined by

$$H_{\max}(X|Y) = \max_y H_{\max}(X|Y = y) ,$$
$$H_{\min}(X|Y) = \min_y H_{\min}(X|Y = y) ,$$

where $H_{\max}(X|Y = y)$ (and $H_{\min}(X|Y = y)$) denote the max-entropy (min-entropy) of a random variable distributed according to the conditional distribution $P_{X|Y=y}$.

In [RW05], max- and min-entropies have been generalized to so-called *smooth max- and min-entropies.* Smooth max- and min-entropies can be expressed in terms of an optimization over events $\mathcal{E}$ with probability at least $1 - \varepsilon$. Let $P_{X\mathcal{E}|Y=y}(x)$ be the probability that $\{X = x\}$ *and* $\mathcal{E}$ occur, conditioned on $Y = y$. We then have

$$H^{\varepsilon}_{\max}(X|Y) = \min_{\mathcal{E}:\Pr(\mathcal{E}) \geq 1-\varepsilon} \max_y \log |\{x : P_{X\mathcal{E}|Y=y}(x) > 0\}| ,$$
$$H^{\varepsilon}_{\min}(X|Y) = \max_{\mathcal{E}:\Pr(\mathcal{E}) \geq 1-\varepsilon} \min_y \min_x (-\log P_{X\mathcal{E}|Y=y}(x)) .$$

Smooth entropies have properties similar to Shannon entropy [RW05].[3] For example, the *chain rule* $H(X|Y) = H(XY) - H(Y)$ translates to[4]

$$
\begin{aligned}
H_{\max}^{\varepsilon+\varepsilon'}(XY) - H_{\max}^{\varepsilon'}(Y) \leq H_{\max}^{\varepsilon}(X|Y) \ , \\
\leq H_{\max}^{\varepsilon_1}(XY) - H_{\min}^{\varepsilon_2}(Y) + \log(1/(\varepsilon - \varepsilon_1 - \varepsilon_2))
\end{aligned}
$$

and

$$
\begin{aligned}
H_{\min}^{\varepsilon_1}(XY) - H_{\max}^{\varepsilon_2}(Y) - \log(1/(\varepsilon - \varepsilon_1 - \varepsilon_2)) \\
\leq H_{\min}^{\varepsilon}(X|Y) \leq H_{\min}^{\varepsilon+\varepsilon'}(XY) - H_{\min}^{\varepsilon'}(Y) \ ,
\end{aligned}
$$

for any $\varepsilon, \varepsilon', \varepsilon_1, \varepsilon_2 > 0$.

Note that these rules also hold conditioned on an additional random variable Z. For instance, we have

$$
H_{\min}^{\varepsilon_1}(XY|Z) - H_{\max}^{\varepsilon_2}(Y|Z) - \log(1/(\varepsilon - \varepsilon_1 - \varepsilon_2)) \leq H_{\min}^{\varepsilon}(X|YZ) \ . \quad (6)
$$

Because $H_{\max}^{\varepsilon_2}(Y|Z) \leq \log|\mathcal{Y}|$, where $\mathcal{Y}$ denotes the alphabet of Y, this inequality is a generalization of the fact that, by conditioning on an additional random variable Y, with high probability, the min-entropy decreases by at most the logarithm of the alphabet size of Y [Cac97, MW97].

2.3 Operational Interpretation

In [SW73] it was shown that the rate at which many independent realizations of X can be compressed is asymptotically equal to $H(X|Y)$ if the decoder is provided with side-information Y. It is easy to see that $H(X|Y)$ can also be interpreted as the rate at which uniform randomness can be extracted from X in such a way that it is independent of Y. In [RW05], these operational interpretations of the Shannon entropy have been generalized to the *single-serving* case, i.e., it was shown that the smooth entropies $H_{\max}^{\varepsilon}$ and $H_{\min}^{\varepsilon}$ quantify compression and randomness extraction, respectively. More precisely, let $H_{\mathrm{comp}}^{\varepsilon}(X|Y)$ be the length of a bit string needed to store *one* instance of X such that X can later be recovered with an error of at most ε using this string and Y. This quantity is then roughly equal to $H_{\max}^{\varepsilon}$, i.e.,

$$
H_{\max}^{\varepsilon}(X|Y) \leq H_{\mathrm{comp}}^{\varepsilon}(X|Y) \leq H_{\max}^{\varepsilon'}(X|Y) + \log(1/(\varepsilon - \varepsilon')) \ .
$$

[3] This is in contrast the the usual, "non-smooth" min- and max-entropies which have many counterintuitive properties that make them less useful in many contexts.

[4] Note that, because of (5), the chain rule $H(X|Y) = H(XY) - H(Y)$ for Shannon entropy can be seen as a special case of the chain rules for smooth max- and min-entropies.

Similarly, let $H^{\varepsilon}_{\mathrm{ext}}(X|Y)$ be the maximum length of a string that can be computed from X, such that this string is uniformly distributed and independent of Y, with an error of at most ε. We then have

$$H^{\varepsilon'}_{\min}(X|Y) - 2\log(1/(\varepsilon - \varepsilon')) \le H^{\varepsilon}_{\mathrm{ext}}(X|Y) \le H^{\varepsilon}_{\min}(X|Y) \ .$$

3 Distributed Randomness Extraction

The *statistical distance* of two random variables X and Y (or two distribution P_X and P_Y) over the same alphabet $\mathcal{U}$ is defined as

$$\Delta(X,Y) := \frac{1}{2}\sum_{u\in\mathcal{U}} \Big| \Pr[X=u] - \Pr[Y=u] \Big| \ .$$

We say that a random variable X *over $\mathcal{X}$ is ε-close to uniform with respect to Y* if $\Delta(P_{XY}, P_U \times P_Y) \le \varepsilon$, where P_U is the uniform distribution over $\mathcal{X}$.

A function $h : \mathcal{S} \times \mathcal{X} \to \{0,1\}^m$ is called a 2*-universal hash function* [CW79], if for all $x_0 \neq x_1 \in \mathcal{X}$ and for S uniform over $\mathcal{S}$, we have

$$\Pr[h(S,x_0) = h(S,x_1)] \le 2^{-m} \ .$$

Lemma 1, first stated in [ILL89] (see also [BBR88]), gives us a bound on the amount of randomness that can be extracted from a random variable X (which might depend on another random variable Y) such that the extracted randomness is (almost) uniform (and independent of Y). It has a wide range of applications, e.g., in cryptology, it can directly be used for *privacy amplification* [BBR88, BBCM95].

Lemma 1 (Leftover hash lemma [ILL89])**.** *Let $\varepsilon > 0$ and let $h : \mathcal{S} \times \mathcal{X} \to \{0,1\}^m$ be a* 2*-universal hash function. For any random variable X over $\mathcal{X}$ satisfying*

$$H_{\min}(X \mid Z) \ge m + 2\log(1/\varepsilon) \ ,$$

and for S uniform over $\mathcal{X}$ and independent of X, the value $h(S,X)$ is ε-close to uniform with respect to (S,Z).

A *distributed* version of the leftover hash lemma has recently been introduced in [Wul07]. It can be applied in the setting where two players independently extract randomness from two *dependent* random variables. Lemma 1 implies that if the lengths of the extracted strings are smaller than the min-entropy of these random variables, then each of them is close to uniform. However, the two strings might still be dependent on each other. Lemma 2 below states that if, in addition, the sum of the lengths of the extracted strings is smaller than the overall min-entropy, then they are almost independent of each other. The obtained bound is optimal.

Lemma 2 (Distributed leftover hash lemma [Wul07])**.** *Let* $\varepsilon > 0$ *and let* $g : \mathcal{S} \times \mathcal{X} \to \{0,1\}^m$ *and* $h : \mathcal{R} \times \mathcal{Y} \to \{0,1\}^n$ *be 2-universal hash functions. For any random variables* X *and* Y *over* $\mathcal{X}$ *and* $\mathcal{Y}$*, respectively, such that*

$$\begin{aligned} H_{\min}(X \mid Z) &\geq m + 2\log(1/\varepsilon) \ , \\ H_{\min}(Y \mid Z) &\geq n + 2\log(1/\varepsilon) \ , \\ H_{\min}(XY \mid Z) &\geq m + n + 2\log(1/\varepsilon) \ , \end{aligned}$$

and for (S, R) *uniform over* $\mathcal{S} \times \mathcal{R}$ *and independent of* (X, Y)*, the pair* $(g(S, X), h(R, Y))$ *is* ε*-close to uniform with respect to* (S, R, Z)*.*

The proof of Lemma 2 is very similar to the proof of the leftover hash lemma (Lemma 1).

Given this lemma, it is easy to generalise it to smooth min-entropy, using simply the definition above and the union bound.

Lemma 3 ("Smoothed" distributed leftover hash lemma)**.** *Let* $\varepsilon > 0$*,* $\varepsilon' \geq 0$*, and let* $g : \mathcal{S} \times \mathcal{X} \to \{0,1\}^m$ *and* $h : \mathcal{R} \times \mathcal{Y} \to \{0,1\}^n$ *be 2-universal hash functions. For any random variables* X *and* Y *over* $\mathcal{X}$ *and* $\mathcal{Y}$*, respectively, such that*

$$\begin{aligned} H^{\varepsilon'}_{\min}(X \mid Z) &\geq m + 2\log(1/\varepsilon) \ , \\ H^{\varepsilon'}_{\min}(Y \mid Z) &\geq n + 2\log(1/\varepsilon) \ , \\ H^{\varepsilon'}_{\min}(XY \mid Z) &\geq m + n + 2\log(1/\varepsilon) \ , \end{aligned}$$

and for (S, R) *uniform over* $\mathcal{S} \times \mathcal{R}$ *and independent of* (X, Y)*, the pair* $(g(S, X), h(R, Y))$ *is* $(\varepsilon + 3\varepsilon')$*-close to uniform with respect to* (S, R, Z)*.*

4 Distributed Compression

Lemma 4 is the single-serving variant of the famous *Slepian-Wolf encoding* [SW73]. The proof is very similar to the proof in [Cov75]. In cryptography, it can be used for so-called *information reconciliation* [BS94]. Unfortunately, the decoding in our schemes is generally not efficient.

Lemma 4 (Single-serving Slepian-Wolf encoding)**.** *Let* $\varepsilon > 0$*,* $\varepsilon' \geq 0$*, and let* X*,* Y*, and* Z *be random variables over* $\mathcal{X}$*,* $\mathcal{Y}$*, and* $\mathcal{Z}$*, respectively. Let* $\mathrm{encx} : \mathcal{R} \times \mathcal{X} \to \{0,1\}^m$ *and* $\mathrm{ency} : \mathcal{S} \times \mathcal{Y} \to \{0,1\}^n$ *be 2-universal hash functions, where*

$$\begin{aligned} m &\geq H^{\varepsilon'}_{\max}(X \mid YZ) + \log(3/\varepsilon) \ , \\ n &\geq H^{\varepsilon'}_{\max}(Y \mid XZ) + \log(3/\varepsilon) \ , \\ m + n &\geq H^{\varepsilon'}_{\max}(XY \mid Z) + \log(3/\varepsilon) \ , \end{aligned}$$

and let (R, S) *be uniform over* $\mathcal{R} \times \mathcal{S}$*. There exists a function* $\mathrm{dec} : \mathcal{R} \times \mathcal{S} \times \{0,1\}^m \times \{0,1\}^n \times \mathcal{Z} \to \mathcal{X} \times \mathcal{Y}$*, such that*

$$\Pr[\mathrm{dec}(R, S, \mathrm{encx}(R, X), \mathrm{ency}(S, Y), Z) \neq (X, Y)] \leq \varepsilon + 3\varepsilon' \ .$$

Proof. Assume first that $\varepsilon' = 0$. Every $x \in \mathcal{X}$ is mapped to any value $e_x \in \{0,1\}^m$ with probability 2^{-m}, and every $y \in \mathcal{Y}$ is mapped to any value $e_y \in \{0,1\}^n$ with probability 2^{-n}. Given a pair (x, y), the corresponding encoding (e_x, e_y) can be decoded if there does not exist another pair (x', y') which is encoded to (e_x, e_y) as well. For every y, there are at most $2^{H_{\max}(X|YZ)}$ different possible values for x'. Hence, the probability that there exists $x' \neq x$ such that (x', y) is mapped to the pair (e_x, e_y) is at most $2^{-m}2^{H_{\max}(X|YZ)} \leq \varepsilon/3$. Similarly, the probability that there exists $y' \neq y$ such that (x, y') is mapped to the pair (e_x, e_y) is at most $2^{-n}2^{H_{\max}(Y|XZ)} \leq \varepsilon/3$. The probability that a pair (x', y') with $x' \neq x$ and $y' \neq y$ is mapped to (e_x, e_y) is at most $2^{-(m+n)}2^{H_{\max}(XY|Z)} \leq \varepsilon/3$. By the union bound, we get an error of at most ε. For $\varepsilon' > 0$, the claim follows from the definition of smooth max-entropy and the union bound. □

5 Distributed One-Way Secret Key Agreement

In this section, we put together the previous results in order to derive our main results. Basically, our protocol for key agreement follows the same steps as "usual" one-way key agreement, namely information reconciliation followed by privacy amplification. The difference is that we now use the distributed versions of these tools. To connect the two steps, we will make use of the chain rule (6) which provides a lower bound on the remaining min-entropy, if some information is leaked to the adversary.

For the following, let X, Y, W, and Z be random variables known to the parties Alice, Bob, and Carol, as well as to the adversary, Eve, respectively. In a *key agreement protocol*, Alice and Bob both send messages, A and B, respectively, to Carol. Then, Alice and Bob compute keys $K_A \in \{0,1\}^{k_A}$ and $K_B \in \{0,1\}^{k_B}$ of length k_A and k_B, respectively. The protocol is said to be ε*-secure* if Carol can guess the pair (K_A, K_B) with probability at least $1 - \varepsilon$, given her information (W, A, B) and, in addition, the pair (K_A, K_B) is uniform with respect to Eve's information (Z, A, B).

Theorem 1. *Let* P_{XYWZ} *be a distribution and let* $\varepsilon > 0$ *and* $\varepsilon' \geq 0$*. For* *any* k_A, k_B *satisfying*

$$\begin{aligned} k_A &\leq H_{\min}^{\varepsilon'}(X \mid Z) - m - 3\log(4/\varepsilon) \ , \\ k_B &\leq H_{\min}^{\varepsilon'}(Y \mid Z) - m - 3\log(4/\varepsilon) \ , \\ k_A + k_B &\leq H_{\min}^{\varepsilon'}(XY \mid Z) - m - 3\log(4/\varepsilon) \end{aligned}$$

where

$$m = \max(H_{\max}^{\varepsilon'}(X \mid YW) + H_{\max}^{\varepsilon'}(Y \mid XW), H_{\max}^{\varepsilon'}(XY \mid W)) + 2\log(1/\varepsilon) + 4 \ ,$$

there exists an $(\varepsilon + 3\varepsilon')$*-secure key agreement protocol generating keys of length* k_A *and* k_B

Proof. We first prove the statement for the case where $\varepsilon' = 0$. Lemma 4 implies that Carol can decode X and Y with a probability of at least $1 - \varepsilon$, if Alice sends a hash value A of length a, and Bob sends a hash value B of length b, where

$$\begin{aligned} a &\geq H_{\max}(X \mid YW) + \log(1/\varepsilon) + 2 \\ b &\geq H_{\max}(Y \mid XW) + \log(1/\varepsilon) + 2 \\ a + b &\geq H_{\max}(XY \mid W) + \log(1/\varepsilon) + 2 \ . \end{aligned}$$

(Note that Alice and Bob additionally need to send the uniform randomness they used for the hashing to Carol. However, for the secrecy considerations below, we can ignore it as it is independent of the rest.) It is easy to see that a and b can always be chosen such that

$$a + b = \max(H_{\max}(X \mid YW) + H_{\max}(Y \mid XW), H_{\max}(XY \mid W)) + 2\log(1/\varepsilon) + 4 \ .$$

The chain rule (6) yields a bound on the amount of uncertainty Eve has over X and Y,

$$\begin{aligned} H_{\min}^{\varepsilon/4}(X \mid ABZ) &\geq H_{\min}(X \mid Z) - a - b - \log(4/\varepsilon) \\ H_{\min}^{\varepsilon/4}(Y \mid ABZ) &\geq H_{\min}(Y \mid Z) - a - b - \log(4/\varepsilon) \\ H_{\min}^{\varepsilon/4}(XY \mid ABZ) &\geq H_{\min}(XY \mid Z) - a - b - \log(4/\varepsilon) \ . \end{aligned}$$

Now, using 2-universal hashing, Alice extracts a key K_A of length k_A from X, and Bob extracts a key K_B of length k_B from Y. (Again, they use additional randomness which they send to Carol as well.) With the choice

$$\begin{aligned} k_A &\leq H_{\min}^{\varepsilon/4}(X \mid ABZ) - 2\log(4/\varepsilon) \\ k_B &\leq H_{\min}^{\varepsilon/4}(Y \mid ABZ) - 2\log(4/\varepsilon) \\ k_A + k_B &\leq H_{\min}^{\varepsilon/4}(XY \mid ABZ) - 2\log(4/\varepsilon) \ . \end{aligned}$$

Lemma 3 guarantees that (K_A, K_B) is ε-close to uniform with respect to the information held by Eve.

The case where $\varepsilon' > 0$ follows from the definition of smooth min- and max-entropies and the union bound. □

As in two-party key-agreement, Alice, Bob and Carol may extract more key if Alice and Bob pre-process their values, by sending parts of of it to Carol, and forgetting parts of it. Let $P_{U_A V_A|X}$ and $P_{U_B V_B|Y}$ be conditional probability distributions. Alice applies $P_{U_A V_A|X}$ on her value X, and Bob $P_{U_B V_B|Y}$ on his value Y. They send V_A and V_B to Carol, and they use U_A and U_B instead of X and Y. Carol has now (W, V_A, V_B), and Eve (Z, V_A, V_B). Applying Theorem 1, we get the following statement.

Corollary 1. *Let P_{XYWZ} be a distribution and let $\varepsilon > 0$ and $\varepsilon' \geq 0$. For any k_A, k_B such that there exist conditional probability distributions $P_{U_B V_B|Y}$ and $P_{U_A V_A|X}$ satisfying*

$$\begin{aligned} k_A &\leq H_{\min}^{\varepsilon'}(U_A \mid Z V_A V_B) - m - 3\log(4/\varepsilon) \ , \\ k_B &\leq H_{\min}^{\varepsilon'}(U_B \mid Z V_A V_B) - m - 3\log(4/\varepsilon) \ , \\ k_A + k_B &\leq H_{\min}^{\varepsilon'}(U_A U_B \mid Z V_A V_B) - m - 3\log(4/\varepsilon) \ , \end{aligned}$$

where

$$\begin{aligned} m = \max(&H_{\max}^{\varepsilon'}(U_A \mid U_B W V_A V_B) + H_{\max}^{\varepsilon'}(U_B \mid U_A W V_A V_B), \\ &H_{\max}^{\varepsilon'}(U_A U_B \mid W V_A V_B)) + 2\log(1/\varepsilon) + 4 \ , \end{aligned}$$

there exists an $(\varepsilon + 3\varepsilon')$-secure key agreement protocol generating keys of length k_A and k_B.

We will now show that Corollary 1 is almost optimal.

Theorem 2. *Let P_{XYWZ} be a distribution and let $\varepsilon > 0$. If there exists a protocol generating keys of length k_A and k_B then there must exist conditional probability distributions $P_{U_B V_B|Y}$ and $P_{U_A V_A|X}$ such that*

$$\begin{aligned} k_A &\leq H_{\min}^{\varepsilon}(U_A \mid Z V_A V_B) - m \ , \\ k_B &\leq H_{\min}^{\varepsilon}(U_B \mid Z V_A V_B) - m \ , \\ k_A + k_B &\leq H_{\min}^{\varepsilon}(U_A U_B \mid Z V_A V_B) - m \ , \end{aligned}$$

where

$$\begin{aligned} m = \max(&H_{\max}^{\varepsilon}(U_A \mid U_B W V_A V_B) + H_{\max}^{\varepsilon}(U_B \mid U_A W V_A V_B), \\ &H_{\max}^{\varepsilon}(U_A U_B \mid W V_A V_B)) \ . \end{aligned}$$

Proof. (Sketch) Let us assume that we have a protocol, where Alice receives the key K_A, and Bob the key K_B. Furthermore, let M_A and M_B be the messages sent by Alice and Bob to Carol.

Since Carol can calculate K_A and K_B with an error of at most ε, we have $H^\varepsilon_{\max}(K_A \mid K_B W M_A M_B) = 0$, $H^\varepsilon_{\max}(K_B \mid K_A W M_A M_B) = 0$ and $H^\varepsilon_{\max}(K_A K_B \mid W V_A V_B) = 0$. Since (K_A, K_B) is ε-close to uniform with respect to (Z, M_A, M_B), we have $H^\varepsilon_{\min}(K_A \mid Z, M_A, M_B) \geq k_A$, $H^\varepsilon_{\min}(K_B \mid Z, M_A, M_B) \geq k_B$, and $H^\varepsilon_{\min}(K_A K_B \mid Z, M_A, M_B) \geq k_A + k_B$. The statement follows now for $(U_A, V_A) := (K_A, M_A)$, and $(U_B, V_B) := (K_B, M_B)$. □

References

[AC93] R. Ahlswede and I. Csiszár. Common randomness in information theory and cryptography – part I: Secret sharing. *IEEE Transactions on Information Theory*, 39(4):1121–1132, 1993.

[BBCM95] C. H. Bennett, G. Brassard, C. Crépeau, and U. Maurer. Generalized privacy amplification. *IEEE Transactions on Information Theory*, 41, 1995.

[BBR88] C. H. Bennett, G. Brassard, and J.-M. Robert. Privacy amplification by public discussion. *SIAM Journal on Computing*, 17(2):210–229, 1988.

[BS94] G. Brassard and L. Salvail. Secret-key reconciliation by public discussion. In *Advances in Cryptology — EUROCRYPT '93*, pages 410–423. Springer-Verlag, 1994.

[Cac97] C. Cachin. *Entropy Measures and Unconditional Security in Cryptography*. PhD thesis, ETH Zurich, Switzerland, 1997.

[CK78] I. Csiszár and J. Körner. Broadcast channels with confidential messages. *IEEE Transactions on Information Theory*, 24:339–348, 1978.

[CN04] I. Csiszár and P. Narayan. Secrecy capacities for multiple terminals. *IEEE Transactions on Information Theory*, 50(12):3047–3061, 2004.

[Cov75] T. Cover. A proof of the data compression theorem of Slepian and Wolf for ergodic sources. *IEEE Transactions on Information Theory*, 21:226–228, 1975.

[CW79] J. L. Carter and M. N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, 18:143–154, 1979.

[ILL89] R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC '89)*, pages 12–24. ACM Press, 1989.

[Mau93] U. Maurer. Secret key agreement by public discussion. *IEEE Transaction on Information Theory*, 39(3):733–742, 1993.

[MW97] U. Maurer and S. Wolf. Privacy amplification secure against active adversaries. In *Advances in Cryptology — CRYPTO '97*, volume 1294 of *Lecture Notes in Computer Science*, pages 307–321. Springer-Verlag, 1997.

[Ren05] R. Renner. *Security of Quantum Key Distribution.* PhD thesis, ETH Zurich, Switzerland, 2005. Available at http://arxiv.org/abs/quant-ph/0512258.

[RW05] R. Renner and S. Wolf. Simple and tight bounds for information reconciliation and privacy amplification. In *Advances in Cryptology — ASIACRYPT 2005*, volume 3788 of *Lecture Notes in Computer Science*, pages 199–216. Springer-Verlag, 2005.

[RWW06] R. Renner, S. Wolf, and J. Wullschleger. The single-serving channel capacity. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT '06)*, 2006.

[Sha48] C. E. Shannon. A mathematical theory of communication. *Bell System Tech. Journal*, 27:379–423, 623–656, 1948.

[SW73] D. Slepian and J. K. Wolf. Noiseless coding of correlated information sources. *IEEE Transactions on Information Theory*, IT-19:471–480, 1973.

[Wul07] J. Wullschleger. Oblivious-transfer amplification. In *Advances in Cryptology — EUROCRYPT '07*, Lecture Notes in Computer Science. Springer-Verlag, 2007.