Linking Information Reconciliation and Privacy Amplification*

(Extended Abstract)

Christian Cachin and Ueli M. Maurer

Institute for Theoretical Computer Science
ETH Zürich
CH-8092 Zürich, Switzerland
E-mail: {cachin,maurer}@inf.ethz.ch

Abstract. Information reconciliation and privacy amplification are important tools in cryptography and information theory. Reconciliation allows two parties knowing correlated random variables, such as a noisy version of the partner's random bit string, to agree on a shared string. Privacy amplification allows two parties sharing a partially secret string, about which an opponent has some partial information, to distill a shorter but almost completely secret key by communicating only over an insecure channel, as long as an upper bound on the opponent's knowledge about the string is known. The relation between these two techniques has not been well understood and it is the purpose of this paper to provide the missing link between these techniques. The results have applications in unconditionally secure key agreement protocols and in quantum cryptography.

1 Introduction

One of the fundamental problems in cryptography is the generation of a shared secret key by two parties, Alice and Bob, not sharing a secret key initially, in the presence of an enemy Eve. One generally assumes that Eve can eavesdrop on the communication between Alice and Bob who are connected only by a public channel. It is easy to see that if this public channel is not assumed to be authentic, then such key agreement is impossible. We therefore assume that any modification or insertion of messages can be detected by Alice and Bob.

This problem can be solved by applying public-key cryptography [8], where one assumes that Eve's computing power is limited. In the recent years, key agreement protocols have been developed that are secure against adversaries with unlimited computing power [1, 10]. The motivation for investigating such protocols is two-fold: First, one avoids having to worry about the generality of a particular computational model, which is of some concern in view of the potential realizability of quantum computers (e.g. [4, 9, 12]). Secondly, and more

* This research was supported by the Swiss National Science Foundation. The full version of this paper has been submitted to the *Journal of Cryptology*.

importantly, no strong rigorous results on the difficulty of breaking a cryptosystem have been proved, and this problem continues to be among the most difficult ones in complexity theory.

Unconditionally secure secret-key agreement [10, 11] takes place in a scenario where Alice, Bob and Eve know the correlated random variables X, Y and Z, respectively, distributed according to some joint probability distribution that may be under partial control of Eve (like for instance in quantum cryptography [1]). One possible scenario considered by Maurer [10] is that X, Y and Z result from a binary random string broadcast by a satellite and received by Alice, Bob and Eve over noisy channels. Secret-key agreement is possible even when Eve's channel is much more reliable than Alice's and Bob's channels.

A key agreement protocol for such a scenario generally consists of three phases:

Advantage Distillation: The purpose of the first phase is to create a random variable W about which either Alice or Bob has more information than Eve. Advantage distillation is only needed when such a W is not immediately available from X and Y, for instance, when Eve's channel is superior in the above satellite scenario. Alice and Bob create W by exchanging messages, summarized as the random variable C, over the public channel.

Information Reconciliation [1, 6]: To agree on a string T with very high probability, Alice and Bob exchange redundant error-correction information U, such as a sequence of parity checks. After this phase, Eve's (incomplete) information about T consists of Z, C and U.

Privacy Amplification [2, 3]: In the final phase, Alice and Bob agree publicly on a compression function G to distill from T a shorter string S about which Eve has only a negligible amount of information. Therefore, S can subsequently be used as a secret key.

Information reconciliation and privacy amplification are fundamental for unconditionally secure key agreement and quantum key distribution.

If after the first phase Alice knows a string about which Bob has more information than Eve, Alice and Bob can choose W to be this string. In other words, using information-theoretic terms, W is a random variable such that $H(W|XC) = 0$ and $H(W|YC) < H(W|ZC)$. In such a case, Bob tries to determine W from Y and the reconciliation string U, which could simply be an error-correction string sent by Alice or could result from an interactive communication with Alice. (Note that $H(U) \geq H(W|YC)$ is a necessary condition.) Hence reconciliation serves to establish $H(W|YCU) \approx 0$ while Eve still has a substantial amount of uncertainty about W: $H(W|ZCU) > 0$. After privacy amplification, $H(S)$ should be as large as possible, and Eve's information about S should be arbitrarily close to zero: $I(S; ZCUG) = H(S) - H(S|ZCUG) \approx 0$. Note that Alice and Bob can both compute S, i.e., $H(S|WG) = 0$.

In the following, let $V = [Z, C]$ summarize Eve's total knowledge about W before reconciliation. For deriving lower bounds on Eve's final information about the secret key S one can either consider a particular value $V = v$ that

Eve knows or one can average over all possible values of V. Note that results for a particular $V = v$, which will be considered in this paper, are stronger than averaging results because they are known to hold for the very instance of the protocol execution. In other words, Eve's information about W is modeled by the probability distribution $P_{W|V=v}$ about which Alice and Bob have some incomplete knowledge. In particular, they know a lower bound on the collision entropy (see below) of the distribution $P_{W|V=v}$ but they do not know v.

It is known [2] that the collision entropy after reconciliation with $U = u$ (i.e., of the distribution $P_{W|V=v,U=u}$) is a lower bound on the size of the secret key that can be distilled safely by privacy amplification. This paper is concerned with understanding the reduction of the collision entropy induced by the side information U, either for a particular value $U = u$, or averaged over all values of U. Although this question is fundamental for any proof in the area of key-agreement protocols, it has previously not been well understood because the behavior of collision entropy is different from that of Shannon entropy with respect to side-information. Existing proofs such as the ingenious Big-Brother argument of [1] work only for particular probability distributions and reconciliation protocols.

The paper is organized as follows. Section 2 reviews privacy amplification and the definition of collision entropy. Section 3 presents upper bounds on the reduction of collision entropy due to side-information for arbitrary probability distributions. Non-interactive reconciliation protocols with uniform and close-to-uniform probability distributions are investigated in Section 4.

2 Review of Privacy Amplification and Collision Entropy

We assume that the reader is familiar with the notion of entropy and the basic concepts of information theory [5]. In privacy amplification, a different and non-standard entropy measure, *collision entropy*, is of central importance [2]. Collision entropy is also known as Rényi entropy of order 2. To distinguish collision entropy from entropy in the sense of Shannon, we will always refer to the latter as *Shannon entropy*. All logarithms in this paper are to the base 2, and entropies are thus measured in bits.

Privacy amplification was introduced by Bennett, Brassard and Robert [3] and investigated further in [2]. Assume Alice and Bob share an n–bit string W about which an eavesdropper Eve has incomplete information characterized by a probability distribution $P_{W|V=v}$ over the n–bit strings, where v denotes the particular value taken on by the random variable V summarizing her side-information. Alice and Bob have some knowledge of the distribution $P_{W|V=v}$, but they do not know exactly what is compromised about their string. Using a public channel, which is totally susceptible to eavesdropping, they wish to agree on a function $g : \{0,1\}^n \to \{0,1\}^r$ such that Eve, despite her partial knowledge about W and complete knowledge of g, almost certainly knows nearly nothing about $g(W)$. This process transforms a partially secret n–bit string W into a highly secret but shorter r–bit string $g(W)$ which can be used as a secret key.

The method for selecting the function g proposed in [3] is to choose it at random from a publicly-known *universal class of hash functions* [7] mapping n-bit strings to r-bit strings.

Bennett, Brassard, Crépeau and Maurer [2] showed that the collision entropy (defined below) of Eve's distribution about W provides a lower bound on the size r of the secret key distillable from W by privacy amplification with a universal hash function.

Definition 1 [2]. Let X be a random variable with alphabet $\mathcal{X}$ and distribution P_X. The *collision probability* $P_c(X)$ of X is defined as the probability that X takes on the same value twice in two independent experiments, i.e.,

$$P_c(X) = \sum_{x \in X} P_X(x)^2.$$

The *collision entropy* of X is defined as the negative logarithm of its collision probability:

$$H_c(X) = -\log P_c(X).$$

For an event $\mathcal{E}$, the *collision entropy of X conditioned on $\mathcal{E}$*, $H_c(X|\mathcal{E})$, is defined naturally as the collision entropy of the conditional distribution $P_{X|\mathcal{E}}$. The *collision entropy conditioned on a random variable*, $H_c(X|Y)$, is defined as the expected value of the conditional collision entropy:

$$H_c(X|Y) = \sum_y P_Y(y)\, H_c(X|Y = y).$$

Equivalently, $H_c(X)$ can be expressed as $H_c(X) = -\log E[P_X(X)]$, where $E[\cdot]$ denotes the expected value. Shannon entropy $H(X)$ can be expressed similarly as $H(X) = -E[\log P_X(X)]$. It follows from Jensen's inequality (see [5], p. 428) that collision entropy is upper bounded by the Shannon entropy:

$$H_c(X) \leq H(X),$$

with equality if and only if P_X is the uniform distribution over $\mathcal{X}$ or a subset of $\mathcal{X}$. Similarly, we have $H(X|Y) \geq H_c(X|Y)$. Note that collision entropy (like Shannon entropy) is always positive.

The following theorem is the main result of [2]:

Theorem 2. *Let X be a random variable on alphabet $\mathcal{X}$ with probability distribution P_X and collision entropy $H_c(X)$. Further, let G be the random variable corresponding to the random choice (with uniform distribution) of a member of a universal class of hash functions from $\mathcal{X} \to \{0,1\}^r$. Then*

$$H(G(X)|G) \geq H_c(G(X)|G) \geq r - \frac{2^{r-H_c(X)}}{\ln 2}.$$

Note that G is a random variable and that the quantity $H(G(X)|G)$ is an average over all choices of the function g. It is possible that $H(g(X)|g) = H(g(X))$ differs from r by a non-negligible amount for some g, but such a g can occur only with negligible probability.

This theorem clearly applies also to conditional probability distributions such as $P_{W|V=v}$ discussed above. If Eve's collision entropy $H_c(W|V=v)$ is known to be at least t and Alice and Bob choose $S = G(W)$ as their secret key, then

$$H_c(S|G, V=v) = H_c(G(W)|G, V=v) \geq r - 2^{r-t}/\ln 2.$$

The key S is indeed virtually secret because $H(S|G, V=v) \geq H_c(S|G, V=v)$ and hence $H(S|G, V=v)$ is arbitrarily close to maximal. More precisely, if $r < t$, then Eve's total information about S decreases exponentially in the excess compression $t - r$.

It should be pointed out that Theorem 2 cannot be generalized to collision entropy conditioned on a random variable, i.e., $H_c(G(W)|GV) \geq r - 2^{r-H_c(W|V)}/\ln 2$ is false in general.

3 The Effect of Side Information on Collision Entropy

As described above, the reconciliation step consists of Alice and Bob exchanging suitable error-correction information U over the public channel. This information decreases Eve's (Shannon and collision) entropy about W. For non-interactive reconciliation, Alice chooses an appropriate error-correction function f and sends $U = f(W)$ to Bob who can then reconstruct W from U and his prior knowledge YC.

The results of this paper will be derived for an arbitrary random variable X with probability distribution P_X and a side-information random variable U jointly distributed with X according to P_{XU}. However, they can just as well be applied to conditional distributions; our intended application is the key agreement scenario mentioned in the introduction, i.e., when P_X and $P_{X|U}$ are replaced by $P_{W|V=v}$ and $P_{W|V=v,U}$, respectively.

In general, giving side-information implies a reduction of entropy. Our goal is to derive upper bounds on the size of this reduction. Giving as side-information the fact that U takes on a particular value u, it is possible for both, Shannon and collision entropies, that the entropy increases or decreases. Moreover, the size of a reduction can be arbitrarily large.

However, the expected reduction (for all values of U) of the Shannon entropy of X by giving U, called the mutual information between X and U, is bounded by $H(U)$:

$$H(X) - H(X|U) = I(X;U) \leq H(U) \tag{1}$$

which follows from the symmetry of $I(X;U)$ and the fact that Shannon entropy (conditional or not) is always positive.

The example below illustrates two facts. First, the reduction of collision entropy implied by giving side-information $U = u$ can exceed the reduction of Shannon entropy, i.e.,

$$H_c(X) - H_c(X|U = u) > H(X) - H(X|U = u)$$

is possible. Second, it shows that the natural generalization of (1) to collision entropy, namely $H_c(X) - H_c(X|U) \leq H_c(U)$, is not true in general. However, Theorem 3 demonstrates that the weaker inequality $H_c(X) - H_c(X|U) \leq H(U)$ is always satisfied.

Example. Let X be a random variable with alphabet $\mathcal{X} = \{a_1, \ldots, a_{10}, b_1, \ldots, b_{10}\}$ distributed according to $P_X(a_i) = 0.01$ and $P_X(b_i) = 0.09$ for $i = 1, \ldots, 10$. We have $H(X) \approx 3.79$ and $H_c(X) \approx 3.61$. Let $f : \mathcal{X} \to \{0, 1\}$ be defined as

$$f(x) = \begin{cases} 0 & \text{if } x \in \{a_1, \ldots, a_9, b_{10}\} \\ 1 & \text{if } x \in \{a_{10}, b_1, \ldots, b_9\} \end{cases}$$

and let $U = f(X)$. Then $H(X|U = 0) \approx 2.58$ and $H_c(X|U = 0) \approx 1.85$. The reduction of collision entropy when given $U = 0$ exceeds the reduction of Shannon entropy, i.e., $H_c(X) - H_c(X|U = 0) \approx 1.76$ whereas $H(X) - H(X|U = 0) \approx 1.21$.

The expected entropy reductions are $H(X) - H(X|U) \approx 0.69$ (in fact $H(X) - H(X|U) = H(U)$ because f is deterministic) and $H_c(X) - H_c(X|U) \approx 0.65$. Note that $H_c(U) \approx 0.50$ and $H_c(X) - H_c(X|U)$ is indeed greater than $H_c(U)$ but less than $H(U)$.

$H(U)$ is not only the maximal expected decrease of Shannon entropy, but $H(U)$ is also an upper bound on the expected decrease of collision entropy, as the following theorem demonstrates. (All theorems will be proved in the full version.)

Theorem 3. *Let X and U be two random variables with alphabets $\mathcal{X}$ and $\mathcal{U}$, respectively. The expected reduction of the collision entropy of X, when given U, does not exceed the Shannon entropy of U, i.e.,*

$$H_c(X) - H_c(X|U) \leq H(U),$$

with equality if and only if U is defined uniquely for each $x \in \mathcal{X}$ and P_U is the uniform distribution over $\mathcal{U}$ or a subset of $\mathcal{U}$.

For a positive-valued random variable X, $E[X] \leq t$ implies that $P[X \geq kt] \leq 1/k$. Hence, according to Theorem 3, the probability that the leaking information $U = u$ decreases collision entropy by more than $kH(U)$ is at most $1/k$, i.e., $P[H_c(X) - H_c(X|U = u) \geq kH(U)] \leq 1/k$. However, such a high probability of partially exposing the string W is unacceptable in a key-agreement scenario. The following theorem provides a much stronger result by showing that the above probability decreases in fact exponentially in k if $H(U)$ is replaced by an expression roughly twice as large.

Theorem 4. *Let X and U be random variables with alphabets $\mathcal{X}$ and $\mathcal{U}$, respectively, and let $s > 0$ be an arbitrary security parameter. With probability at least $1 - 2^{-s}$, U takes on a value u for which*

$$H_c(X) - H_c(X|U = u) \leq 2 \log |\mathcal{U}| + 2s.$$

Remark. The statement of the theorem is equivalent to

$$\sum_{u:\, H_c(X) - H_c(X|U=u) \leq 2 \log |\mathcal{U}| + 2s} P_U(u) \geq 1 - 2^{-s}.$$

Equivalently, but less intuitively, we can write

$$P[H_c(X) - H_c(X|U = u) \leq 2 \log |\mathcal{U}| + 2s] \geq 1 - 2^{-s},$$

if we view $H_c(X|U = u)$ as a function of u. (Note that this probability is defined although u could be a function of x, and that $H_c(X)$ is a constant that does not depend on x or u.)

Because of its importance we restate Theorem 4 for the key-generation scenario, replacing P_X by $P_{W|V=v}$, with the side-information consisting of k bits, for instance k parity checks of W when W is an n-bit string.

Corollary 5. *Let W be a random variable with alphabet $\mathcal{W}$, let v and u be particular values of the random variables V and U, correlated with W, and let $s > 0$ be a given security parameter. Then, with probability at least $1 - 2^{-s}$, U takes on a value u such that the decrease in collision entropy by giving u, $H_c(W|V = v) - H_c(W|V = v, U = u)$, is at most $2k + 2s$.*

4 Almost Uniform Distributions

As shown above, giving side information of the form $U = u$ can reduce the collision entropy by an arbitrary amount, although the probability that this happens is bounded by Theorem 4. In this section we derive better bounds on the reduction for non-interactive reconciliation and special probability distributions. It is easy to see that for uniform distributions and deterministic side-information $U = f(W)$, the reduction of collision entropy depends only on the size of the preimage of $u = f(x)$:

Lemma 6 [2]. *Let X be a random variable with alphabet $\mathcal{X}$, let $f : \mathcal{X} \to \mathcal{U}$ be an arbitrary function taking on values in a given set $\mathcal{U}$, let U be defined as $U = f(X)$, and set $\mathcal{X}_u = \{x \in \mathcal{X} : f(x) = u\}$. If X is distributed uniformly over $\mathcal{X}$, then*

$$H_c(X) - H_c(X|U = u) = \log \frac{|\mathcal{X}|}{|\mathcal{X}_u|}.$$

In particular, if f is symmetric (i.e., $|\mathcal{X}_u|$ is the same for all $u \in \mathcal{U}$), knowledge of $U = u$ reduces the collision entropy by $\log |\mathcal{U}|$.

Theorem 7 states a bound on the reduction of collision entropy for almost uniform distributions.

Theorem 7. *For given $\alpha > 1$ and $\beta > 1$, let X be a random variable with alphabet $\mathcal{X}$ and probability distribution P_X such that $\frac{1}{\alpha|\mathcal{X}|} \leq P_X(x) \leq \frac{\beta}{|\mathcal{X}|}$ for all $x \in \mathcal{X}$. Define f, U and $\mathcal{X}_u$ as in Lemma 6. Then*

$$H_c(X) - H_c(X|U=u) \leq \log \frac{|\mathcal{X}|}{|\mathcal{X}_u|} + 4\log\alpha + 2\log\beta.$$

In particular, if f is symmetric, then $H_c(X) - H_c(X|U=u) \leq \log|\mathcal{U}| + 4\log\alpha + 2\log\beta$.

This result can be applied to the important class of scenarios where a certain random experiment is repeated independently a large number of times. For example, W could be the result of receiving independently generated bits over a memoryless channel, as in the satellite scenario mentioned earlier. A fundamental theorem of information theory [5] states that in such a scenario all occurring sequences can be divided into a typical set and a non-typical set, where the probability that a randomly selected sequence of length n lies in the typical set approaches 1 for all sufficiently large n. Furthermore, all sequences in the typical set are almost equally probable. As will be shown in the final version of the paper, this crucial observation allows us to bound the decrease of collision entropy for the sequences in the typical set by Theorem 7, leading to a result similar to Lemma 6.

5 Conclusions

The described link between information reconciliation and privacy amplification for unconditionally secure secret-key agreement can be summarized as follows. Assume that Alice knows a random variable W and that Bob and Eve have partial knowledge about W, characterized by the random variables W' and V, respectively. These random variables could for instance result from the described satellite scenario with W and W' being functions of $[X, C]$ and $[Y, C]$, respectively, and with $V = [Z, C]$. In order to state the results in the strongest possible form we consider a particular value $V = v$ held by Eve rather than the average over all values of V.

When V gives less information than W' about W, i.e., $H(W|V) > H(W|W')$, and a lower bound $t > 0$ on the collision entropy of Eve's probability distribution of W is known, i.e., $H_c(W|V=v) \geq t$, then Alice and Bob can generate a shared secret key S as follows. Alice and Bob exchange error-correcting information U consisting of $k > H(W|W')$ bits over the public channel such that Bob can reconstruct W, i.e., $H(W|W'U) \approx 0$. Eve gains additional knowledge about W by seeing $U = u$. However, Corollary 5 shows that with probability at least $1 - 2^{-s}$ (over all values of U) where the security parameter s can be chosen arbitrarily, her collision entropy is bounded from below by $H_c(W|V=v, U=$

$u) \geq t - 2k - 2s$. Using privacy amplification, Alice and Bob can now generate an r-bit secret key S, where r has to be chosen smaller than $t - 2k - 2s$ and Eve's total information about S is exponentially small in $t - 2k - 2s - r$, namely less than $2^{r-(t-2k-2s)}/\ln 2$ bits.

The main advantage of Theorem 4 is that it applies to any distribution and any reconciliation protocol whereas previously obtained results held only for particular distributions and protocols. However, as was demonstrated in Section 4, a larger secret key than suggested by Theorem 4 can be obtained by Alice and Bob for special distributions.

Acknowledgement

It is a pleasure to thank Charles Bennett, Gilles Brassard, Claude Crépeau, and Martin Gander for interesting discussions.

References

1. C. H. Bennett, F. Bessette, G. Brassard, L. Salvail, and J. Smolin, "Experimental quantum cryptography," *Journal of Cryptology*, vol. 5, no. 1, pp. 3–28, 1992.
2. C. H. Bennett, G. Brassard, C. Crépeau, and U. M. Maurer, "Generalized privacy amplification." Preprint, 1994.
3. C. H. Bennett, G. Brassard, and J.-M. Robert, "Privacy amplification by public discussion," *SIAM Journal on Computing*, vol. 17, pp. 210–229, Apr. 1988.
4. E. Bernstein and U. Vazirani, "Quantum complexity theory," in *Proc. 25th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 11–20, 1993.
5. R. E. Blahut, *Principles and Practice of Information Theory.* Reading, MA: Addison-Wesley, 1987.
6. G. Brassard and L. Salvail, "Secret-key reconciliation by public discussion," in *Advances in Cryptology — EUROCRYPT '93* (T. Helleseth, ed.), vol. 765 of *Lecture Notes in Computer Science*, pp. 410–423, Springer-Verlag, 1994.
7. J. L. Carter and M. N. Wegman, "Universal classes of hash functions," *Journal of Computer and System Sciences*, vol. 18, pp. 143–154, 1979.
8. W. Diffie and M. E. Hellman, "New directions in cryptography," *IEEE Transactions on Information Theory*, vol. 22, pp. 644–654, Nov. 1976.
9. S. Lloyd, "A potentially realizable quantum computer," *Science*, vol. 261, pp. 1569–1571, 1993.
10. U. M. Maurer, "Secret key agreement by public discussion from common information," *IEEE Transactions on Information Theory*, vol. 39, pp. 733–742, May 1993.
11. U. M. Maurer, "The strong secret key rate of discrete random triples," in *Communications and Cryptography: Two Sides of One Tapestry* (R. E. Blahut *et al.*, eds.), Kluwer, 1994.
12. P. W. Shor, "Algorithms for quantum computation." Submitted to FOCS'94, 1994.