On Recent Advances in Key Derivation via the Leftover Hash Lemma

Maciej Skorski

Cryptology and Data Security Group, University of Warsaw
maciej.skorski@gmail.com

Abstract

Barak et al. showed how to significantly reduce the entropy loss, which is necessary in general, in the use of the Leftover Hash Lemma (LHL) to derive a secure key for many important cryptographic applications. If one wants this key to be secure against any additional *short* leakage, then the min-entropy of the source used with the LHL must be appropriately bigger (roughly by the length of the supposed leakage). Recently, Berens came up with a notion of collision entropy that is much weaker than min-entropy and allows proving a version of the LHL with leakage robustness but without any entropy saving. We combine both approaches and extend the results of Barak et. al to Beren's collision entropy. Summarizing, we obtain a version of the LHL with optimized entropy loss, leakage robustness and weak entropy requirements.

1998 ACM Subject Classification F.1.3 Complexity Measures and Classes, K.6.5 Security and Protection

Keywords and phrases Key derivation, Leftover Hash Lemma, leakage resilient cryptography

Digital Object Identifier 10.4230/OASIcs.ICCSW.2014.83

1 Introduction

1.1 Randomness extractors

Analyzing security of cryptographic primitives one usually assumes an access to *perfect randomness*. However in practice we are limited to imperfect sources, which have a lot of randomness (measured by entropy[1]) but exhibit some patterns, bias or correlations between particular bits in generated sequences. As an entropy source one can use biometric data (like fingerprints), data collected from user-application interaction (mouse movements, typing the keyboard, timing events and others) and even physical sources (thermal noise, nuclear decay). Having collected "noisy" data with enough entropy, in the post-processing phase one "extracts" pure randomness by dedicated procedures called *randomness extractors*. One needs to stress that in general these procedures, in addition to a high-entropy source, require some amount of pure randomness used as a "catalyst"[2], referred to as the *seed*.

1.2 Leftover Hash Lemma

The famous Leftover Hash Lemma [4] states that universal hash functions are good extractors: if an n-bit source X has at least $\ell = k + 2\log(1/\epsilon)$ bits of min-entropy[3] then the distribution

[1] In information theory, the most popular notion is Shannon Entropy. In the context of extracting randomness one typically uses min-entropy or collision entropy.
[2] Since obtaining true random bits might be hard and expensive, a lot of research is focused on deterministic extractors, which don't need to be seeded. However, they work only for very limited class of sources.
[3] Which means that no adversary can guess the output of X with probability better than $2^{-\ell}$.

Imperial College Computing Student Workshop (ICCSW'14).
Editors: Rumyana Neykova and Nicholas Ng; pp. 83–90

OpenAccess Series in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

$H(X), H$, where H is a randomly chosen universal hash function from n to k bits (the seed), is ϵ-close to the k-bit uniform distribution, even if H is published.

Entropy loss and RT bound. Note that the Leftover Hash Lemma requires the significantly bigger entropy on its input comparing to what it extracts. More specifically, we sacrifice $L = 2\log(1/\epsilon)$ bits of entropy in order to obtain an output of quality ϵ. The result of Radhakrishnan and Ta-Shma [5], called the RT-bound, shows that this loss is necessary for any extractor. In this sense, the LHL as an extractor achieves the optimal entropy loss.

Importance of optimizing the LHL. Although the loss of $2\log(1/\epsilon)$ bits might seem to be negligible from an asymptotic point of view, however in some important applications it affects the efficiency (for instance the Diffie-Hellmann key exchange). Thus, minimizing it is important for efficiency. Similarly, shorter seeds than that one of the LHL are desired.

1.3 Key derivation: ideal and real settings

For any cryptographic primitive (like an encryption scheme or a signature), which uses randomness R to derive its secure key, we compare two different settings:

(a) ideal: R is *perfectly* uniform and independent of any attacker's side information Z
(b) real settings: the key owner has only an *imperfect* entropy source X and uses the key extracted (in our case: by hashing) from X as R. In addition, an attacker may have some *side information* Z correlated with X.

The security of the primitive is parametrized by ϵ, which is the success probability or the advantage of an attacker with certain resources. The LHL implies that if the security is ϵ for uniform R, then the same application keyed with a random hash of X is ϵ'-secure, where

$$\epsilon' \leqslant \epsilon + \sqrt{2^{-L}} \tag{1}$$

and the entropy loss L is the difference between the entropy of X given Z (suitably defined) and the length of the hashing output (the length of the extracted key).

1.4 Side leakage and chain rules

In the context of leakage, we want the guaranteed security not to degrade much when some extra but short information is revealed to an attacker. For an entropy notion $\mathbf{H}$ and two correlated random variables X, Z, the chain rule is an inequality of the following form

$$\mathbf{H}(X|Z) \geqslant \mathbf{H}(X) - C \cdot |Z| \tag{2}$$

where $|Z| = \log|\mathrm{supp}(Z)|$ is the length of Z and C is some constant (ideally $C = 1$).

Applications of chain rules. Many information-theoretic notions of entropy satisfies the property (2). In fact, it is what one expects from a "good" notion of entropy and is often used in security proofs. Typically, high entropy corresponds to high security. Thus, Equation (2) is often used to prove security in the presence of leakage.

Example: Guessing Probability. The min-entropy of a random variable X given Z is defined as the negative logarithm of the maximal probability that an attacker can guess the output of X given only Z. It is known that it satisfies the inequality (2) with $C = 1$. In particular, if X is 128-bit uniform key, then any adversary given 10 bits of extra information, can guess X correctly with probability at most 2^{-118}, comparing to original 2^{-128}.

Example: leakage-robustness of the LHL. From Equation (1) it follows that to derive a secure key of required length we need to guarantee that $\mathbf{H}(X|Z)$, where $\mathbf{H}$ is the min-entropy, is big enough. Let us say that we have an application whose security for the uniform k-bit key is ϵ and that we want to derive a k-bit key which guarantees the security (with no extra side information) 2ϵ. For this we need $\mathbf{H}(X) \geqslant k + 2\log(1/\epsilon)$. Suppose now that we look at the security against *any* leakage Z of length $m = 20$. The chain rule for min-entropy (2) implies that the entropy loss $L = \mathbf{H}(X|Z) - k$ is *not smaller* than $\mathbf{H}(X) - m - k$. Combining this with (1) we see that the security is now $\epsilon + \epsilon \cdot 2^{m/2} \approx 2^{10}\epsilon$. Alternatively, we could start with $m = 20$ bits of entropy more and achieve the security 2ϵ. The chain rule is *necessary* to derive keys which remain secure against any bounded leakage.

1.5 Improvements in key derivation by the LHL

Reducing the entropy loss . The RT bound shows that it is impossible to avoid losing $2\log(1/\epsilon)$ bits of entropy if we want the extracted output to be ϵ-indistinguishable from uniform (i.e. ϵ-close) by *all* statistical tests. However, in cryptographic applications we are interested only in *very special* tests, corresponding to the definition of the security of the application. This suggests that one can overcome the RT bound in specific situations. Indeed, Barak et al. proved in [1] the stronger version of LHL, which essentially says that for many scenarios the "ideal" security ϵ and "real" security ϵ' are related by

$$\epsilon' \leqslant \epsilon + \sqrt{\epsilon 2^{-L}} \tag{3}$$

where L is the entropy loss. This shows that one obtains roughly the same level of security with L reduced by half up to $L = \log(1/\epsilon)$. The result is valid for the broad class of cryptographic applications, including unpredictability applications (for example, one-way functions) and some indistinguishability applications (the so called squared-friendly applications [1, 3], like weak pseudo-random functions or stateless chosen plaintext attack secure encryption).

Reducing the seed length. It is known that seed for LHL must grow linearly with the number of extracted bits [7]. In [1] the authors also observed that in some cases the length of the seed in the LHL can be reduced. The natural idea of expanding a shorter seed works either for small number of bits or in minicrypt [1].

Relaxing the entropy assumptions. The assumptions in the Leftover Hash Lemma are typically formulated in terms of min-entropy. However, it gives the same security guarantees when the upper bound on the probability of guessing X (which corresponds to min-entropy) is replaced by the same bound on the *collision probability* of X [4]. Since the collision probability is typically much bigger than the guessing probability, this means that from many sources we can *extract more* bits than it is guaranteed by the min-entropy bounds[5]. This result can be extended into the conditional case (when there is side information Z correlated with X):

(a) The generalized LHL of Barak et al. [1]. The min-entropy requirement can be replaced by the upper bound on the collision probability of X given Z, as observed by the authors. Unfortunately, the latter does not guarantee the leakage-robustness.

[4] This fact is actually intuitive, as the hash functions are, by definition, collision resistant and one can expect that the entropy notions based on the collision probability fit well to that settings

[5] Note that this does not contradict to the RT-bound, because that counterexample is a flat distribution for which the collision and guessing probability coincide (and thus, is a very special case)

Table 1 Improvements of the Leftover Hash Lemma

Statement	Standard			Generalized		
Assumptions	min-entropy	collision prob.	collision entropy	min-entropy	collision prob.	collision entropy
Reduced entropy loss	No	No	No	Yes	Yes	?
Reduced seed length	No	No	No	Yes	Yes	?
Side leakage robustness	Yes	No	Yes	Yes	No	?
Minimal entropy requirements	No	No	?	No	?	?

(b) The standard formulation of the LHL given for the appropriate notion of conditional collision entropy [2]. This notion has two big advantages: is the weakest among other notions proposed for the collision entropy (weaker than conditional collision probability!) and it satisfies the chain rule, guaranteeing leakage robustness.

The ideal version of the LHL - issues. On the one hand the generalized LHL allows reducing the entropy loss, the seed length and also handling side information, provided that we quantify randomness by min-entropy. On the other hand, we know how to minimize the entropy requirements for the standard LHL. This landscape is summarized in Table 1 below.

The discussion leads to the natural question about the existence of the "ideal" LHL:

Question: *Does there exist a variant of the LHL which simultaneously captures the following advantages: reducing the entropy loss, reducing the seed length, leakage robustness and possibly minimal entropy assumptions?*

1.6 Our results

Summary. We answer this question affirmatively. As a first contribution we show that that the generalized LHL is not guaranteed to be leakage-robust with low-collision-probability sources. Second, we show that this can be fixed with a "correct" notion of collision entropy.

Conditional collision probability is not leakage-robust. We show (Section 4) that the generalized LHL [1] does not guarantee security for the key derived from a source of low collision probability against leakages of even one extra bit!

The generalized LHL works with the conditional collision entropy. We extend the generalized LHL [1] to work with Beren's entropy [2] (Section 5). It gives more security (because of the weaker entropy notion) and leakage robustness (because of the chain rule).

Applications - saving entropy. For a *low-collision-probability* source X, from which we want to derive a key ϵ-secure and secure against *arbitrary* one-bit leakage Z we save even $\log(1/\epsilon) - \mathcal{O}(1)$ bits of entropy comparing to what follows from [1]. Indeed, the only way to apply the statement of Barak et al. would be to convert first the "entropy" in X into min-entropy (because there is no chain rule for collision probability). By "entropy smoothing" [6], this can be achieved with loss of $\log(1/\epsilon)$ bits. In our approach this is not necessary.

2 Preliminaries

Notation. By U_k we denote the uniform distribution over $\{0,1\}^k$ and, more generally, by U_S we denote the distribution uniform over a set S. By $\mathrm{supp}(X)$ we denote the support of the distribution X.

Distinguishing Advantage and Statistical Distance. For two distributions X, Y defined on the same set we define $\Delta_{\mathrm{D}}(X;Y) = \mathbf{E}\,\mathrm{D}(X) - \mathbf{E}\,\mathrm{D}(Y)$ as the advantage of a function (attacker) D in distinguishing X and Y. The statistical distance is defined by $\mathrm{SD}(X;Y) = \frac{1}{2}\sum_x |\Pr(X=x) - \Pr(Y=x)|$, we also denote for shortness $\mathrm{SD}(X;Y|Z) = \mathrm{SD}(X,Z;Y,Z)$ and $\Delta_{\mathrm{D}}(X;Y|Z) = \Delta_{\mathrm{D}}(X,Z;Y,Z)$. We have $\mathrm{SD}(X,Y) = \max_{\mathrm{D}} \Delta_{\mathrm{D}}(X,Y)$, where the maximum is taken over all boolean functions D (possibly probabilistic).

Entropy notions and hash functions . Here we provide necessary entropy definitions.

▶ **Definition 1** (Min-entropy). The min-entropy of X given Z is defined as

$$\widetilde{\mathbf{H}}_\infty(X|Z) = -\log\left(\mathop{\mathbf{E}}_{z\leftarrow Z}\left[\max_x \Pr(X=x|Z=z)\right]\right). \tag{4}$$

▶ **Definition 2** (Collision-probability). The collision probability of X given Z is given by

$$\mathrm{CP}(X|Z) = \mathop{\mathbf{E}}_{z\leftarrow Z}\left(\sum_x \Pr(X=x|Z=z)^2\right) = \mathop{\mathbf{E}}_{z\leftarrow Z}\mathrm{CP}(X|_{Z=z}). \tag{5}$$

▶ **Definition 3** (Collision-entropy [2]). The collision entropy of X given Z is equal to

$$\mathbf{H}_2(X|Z) = -\log\left(\mathop{\mathbf{E}}_{z\leftarrow Z}\sqrt{\mathrm{CP}(X|_{Z=z})}\right)^2. \tag{6}$$

All these three definitions are related as follows:

▶ **Lemma 4.** *For any joint distribution X, Z we have*

$$\mathbf{H}_2(X|Z) \geqslant -\log \mathrm{CP}(X|Z) \geqslant \widetilde{\mathbf{H}}_\infty(X|Z) \tag{7}$$

▶ **Definition 5** (Almost universal families). A family $\mathcal{H}$ of functions $h : \mathcal{X} \to \{0,1\}^k$ is called γ-universal hash family, if for any $x_1, x_2 \in \mathcal{X}$, $x_1 \neq x_2$ we have $\Pr_{h\leftarrow\mathcal{H}}[h(x_1) = h(x_2)] \leqslant \gamma$. If $\gamma = 2^{-k}$ then we say that $\mathcal{H}$ is universal.

3 Leftover Hash Lemma

3.1 Standard LHL

Below we formulate the Leftover Hash Lemma for the conditional min-entropy.

▶ **Theorem 6** (The LHL [4]). *Let (X,Z) be a joint distribution on $\mathcal{X}\times\mathcal{Z}$, let $\mathcal{H} = \{h :\ \mathcal{X} \to \{0,1\}^k\}$ be a $\frac{1+\gamma}{2^k}$-universal family and let H be a random member of $\mathcal{H}$. Then we have*

$$\mathrm{SD}(H(X); U_k|H, Z) \leqslant \frac{1}{2}\cdot\sqrt{2^{-L}+\gamma}, \tag{8}$$

where $L = \widetilde{\mathbf{H}}_\infty(X|Z) - k$ is the entropy loss.

For universal hashing we need $L \approx 2\log(1/\epsilon)$ to make the statistical distance smaller than ϵ.

Entropy requirements for the standard LHL. It is well known that in Theorem 6 one can use collision probability instead of min-entropy. More precisely, we have

$$\mathrm{SD}(H(X);U_k|H,Z) \leqslant \frac{1}{2}\cdot\sqrt{2^k\mathrm{CP}(X|Z)+\gamma}. \tag{9}$$

Since $\mathrm{CP}(X|Z) \leqslant 2^{-\widetilde{\mathbf{H}}_\infty(X|Z)}$ this implies the bound in Equation (8). Berens [2] observed that even weaker assumption is enough. Namely, we have

$$\mathrm{SD}(H(X);U_k|H,Z) \leqslant \frac{1}{2}\cdot\sqrt{2^{k-\mathbf{H}_2(X|Z)}+\gamma}. \tag{10}$$

By Theorem 4 this implies both Equation (9) and Equation (8).

3.2 Generalized LHL

We start with the inequality that can be thought of as an abstract formulation of the LHL:

▶ **Lemma 7.** *[1] Let (Y,Z) be a joint distribution on $\mathcal{Y}\times\mathcal{Z}$ and let U be independent and uniform on $\mathcal{Y}$. Then for all real-valued functions* D *on $\mathcal{Y}\times\mathcal{Z}$, we have*

$$\mathbf{E}\,\mathrm{D}(Y,Z) - \mathbf{E}\,\mathrm{D}(U,Z) \leqslant \sqrt{\mathrm{Var}\,\mathrm{D}(U_{\mathcal{Y}},Z)}\cdot\sqrt{|\mathcal{Y}|\cdot\mathrm{CP}(Y|Z)-1}. \tag{11}$$

The generalized Leftover Hash Lemma is a special case of this result, where $Y=(H(X),H)$ for a randomly chosen hash function H. We need only one simple fact (we omit the proof):

▶ **Lemma 8.** *Let $\mathcal{H}=\{h:\ \mathcal{X}\to\{0,1\}^k\}$ be a $\frac{1+\gamma}{2^k}$-universal family and let H be a random member of $\mathcal{H}$. Then* $\mathrm{CP}(H(X),H) \leqslant \left(\mathrm{CP}(X)+2^{-k}(1+\gamma)\right)/|\mathcal{H}|$.

By Theorem 7 and Theorem 8 one obtains the following generalization of Theorem 6:

▶ **Theorem 9** (Generalized LHL [1])**.** *Let (X,Z) be a joint distribution on $\mathcal{X}\times\mathcal{Z}$, let $\mathcal{H}=\{h:\ \mathcal{X}\to\{0,1\}^k\}$ be $\frac{1+\gamma}{2^k}$-universal and H be a random member of $\mathcal{H}$. Then*

$$\mathbf{E}\,\mathrm{D}(H(X),H,Z) - \mathbf{E}\,\mathrm{D}(U_k,H,Z) \leqslant \sqrt{\mathrm{Var}\,\mathrm{D}(U_k,U_{\mathcal{H}},Z)}\cdot\sqrt{2^k\mathrm{CP}(X|Z)+\gamma}, \tag{12}$$

for any real valued function D *on $\{0,1\}^k\times\mathcal{H}\times\mathcal{Z}$. In particular, if $L=\widetilde{\mathbf{H}}_\infty(X|Z)-k$ is the entropy loss then we obtain*

$$\mathbf{E}\,\mathrm{D}(H(X),H,Z) - \mathbf{E}\,\mathrm{D}(U,H,Z) \leqslant \sqrt{\mathrm{Var}\,\mathrm{D}(U_k,U_{\mathcal{H}},Z)}\cdot\sqrt{2^{-L}+\gamma}. \tag{13}$$

▶ Remark. Note that we recover the standard LHL in Equation (8) by applying the above theorem to $\mathrm{D}'=\mathrm{D}-\frac{1}{2}$ where D is $[0,1]$-valued and taking the maximum over D.

3.3 Applications of the generalized LHL

The bound in Theorem 9 introduces the factor $\mathrm{Var}\,\mathrm{D}(U_k,U_{\mathcal{H}},Z)$ depending only on the distinguisher D. If we only want the extracted key $H(X)$ to be (almost) as good as the uniform key U for a *restricted* class of distinguishers[6] D then Equation (13) might offer a significant improvement over Equation (8).

[6] For example, in secure unpredictability applications we assume that D almost always outputs 0.

Security games. In the security game an attacker tries to win against a challenger which uses a key r. To cover the case when the key is extracted by a hash function h and there is some side information z, we assume that the attacker is given also h and z. By $\mathrm{Win}_{\mathcal{A}}(r,h,z)$ we denote the probability that the attacker $\mathcal{A}$ wins the game given h,z and challenged on r.

Unpredictability vs indistinguishability applications. The security requires the advantage of the winning probability over the "trivial strategy"[7] to be small.

▶ **Definition 10** (Security of applications, [1]). Let $c=0$ for an unpredictability and $c=\frac{1}{2}$ for indistinugishability application. The application is (T,ϵ)-secure in the ideal model if $\mathbf{E}\,\mathrm{Win}_{\mathcal{A}}(U,H,Z) \leqslant c+\epsilon$ for all attackers $\mathcal{A}$ with resources (running time, number of oracle questions...) less than T. The application is (T,ϵ')-secure in the real model if $\mathbf{E}\,\mathrm{Win}_{\mathcal{A}}(H(X),H,Z) \leqslant c+\epsilon'$ for all attackers $\mathcal{A}$ with resources less that T.

▶ Remark. In the "ideal" scenario the values of h and z do not help the attacker. That is, for the security in the ideal model it is enough to assume that the winning probability is smaller than $c+\epsilon$ for any $\mathcal{A}$ challenged on r which does not know h and z.

Define $\mathrm{D}(r,h,z)$ to 1 if $\mathcal{A}(h,z)$ wins when challenged on r, and 0 otherwise. Clearly we have $\mathrm{Win}_{\mathcal{A}}(H(X),H,Z)-c-(\mathrm{Win}_{\mathcal{A}}(U,H,Z)-c)=\Delta_{\mathrm{D}}(H(X);U|H,Z)$. Theorem 9 implies

$$\epsilon' \leqslant \epsilon + \sqrt{\mathrm{Var}\left(\mathrm{Win}_{\mathcal{A}}(U_k,U_{\mathcal{H}},Z)\right)}\cdot\sqrt{2^{-L}+\gamma}.$$

If the variance term is not bigger than ϵ, we see that the use of the extracted key and the use of the ideally random key are (roughly) equally secure when $L=\log(1/\epsilon)$.

Reducing the entropy loss by bounding the variance term. For the variance term to be small, the adversary's winning probability in the ideal setting must be concentrated around its mean. This is always the case of unpredictability because then we have $\mathrm{Var}\left(\mathrm{Win}_{\mathcal{A}}(U,H,Z)\right) \leqslant \mathbf{E}\,\mathrm{Win}_{\mathcal{A}}(U,H,Z) \leqslant \epsilon$. However, the case of indistinguishability is more subtle. For example, it might happen that the adversary always wins on one half of the keys and always fails on the second half. Then the variance is equal to $\frac{1}{2}$ and we do not get any improvement in the entropy loss. For more details we refer the reader to [1].

4 The generalized LHL and collision probability: no robustness

We show that the "natural" use of the collision probability to define the conditional collision entropy leads to the notion for which any reasonable chain rule fails.

▶ **Lemma 11** (Chain Rule fails for collision probability). *For every k there exist random variables $X\in\{0,1\}^{k+1}$ and $Z\in\{0,1\}$ such that $\mathrm{CP}(X)=2^{-k}$ but $\mathrm{CP}(X|Z)\geqslant 2^{-\frac{k+1}{2}}$. Thus, if we define $\mathbf{H}'_2(X|Z)=-\log\mathrm{CP}(X|Z)$, then for some X and a bit Z we have $\mathbf{H}'_2(X)=k$ but $\mathbf{H}'_2(X|Z)\approx\frac{k}{2}$.*

Proof. Let $p=2^{-\frac{k+1}{2}}$ and $q=1-p$. Fix a point $x_0\in\{0,1\}^n$ and let $S\subset\{0,1\}^n$ be a set such that $|S|=2^{k+1}q^2$ and does not contain x_0. Let $Z=0$ with probability p and 1 with probability q, and let $X|_{Z=0}$ be the point mass at x_0 and let $X|_{Z=1}$ be uniform on S. Then $\mathrm{CP}(X)=p^2+q^2/\left(2^{k+1}q^2\right)=2^{-k}$ and $\mathrm{CP}(X|Z)=p\cdot 1+\frac{q}{(2^{k+1}q^2)}=\frac{1}{2^{\frac{k+1}{2}}-1}\geqslant 2^{-\frac{k+1}{2}}$. ◀

[7] In unpredictability games, an adversary needs to guess a long string so the trivial guess succeeds with the probability close to $c=0$. In indistinguishability games he needs to guess a bit, thus the trivial guess wins with probability $c=\frac{1}{2}$.

5 Generalized LHL works with conditional collision entropy

▶ **Theorem 12.** *Let X be an n-bit random variable, $\mathcal{H}$ be a $\frac{1+\gamma}{2^k}$-universal family from n to k bits and H be a random member of $\mathcal{H}$. Suppose that we have an application which is (T,ϵ)-secure in the ideal model. Then the same application is (T',ϵ')-secure in the real model where $\epsilon' \leqslant \sqrt{\epsilon 2^{k-\mathbf{H}_2(X|Z)}} + \sqrt{\epsilon\gamma}$ and $T' \approx T$ (against non-uniform attackers).*

Proof. Define $\mathrm{D}(r,h,z) = \mathrm{Win}_A(r,h,z) - c$. For every fixed z, by Theorem 7 and Theorem 8

$$\mathbf{E}\,\mathrm{D}(H(X|_{Z=z},H,z) - \mathbf{E}\,\mathrm{D}(U,H,z) \leqslant \sqrt{\mathrm{Var}\,\mathrm{D}(U,H,z)} \cdot \sqrt{2^k \mathrm{CP}\,(X|_{Z=z}) + \gamma} \tag{14}$$

Since H is independent of X we have $\mathrm{Var}\,\mathrm{D}(U,H,z) \leqslant \max_h \mathrm{Var}\,\mathrm{D}(U,h,z)$ which is smaller by ϵ by the assumptions (applied to D with fixed h and z). Hence,

$$\begin{aligned}\mathbf{E}\,\mathrm{D}(H(X|_{Z=z},H,z) - \mathbf{E}\,\mathrm{D}(U,H,z) &\leqslant \sqrt{\epsilon} \cdot \sqrt{2^k \mathrm{CP}\,(X|_{Z=z}) + \gamma} \\ &\leqslant \sqrt{\epsilon} \cdot \sqrt{2^k \mathrm{CP}\,(X|_{Z=z})} + \sqrt{\epsilon\gamma}. \end{aligned} \tag{15}$$

Since $\mathbf{E}_{z\leftarrow Z} \sqrt{\mathrm{CP}(X|_{Z=z})} = 2^{-\frac{1}{2}\mathbf{H}_2(X|Z)}$, the result follows. ◀

6 Conclusion

Our result extend all the results of [1] for the case of collision entropy, which is much less restrictive and still provides the leakage robustness because of the chain rule. Moreover, our results strongly support the belief that this notion of collision entropy is the "right" one.

References

1 Boaz Barak, Yevgeniy Dodis, Hugo Krawczyk, Olivier Pereira, Krzysztof Pietrzak, Francois-Xavier Standaert, and Yu Yu. Leftover hash lemma, revisited. Cryptology ePrint Archive, Report 2011/088, 2011. `http://eprint.iacr.org/`.

2 Stefan Berens. Conditional renyi entropy. Master's thesis, Mathematisch Instituut, Universiteit Leiden, 2013.

3 Yevgeniy Dodis and Yu Yu. Overcoming weak expectations. In Amit Sahai, editor, *Theory of Cryptography*, volume 7785 of *Lecture Notes in Computer Science*, pages 1–22. Springer Berlin Heidelberg, 2013.

4 Johan Hastad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.

5 Jaikumar Radhakrishnan and Amnon Ta-Shma. Bounds for dispersers, extractors, and depth-two superconcentrators. *SIAM Journal on Discrete Mathematics*, 13:2000, 2000.

6 R. Renner and S. Wolf. Smooth Renyi entropy and applications. In *International Symposium on Information Theory, 2004. ISIT 2004. Proceedings.*, page 232. IEEE, 2004.

7 D.R. Stinson. Universal hashing and authentication codes. *Designs, Codes and Cryptography*, 4(3):369–380, 1994.