

A trivial debiasing scheme for Helper Data Systems

Boris Škorić

b.skoric@tue.nl

Abstract

We introduce a debiasing scheme that solves the more-noise-than-entropy problem which can occur in Helper Data Systems when the source is very biased. We perform a condensing step, similar to Index Based Syndrome coding, that reduces the size of the source space in such a way that some source entropy is lost while the noise entropy is greatly reduced. In addition, our method allows for even more entropy extraction by means of a ‘spamming’ technique. Our method outperforms solutions based on the one-pass von Neumann algorithm.

1 Introduction

1.1 Helper Data Systems

The past decade has seen a lot of interest in the field of *security with noisy data*. In several security applications it is necessary to reproducibly extract secret data from noisy measurements on a physical system. One such application is read-proof storage of cryptographic keys using Physical Unclonable Functions (PUFs) [17, 18, 16, 4, 14]. Another application is the privacy-preserving storage of biometric data.

Storage of keys in nonvolatile digital memory can often be considered insecure because of the vulnerability to physical attacks. (For instance, fuses can be optically inspected with a microscope; flash memory may be removed and read out.) PUFs provide an alternative way to store keys, namely in analog form, which allows the designer to exploit the inscrutability of analog physical behavior. Keys stored in this way are sometimes referred to as Physically Obfuscated Keys (POKs) [10]. In both the biometrics and the PUF/POK case, one faces the problem that some form of error correction has to be performed, but under the constraint that the redundancy data, which is visible to attackers, does not endanger the secret extracted from the physical measurement. This problem is solved by a special security primitive, the *Helper Data System* (HDS). A HDS in its most general form is shown in Fig. 1. The **Gen** procedure takes as input a measurement X . It outputs a secret S and (public) Helper Data W . The helper data is stored. In the reconstruction phase, a fresh measurement X' is obtained. Typically X' is a noisy version of X , close to X (in terms of e.g. Euclidean distance or Hamming distance) but not necessarily identical. The **Rec** (reconstruction) procedure takes X' and W as input. It outputs \hat{S} , an estimate of S . If X' is not too noisy then $\hat{S} = S$.

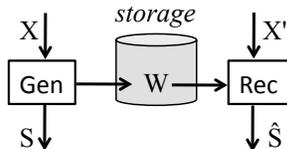


Figure 1: *Data flow in a generic Helper Data System.*

Two special cases of the general HDS are the Secure Sketch (SS) and the Fuzzy Extractor (FE) [9]. The Secure Sketch has $S = X$ (and $\hat{S} = \hat{X}$, an estimator for X). If X is not uniformly distributed, then S is not uniform. The SS is suitable for privacy-preserving biometrics, where high entropy of S (given W) is required, but not uniformity. The Fuzzy Extractor is required to have a (nearly) uniform S given W . The FE is typically used for extracting keys from PUFs and POKs.

1.2 The Code Offset Method (COM)

By way of example we briefly present the Code Offset Method [3, 12, 9, 8, 19], the oldest and most well-known HDS. The COM makes use of a linear code \mathcal{C} with k -bit messages and n -bit codewords. The syndrome function is denoted as $\mathbf{Syn} : \{0, 1\}^n \rightarrow \{0, 1\}^{n-k}$. The code supports a syndrome decoding algorithm $\mathbf{SynDec} : \{0, 1\}^{n-k} \rightarrow \{0, 1\}^n$ that maps syndromes to error patterns. Use is also made of a key derivation function $\mathbf{KeyDeriv} : \{0, 1\}^n \times \{0, 1\}^* \rightarrow \{0, 1\}^\ell$, with $\ell \leq k$. Below we show how the COM can be used as a FE for either uniform or non-uniform source X .

Enrollment:

The enrollment measurement gives $X \in \{0, 1\}^n$. The helper data is computed as $W = \mathbf{Syn} X$. The key is computed as $S = \mathbf{KeyDeriv}(X, R)$, where R is (optional) public randomness. The W and R are stored.

Reconstruction:

A fresh measurement of the PUF gives $X' \in \{0, 1\}^n$. The estimator for X is computed as $\hat{X} = X' \oplus \mathbf{SynDec}(W \oplus \mathbf{Syn} X')$ and the reconstructed key is $\hat{S} = \mathbf{KeyDeriv}(\hat{X}, R)$.

1.3 The problem of bias

After seeing the helper data W as specified above, an attacker has uncertainty $\mathbf{H}(X|W)$ about X .¹ As W is a function of X , we have $\mathbf{H}(X|W) = \mathbf{H}(X|\mathbf{Syn} X) = \mathbf{H}(X) - \mathbf{H}(\mathbf{Syn} X)$. Let us consider the simplest possible noise model, the Binary Symmetric Channel (BSC). In the case of a BSC with bit error rate β , the code's redundancy has to satisfy $n - k \geq nh(\beta)$ in order for the error-correction to work. (Here h denotes the binary entropy function, $h(\beta) = -\beta \log \beta - (1 - \beta) \log(1 - \beta)$, with 'log' the base 2 logarithm). More generally, $\mathbf{H}(\mathbf{Syn} X)$ has at least to be equal to the entropy of the noise. Hence the COM helper data in the BSC case leaks at least $nh(\beta)$ bits of information about X . This becomes a problem when X itself does not have much entropy, which occurs for instance if the bits in X are highly biased [11, 13]. Note that the problem is not fundamental: the secrecy capacity is $nh(p + \beta - 2p\beta) - nh(\beta)$, which is positive.

A solution was proposed by Maes et al. [15]. Their approach is to combine debiasing with error correction. For the debiasing they use the von Neumann algorithm in a single-pass or multi-pass manner. Their helper data comprises a selection 'mask' that helps the reconstruction algorithm to identify the locations where von Neumann should yield output.

In this paper we follow a simpler approach similar to the Index Based Syndrome (IBS) method [20]. In IBS the helper data consists of an ordered list of pointers to locations in X ; the content of X in those locations together forms a codeword. IBS avoids the leakage caused by publishing error-correcting-code syndromes.

1.4 Contributions and outline

We introduce an alternative solution to the bias problem in helper data systems. We follow the *condense-then-fuzzy-extract* philosophy proposed by Canetti et al. [6] as one of the available options when faced with 'more noise than entropy' scenarios. Condensing means mapping the source variable X to a smaller space such that most of the entropy of X is retained, but the noise entropy is greatly reduced. Our way of condensing the source is to restrict X to the bit positions \mathcal{U} where $\mathcal{U} \subset \{1, \dots, n\}$ is a random subset containing all the rare symbols. The set \mathcal{U} becomes part of the helper data.

Our \mathcal{U} bears some similarity to the von Neumann mask in [15], but there are important differences. (i) The size of \mathcal{U} is tunable. (ii) We can extract source information based on the legitimate party's ability to distinguish \mathcal{U} from fake instances of \mathcal{U} when a 'spamming' technique similar to [19] is applied.

The outline of this paper is as follows. Section 2 gives the details of the scheme. Section 3 presents an analysis of the extractable entropy and of the practicality of the 'spamming' option. Section 4 summarises and suggests future work.

¹The notation \mathbf{H} stands for Shannon entropy. For information-theoretic concepts see e.g. [7].

2 Debiasing based on subset selection

We will use the following notation. The set $\{1, \dots, n\}$ is written as $[n]$. The notation $X_{\mathcal{U}}$ means X restricted to the positions specified in \mathcal{U} . Set difference is written as \setminus . A string consisting of n zeroes is written as 0^n . The Hamming weight of X is denoted as $w(X)$. We will consider a source $X \in \{0, 1\}^n$ made up of i.i.d. bits X_i following a Bernoulli distribution with parameter p , i.e. $\Pr[X_i = 1] = p$. Without loss of generality we take $p \in (0, \frac{1}{2})$. In particular we are interested in the case $p < \beta$ where direct application of the COM fails. The notation ‘log’ stands for the base-2 logarithm. Information distance (Kullback-Leibler divergence) is denoted as $D(p||q) = p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}$ for $p, q \in (0, 1)$.

2.1 The scheme

Below we present a barebones version of our proposed scheme. We omit details concerning the protection of the stored data. There are well known ways to protect helper data, using either Public Key Infrastructure or one-way functions [5]. We also omit details that have to do with the verification of the reconstructed key. These details are trivially added.

System setup

The following system parameters are fixed. An integer u satisfying $np < u < n$, representing the size of \mathcal{U} ; a list length L ; a pseudorandom generator f that takes as input a seed σ and a counter j , and outputs a subset $f(\sigma, j) \subset [n]$ such that $|f(\sigma, j)| = u$; a Secure Sketch (**Gen**, **Rec**) that acts on a source in $\{0, 1\}^u$ and is able to handle bit error rate β ; a key derivation function $\text{KDF} : \{0, 1\}^u \times [L] \rightarrow \{0, 1\}^\ell$. All these parameters are public.

Enrollment

- E1. Measure $X \in \{0, 1\}^n$.
- E2. Draw a random subset $\mathcal{U} \subset [n]$ of size u such that $X_{[n] \setminus \mathcal{U}} = 0^{n-u}$.
- E3. Compute $Y = X_{\mathcal{U}}$ and $W = \text{Gen}(Y)$.
- E4. Draw a random seed σ .
- E5. Draw a random $z \in [L]$. Determine a permutation $\pi : [n] \rightarrow [n]$ such that² $\pi(\mathcal{U}) = f(\sigma, z)$.
- E6. Derive the secret key as $S = \text{KDF}(Y, z)$.
- E7. Store σ, π, W .

Reconstruction

- R1. Read σ, π, W .
- R2. Measure $X' \in \{0, 1\}^n$.
- R3. Construct a set $\mathcal{M}' = \{i \in [n] : X'_i = 1\}$. Compute $\mathcal{M} = \pi(\mathcal{M}')$.
- R4. Construct the list $\mathcal{L} = (f(\sigma, j))_{j=1}^L$.
- R5. Determine the index $\hat{z} \in [L]$ for which $\mathcal{L}_{\hat{z}}$ has the largest overlap with \mathcal{M} .
- R6. Compute $\hat{\mathcal{U}} = \pi^{\text{inv}}(\mathcal{L}_{\hat{z}})$ and $\hat{Y} = \text{Rec}(X'_{\hat{\mathcal{U}}}, W)$.
- R7. Reconstruct the key as $\hat{S} = \text{KDF}(\hat{Y}, \hat{z})$.

There is a small probability that $w(X) > u$ occurs, which would cause a problem in step E2. However, we will typically consider $u \geq 2np$, in which case the probability that step E2 fails is exponentially small.

² $\pi(\mathcal{U})$ means π applied to each element of \mathcal{U} individually.

2.2 Explanation of the scheme

The effect of steps E4,E5 is to create a list of \mathcal{U} -candidates, of which only entry z is correct. To an attacker (who knows u but does not know X or X') the L candidates are indistinguishable.³ Steps R3–R5 allow for a quick search to identify the index z of the correct \mathcal{U} -candidate. Note that the reconstruction algorithm compares a permuted \mathcal{M} to \mathcal{L} -entries instead of \mathcal{M} to permuted \mathcal{L} -entries; this improves speed. To further optimize for speed, steps R4 and R5 can be combined to select good z values on the fly as soon as a new \mathcal{L}_j is generated.

Note that extremely fast pseudorandom generators exist which spew out more than 8 bits per clock cycle [1, 2]. This makes it practical to work with large values of L , as long as not too many plausible z -candidates are generated. (See Section 3.3).

We did not explicitly specify how to map pseudorandom numbers to a size- u subset of $[n]$. A very straightforward algorithm would be to keep uniformly picking locations in $[n]$, until u different ones have been found.

We did not specify an algorithm for determining the permutation π , nor did we specify in which form π is stored. These are minor details and have no impact on the overall efficiency of the scheme, since steps E5 and R3 are performed only once. The computational bottleneck is R4,R5. Note that inputting z into the key derivation function increases the entropy of S by $\log L$ bits.

3 Analysis

3.1 Entropy after condensing

The Hamming weight of X carries very little information. Let's assume for the moment that $T = w(X) \in \{0, \dots, u\}$ is known to the attacker⁴, just to simplify the analysis.

Even if the attacker knows t and \mathcal{U} (i.e. z), there are $\binom{u}{t}$ equally probable possibilities for Y . Hence

$$\mathsf{H}(Y|Z = z, T = t) = \log \binom{u}{t} > uh\left(\frac{t}{u}\right) - \frac{1}{2} \log \frac{t(u-t)}{u} - \log \frac{e^2}{\sqrt{2\pi}}. \quad (1)$$

The inequality follows from Stirling's approximation.

As (1) does not depend on z , the entropy $\mathsf{H}(Y|T = t)$ is also given by (1). A lower bound on $\mathsf{H}(Y|T)$ is obtained by taking the expectation over t . This turns out to be rather messy, since the distribution of t is a truncated binomial. (It is given that $t \leq u$, while originally $w(X) \in \{0, \dots, n\}$). As t equals approximately np on average, the result is $\mathsf{H}(Y|T) \approx uh\left(\frac{np}{u}\right)$. A more precise lower bound is given in Theorem 1 below.

Theorem 1 *Let δ be defined as*

$$\delta = \min\{e^{-2np^2\left(\frac{u}{np}-1\right)^2}, e^{-np\frac{1}{3}\left(\frac{u}{np}-1\right)}, e^{-nD\left(\frac{u}{n}\|p\right)}\}. \quad (2)$$

Let $\delta < p$. Let $u \geq 2np/(1-\delta)$. Then the entropy of Y given T can be lowerbounded as

$$\begin{aligned} \mathsf{H}(Y|T) &> uh\left(\frac{np-n\delta}{u}\right) - \frac{1}{2} \log \frac{np}{1-\delta} - \log \frac{e^2}{\sqrt{2\pi}} + \frac{1}{2\ln 2} \frac{np-n\delta}{u} \\ &\quad - u \left\{ \frac{1-p}{np} + \frac{2\delta}{p} + \delta + \frac{\delta^3}{p^3} \right\} (1-\delta)^{-1} \left(1 - \frac{\delta}{p}\right)^{-2}. \end{aligned} \quad (3)$$

Proof: The proof is rather tedious and can be found in Appendix A. \square

The entropy of Y is obtained as follows,

$$\mathsf{H}(Y) = \mathsf{H}(YT) = \mathsf{H}(T) + \mathsf{H}(Y|T). \quad (4)$$

The $\mathsf{H}(T)$ is the entropy of the *truncated* binomial distribution.

³Given the i.i.d. model for creating X , the set \mathcal{U} itself is uniformly random. If we want a different model for X , e.g. with asymmetries between the positions, then \mathcal{L} will need to be generated in a way that follows the statistics of X .

⁴In the security analysis we denote random variables using capitals, and their numerical realisations in lowercase.

Theorem 2 Let $q_t = \binom{n}{t} p^t (1-p)^{n-t}$ denote the probabilities in the full binomial distribution. Let δ be defined as in Theorem 1.

$$H(T) > \log \frac{2\pi}{e} \sqrt{np(1-p)} + \log(1-\delta) - (n-u)q_u \log \frac{1}{q_u}. \quad (5)$$

Proof: See Appendix B. □

Theorem 3 Let δ be defined as in Theorem 1. Let $\delta < p$. Let $u \geq 2np/(1-\delta)$.

$$H(Y) > uh\left(\frac{np-n\delta}{u}\right) + \frac{1}{2\ln 2} \frac{np-n\delta}{u} - \log \frac{e^3}{(2\pi)^{3/2}} + \frac{1}{2} \log(1-p) + \frac{3}{2} \log(1-\delta) \\ - (n-u)q_u \log \frac{1}{q_u} - u \left\{ \frac{1-p}{np} + \frac{2\delta}{p} + \delta + \frac{\delta^3}{p^3} \right\} (1-\delta)^{-1} (1-\frac{\delta}{p})^{-2}. \quad (6)$$

Proof: Follows directly from Theorems 1 and 2 by adding $H(T) + H(Y|T)$. □

The complicated expression (6) can be well approximated by $uh(np/u) - u/np$.

Note that the difference between the original source entropy $H(X) = nh(p)$ and the condensed form $H(Y) \approx uh(np/u) - u/np$ is considerable. For example, setting $n = 640$, $p = 0.1$, $u = 128$ yields $H(Y) \approx 126$ and $H(X) \approx 300$. A small part of this huge difference can be regained using the trick with the entropy of Z . The practical aspects are discussed in Section 3.3.

Note also that our scheme outperforms simple von Neumann debiasing a factor of at least 2. The von Neumann algorithm takes $n/2$ pairs of bits; each pair has a probability $2p(1-p)$ of generating a (uniform) output bit; hence the extracted entropy is $np(1-p) < np$. In our scheme, if we set $u \approx 2np$ we get $H(Y) \approx 2np - 2$. Furthermore, $H(Y)$ is an increasing function of u .

3.2 Fuzzy Extraction after condensing

The Code Offset Method applied to $Y \in \{0,1\}^u$ leaks at least $uh(\beta)$ bits of information about Y . In case of a ‘perfect’ error-correcting code, the length of a noise-robust key reconstructed with the COM is $H(Y) - uh(\beta)$. In Fig. 2 we plot $H(Y) - uh(\beta)$ for some example parameter settings. (Note that in all cases shown here $\beta \geq p$ holds; the COM acting on the original source X would be unable to extract any entropy.) Clearly there is an optimum (n, u) for given (p, β) .

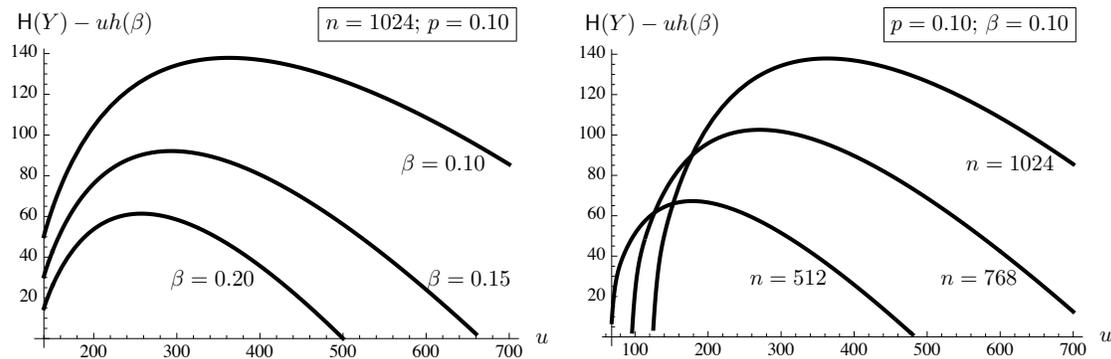


Figure 2: The key length $H(Y) - uh(\beta)$ that the Code Offset Method can reproducibly extract after the condensing step, plotted as a function of u for various parameter values. For $H(Y)$ the bound in Theorem 3 is used.

3.3 The list size L

We briefly discuss how large L can be made before the reconstruction procedure sketched in Section 2 starts to produce too many candidates for z . We define $\tilde{p} = p + \beta - 2p\beta$.

On the one hand there is the number of ‘1’ symbols in $X'_{\mathcal{U}}$ for the **correct** \mathcal{U} . The number of 1’s in $X_{\mathcal{U}}$ is on average np . Of these, on average $np(1 - \beta)$ will be a ‘1’ in X' . The $(u - np)$ zeroes in $X_{\mathcal{U}}$ will generate on average $(u - np)\beta$ 1’s in X' . Hence the number of 1’s in $X'_{\mathcal{U}}$ is expected to be around $np(1 - \beta) + (u - np)\beta = np + (u - 2np)\beta$.

On the other hand there is the number of 1’s for **incorrect** \mathcal{U} -candidates. The number of 1’s in X' is approximately $n\tilde{p}$. We pretend that $n\tilde{p}$ is integer, for notational simplicity. We denote by $A \in \{0, \dots, n\tilde{p}\}$ the number of 1’s in $X'_{\mathcal{V}}$ for a *randomly chosen* subset \mathcal{V} , with $|\mathcal{V}| = u$. The A follows a hypergeometric probability distribution

$$\Pr[A = a] = \frac{\binom{n\tilde{p}}{a} \binom{n-n\tilde{p}}{u-a}}{\binom{n}{u}} = \frac{\binom{u}{a} \binom{n-u}{n\tilde{p}-a}}{\binom{n}{n\tilde{p}}}. \quad (7)$$

The first expression looks at the process of selecting u out of n positions with exactly a 1’s hitting the $n\tilde{p}$ existing 1’s in X' ; the second expression looks at the process of selecting $n\tilde{p}$ positions in X' such that exactly a of them lie in \mathcal{V} . We have $\mathbb{E}_a a = u\tilde{p}$ and $\mathbb{E}_a (a - u\tilde{p})^2 = u\tilde{p}(1 - \tilde{p})(n - u)/(n - 1) < u\tilde{p}$. In other words, a is sharply peaked around $u\tilde{p}$.

We can put a threshold θ somewhere in the gap between $u\tilde{p}$ and $np + (u - 2np)\beta$, and declare a \mathcal{U} -candidate \mathcal{V} to be bad if the Hamming weight of $X'_{\mathcal{V}}$ is lower than θ . Let’s set $\theta = np + (u - 2np)\beta - c \cdot \sqrt{n\tilde{p}}$ with c sufficiently large to avoid false negatives (i.e. missing the correct \mathcal{U}). In a way analogous to (2), we can bound the false positive probability as

$$\Pr[A \geq \theta] < \min\{e^{-2u\tilde{p}^2(\theta/u\tilde{p}-1)^2}, e^{-u\tilde{p}\frac{1}{3}(\theta/u\tilde{p}-1)}, e^{-uD(\frac{\theta}{u}||\tilde{p})}\}. \quad (8)$$

These bounds are obtained by applying (2) and replacing $u \rightarrow \theta$, $n \rightarrow u$, $p \rightarrow \tilde{p}$. Fig. 3 shows how many bits of entropy ($\approx -\log \Pr[A > \theta]$) can be obtained from z without running into false positives in step R5. To extract b bits of entropy, a list length $L = 2^b$ is needed. Fig. 3 serves just to show the orders of magnitude and is by no means an exhaustive treatment of the whole parameter space n, p, β, u, θ . We remark that the curves depend quite strongly on the threshold parameter c .

It is important to note that it is perfectly possible to make L extremely large. Then many false positives occur, but this is not a fundamental problem. It requires extra work: one key reconstruction and one verification (e.g. of a key hash) per false positive. Depending on the available n , the computing platform etc. this may be a viable option.

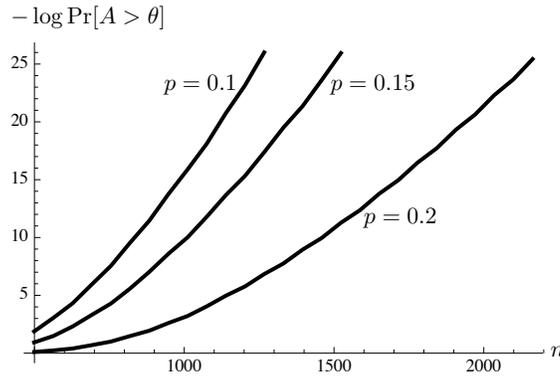


Figure 3: Plots of $-\log \Pr[A > \theta]$ as a function of n for $\beta = p$, $u = 2.25np$, and $\theta = np + (u - 2np)\beta - c\sqrt{n\tilde{p}}$ with $c = 4$. The $\Pr[A > \theta]$ was obtained by numerical summation of (7).

4 Summary

We have introduced a method for source debiasing that can be used in Helper Data Systems to solve the ‘more noise than entropy’ problem. Our method applies the condense-then-fuzzy-extract idea [6] in a particularly simple way: the space $\{0,1\}^n$ is condensed to $\{0,1\}^u$ in such a way that all the rare symbols are kept; meanwhile the noise entropy is reduced from $nh(\beta)$ to $uh(\beta)$. Theorem 3 gives a lower bound on the retained entropy $H(Y)$. Furthermore, there is the option of extracting additional entropy from the index z , which points at the real subset \mathcal{U} among the fakes. Even in its bare form, without the fake subsets, our method outperforms basic von Neumann debiasing by factor of at least 2.

Fig. 2 shows that after the condensing step the Code Offset Method can extract significant entropy in a situation where the bare COM fails. It also shows the tradeoff between the reduction of source entropy and noise entropy as u varies.

In Section 3.3 we did a very preliminary analysis of the practicality of extracting information from the index z . More work is needed to determine how this works out for real-world parameter values n, p, β and to see how the computations in steps R4 and R5 can be optimised for speed. Another topic for future work is to study different noise models such as asymmetric noise.

Acknowledgments

We thank Sebastian Verschoor and Frans Willems for useful discussions.

A Proof of Theorem 1

We start with a number of definitions and supporting lemmas. We define the binomial probability $q_t = \binom{n}{t} p^t (1-p)^{n-t}$. We define $\Delta = \Pr[w(X) > u] = \sum_{t=u+1}^n q_t$ and $\pi_t = q_t / (1 - \Delta)$ for $t \leq u$, such that the vector (π_t) is the probability distribution of t from the attacker’s point of view. The notation \mathbb{E}_t will refer to the binomial distribution $(q_t)_{t=0}^n$, while $\tilde{\mathbb{E}}$ will refer to the truncated binomial $(\pi_t)_{t=0}^u$.

Lemma 1 *Let $k \geq 1, n \geq 2$.*

$$nh\left(\frac{k}{n}\right) - \frac{1}{2} \log \frac{k(n-k)}{n} - \log \frac{e^2}{\sqrt{2\pi}} \leq \log \binom{n}{k} \leq nh\left(\frac{k}{n}\right) - \frac{1}{2} \log \frac{k(n-k)}{n} - \log \frac{2\pi}{e}. \quad (9)$$

Proof: Follows directly from Stirling’s approximation. \square

Lemma 2 *Let $u > 2np$.*

$$\Delta \leq \delta \quad \text{with} \quad \delta \stackrel{\text{def}}{=} \min\{e^{-2np^2(\frac{u}{np}-1)^2}, e^{-np\frac{1}{3}(\frac{u}{np}-1)}, e^{-nD(\frac{u}{n}||p)}\}. \quad (10)$$

Proof: These are the inequalities of Hoeffding, Chernoff, Chernoff-Hoeffding respectively. The listed form of the Chernoff inequality needs $u > 2np$. \square

Lemma 3

$$\tilde{\mathbb{E}}_t t > n(p - \delta). \quad (11)$$

Proof:

$$\tilde{\mathbb{E}}_t t = \frac{1}{1-\Delta} \sum_{t=0}^u q_t t > \sum_{t=0}^u q_t t = np - \sum_{t=u+1}^n q_t t > np - \sum_{t=u+1}^n q_t n = np - n\Delta \geq np - n\delta. \quad (12)$$

In the last step we used Lemma 2. \square

Lemma 4

$$\tilde{\mathbb{E}}_t t < \frac{np}{1-\delta}. \quad (13)$$

Proof:

$$\tilde{\mathbb{E}}_t t = \frac{1}{1-\Delta} \sum_{t=0}^u q_t t < \frac{1}{1-\Delta} \sum_{t=0}^n q_t t = \frac{np}{1-\Delta} \leq \frac{np}{1-\delta}. \quad (14)$$

In the last step we used Lemma 2. \square

Lemma 5

$$\tilde{\mathbb{E}}_t t^2 < \frac{(np)^2 + np(1-p)}{1-\delta}. \quad (15)$$

Proof:

$$\tilde{\mathbb{E}}_t t^2 = \frac{1}{1-\Delta} \sum_{t=0}^u q_t t^2 < \frac{1}{1-\Delta} \mathbb{E}_t t^2 = \frac{(np)^2 + np(1-p)}{1-\Delta} \leq \frac{(np)^2 + np(1-p)}{1-\delta}. \quad (16)$$

In the last step we used Lemma 2. \square

Lemma 6 *Let $p \in [0, 1]$. Let $r \in (0, \frac{1}{2}]$. Then it holds that*

$$h(p) \geq \Omega_r(p) \quad , \quad \Omega_r(p) = h(r) + (p-r)h'(r) - \frac{(p-r)^2}{r^2} [h(r) - rh'(r)]. \quad (17)$$

The expression $h(r) - rh'(r)$ is a nonnegative increasing function on $r \in (0, \frac{1}{2}]$.

Proof: Ω_r is a parabola constructed such that $\Omega_r(0) = 0$, $\Omega_r(r) = h(r)$ and $\Omega_r'(r) = h'(r)$. The property $h(p) \geq \Omega_r(p)$ is verified by visual inspection. We define $g(r) = h(r) - rh'(r)$. We have $\lim_{r \rightarrow 0} g(r) = 0$. Furthermore $\frac{d}{dr} g(r) = -rh''(r) > 0$, which proves that g is increasing. Together with $g(0) = 0$ that implies that $g(r)$ is nonnegative on the given interval. \square

Lemma 7 *Let $\delta < p$ and $u \geq 2np/(1-\delta)$.*

$$\tilde{E}_t h\left(\frac{t}{u}\right) > h\left(\frac{np-n\delta}{u}\right) - \left\{ \frac{1-p}{np} + \frac{2\delta}{p} + \delta + \frac{\delta^3}{p^3} \right\} (1-\delta)^{-1} \left(1 - \frac{\delta}{p}\right)^{-2}. \quad (18)$$

Proof: We start from Lemma 6 to expand $h(\frac{t}{u})$ around $t = \tilde{\mathbb{E}}_t t$,

$$h\left(\frac{t}{u}\right) \geq h\left(\frac{\tilde{\mathbb{E}}_t t}{u}\right) + \text{linear} - \frac{(t - \tilde{\mathbb{E}}_t t)^2}{(\tilde{\mathbb{E}}_t t)^2} [h(\frac{\tilde{\mathbb{E}}_t t}{u}) - \tilde{\mathbb{E}}_t [\frac{t}{u}] h'(\frac{\tilde{\mathbb{E}}_t t}{u})]. \quad (19)$$

When we take the expectation $\tilde{\mathbb{E}}_t$, the term linear in $t - \tilde{\mathbb{E}}_t t$ disappears,

$$\tilde{\mathbb{E}}_t h\left(\frac{t}{u}\right) \geq h\left(\frac{\tilde{\mathbb{E}}_t t}{u}\right) - \left[\frac{\tilde{\mathbb{E}}_t t^2}{(\tilde{\mathbb{E}}_t t)^2} - 1 \right] \left[h\left(\frac{\tilde{\mathbb{E}}_t t}{u}\right) - \tilde{\mathbb{E}}_t \left[\frac{t}{u} \right] h' \left(\frac{\tilde{\mathbb{E}}_t t}{u} \right) \right]. \quad (20)$$

We use $\tilde{\mathbb{E}}_t t < u/2$ to bound the second occurrence of $h(\frac{\tilde{\mathbb{E}}_t t}{u})$ as $h(\dots) < 1$ and to use $h' > 0$. For the first occurrence of $h(\frac{\tilde{\mathbb{E}}_t t}{u})$ we use that h is an increasing function and apply Lemma 3.

$$\tilde{\mathbb{E}}_t h\left(\frac{t}{u}\right) > h\left(\frac{np-n\delta}{u}\right) - \left[\frac{\tilde{\mathbb{E}}_t t^2}{(\tilde{\mathbb{E}}_t t)^2} - 1 \right]. \quad (21)$$

Finally we bound $\tilde{\mathbb{E}}_t t^2$ using Lemma 5 and we bound $\tilde{\mathbb{E}}_t t$ using Lemma 3. \square

Lemma 8

$$\tilde{\mathbb{E}}_t \log \binom{u}{t} > u \tilde{\mathbb{E}}_t h\left(\frac{t}{u}\right) - \log \frac{e^2}{\sqrt{2\pi}} - \frac{1}{2} \log \frac{np}{1-\delta} + \frac{1}{2 \ln 2} \frac{np - n\delta}{u}. \quad (22)$$

Proof: We start from Lemma 1 and write

$$\log \binom{u}{t} \geq uh\left(\frac{t}{u}\right) - \frac{1}{2} \log \frac{t(u-t)}{u} - \log \frac{e^2}{\sqrt{2\pi}}. \quad (23)$$

We expand the second term as $-\frac{1}{2} \log \frac{t(u-t)}{u} = -\frac{1}{2} \log t - \frac{1}{2} \log(1 - \frac{t}{u}) > -\frac{1}{2} \log t + \frac{1}{2 \ln 2} \frac{t}{u}$. Then we apply $\tilde{\mathbb{E}}_t$. Jensen followed by Lemma 4 gives $\tilde{\mathbb{E}}_t \log t \leq \log \tilde{\mathbb{E}}_t t < \log \frac{np}{1-\delta}$. Lemma 3 gives $\mathbb{E}_t \frac{t}{u} > \frac{np-n\delta}{u}$. \square

With all these lemmas we can finally prove Theorem 1. From (1) we have

$$\mathbb{H}(Y|T) = \tilde{\mathbb{E}}_t \mathbb{H}(Y|T=t) = \tilde{\mathbb{E}}_t \log \binom{u}{t}. \quad (24)$$

We apply Lemma 8, and then Lemma 7 to lowerbound $\tilde{\mathbb{E}}_t h\left(\frac{t}{u}\right)$.

B Proof of Theorem 2

We begin by lowerbounding the entropy of the un-truncated distribution $(q_t)_{t=0}^n$.

Lemma 9

$$\mathbb{H}(\mathbf{q}) \geq \log \frac{2\pi}{e} \sqrt{np(1-p)}. \quad (25)$$

Proof:

$$\begin{aligned} \mathbb{H}(\mathbf{q}) &= -\mathbb{E}_t \log \left[\binom{n}{t} p^t (1-p)^{n-t} \right] = (-\mathbb{E}_t t) \log p - (n - \mathbb{E}_t t) \log(1-p) - \mathbb{E}_t \log \binom{n}{t} \\ &= nh(p) - \mathbb{E}_t \log \binom{n}{t} \geq nh(p) - \log \binom{n}{\mathbb{E}_t t} = nh(p) - \log \binom{n}{np}. \end{aligned} \quad (26)$$

In the last line we used Jensen's inequality for the function $\log \binom{n}{t}$. We use Lemma 1 to upperbound $\log \binom{n}{np}$. \square

Now we write

$$\begin{aligned} \mathbb{H}(T) &= \tilde{\mathbb{E}}_t \log \frac{1-\Delta}{q_t} = \log(1-\Delta) + \frac{1}{1-\Delta} \sum_{t=0}^u q_t \log \frac{1}{q_t} \\ &> \log(1-\delta) + \sum_{t=0}^u q_t \log \frac{1}{q_t} = \log(1-\delta) + \mathbb{H}(\mathbf{q}) - \sum_{t=u+1}^n q_t \log \frac{1}{q_t} \\ &> \log(1-\delta) + \mathbb{H}(\mathbf{q}) - \sum_{t=u+1}^n q_u \log \frac{1}{q_u} = \log(1-\delta) + \mathbb{H}(\mathbf{q}) - (n-u)q_u \log \frac{1}{q_u}. \end{aligned} \quad (27)$$

In the last inequality we used that $q_t \log \frac{1}{q_t}$ is a decreasing function of t . Finally we lowerbound $\mathbb{H}(\mathbf{q})$ with Lemma 9.

References

- [1] <http://gjrands.sourceforge.net/>.
- [2] <http://www.digicortex.net/node/22>.

- [3] C.H. Bennett, G. Brassard, C. Crépeau, and M. Skubiszewska. Practical quantum oblivious transfer. In *CRYPTO*, pages 351–366, 1991.
- [4] C. Böhm and M. Hofer. *Physical Unclonable Functions in Theory and Practice*. Springer, 2013.
- [5] X. Boyen. Reusable cryptographic fuzzy extractors. In *ACM Conference on Computer and Communications Security*, pages 82–91, 2004.
- [6] R. Canetti, B. Fuller, O. Paneth, L. Reyzin, and A. Smith. Reusable fuzzy extractors for low-entropy distributions. In *Eurocrypt 2016*, 2016. eprint.iacr.org/2014/243.
- [7] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., second edition, 2005.
- [8] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. Fuzzy Extractors: how to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008.
- [9] Y. Dodis, M. Reyzin, and A. Smith. Fuzzy Extractors: How to generate strong keys from biometrics and other noisy data. In *Eurocrypt 2004*, volume 3027 of *LNCS*, pages 523–540. Springer-Verlag, 2004.
- [10] B. Gassend. Physical Random Functions. Master’s thesis, Massachusetts Institute of Technology, 2003.
- [11] T. Ignatenko and F.M.J. Willems. Information leakage in fuzzy commitment schemes. *IEEE Transactions on Information Forensics and Security*, 5(2):337–348, 2010.
- [12] A. Juels and M. Wattenberg. A fuzzy commitment scheme. In *ACM Conference on Computer and Communications Security (CCS) 1999*, pages 28–36, 1999.
- [13] P. Koeberl, J. Li, A. Rajan, and W. Wu. Entropy loss in PUF-based key generation schemes: The repetition code pitfall. In *IEEE International Symposium on Hardware-Oriented Security and Trust (HOST) 2014*, pages 44–49, 2014.
- [14] R. Maes. *Physically Unclonable Functions: Constructions, Properties and Applications*. Springer, 2013.
- [15] R. Maes, V. van der Leest, E. van der Sluis, and F.M.J. Willems. Secure key generation from biased PUFs. In *Cryptographic Hardware and Embedded Systems (CHES) 2015*, volume 9293 of *LNCS*, pages 517–534. Springer, 2015.
- [16] A.-R. Sadeghi and D. Naccache, editors. *Towards hardware-intrinsic security*. Springer, 2010.
- [17] P. Tuyls, G.-J. Schrijen, B. Škorić, J. van Geloven, R. Verhaegh, and R. Wolters. Read-proof hardware from protective coatings. In *Cryptographic Hardware and Embedded Systems (CHES) 2006*, volume 4249 of *LNCS*, pages 369–383. Springer-Verlag, 2006.
- [18] P. Tuyls, B. Škorić, and T. Kevenaar. *Security with Noisy Data: Private Biometrics, Secure Key Storage and Anti-Counterfeiting*. Springer, London, 2007.
- [19] B. Škorić and N. de Vreede. The Spammed Code Offset Method. *IEEE Transactions on Information Forensics and Security*, 9(5):875–884, May 2014.
- [20] M.-D. Yu and S. Devadas. Secure and robust error correction for Physical Unclonable Functions. *IEEE Design & Test of Computers*, 27(1):48–65, 2010.