

Efficient Fuzzy Extraction of PUF-Induced Secrets: Theory and Applications*

Jeroen Delvaux^{1,2}, Dawu Gu², Ingrid Verbauwhede¹,
Matthias Hiller³ and Meng-Day (Mandel) Yu^{4,1,5}

¹ KU Leuven, ESAT/COSIC and iMinds,
Kasteelpark Arenberg 10, B-3001 Leuven, Belgium
{jeroen.delvaux, ingrid.verbauwhede}@esat.kuleuven.be

² Shanghai Jiao Tong University, CSE/LoCCS,
800 Dongchuan Road, Shanghai 200240, China
dwgu@sjtu.edu.cn

³ Chair of Security in Information Technology,
Technical University of Munich, Germany
matthias.hiller@tum.de

⁴ Verayo Inc., USA
myu@verayo.com

⁵ CSAIL, MIT, USA

Abstract. The device-unique response of a *physically unclonable function* (PUF) can serve as the root of trust in an embedded cryptographic system. *Fuzzy extractors* transform this noisy non-uniformly distributed secret into a stable high-entropy key. The overall efficiency thereof, typically depending on error-correction with a binary $[n, k, d]$ block code, is determined by the universal and well-known $(n - k)$ bound on the min-entropy loss. We derive new considerably tighter bounds for PUF-induced distributions that suffer from, e.g., bias or spatial correlations. The bounds are easy-to-evaluate and apply to large non-trivial codes, e.g., BCH and Reed-Muller codes. Apart from an inherent reduction in implementation footprint, the newly developed theory also facilitates the analysis of state-of-the-art error-correction methods for PUFs. As such, we debunk the reusability claim of the reverse fuzzy extractor. Moreover, we provide proper quantitative motivation for debiasing schemes, as this was missing in the original proposals.

Keywords: fuzzy extractor, secure sketch, min-entropy, physically unclonable function, coding theory

* This manuscript comprehends an extended version of our prior CHES 2016 work. The most notable changes are as follows. First, the equivalence among secure sketch constructions is proven in Appendix A. Second, Table 1 illustrates the reduction in implementation footprint attributed to the newly developed bounds. Third, the IBS and von Neumann debiasing schemes are fully specified in order to make the manuscript self-sustaining. More importantly, Table 2 evaluates the performance of the latter schemes, enabling a comparison with Table 1. Fourth, the upper bound on the residual min-entropy of a biased distribution in Section 3.3 has been improved for non-perfect codes.

1 Introduction

Cryptography relies on reproducible uniformly distributed secret keys. Obtaining affordable physically secure key-storage in embedded non-volatile memory is hard though. Harvesting entropy from *physically unclonable functions* (PUFs) comprehends an alternative that lowers the vulnerability during the power-off state. Unfortunately, PUF responses are corrupted by noise and non-uniformities are bound to occur. A *fuzzy extractor* [14] provides an *information-theoretically secure* mechanism to convert PUF responses into high-quality keys. The essential building block for handling noisiness is the *secure sketch*, providing error-correction with most frequently a binary $[n, k, d]$ block code. Associated public helper data reveals information about the PUF response though; the system provider should hence quantify how much min-entropy remains. So far, the conservative $(n - k)$ upper bound on the min-entropy loss has been applied. Unfortunately, the residual min-entropy is underestimated, implying that more PUF response bits than necessary have to be used. Expensive die area is hence blocked by PUF circuits that are not strictly required to obtain the desired security level, i.e., symmetric key length.

1.1 Contribution

The novelty of our work is twofold:

- First, we derive new bounds on the secure sketch min-entropy loss for PUF-induced distributions with practical relevance. Our bounds are considerably tighter than the well-known $(n - k)$ formula, hereby improving the implementation efficiency of PUF-based key generators. The discrepancy is showcased for two predominant PUF imperfections, i.e., biased and spatially correlated response bits. It is important to note that a variety of commonly used code is covered, e.g., BCH, Golay, and Reed-Muller codes, regardless of their algebraic complexity. Furthermore, a large variety of distributions could be supported. Therefore, our scope reaches considerably further than related work in [11, 28], focussing on simple repetition codes and biased distributions only. As in the latter works, our bounds are easy-to-evaluate and able to support large codes.
- Second, the newly developed theory is applied to state-of-the-art error-correction methods for PUFs. As such, we reveal a fundamental flaw in the reverse fuzzy extractor, proposed by Van Herrewege et al. [36] at Financial Crypto 2012. The latter lightweight primitive is gaining momentum and has also been adopted in the CHES 2015 protocol of Aysu et al. [2]. We debunk the main security claim that repeated helper data exposure does not result in additional min-entropy loss. Furthermore, we contribute to the motivation of debiasing schemes such as the *index-based syndrome* (IBS) proposal of Yu et al. [40], and the CHES 2015 proposal of Maes et al. [28]. The latter proposals assume that a stand-alone sketch cannot handle biased distributions. We eliminate the need for an educated guess that originates from the

extrapolation of repetition code insights and/or the application of the overly conservative $(n - k)$ bound.

1.2 Organization

The remainder of this manuscript is organized as follows. Section 2 introduces notation and preliminaries. Section 3 derives new tight bounds on the secure sketch min-entropy loss. Section 4 elaborates applications of the newly developed theory. Section 5 concludes the work.

2 Preliminaries

2.1 Notation

Binary vectors are denoted with a bold lowercase character, e.g., $\mathbf{x} = (x_1 \ x_2 \ x_3)$. All vectors are row vectors. All-zeros and all-ones vectors are denoted with $\mathbf{0}$ and $\mathbf{1}$ respectively. Binary matrices are denoted with a bold uppercase character, e.g., \mathbf{H} . A random variable and its corresponding set of outcomes are denoted with an uppercase *italic* and calligraphic character respectively, e.g., X and \mathcal{X} . Variable assignment is denoted with an arrow, e.g., $\mathbf{x} \leftarrow X$. Custom-defined procedure names are printed in a sans-serif font, e.g., Hamming weight $\text{HW}(\mathbf{x})$ and Hamming distance $\text{HD}(\mathbf{x}, \tilde{\mathbf{x}})$. The probability of an event A is denoted as $\mathbb{P}(A)$. The expected value of a function $g(X)$ of random variable X is denoted as $\mathbb{E}_{x \leftarrow X}[g(X)]$. The *probability density function* and *cumulative distribution function* of a standard normal distribution $N(0, 1)$ are denoted as $f_{\text{norm}}(\cdot)$ and $F_{\text{norm}}(\cdot)$ respectively. For a binomial distribution $B(n, p)$ with n trials and success probability p , we use $f_{\text{bino}}(\cdot; n, p)$ and $F_{\text{bino}}(\cdot; n, p)$ respectively.

2.2 Min-Entropy Definitions

The *min-entropy* of a random variable X is as defined in (1). Consider now a pair of possibly correlated random variables: X and P . The *conditional min-entropy* [14] of X given P is as defined in (2). Terms with $\mathbb{P}(P = p) = 0$ are evaluated as 0. Both definitions quantify the probability that an attacker guesses a secret $x \leftarrow X$ first time right, on a logarithmic scale. We emphasize that min-entropy is a more conservative notion than Shannon entropy and therefore often preferred within cryptology.

$$\mathbb{H}_{\infty}(X) = -\log_2\left(\max_{x \in \mathcal{X}} \mathbb{P}(X = x)\right). \quad (1)$$

$$\tilde{\mathbb{H}}_{\infty}(X|P) = -\log_2\left(\mathbb{E}_{p \leftarrow P}\left[\max_{x \in \mathcal{X}} \mathbb{P}((X = x)|(P = p))\right]\right). \quad (2)$$

2.3 Physically Unclonable Functions

A prominent category of PUFs, suitable for key generation in particular, consists of an array of identically designed cells. Each cell produces a single bit, or occasionally a few bits. This includes memory-based designs, such as the SRAM PUF [20], as well as the coating PUF [33] and a subset of the large number of ring oscillator-based designs, e.g., [38]. The most prominent entropy-degrading effects for such PUFs are bias and spatial correlations. Bias comprehends an imbalance between the number of zeros and ones. Spatial correlations implicate that neighboring cells might influence each other.

The analysis of error-correction methods for PUFs is greatly facilitated by having a generic yet accurate noise model. We describe a parameterized probability distribution for the evaluation of individual PUF response bits \tilde{x}_i , with $i \in [1, n]$. Experimental validation on various PUF circuits, e.g., in [26, 13], labelled the model as accurate. Two random variables are incorporated in (3). First, the normalized manufacturing variability $V_i \sim N(0, 1)$, drawn once for each response bit \tilde{x}_i . Second, additive noise $V_{\text{noise},i}^{(j)} \sim N(0, \sigma_{\text{noise}})$, with standard deviation σ_{noise} a fixed parameter, and drawn for each evaluation $j \in [1, n_{\text{runs}}]$ of \tilde{x}_i . Threshold v_{thres} is a fixed parameter; a nonzero value implies bias. Spatial correlations can be incorporated via a multivariate normal distribution $V \sim N(\mathbf{0}, \Sigma)$, with Σ the symmetric $n \times n$ covariance matrix.

$$\tilde{x}_i^{(j)} = \begin{cases} 1 & \text{if } (v_i + v_{\text{noise},i}^{(j)}) > v_{\text{thres}}, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Error rates are defined with respect to a reference response. For ease of analysis, we consider the response bits x_i obtained by thresholding $v_i > v_{\text{thres}}$ as a reference. In practice, these nominal values can be approximated via a majority vote among noisy replicas \tilde{x}_i , possibly accelerated via circuit techniques [6, 40]. Bias parameter b , defined as the probability $\mathbb{P}(x_i = 1)$, equals $F_{\text{norm}}(-v_{\text{thres}})$. Zero bias corresponds to $b = 0.5$. The error rate $p_{\text{error},i} \in (0, \frac{1}{2}]$ of a response bit \tilde{x}_i with respect to its reference, i.e., the probability $\mathbb{P}(x_i \neq \tilde{x}_i)$, equals $F_{\text{norm}}(-|v_i - v_{\text{thres}}|/\sigma_{\text{noise}})$.

2.4 Secure Sketch and Fuzzy Extractor Definitions

Secure sketches operate on a metric space \mathcal{X} with distance function dist . For PUFs, we can restrict our attention to binary vectors $\mathbf{x} \in \{0, 1\}^{1 \times n}$ and the Hamming distance HD therebetween. An attacker knows the probability distribution of $\mathbf{x} \leftarrow X$. Consider a noisy version $\tilde{\mathbf{x}}$ of sample \mathbf{x} . A *secure sketch* [14] is a pair of efficient and possibly randomized procedures: the sketching procedure $\mathbf{p} \leftarrow \text{SSGen}(\mathbf{x})$, generating helper data $\mathbf{p} \in \mathcal{P}$, and the recovery procedure $\hat{\mathbf{x}} \leftarrow \text{SSRep}(\tilde{\mathbf{x}}, \mathbf{p})$. There are two defining properties:

- *Correctness.* If $\text{HD}(\mathbf{x}, \tilde{\mathbf{x}}) \leq t$, correctness of reconstruction is guaranteed, i.e., $\hat{\mathbf{x}} = \mathbf{x}$. If $\text{HD}(\mathbf{x}, \tilde{\mathbf{x}}) > t$, there is no guarantee whatsoever.

- *Security.* Given a certain lower bound h_{in} on the ingoing min-entropy, i.e., $\mathbb{H}_{\infty}(X) \geq h_{\text{in}}$, a corresponding lower bound h_{out} on the residual min-entropy, i.e., $\tilde{\mathbb{H}}_{\infty}(X|P) \geq h_{\text{out}}$, can be imposed. Often, but not necessarily, this condition can be satisfied regardless of h_{in} . Or stated otherwise, there is a certain upper bound on the min-entropy loss $\Delta\mathbb{H}_{\infty} = \mathbb{H}_{\infty}(X) - \tilde{\mathbb{H}}_{\infty}(X|P)$.

A slightly modified notion brings us to the *fuzzy extractor* [14]. Output $\mathbf{k} \in \mathcal{K}$ is then required to be nearly-uniform, given observation of $\mathbf{p} \leftarrow P$, and is therefore suitable as a secret key. There is a proven standard method to craft a fuzzy extractor from a secure sketch. In particular, a *randomness extractor* could derive a key from the secure sketch output, i.e., $\mathbf{k} \leftarrow \text{Ext}(\mathbf{x})$. *Universal hash functions* [9] are good randomness extractors, according to the (*generalized*) *leftover hash lemma* [16, 3]. Unfortunately, their min-entropy loss is quite substantial. In practice, key generators therefore often rely on a cryptographic hash function that is assumed to behave as a *random oracle*. The latter idealized heuristic results in zero min-entropy loss.

2.5 Coding Theory

A *binary code* \mathcal{C} is a bijection from a message space \mathcal{M} to a codeword space $\mathcal{W} \subseteq \{0, 1\}^{1 \times n}$. The *minimum distance* d is the minimum number of bits in which any two distinct codewords differ. A procedure $\mathbf{w} \leftarrow \text{Encode}(\mathbf{m})$ maps a message $\mathbf{m} \in \mathcal{M}$ to a codeword $\mathbf{w} \in \mathcal{W}$. A procedure $\hat{\mathbf{w}} \leftarrow \text{Correct}(\tilde{\mathbf{w}})$ corrects up to $t = \lfloor \frac{d-1}{2} \rfloor$ errors for any noise-corrupted codeword $\tilde{\mathbf{w}} = \mathbf{w} \oplus \mathbf{e}$, with $\text{HW}(\mathbf{e}) \leq t$. An extended procedure $\hat{\mathbf{m}} \leftarrow \text{Decode}(\tilde{\mathbf{w}})$ returns the corresponding message instead. Equation (4) expresses the Hamming bound [24]. The equality holds for *perfect codes* only, implicating that any vector in $\{0, 1\}^{1 \times n}$ is within distance t of a codeword. All other codes are subject to the inequality.

$$\sum_{i=0}^t \binom{n}{i} |\mathcal{M}| \leq 2^n. \quad (4)$$

A binary $[n, k, d]$ *block code* \mathcal{C} restricts the message length $k = \log_2(|\mathcal{M}|)$ to an integer. For a linear block code, any linear combination of codewords is again a codeword. A $k \times n$ *generator matrix* \mathbf{G} , having full rank, can then implement the encoding procedure, i.e., $\mathbf{w} = \mathbf{m} \cdot \mathbf{G}$. For any translation $\mathbf{t} \in \{0, 1\}^{1 \times n}$ and linear code \mathcal{C} , the set $\{\mathbf{t} \oplus \mathbf{w} : \mathbf{w} \in \mathcal{W}\}$ is referred to as a *coset*. Two cosets are either disjoint or coincide. Therefore, the vector space $\{0, 1\}^{1 \times n}$ is fully covered by 2^{n-k} cosets, referred to as the *standard array*. The minimum weight vector \mathbf{e} in a coset is called the *coset leader*. In case of conflict, i.e., a common minimum $\text{HW}(\mathbf{e}) > t$, an arbitrary leader can be selected. The minimum distance d of a linear code equals the minimum Hamming weight of its nonzero codewords. A linear code \mathcal{C} is *cyclic* if every circular shift of a codeword is again a codeword belonging to \mathcal{C} .

2.6 The Code-Offset Secure Sketch

Several secure sketch constructions rely on a binary code \mathcal{C} . For ease of understanding, we focus on the code-offset method of Dodis et al. [14] exclusively. Nevertheless, equivalencies in Appendix A prove that all results in this manuscript apply to six other constructions equally well. The code \mathcal{C} that instantiates the code-offset method in Fig. 1 is not necessarily linear. Even more, it is not required to be a block code either. Linear codes (BCH, Golay, repetition, etc.) remain the most frequently used though due to their efficient decoding algorithms [24]. Correctness of reconstruction is guaranteed if $\text{HD}(\mathbf{x}, \tilde{\mathbf{x}}) \leq t$, with t the error-correcting capability of the code.

$\mathbf{p} \leftarrow \text{SSGen}(\mathbf{x})$	$\hat{\mathbf{x}} \leftarrow \text{SSRep}(\tilde{\mathbf{x}}, \mathbf{p})$
Random $\mathbf{w} \in \mathcal{C}$	$\tilde{\mathbf{w}} \leftarrow \tilde{\mathbf{x}} \oplus \mathbf{p} = \mathbf{w} \oplus \mathbf{e}$
$\mathbf{p} \leftarrow \mathbf{x} \oplus \mathbf{w}$	$\hat{\mathbf{x}} \leftarrow \mathbf{p} \oplus \text{Correct}(\tilde{\mathbf{w}})$

Fig. 1. The code-offset secure sketch, having an n -bit reference input \mathbf{x} .

Min-entropy loss can be understood as a *one-time pad* imperfection. Sketch input \mathbf{x} is masked with a random codeword \mathbf{w} , i.e., an inherent entropy deficiency: $\mathbb{H}_\infty(W) = \log_2(|\mathcal{M}|) < n$. For linear codes in particular, we highlight a convenient interpretation using cosets. Helper data \mathbf{p} then reveals in which coset reference \mathbf{x} resides. It can be seen easily that \mathbf{p} is equal to a random vector in the same coset as \mathbf{x} . The residual min-entropy in (2) hence reduces to (5) for linear codes, with \mathbf{e} a coset leader. We emphasize that the min-entropy loss $\Delta\mathbb{H}_\infty$ does not depend on the decoding method, simply because the helper data is not affected. For $[n, k, d]$ block codes in particular, the well-known upper bound $\Delta\mathbb{H}_\infty \leq (n - k)$ holds, as proven in [14]. More generally, this extends to $\Delta\mathbb{H}_\infty \leq n - \log_2(|\mathcal{M}|)$.

$$\tilde{\mathbb{H}}_\infty(X|P) = -\log_2\left(\mathbb{E}_{\mathbf{e} \leftarrow E} \left[\max_{\mathbf{w} \in \mathcal{W}} \mathbb{P}((X = \mathbf{e} \oplus \mathbf{w}) | (E = \mathbf{e})) \right]\right). \quad (5)$$

2.7 Repeated Execution of a Concatenated Code

Optimized fuzzy extractors often rely on a concatenated code $\mathcal{C}_2 \circ \mathcal{C}_1$ that processes z non-overlapping blocks of PUF response bits independently [7, 29]. The inner code \mathcal{C}_2 is a small $[n_2, k_2 = 1, d_2 = n_2]$ repetition code, with n_2 odd, allowing to support a high bit error rate $\mathbb{E}_{v_i \leftarrow V_i} [p_{\text{error}, i}]$. A large $[n_1, k_1, d_1]$ outer code \mathcal{C}_1 , e.g., a BCH code [24], is faced with a considerably lower bit error rate so that its min-entropy loss can be relatively small. The size of \mathcal{C}_1 is nevertheless limited due to the implementation footprint of `Correct`. PUF response \mathbf{x} hence needs to be partitioned in z blocks in order to generate a key \mathbf{k} of sufficient length. We encapsulate the operation $z \times [n_2, k_2, d_2] \circ [n_1, k_1, d_1]$ in a single umbrella block code with $[n = z \cdot n_1 \cdot n_2, k = z \cdot k_1, d = d_1 \cdot d_2]$ and $t = t_1(t_2 + 1) \leq \lfloor \frac{d-1}{2} \rfloor$.

3 Tight Bounds on the Min-Entropy Loss

Currently, secure sketch implementations rely on the $(n - k)$ upper bound on the min-entropy loss, e.g., [29]. Unfortunately, this leads to an overly conservative design when instantiating security parameters accordingly. We develop a graphical framework that produces tight bounds on $\tilde{\mathbb{H}}_\infty(X|P)$ for typical PUF-induced distributions. The critical *first-order* effects of bias and spatial correlations are captured. Both lower and upper bounds are supported. The lower bounds are of primary interest for a conservative system provider, entertaining the worst-case scenario. We considerably improve upon the $(n - k)$ bound, i.e., the leftmost inequality in (6). We also improve upon the rather trivial upper bounds [14] that comprehend the rightmost inequality in (6).

$$\underbrace{\max(\mathbb{H}_\infty(X) - (n - \log_2(|\mathcal{M}|)), 0)}_{\text{worst-case}} \leq \tilde{\mathbb{H}}_\infty(X|P) \leq \underbrace{\min(\log_2(|\mathcal{M}|), \mathbb{H}_\infty(X))}_{\text{best-case}}. \quad (6)$$

Our lower and upper bounds combined define a relatively narrow interval in which the exact value of $\tilde{\mathbb{H}}_\infty(X|P)$ is enclosed. We considerably extend related work in [11, 28] as follows. First, we cover a variety of codes, regardless of their algebraic complexity. Prior work focussed on repetition codes only. Although frequently used as the inner code of a concatenated code [7], full-fledged key generators [29] typically rely on non-trivial codes, e.g., BCH codes [24]. Second, our techniques may be applied to a variety of distributions, while prior work covered biased distributions only. Our bounds remain easy-to-evaluate and are able to handle large codes. Although derived for the code-offset sketch of Dodis et al. [14] in particular, Appendix A establishes the equivalence with six other constructions.

3.1 Distributions

Our work is generic in the sense that a large variety of distributions X could be covered. We only require that $\mathcal{X} = \{0, 1\}^{1 \times n}$ can be partitioned in a limited number of subsets φ_j , with $j \in [1, n_{\text{sets}}]$, so that all elements of φ_j have the same probability of occurrence q_j . Formally, $\mathbb{P}(X = \mathbf{x}) = q_j$ if and only if $\mathbf{x} \in \varphi_j$. These probabilities are strictly monotonically decreasing, i.e., $q_j > q_{j+1}$, with $j \in [1, n_{\text{sets}} - 1]$. Occasionally, $q_{n_{\text{sets}}} = 0$. The ingoing min-entropy is easily computed as $\mathbb{H}_\infty(X) = -\log_2(q_1)$.

We determine bounds on $\tilde{\mathbb{H}}_\infty(X|P)$. The runtime of the corresponding algorithms is roughly proportional to n_{sets} . The crucial observation is that a small n_{sets} might suffice to capture realistic PUF models. Below, we describe a parameterized distribution X for both biased and spatially correlated PUFs. Both distributions are to be considered as proof-of-concept models, used in showcasing the feasibility of a new research direction. In case a given PUF is not approximated accurately enough, one can opt for an alternative and possibly more

complicated *second-order* distribution. As long as n_{sets} is limited, bounds can be evaluated in milliseconds-minutes on a standard desktop computer.

- *Biased distribution.* We assume response bits to be independent and identically distributed (i.i.d.) so that $\mathbb{P}(X_i = 1) = b$, with $i \in [1, n]$ and a real-valued $b \in [0, 1]$. For $b = \frac{1}{2}$, this corresponds to a uniform distribution. The latter bias model comprehends a very popular abstraction in PUF literature. The min-entropy loss of various other helper data methods has been analyzed as such, e.g., *soft-decision decoding* [27, 11] as well as *IBS* [40, 18] and von Neumann [28, 35] debiasing. Therefore, our results enable adequate comparison with related methods, all using a common baseline distribution.
- *Correlated distribution.* We assume response bits to be distributed so that $\mathbb{P}(X_i = X_{i+1}) = c$, with $i \in [1, n-1]$ and a real-valued $c \in [0, 1]$. This extends to (7) for larger neighborhoods. There is no bias, i.e., $\mathbb{P}(X_i = 1) = \frac{1}{2}$. For $c = \frac{1}{2}$, the latter model corresponds to a uniform distribution. Although spatial correlations are frequently encountered in experimental work, e.g., byte-level dependencies for the SRAM PUFs in [17, 2], these are often neglected in information theoretic work due to their complexity. We hope that our results may help turn the tide on this.

$$c_{i,j} = \mathbb{P}(X_i = X_j) = \sum_{u=0}^{\lfloor |i-j|/2 \rfloor} \mathbf{f}_{\text{bino}}(2u; |i-j|, 1-c), \quad \text{with } i, j \in [1, n]. \quad (7)$$

Fig. 2 specifies the subsets φ_j for both distributions. For the biased distribution, we partition according to $\text{HW}(\mathbf{x})$. This corresponds to a binomial distribution with $j-1$ successes for n Bernoulli trials, each having success probability $b_\star = \min(b, 1-b)$. For the correlated distribution, we partition according to $\sum_{i=1}^{n-1} \text{HD}(x_i, x_{i+1})$, i.e., the number of transitions in \mathbf{x} . Inputs in subset φ_j exhibit $j-1$ transitions and obey either one out of two forms, i.e., $\mathbf{x} = (\mathbf{0} \parallel \mathbf{1} \parallel \mathbf{0} \parallel \dots)$ and $\mathbf{x} = (\mathbf{1} \parallel \mathbf{0} \parallel \mathbf{1} \parallel \dots)$. A related observation is that if $\mathbf{x} \in \varphi_j$, then so is its ones' complement, i.e., $\bar{\mathbf{x}} \in \varphi_j$. This explains the factors 2 and $\frac{1}{2}$ everywhere. Set size $|\varphi_j|$ is further determined with *stars and bars* combinatorics [15]. In particular, we separate n indistinguishable stars into j distinguishable bins by adding $j-1$ out of $n-1$ bars.

We treat the degenerate case $b = c = \frac{1}{2}$, i.e., a uniform distribution, separately. There is only one set then. Formally, $n_{\text{sets}} = 1$, $|\varphi_1| = 2^n$, and $q_1 = 1/2^n$. As proven by Reyzin [30], the min-entropy loss of a secure sketch is maximal for a uniformly distributed input, making this a case of special interest.

3.2 Generic Bounds

We derive generic bounds that can be applied to any distribution X with n_{sets} limited. Equation (8) holds for the code-offset construction of Dodis et al. [14], given that a codeword is selected fully at random during enrollment.

$$\mathbb{P}((P = \mathbf{p}) | (X = \mathbf{x})) = \begin{cases} 1/|\mathcal{M}|, & \text{if } \exists \mathbf{w} : \mathbf{p} = \mathbf{x} \oplus \mathbf{w} \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

j	$ \varphi_j $	q_j	j	$ \varphi_j $	q_j
1	1	$(1 - b_\star)^n$	1	2	$\frac{1}{2}(1 - c_\star)^{n-1}$
2	n	$b_\star(1 - b_\star)^{n-1}$	2	$2(n-1)$	$\frac{1}{2}c_\star(1 - c_\star)^{n-2}$
...
j	$\binom{n}{j-1}$	$(b_\star)^{j-1}(1 - b_\star)^{n-j+1}$	j	$2\binom{n-1}{j-1}$	$\frac{1}{2}(c_\star)^{j-1}(1 - c_\star)^{n-j}$
...
n	n	$(b_\star)^{n-1}(1 - b_\star)$	$n-1$	$2(n-1)$	$\frac{1}{2}(c_\star)^{n-2}(1 - c_\star)$
$n+1$	1	$(b_\star)^n$	n	2	$\frac{1}{2}(c_\star)^{n-1}$

Fig. 2. Subsets φ_j for a biased and correlated distribution X , left and right respectively. We define $b_\star = \min(b, 1 - b)$ and $c_\star = \min(c, 1 - c)$.

Equation (9) applies Bayes' rule to the definition of conditional min-entropy in (2) and fills in (8). The 0 case is resolved by switching variables for the max operator. A direct exhaustive evaluation of the resulting formula requires up to $2^n |\mathcal{M}|$ operations.

$$\begin{aligned} \tilde{\mathbb{H}}_\infty(X|P) &= -\log_2 \left(\sum_{\mathbf{p} \in \mathcal{P}} \cancel{\mathbb{P}(P = \mathbf{p})} \max_{\mathbf{x} \in \mathcal{X}} \frac{\mathbb{P}(X = \mathbf{x}) \cancel{\mathbb{P}((P = \mathbf{p})|(X = \mathbf{x}))}}{\cancel{\mathbb{P}(P = \mathbf{p})}} \right) \\ &= -\log_2 \left(\frac{1}{|\mathcal{M}|} \sum_{\mathbf{p} \in \mathcal{P}} \max_{\mathbf{w} \in \mathcal{W}} \mathbb{P}(X = \mathbf{p} \oplus \mathbf{w}) \right). \end{aligned} \quad (9)$$

For linear codes, the workload can be reduced substantially. With a similar derivation as before, we rewrite (5) as shown in (10). Up to 2^n operations suffice. Nevertheless, direct evaluation is only feasible for small codes. We emphasize that our bounds are able to handle large codes, as is typically the case for a practical key generator.

$$\tilde{\mathbb{H}}_\infty(X|P) = -\log_2 \left(\sum_{\mathbf{e} \in \mathcal{E}} \max_{\mathbf{w} \in \mathcal{W}} \mathbb{P}(X = \mathbf{e} \oplus \mathbf{w}) \right). \quad (10)$$

Equation (9) iterates over all \mathbf{p} 's and selects each time the most likely \mathbf{x} that is within range, via the addition of a codeword $\mathbf{w} \in \mathcal{W}$. We now reverse the roles, as shown in Fig. 3. We iterate over all \mathbf{x} 's, from most likely to least likely, i.e., from φ_1 to φ_J . Within a certain φ_j , the order of the \mathbf{x} 's may be chosen arbitrarily. Subsequently, we assign \mathbf{p} 's to each \mathbf{x} , as represented by the black squares, until the set \mathcal{P} of size 2^n is depleted. For each assigned \mathbf{p} , we assume that the corresponding \mathbf{x} is the most likely vector, according to (9). Let $s_j^{\mathbf{p}}$ denote the number of black squares assigned to set φ_j . The residual min-entropy is then easily computed as in (11).

$$\tilde{\mathbb{H}}_\infty(X|P) = -\log_2 \left(\frac{1}{|\mathcal{M}|} \sum_{j=1}^J s_j^{\mathbf{p}} q_j \right). \quad (11)$$

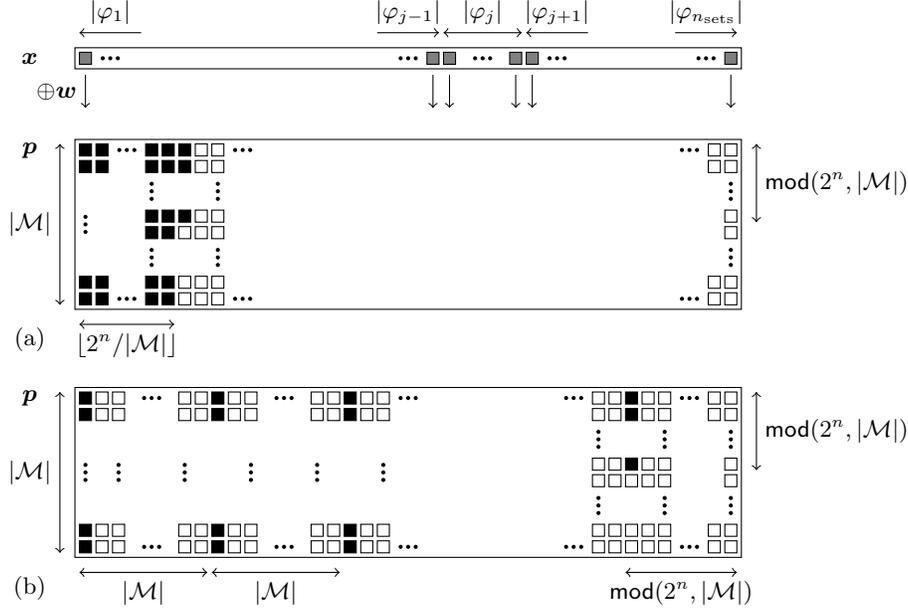


Fig. 3. Reversal of the roles in (9). (a) A lower bound on $\tilde{\mathbb{H}}_\infty(X|P)$. (b) An upper bound on $\tilde{\mathbb{H}}_\infty(X|P)$. Black squares represent terms that contribute to $\tilde{\mathbb{H}}_\infty(X|P)$, one for each $\mathbf{p} \in \mathcal{P}$. White squares represent non-contributing terms, overruled by the max operator. In general, there are few black squares but many white squares, 2^n versus $(|\mathcal{M}| - 1)2^n$ to be precise. For block codes, i.e., $|\mathcal{M}| = 2^k$, the last column of black squares is completely filled.

Both linear and non-linear codes are supported by former graphical representation. Nevertheless, we elaborate linear codes as a special case due to their practical relevance. Fig. 4 swaps the order of iteration in (10). Only one row suffices, i.e., each column of helper data vectors \mathbf{p} in Figure 3 is condensed to a single square. Black and white squares are now assigned to cosets, as represented by their coset leaders \mathbf{e} . Let s_j^e denote the number of black squares assigned to set φ_j . The residual min-entropy is then easily computed as in (12), hereby dropping denominator $|\mathcal{M}|$ compared to (11), given that $s_j^{\mathbf{p}} = 2^k \cdot s_j^e$.

$$\tilde{\mathbb{H}}_\infty(X|P) = -\log_2 \left(\sum_{j=1}^J s_j^e q_j \right). \quad (12)$$

In the worst-case scenario, the most likely \mathbf{x} 's all map to unique \mathbf{p} 's, without overlap, resulting in a lower bound on $\tilde{\mathbb{H}}_\infty(X|P)$. For a linear code, this would be the case if the first 2^{n-k} \mathbf{x} 's all belong to different cosets. In the best-case scenario, our sequence of \mathbf{x} 's exhibits maximum overlap in terms of \mathbf{p} , resulting in an upper bound on $\tilde{\mathbb{H}}_\infty(X|P)$. For a linear code, this would be the case if the first 2^k \mathbf{x} 's all map to the same coset, and this repeated for all 2^{n-k} cosets.

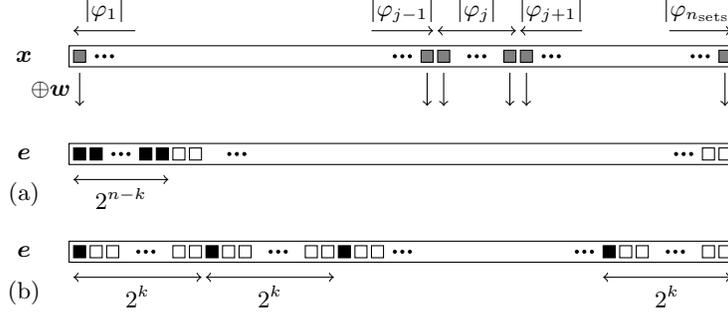


Fig. 4. Reversal of the roles in (10), as applied to linear codes. (a) A lower bound on $\tilde{\mathbb{H}}_\infty(X|P)$. (b) An upper bound on $\tilde{\mathbb{H}}_\infty(X|P)$. Black squares represent terms that contribute to $\tilde{\mathbb{H}}_\infty(X|P)$, one for each $e \in \mathcal{E}$. White squares represent non-contributing terms, overruled by the max operator.

Algorithms 1 and 2 comprehend a literal transcript of Fig. 3 and compute the lower bound and upper bound respectively. Auxiliary variables s^p and s^x accumulate black and gray squares respectively. To maintain generality, we abstain from special case algorithms for linear codes, although it would result in a few simplifications.

Algorithm 1: BoundWorstCase	Algorithm 2: BoundBestCase
<p>Input: List $\langle \varphi_j , q_j \rangle$</p> <p>Output: Lower bound on $\tilde{\mathbb{H}}_\infty(X P)$</p> <p>$j, q, s^p \leftarrow 0$</p> <p>while $s^p < 2^n$ do</p> <div style="border-left: 1px solid black; border-right: 1px solid black; padding: 0 10px;"> <p>$j \leftarrow j + 1$</p> <p>$s_j^p \leftarrow \min(\varphi_j \mathcal{M} , 2^n - s^p)$</p> <p>$s^p \leftarrow s^p + s_j^p$</p> <p>$q \leftarrow q + s_j^p \cdot q_j$</p> </div> <p>$\tilde{\mathbb{H}}_\infty(X P) \leftarrow -\log_2(q/ \mathcal{M})$</p>	<p>Input: List $\langle \varphi_j , q_j \rangle$</p> <p>Output: Upper bound on $\tilde{\mathbb{H}}_\infty(X P)$</p> <p>$j, q, s^p, s^x \leftarrow 0$</p> <p>while $s^p < 2^n$ do</p> <div style="border-left: 1px solid black; border-right: 1px solid black; padding: 0 10px;"> <p>$j \leftarrow j + 1$</p> <p>$s^x \leftarrow s^x + \varphi_j$</p> <p>$s_j^p \leftarrow \lceil (s^x - s^p) / \mathcal{M} \rceil \mathcal{M}$</p> <p>$s_j^p \leftarrow \min(\max(s_j^p, 0), 2^n - s^p)$</p> <p>$s^p \leftarrow s^p + s_j^p$</p> <p>$q \leftarrow q + s_j^p \cdot q_j$</p> </div> <p>$\tilde{\mathbb{H}}_\infty(X P) \leftarrow -\log_2(q/ \mathcal{M})$</p>

Algorithms 1 and 2 may now be applied to a variety of distributions. For a uniform distribution, the lower and upper bound both evaluate to $\tilde{\mathbb{H}}_\infty(X|P) = \log_2(|\mathcal{M}|)$, regardless of other code specifics. Or simply k , for block codes in particular. The min-entropy loss is hence exactly $(n-k)$ bits, given that $\mathbb{H}_\infty(X) = n$. Reyzin's proof [30] therefore implies that the general-purpose $(n-k)$ bound can-

not be tightened any further. Although results are fairly presentable already for the biased and correlated distributions, we further tighten these bounds first.

3.3 Tighter Bounds

Tighter bounds can be obtained by leveraging code and distribution properties more effectively. Algorithms 3 and 4 generalize Algorithms 1 and 2 respectively. In the former case, an additional input imposes an upper bound on the accumulated number of black squares, i.e., $\forall j, (s_1^p + s_2^p + \dots + s_j^p) \leq (u_1^p + u_2^p + \dots + u_j^p)$. In the latter case, an additional input imposes a lower bound on the accumulated number of black squares, i.e., $\forall j, (s_1^p + s_2^p + \dots + s_j^p) \geq (l_1^p + l_2^p + \dots + l_j^p)$. We now provide several examples.

Algorithm 3: BoundWorstCase2	Algorithm 4: BoundBestCase2
<p>Input: List $\langle \varphi_j , q_j, u_j^p \rangle$</p> <p>Output: Lower bound on $\tilde{\mathbb{H}}_\infty(X P)$</p> <p>$j, q, s^p, u^p \leftarrow 0$</p> <p>while $s^p < 2^n$ do</p> <div style="margin-left: 20px;"> <p>$j \leftarrow j + 1$</p> <p>$u^p \leftarrow u^p + u_j^p$</p> <p>$s_j^p \leftarrow \min(\varphi_j \mathcal{M} , u^p - s^p)$</p> <p>$s_j^p \leftarrow \min(s_j^p, 2^n - s^p)$</p> <p>$s^p \leftarrow s^p + s_j^p$</p> <p>$q \leftarrow q + s_j^p \cdot q_j$</p> </div> <p>$\tilde{\mathbb{H}}_\infty(X P) \leftarrow -\log_2(q/ \mathcal{M})$</p>	<p>Input: List $\langle \varphi_j , q_j, l_j^p \rangle$</p> <p>Output: Upper bound on $\tilde{\mathbb{H}}_\infty(X P)$</p> <p>$j, q, s^p, s^x, l^p \leftarrow 0$</p> <p>while $s_{1:j}^p < 2^n$ do</p> <div style="margin-left: 20px;"> <p>$j \leftarrow j + 1$</p> <p>$s^x \leftarrow s^x + \varphi_j$</p> <p>$l^p \leftarrow l^p + l_j^p$</p> <p>$s_j^p \leftarrow \lceil (s^x - s^p) / \mathcal{M} \rceil \mathcal{M}$</p> <p>$s_j^p \leftarrow \max(s_j^p, l^p - s^p, 0)$</p> <p>$s_j^p \leftarrow \min(s_j^p, 2^n - s^p)$</p> <p>$s^p \leftarrow s^p + s_j^p$</p> <p>$q \leftarrow q + s_j^p \cdot q_j$</p> </div> <p>$\tilde{\mathbb{H}}_\infty(X P) \leftarrow -\log_2(q/ \mathcal{M})$</p>

Worst-Case Bounds We further tighten the lower bound on $\tilde{\mathbb{H}}_\infty(X|P)$ for the correlated distribution. The improvement applies to linear codes that have the all-ones vector $\mathbf{1}$ of length n as a codeword. This includes Reed-Muller codes of any order [24]. This also includes many BCH, Hamming and repetition codes, on the condition that these are cyclic and having d odd, as easily proven hereafter. Consider an arbitrary codeword with Hamming weight d . XORing all 2^n circular shifts of this codeword results in the all-ones codeword, which ends the proof. As mentioned before, each set φ_j of the correlated distribution can be partitioned in pairs $\{\mathbf{x}, \bar{\mathbf{x}}\}$, with $\bar{\mathbf{x}}$ the ones' complement of \mathbf{x} . Paired inputs belong to the same coset, i.e., maximum overlap in terms of helper data \mathbf{p} . Therefore, we impose the cumulative upper bound in (13).

$$u_j^p = |\mathcal{M}| \frac{|\varphi_j|}{2} = 2^{k-1} |\varphi_j|. \quad (13)$$

For instance, consider linear/cyclic $[n, k = 1, d = n]$ repetition codes, i.e., having generator matrix $\mathbf{G} = \mathbf{1}$, with n odd. Algorithms `BoundWorstCase2` and `BoundBestCase` then converge to the exact result $\tilde{\mathbb{H}}_\infty(X|P) = 1$, not depending on parameter c . This is the best-case scenario, given the universal bound $\tilde{\mathbb{H}}_\infty(X|P) \leq k$. Fig. 5 illustrates the former with squares for $n = 5$. The result also holds if the repetition code is neither linear/cyclic nor odd. As long as $\mathbf{w}_1 \oplus \mathbf{w}_2 = \mathbf{1}$, the elements of each φ_j can be paired into cosets. Although the term coset is usually preserved for linear codes, translations of a non-linear repetition code are either disjoint or coincide and still partition the space $\{0, 1\}^{1 \times n}$. As a side note, the result offers another [11] refutation of the *repetition code pitfall* of Koeberl et al. [23], a work that overlooks that $(n - k)$ is an upper bound only.

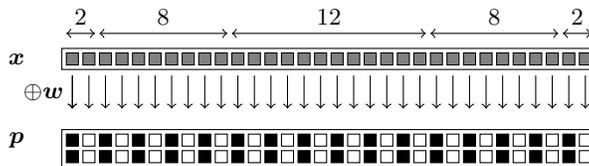


Fig. 5. The exact residual min-entropy $\tilde{\mathbb{H}}_\infty(X|P)$ for the correlated distribution and an $[n = 5, k = 1, d = 5]$ repetition code.

Best-Case Bounds We improve the upper bound on $\tilde{\mathbb{H}}_\infty(X|P)$ for both the biased and correlated distribution. In particular, we take minimum distance d into account. The main insight is that two slightly differing inputs $\mathbf{x}_u \neq \mathbf{x}_v$ do not overlap in terms of helper data \mathbf{p} . More precisely, if $\text{HD}(\mathbf{x}_u, \mathbf{x}_v) \in [1, d - 1]$, then $\{\mathbf{x}_u \oplus \mathbf{w} : \mathbf{w} \in \mathcal{W}\} \cap \{\mathbf{x}_v \oplus \mathbf{w} : \mathbf{w} \in \mathcal{W}\} = \emptyset$. For the biased distribution, the following holds: $\text{HD}(\mathbf{x}_u, \mathbf{x}_v) \in [1, d - 1]$ if $\mathbf{x}_u \neq \mathbf{x}_v$ and $\mathbf{x}_u, \mathbf{x}_v \in (\varphi_1 \cup \varphi_2 \cup \dots \cup \varphi_{t+1})$. Or stated otherwise, the elements of the first $t + 1$ sets all result in unique \mathbf{p} 's. Therefore, we can impose the constraint given in (14). Fig. 6 depicts the squares.

$$l_j^{\mathbf{p}} = \begin{cases} |\varphi_j| |\mathcal{M}|, & \text{if } j \in [1, t + 1] \\ 0, & \text{otherwise} \end{cases}. \quad (14)$$

There is an interesting observation for perfect codes in particular. As clear from the Hamming bound in (4), all unique \mathbf{p} 's are covered by the first $t + 1$ sets exclusively. `BoundWorstCase` and `BoundBestCase2` hence produce the same output, implying that the residual min-entropy is evaluated exactly, as further simplified in (15). Delvaux et al. [11] derived the same formula for $[n, k = 1, d = n]$ repetition codes with n odd. The scope of their result is hence extended from perfect repetition codes to perfect codes in general. As a side note, the formula was originally adopted to debunk the aforementioned *repetition code pitfall* [23].

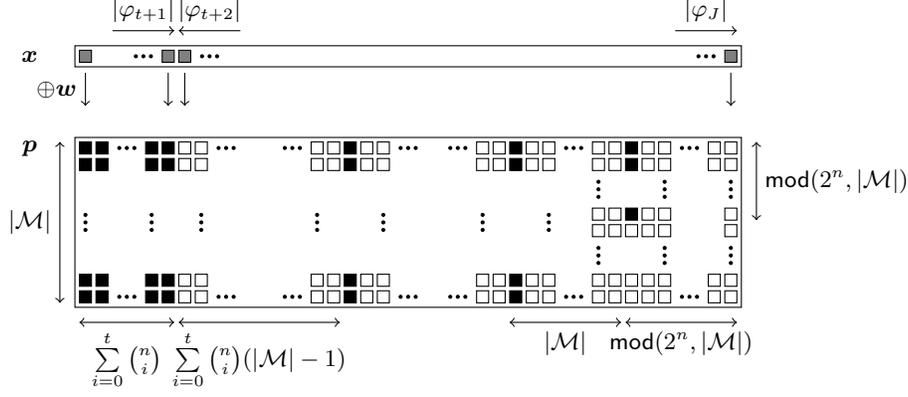


Fig. 6. A tightened upper bound on $\tilde{\mathbb{H}}_\infty(X|P)$ for the biased distribution, hereby making use of (14).

Maes et al. [28] later presented a similar contribution at CHES 2015, differing in its use of Shannon entropy rather than min-entropy.

$$\tilde{\mathbb{H}}_\infty(X|P) = -\log_2\left(\sum_{j=1}^{t+1} |\varphi_j| \cdot q_j\right) = -\log_2(\mathbf{F}_{\text{bino}}(t; n, \min(b, 1-b))). \quad (15)$$

For codes that do not happen to be perfect, there is still margin for improvement. We inject some promising thoughts but abstain from numerical results later-on. Consider a linear code of which the Hamming weight distribution of the coset leaders e is well-understood. Let $|\mathcal{E}_h|$ denote the number of cosets such that $h = \text{HW}(e)$. Clearly, $|\mathcal{E}_h| = \binom{n}{h}$ for $h \in [0, t]$. Our interest concerns $|\mathcal{E}_h|$ for $h > t$, all of which are exactly known in the ideal case, as in [10] for certain BCH codes. The largest h for which $|\mathcal{E}_h| > 0$ is also referred to as the *covering radius* h_{cr} of the code. For a bias $b < \frac{1}{2}$, (16) comprehends the exact residual min-entropy. The latter expression extends to $b > \frac{1}{2}$ in case the all-ones vector $\mathbf{1}$ is a codeword. This includes Reed-Muller codes as well as cyclic codes with d odd, as has been argued earlier-on. If only bounds on $|\mathcal{E}_h|$ and/or h_{cr} are known, one might still be able to further tighten the bounds on $\tilde{\mathbb{H}}_\infty(X|P)$ correspondingly.

$$\tilde{\mathbb{H}}_\infty(X|P) = -\log_2\left(\frac{1}{|\mathcal{M}|} \sum_{h=0}^{h_{\text{cr}}} |\mathcal{E}_h| \cdot |\mathcal{M}| \cdot q_{h+1}\right) = -\log_2\left(\sum_{h=0}^{h_{\text{cr}}} |\mathcal{E}_h| \cdot q_{h+1}\right). \quad (16)$$

For instance, consider $[n, k = 1, d = n]$ repetition codes with n even. These form the non-perfect and therefore less popular counterpart of n odd. Inputs \mathbf{x} belonging to φ_j and φ_{n+2-j} are still paired in order to form the cosets. Unlike n odd, there is a central set φ_{t+2} that contains both members of each pair.

Therefore, $h_{cr} = t + 1$ and $|\mathcal{E}_{t+1}| = |\varphi_{t+2}|/2$. As argued before, the operational principles of cosets extend to non-linear repetition codes. Fig. 7 depicts the squares for $n = 4$. Equation (17) evaluates the residual min-entropy.

$$\tilde{\mathbb{H}}_{\infty}(X|P) = -\log_2\left(\mathbb{F}_{\text{bino}}(t; n, \min(b, 1 - b)) + \frac{1}{2}\binom{n}{\frac{n}{2}}(b(1 - b))^{\frac{n}{2}}\right). \quad (17)$$

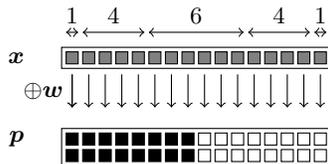


Fig. 7. The exact residual min-entropy $\tilde{\mathbb{H}}_{\infty}(X|P)$ for the biased distribution and an $[n = 4, k = 1, d = 4]$ repetition code.

Also for the correlated distribution, distance d might be incorporated to tighten the upper bound on $\tilde{\mathbb{H}}_{\infty}(X|P)$. First of all, we assign $|\mathcal{M}|$ unique \mathbf{p} 's to one out of two elements in φ_1 . For ease of understanding, assume $\mathbf{x} = \mathbf{0}$, comprehending the first case in (18). For each set φ_j , with $j \in [2, n]$, we then count the number of inputs $\mathbf{x} \in \varphi_j$ such that $h = \text{HW}(\mathbf{x}) \leq t$. The latter constraint guarantees all assigned \mathbf{p} 's to be unique. We distinguish between two forms, $\mathbf{x} = (\mathbf{0} \|\mathbf{1} \|\mathbf{0} \|\dots)$ and $\mathbf{x} = (\mathbf{1} \|\mathbf{0} \|\mathbf{1} \|\dots)$, resulting in two main terms. For each form, we apply *stars and bars* combinatorics twice. In particular, we assign h indistinguishable stars, i.e., ones, to distinguishable bins and independently also for $n - h$ zeros. Note that $l_j^{\mathbf{p}} = 0$ for $j > 2t + 1$. To ensure formula correctness, one may verify numerically that $l_1^{\mathbf{p}} + l_2^{\mathbf{p}} + \dots + l_{2t+1}^{\mathbf{p}}$ equals the left hand side of the Hamming bound in (4).

$$l_j^{\mathbf{p}} = \begin{cases} |\mathcal{M}|, & \text{if } j = 1 \\ |\mathcal{M}| \left(\sum_{h=\lfloor j/2 \rfloor}^t \binom{h-1}{\lfloor j/2 \rfloor - 1} \binom{n-h-1}{\lfloor j/2 \rfloor - 1} \right. \\ \quad \left. + \sum_{h=\lceil j/2 \rceil}^t \binom{h-1}{\lfloor j/2 \rfloor - 1} \binom{n-h-1}{\lfloor j/2 \rfloor - 1} \right), & \text{otherwise.} \end{cases} \quad (18)$$

3.4 Numerical Results

Fig. 8 presents numerical results for various BCH codes. We focus on small codes, as these allow for an exact exhaustive evaluation of the residual min-entropy using (9) and/or (10). As such, the tightness of various bounds can be assessed adequately. Fig. 8(d) nevertheless demonstrates that our algorithms support large codes equally well, in compliance with a practical key generator. Note that only half of the bias interval $b \in [0, 1]$ is depicted. The reason is that

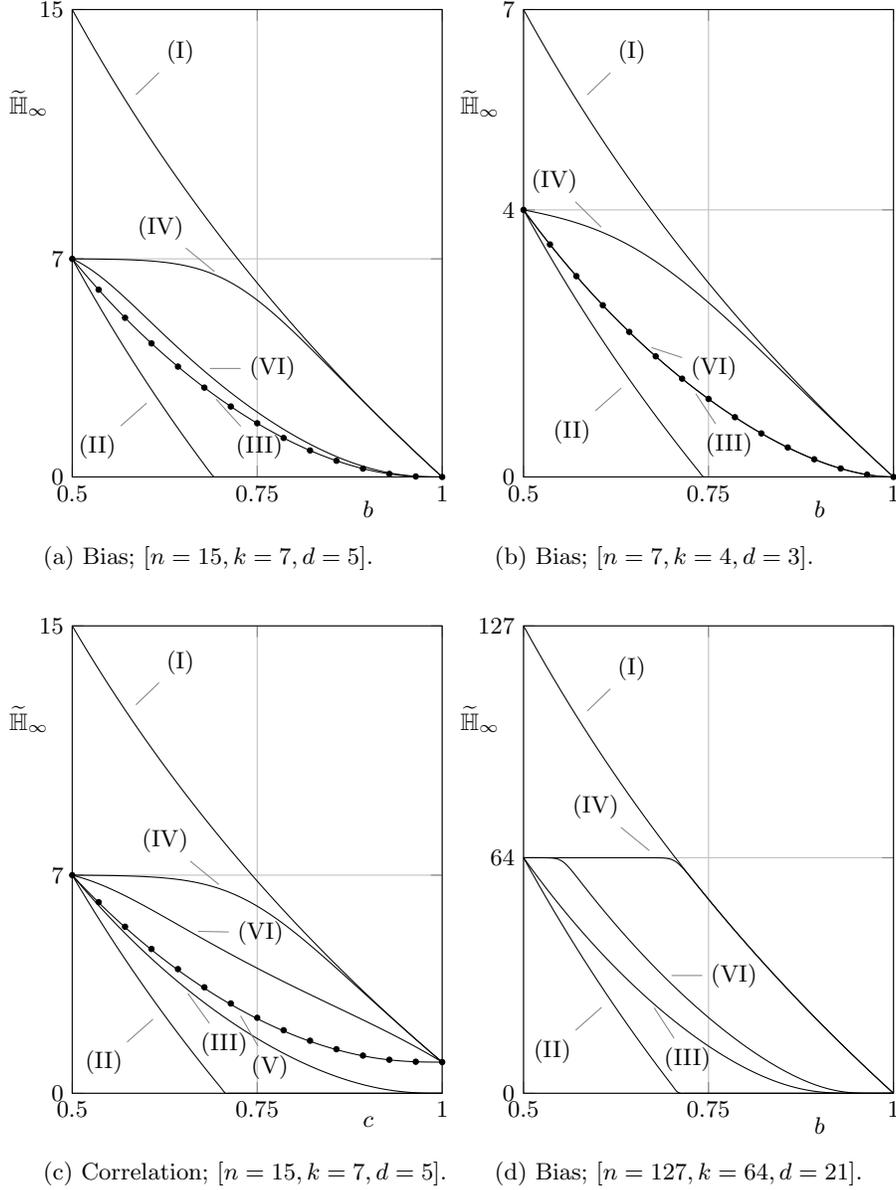


Fig. 8. The secure sketch min-entropy loss for various BCH codes. Dots correspond to an exact exhaustive evaluation of (9)/(10). The legend of the curves is as follows. (I) The ingoing min-entropy $\mathbb{H}_\infty(X) = -\log_2(q_1)$. (II) The lower bound $\tilde{\mathbb{H}}_\infty(X|P) = \max(\mathbb{H}_\infty(X) - (n - k), 0)$. (III) The lower bound on $\tilde{\mathbb{H}}_\infty(X|P)$ according to **BoundWorstCase**. (IV) The upper bound on $\tilde{\mathbb{H}}_\infty(X|P)$ according to **BoundBestCase**. (V) The lower bound on $\tilde{\mathbb{H}}_\infty(X|P)$ according to **BoundWorstCase2**. (VI) The upper bound on $\tilde{\mathbb{H}}_\infty(X|P)$ according to **BoundBestCase2**.

all curves mirror around the vertical axis of symmetry $b = \frac{1}{2}$. The same holds for the correlated distribution with parameter c .

Especially the lower bounds perform well, which benefits a conservative system provider. The best lower bounds in Figs. 8(a), (b) and (c) visually coincide with the exact result. The gap with the $(n - k)$ bound is the most compelling around $b, c \approx 0.7$, where the corresponding curves hit the horizontal axis $\mathbb{H}_\infty(X|P) = 0$. Also our upper bounds are considerably tighter than their more general alternatives in (6). Nevertheless, the latter bounds remain open for further improvement, with the exception of Fig. 8(b). An $[n = 7, k = 4, d = 3]$ code is perfect and lower and upper bounds then converge to the exact result for a biased distribution.

Table 1 quantifies the reduction in implementation footprint for a fuzzy extractor that produces a 128-bit key from a biased PUF. A concatenated code $\mathcal{C}_2 \circ \mathcal{C}_1$ is applied to z non-overlapping blocks of PUF response bits independently. We consider all 70 BCH codes \mathcal{C}_1 with $n_2 \leq 255$ and all 7 repetition codes \mathcal{C}_2 with $n_2 \leq 13$ and n_2 odd. The degenerate case $n_2 = 1$ ensures that our search space of 490 codes includes stand-alone BCH codes. Given a bias b and an expected bit error rate $\mathbb{E}_{v_i \leftarrow V_i}[p_{\text{error},i}]$, we retain the code that minimizes the number of PUF response bits n while satisfying the following two constraints. First, the residual min-entropy $\tilde{\mathbb{H}}_\infty(X|P) \geq 128$. Due to i.i.d. response bits, algorithm `BoundWorstCase` can be applied to an $[n_1 \cdot n_2, k_1, d_1 \cdot d_2]$ umbrella code and the residual min-entropy thereof is multiplied by z . A second constraint states that the expected device failure rate $\mathbb{E}_{v \leftarrow V}[P_{\text{fail}}] \leq 10^{-6}$. Due to i.i.d. response bits, we easily compute $\mathbb{E}_{v \leftarrow V}[P_{\text{fail}}] = 1 - (\text{F}_{\text{bino}}(t_1; n_1, 1 - \text{F}_{\text{bino}}(t_2; n_2, \mathbb{E}_{v_i \leftarrow V_i}[P_{\text{error}}])))^z$ [11].

According to the $(n - k)$ bound, a modest bias is highly detrimental already. Most notably, for $b = 0.56$, there is no code within the search space that satisfies all the design constraints. According to the newly derived bound, PUFs with a considerable bias can be supported. We emphasize that a carefully balanced custom-designed PUF tends to have a low bias. Notable cases of a high bias can typically be attributed to an asymmetry in either the PUF circuit or its lay-out, e.g., the D flip-flop PUF in [34] with $b > 0.7$. For low-bias PUFs, with $b \in [0.42, 0.58]$, a stand-alone secure sketch turns out to be competitive with state-of-the-art debiasing schemes [40, 18, 34, 35, 28].

3.5 Experimental Procedures

Fig. 9 depicts the newly developed theory in a larger practical context, as experienced by system providers. Conventionally, the min-entropy $\mathbb{H}_\infty(X)$ is estimated from a series of experimentally measured PUF responses \mathbf{x} , with subsequent application of the $(n - k)$ bound. There is no golden standard procedure for the former step though. Compression algorithms are occasionally applied; the main drawback is that these produce an upper bound on $\mathbb{H}_\infty(X)$ rather than a lower bound. A frequently used procedure is the estimation of a distribution X . For top-quality PUFs, the distribution is often assumed to be uniform. Various statistical tests that detect non-uniformities, e.g., the inter-distance metric [25] and

	b	$\mathbb{E}[P_{\text{error}}]$	$z \times [n_2, k_2, d_2] \circ [n_1, k_1, d_1]$	$\tilde{\mathbb{H}}_{\infty}(X P)$	PUF size n	$\mathbb{E}[P_{\text{fail}}]$
$n - k$ bound	0.50	$\approx 10.0\%$	$2 \times [5, 1, 5] \circ [127, 64, 21]$	128	1270	$\approx 3.26\text{E}-8$
	0.52	$\approx 10.0\%$	$3 \times [3, 1, 3] \circ [255, 87, 53]$	≈ 131.1	2295	$\approx 1.44\text{E}-8$
	0.54	$\approx 9.96\%$	$10 \times [5, 1, 5] \circ [255, 155, 27]$	≈ 134.4	12750	$\approx 5.56\text{E}-7$
	0.56	$\approx 9.90\%$	No code within the search space satisfies the constraints.			
new bound	0.50	$\approx 10.0\%$	$2 \times [5, 1, 5] \circ [127, 64, 21]$	128	1270	$\approx 3.26\text{E}-8$
	0.52	$\approx 10.0\%$	$1 \times [5, 1, 5] \circ [255, 163, 25]$	≈ 134.3	1275	$\approx 4.27\text{E}-7$
	0.54	$\approx 9.96\%$	$2 \times [3, 1, 3] \circ [255, 99, 47]$	≈ 132.5	1530	$\approx 5.35\text{E}-7$
	0.56	$\approx 9.90\%$	$3 \times [3, 1, 3] \circ [255, 87, 53]$	≈ 131.3	2295	$\approx 9.90\text{E}-9$
	0.58	$\approx 9.81\%$	$2 \times [5, 1, 5] \circ [255, 163, 25]$	≈ 130.0	2550	$\approx 4.85\text{E}-7$
	0.60	$\approx 9.71\%$	$3 \times [5, 1, 5] \circ [255, 155, 27]$	≈ 129.5	3825	$\approx 6.96\text{E}-8$
	0.62	$\approx 9.58\%$	$4 \times [5, 1, 5] \circ [255, 163, 25]$	≈ 130.4	5100	$\approx 4.42\text{E}-7$
	0.64	$\approx 9.42\%$	$10 \times [3, 1, 3] \circ [255, 99, 47]$	≈ 132.8	7650	$\approx 3.87\text{E}-7$
0.66	$\approx 9.24\%$	$17 \times [3, 1, 3] \circ [255, 99, 47]$	≈ 129.7	13005	$\approx 3.28\text{E}-7$	

Table 1. The implementation footprint of a practical fuzzy extractor, using the $(n - k)$ bound and the BoundWorstCase algorithm respectively. We assume i.i.d. response bits with bias b and noise component $\sigma_{\text{noise}} = 0.325$. Error-correction relies on the concatenation of a BCH code \mathcal{C}_1 and a repetition code \mathcal{C}_2 , with size $n_1 \in \{7, 15, 31, 63, 127, 255\}$ and $n_2 \in \{1, 3, 5, 7, 9, 11, 13\}$ respectively. Each row specifies the concatenated code that minimizes the number of PUF response bits n while satisfying the constraints, i.e., a residual min-entropy $\tilde{\mathbb{H}}_{\infty}(X|P) \geq 128$ and an expected failure rate $\mathbb{E}[p_{\text{fail}}] \leq 10^{-6}$. The helper data size of the code-offset sketch equals n also and is hence minimized simultaneously.

the NIST test suite [31], may provide some reassurance. Inspection of local bias and correlation effects often indicates the need for a more flawed distribution though [17, 2].

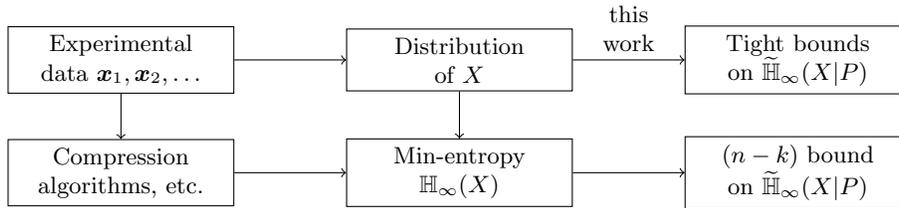


Fig. 9. Procedures for estimating the initial and residual min-entropy of an array-based PUF. The starting point is the experimental read-out of the PUF response \mathbf{x} of one or more fabricated devices. Among several alternatives, estimating the distribution of X comprehends a well-established technique for determining the initial min-entropy $\mathbb{H}_{\infty}(X)$. As elaborated in this work, it allows for tighter bounds on the residual min-entropy $\tilde{\mathbb{H}}_{\infty}(X|P)$, compared to the conventional $(n - k)$ formula.

As emphasized earlier, the biased and correlated distribution in this work are to be understood as 1-parameter proof-of-concept models. Experimental measurements should be performed in order to select the most suitable distribution. Nevertheless, for array-based PUFs, the distribution is expected to be representable in terms of local bias and/or correlation effects. As long as the number of sets n_{sets} of the distribution is limited, bounds can be produced.

4 Applications

The newly developed theory of Section 3 facilitates the design and analysis of error-correction methods for PUFs, as exemplified in twofold manner. First, we point out a fundamental security flaw in the reverse fuzzy extractor [36]. Second, we provide a motivational framework for debiasing schemes [40, 18, 34, 35, 28].

4.1 A Fundamental Security Flaw in Reverse Fuzzy Extractors

The reverse fuzzy extractor, as proposed by Van Herrewege et al. [36] at Financial Crypto 2012, improves the lightweight perspectives of PUF-based authentication protocols. The construction was therefore also adopted in the CHES 2015 protocol of Aysu et al. [2]. Instead of a single helper data exposure only, $\mathbf{p} \leftarrow \text{SSGen}(\tilde{\mathbf{x}})$ is regenerated and transferred with each protocol run by a resource-constrained PUF-enabled device. A receiving resource-rich server, storing reference response \mathbf{x} , can hence reconstruct $\tilde{\mathbf{x}} \leftarrow \text{SSRec}(\mathbf{x}, \mathbf{p})$ and establish a shared secret as such. The footprint of the device is reduced due to the absence of the heavyweight SSRec procedure.

We debunk the main security claim that repeated helper data exposure does not result in additional min-entropy loss. The revealed flaw is attributed to the misuse of a reusability proof of Boyen [8]. For the code-offset sketch with linear codes, the exposure of $\mathbf{p}^{(1)} \leftarrow \text{SSGen}(\mathbf{x})$ and $\mathbf{p}^{(2)} \leftarrow \text{SSGen}(\mathbf{x} \oplus \mathbf{e})$, with perturbation \mathbf{e} known and fully determined by the attacker, is provably equivalent. The latter helper data reveals that \mathbf{x} belongs to an identical coset $\{\mathbf{p}^{(1)} \oplus \mathbf{w} : \mathbf{w} \in \mathcal{W}\} = \{\mathbf{p}^{(2)} \oplus \mathbf{e} \oplus \mathbf{w} : \mathbf{w} \in \mathcal{W}\}$. However, perturbation \mathbf{e} is determined by PUF noisiness rather than by the attacker and its release hence reveals new information.

Given a sequence of protocol runs, the attacker can approximate all individual bit error rates, i.e., $p_{\text{error},i}$ with $i \in [1, n]$, as well as the coset to which reference \mathbf{x} belongs. For this purpose, the attacker collects helper data $\mathbf{p}^{(j)} \leftarrow \text{SSGen}(\mathbf{x} \oplus \mathbf{e}^{(j)})$, with $j \in [1, n_{\text{runs}}]$. The difference vector among each pair of noisy responses can be recovered as long as its Hamming weight does not exceed t ; consider a non-redundant set $(\mathbf{e}^{(1)} \oplus \mathbf{e}^{(j)})$ with $j \in [2, n_{\text{runs}}]$. For $n_{\text{runs}} \rightarrow \infty$, the estimates in (19) converge to their exact counterpart.

$$\lim_{n_{\text{runs}} \rightarrow \infty} \left[\begin{array}{l} (0, p_i), \\ (1, 1 - p_i), \end{array} \begin{array}{l} \text{if } p_i < 1/2, \\ \text{otherwise} \end{array} \right] = (e_i^{(1)}, p_{\text{error},i}) \quad \text{with } p_i = \sum_{j=2}^{n_{\text{runs}}} \frac{e_i^{(1)} \oplus e_i^{(j)}}{n_{\text{runs}} - 1}. \quad (19)$$

Exposure of $p_{\text{error},i} = \mathbf{F}_{\text{norm}}(-|v_i - v_{\text{thres}}|/\sigma_{\text{noise}})$ implies knowledge of threshold discrepancy $|v_i - v_{\text{thres}}|$. The residual min-entropy of reference response X is captured by (20).

$$\begin{aligned} \tilde{\mathbb{H}}_{\infty}(X|(P_{\text{SS}}, P_{\text{error},1}, \dots, P_{\text{error},n})) &= -\log_2\left(\mathbb{E}_{\mathbf{v} \leftarrow V} \left[\frac{\max_{\mathbf{w} \in \mathcal{W}} \mathbb{P}(V = \mathbf{v}_{\mathbf{w}})}{\sum_{\mathbf{w} \in \mathcal{W}} \mathbb{P}(V = \mathbf{v}_{\mathbf{w}})} \right]\right), \\ &\text{with } \mathbf{v}_{\mathbf{w},i} = v_{\text{thres}} + (1 - 2w_i)(v_i - v_{\text{thres}}) \text{ and } i \in [1, n]. \end{aligned} \quad (20)$$

Fig. 10 quantifies the residual min-entropy of X with the exclusion and inclusion of revealed bit error rates $p_{\text{error},i}$ respectively. In the latter case, we rely on a Monte Carlo evaluation of (20), as enabled by choosing a small $[n = 15, k = 7, d = 5]$ BCH code, given that an analytical approach is not so very straightforward. For both the biased and correlated distribution, it turns out that repeated helper data exposure results in additional min-entropy loss.

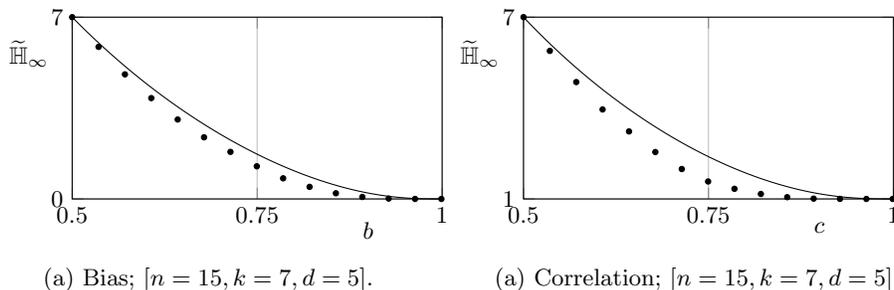


Fig. 10. The additional min-entropy loss attributed to revealed bit error rates. Solid lines represent $\tilde{\mathbb{H}}_{\infty}(X|P_{\text{SS}})$, as computed with `BoundWorstCase2`; Fig. 8 confirms the visual overlap with the exact result. Dots include revealed bit error rates, hereby relying on Monte Carlo evaluations of size 10^6 , i.e., the number of samples $\mathbf{v} \leftarrow V$.

The crucial insight for the biased distribution is that majority and minority bits tend to exhibit lower and higher error rates respectively. Note that $\mathbb{E}_{\mathbf{v}_i \leftarrow V_i}[P_{\text{error},i}|X_i = 1] < \mathbb{E}_{\mathbf{v}_i \leftarrow V_i}[P_{\text{error},i}|X_i = 0]$ if $b > \frac{1}{2}$ and vice versa otherwise. In terms of unanticipated min-entropy loss, the situation is identical to the *soft-decision decoding* scheme of Maes et al. [27]. As pointed out by Delvaux et al. [11], the attacker obtains the bit-specific bias $b_i = \mathbb{P}(X_i = 1|P_{\text{error},i} = p_{\text{error},i})$ in (21), which is more informative than $b = \mathbb{P}(X_i = 1)$.

$$b_i = \frac{\mathbf{f}_{\text{norm}}(v_{\text{thres}} + |v_i - v_{\text{thres}}|)}{\mathbf{f}_{\text{norm}}(v_{\text{thres}} + |v_i - v_{\text{thres}}|) + \mathbf{f}_{\text{norm}}(v_{\text{thres}} - |v_i - v_{\text{thres}}|)}. \quad (21)$$

The crucial insight for the correlated distribution is that correlation among V_i and V_j , with $i, j \in [1, n]$, implies correlation among $P_{\text{error},i}$ and $P_{\text{error},j}$. We

compute $\mathbb{P}(V = \mathbf{v}_w) = f_{\text{norm}}(\mathbf{v}_w, \mathbf{0}, \boldsymbol{\Sigma})$ in (20) with $\Sigma_{i,j} = \sin(\pi(c_{i,j} - 1/2))$ and $c_{i,j}$ defined in (7). The latter relation can be proven by integrating (22) in polar coordinates. The diagonal elements $\Sigma_{i,i} = 1$.

$$c_{i,j} = 2 \int_0^\infty \int_0^\infty f_{\text{norm}}\left((v_i \ v_j); (0 \ 0), \begin{pmatrix} 1 & \Sigma_{i,j} \\ \Sigma_{i,j} & 1 \end{pmatrix}\right) dv_i dv_j. \quad (22)$$

The revealed flaw differs from existing attacks by Delvaux et al. [12] and Becker [4] that apply to the original protocol [36] exclusively. The latter attacks comprehend the modeling of the highly correlated arbiter PUF via repeated helper data exposure; a preemptive fix can be found in the PhD thesis of Maes [25]. The newly revealed flaw is more fundamentally linked to the reverse fuzzy extractor primitive and applies to all existing protocols so far [36, 25, 2]. Observe in Fig. 10 that the overly conservative $(n - k)$ bound would compensate for the additional unanticipated min-entropy loss. However, this somewhat defeats the purpose in light of the original lightweight intentions, and this observation might not necessarily hold for every possible distribution. Further theoretical work may determine to which extent and at which cost reverse fuzzy extractors can be repaired. A potential fix already exists for biased distributions, as illustrated later-on.

4.2 Motivation for Debiasing Schemes

Debiasing schemes transform a biased PUF-induced distribution into a uniform distribution. A considerable fraction of the response bits is discarded in order to restore the balance between 0 and 1. Indices of retained bits are stored as helper data. A subsequent secure sketch, known to have an exact min-entropy loss of $(n - k)$ bits for uniform inputs, still corrects the errors. A first debiasing proposal is the *index-based syndrome* (IBS) scheme of Yu et al. [40], further generalized by Hiller et al. [18]. Second, several variations of the von Neumann debiasing algorithm can be applied. This was first proposed by van der Leest et al. [34], and later also by Van Herrewege in his PhD thesis [35]. Most recently, Maes et al. [28] presented an optimization of the von Neumann algorithm that applies to repetition codes in particular.

The generalized IBS debiasing scheme [18] in Fig. 11 locally rearranges the order of PUF response bits x so that a randomly chosen secret \mathbf{y} is reproduced. Although Y could be uniform over $\{0, 1\}^{1 \times n}$, a joint optimization with the subsequent secure sketch limits its set of outcomes to the codewords of a concatenated code $\mathcal{C}_2 \circ \mathcal{C}_1$. Response \mathbf{x} is partitioned in blocks of size n_{index} , each reproducing a codeword of the embedded $[n_2, k_2 = 1, d_2 = n_2]$ repetition code with $n_2 \leq n_{\text{index}}$. Helper data pointers are chosen so that the reproduction is as reliable as possible. This requires the estimation of individual bit error rates p_{error} , allowing to favor the selection of the most reliable zeros and ones. The codewords of the repetition code are approximately balanced in terms of Hamming weight, e.g., alternating patterns $\mathcal{W} = \{(0 \ 1 \ 0 \ 1 \ 0), (1 \ 0 \ 1 \ 0 \ 1)\}$ for $n_2 = 5$. Compared to the secure sketch in Fig. 1, the reconstruction reduces to $\hat{\mathbf{y}} \leftarrow \text{Correct}(\hat{\mathbf{y}})$, i.e., no

helper data is needed in addition to the index pointers. One may also output the uniform k_1 -bit secret $\text{Decode}(\tilde{\mathbf{y}})$.

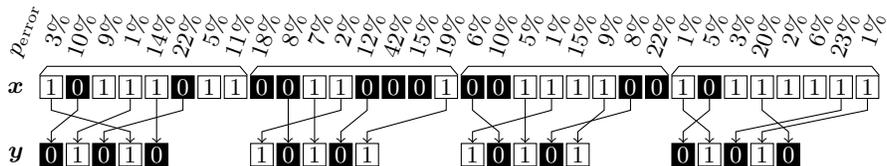


Fig. 11. The generalized IBS debiasing scheme. PUF response bits x are locally shuffled in order to reproduce the randomly chosen codeword \mathbf{y} of a concatenated code. In alignment with an $[n_2 = 5, k_2 = 1, d_2 = 5]$ repetition code, the three most reliable zeros and two most reliable ones, or vice versa depending on the repetition codeword, are selected within partitions of size $n_{\text{index}} = 8$. When all nominal ones within a certain partition are depleted, the least reliable zeros serve as a replacement, and vice versa.

The von Neumann debiasing schemes [34, 35, 28] in Fig. 12 partition the biased PUF response \mathbf{x} into pairs of bits. According to the original algorithm, the first bit of pairs 01 and 10 is retained, while pairs 00 and 11 are discarded as a whole. The number of retained bits obeys a binomial distribution $B(\lfloor n/2 \rfloor, 2b(1-b))$. A second pass of the algorithm on the decimated discarded pairs increases the expected number of retained bits. Three or more passes can be performed, but the gain in retention ratio drops sharply with the number of passes. The outgoing string \mathbf{y} is uniformly distributed and fed into a secure sketch. Maes et al. [28] improved the retention ratio for concatenated codes that embed an $[n_2, k_2 = 1, d_2 = n_2]$ repetition code with n_2 even. Undecimated sequences can be retained as a whole, given that \mathbf{y} is shuffled so that each undecimated sequence remains within the boundaries of a single repetition code. There is no additional min-entropy loss, i.e., a repetition code reveals all pairwise equalities among its corresponding response bits anyway. Note however that for three passes, $n_2 \geq 8$ already.

Table 2 quantifies the implementation footprint for a fuzzy extractor that produces a 128-bit key from a biased PUF. For IBS, the expected failure rate for reconstructing repetition codewords, i.e., $\mathbb{E}_{\mathbf{v} \leftarrow V}[P_{\text{fail}, \mathcal{C}_2}]$, is approximated via Monte Carlo simulations of size 10^6 . An exact evaluation via joint order statistics is not so very straightforward [11]. For the von Neumann schemes, we use the exact formula in (23), which incorporates a failure probability of $1/2$ whenever $n_2/2$ errors are detected. Note that $\mathbb{E}_{\mathbf{v} \leftarrow V}[P_{\text{fail}}] = 1 - (\mathbb{F}_{\text{bino}}(t_1; n_1, \mathbb{E}_{\mathbf{v} \leftarrow V}[P_{\text{fail}, \mathcal{C}_2}]))^z$. A complication for the von Neumann schemes is that the length of \mathbf{y} varies with \mathbf{x} . Therefore a yield is defined, i.e., the probability that sufficient bits can be provided for the subsequent secure sketch. An exact analytical evaluation of the retention ratio is computationally intensive from 3 passes onwards [28], so we rely on Monte Carlo simulations of size 10^6 instead.

b	$\mathbb{E}[P_{\text{error}}]$	$\mathbb{E}[P_{\text{error}} 0]$	$\mathbb{E}[P_{\text{error}} 1]$	Parameters	Retention	$z \times [n_2, k_2, d_2] \circ [n_1, k_1, d_1]$	$\tilde{\mathbb{H}}_{\infty}(X P)$	PUF size n	$\mathbb{E}[p_{\text{fail}, c_2}]$	$\mathbb{E}[P_{\text{fail}}]$
0.50	$\approx 10.0\%$	$\approx 10.0\%$	$\approx 10.0\%$	$n_{\text{index}} = 7$	$\approx 71.4\%$	$2 \times [5, 1, 5] \circ [127, 64, 21]$	128	1778	$\approx 1.01\text{E}-2$	$\approx 1.70\text{E}-7$
0.54	$\approx 9.96\%$	$\approx 10.6\%$	$\approx 9.40\%$	$n_{\text{index}} = 7$	$\approx 71.4\%$	$2 \times [5, 1, 5] \circ [127, 64, 21]$	128	1778	$\approx 1.12\text{E}-2$	$\approx 4.57\text{E}-7$
0.58	$\approx 9.81\%$	$\approx 11.2\%$	$\approx 8.79\%$	$n_{\text{index}} = 7$	$\approx 71.4\%$	$1 \times [5, 1, 5] \circ [255, 131, 37]$	131	1785	$\approx 1.41\text{E}-2$	$\approx 6.46\text{E}-9$
0.62	$\approx 9.58\%$	$\approx 11.8\%$	$\approx 8.18\%$	$n_{\text{index}} = 8$	62.5%	$2 \times [5, 1, 5] \circ [127, 64, 21]$	128	2032	$\approx 1.17\text{E}-2$	$\approx 7.09\text{E}-7$
0.66	$\approx 9.24\%$	$\approx 12.5\%$	$\approx 7.56\%$	$n_{\text{index}} = 8$	62.5%	$1 \times [5, 1, 5] \circ [255, 131, 37]$	131	2040	$\approx 1.83\text{E}-2$	$\approx 3.59\text{E}-7$
0.70	$\approx 8.80\%$	$\approx 13.2\%$	$\approx 6.92\%$	$n_{\text{index}} = 9$	$\approx 77.8\%$	$1 \times [7, 1, 7] \circ [255, 131, 37]$	131	2295	$\approx 1.90\text{E}-2$	$\approx 6.27\text{E}-7$
0.74	$\approx 8.24\%$	$\approx 13.9\%$	$\approx 6.27\%$	$n_{\text{index}} = 11$	$\approx 81.8\%$	$1 \times [9, 1, 9] \circ [255, 131, 37]$	131	2805	$\approx 1.62\text{E}-2$	$\approx 5.72\text{E}-8$
0.78	$\approx 7.57\%$	$\approx 14.6\%$	$\approx 5.58\%$	$n_{\text{index}} = 13$	$\approx 84.6\%$	$1 \times [11, 1, 11] \circ [255, 131, 37]$	131	3315	$\approx 1.65\text{E}-2$	$\approx 7.32\text{E}-8$
0.82	$\approx 6.76\%$	$\approx 15.4\%$	$\approx 4.85\%$	$n_{\text{index}} = 16$	$\approx 68.8\%$	$1 \times [11, 1, 11] \circ [255, 131, 37]$	131	4080	$\approx 1.66\text{E}-2$	$\approx 7.57\text{E}-8$
0.86	$\approx 5.80\%$	$\approx 16.4\%$	$\approx 4.07\%$	$n_{\text{index}} = 16$	$\approx 81.3\%$	$2 \times [13, 1, 13] \circ [255, 71, 59]$	142	8160	$\approx 3.57\text{E}-2$	$\approx 2.85\text{E}-8$
0.90	$\approx 4.64\%$	$\approx 17.5\%$	$\approx 3.21\%$	$n_{\text{index}} = 16$	$\approx 81.3\%$	$3 \times [13, 1, 13] \circ [255, 45, 87]$	135	12240	$\approx 7.51\text{E}-2$	$\approx 6.42\text{E}-7$
0.50	$\approx 10.0\%$	$\approx 10.0\%$	$\approx 10.0\%$		$\approx 83.4\%$	$4 \times [8, 1, 8] \circ [63, 36, 11]$	144	2418	$\approx 2.73\text{E}-3$	$\approx 9.85\text{E}-8$
0.54	$\approx 9.96\%$	$\approx 10.6\%$	$\approx 9.40\%$		$\approx 81.6\%$	$4 \times [8, 1, 8] \circ [63, 36, 11]$	144	2471	$\approx 2.72\text{E}-3$	$\approx 9.73\text{E}-8$
0.58	$\approx 9.81\%$	$\approx 11.2\%$	$\approx 8.79\%$	3 passes	$\approx 77.0\%$	$4 \times [8, 1, 8] \circ [63, 36, 11]$	144	2617	$\approx 2.71\text{E}-3$	$\approx 9.37\text{E}-8$
0.62	$\approx 9.58\%$	$\approx 11.8\%$	$\approx 8.18\%$		$\approx 70.7\%$	$3 \times [10, 1, 10] \circ [63, 45, 7]$	135	2675	$\approx 8.70\text{E}-4$	$\approx 9.81\text{E}-7$
0.66	$\approx 9.24\%$	$\approx 12.5\%$	$\approx 7.56\%$	multi-out ($n_2 \geq 8$)	$\approx 63.6\%$	$3 \times [10, 1, 10] \circ [63, 45, 7]$	135	2971	$\approx 8.52\text{E}-4$	$\approx 9.05\text{E}-7$
0.70	$\approx 8.80\%$	$\approx 13.2\%$	$\approx 6.92\%$		$\approx 56.2\%$	$3 \times [10, 1, 10] \circ [63, 45, 7]$	135	3365	$\approx 8.29\text{E}-4$	$\approx 8.12\text{E}-7$
0.74	$\approx 8.24\%$	$\approx 13.9\%$	$\approx 6.27\%$		$\approx 48.6\%$	$3 \times [10, 1, 10] \circ [63, 45, 7]$	135	3885	$\approx 8.00\text{E}-4$	$\approx 7.06\text{E}-7$
0.78	$\approx 7.57\%$	$\approx 14.6\%$	$\approx 5.58\%$	retention	$\approx 41.4\%$	$3 \times [10, 1, 10] \circ [63, 45, 7]$	135	4567	$\approx 7.65\text{E}-4$	$\approx 5.91\text{E}-7$
0.82	$\approx 6.76\%$	$\approx 15.4\%$	$\approx 4.85\%$	yield 99%	$\approx 33.5\%$	$3 \times [10, 1, 10] \circ [63, 45, 7]$	135	5650	$\approx 7.23\text{E}-4$	$\approx 4.72\text{E}-7$
0.86	$\approx 5.80\%$	$\approx 16.4\%$	$\approx 4.07\%$		$\approx 26.1\%$	$3 \times [10, 1, 10] \circ [63, 45, 7]$	135	7237	$\approx 6.73\text{E}-4$	$\approx 3.55\text{E}-7$
0.90	$\approx 4.64\%$	$\approx 17.5\%$	$\approx 3.21\%$		$\approx 18.5\%$	$3 \times [10, 1, 10] \circ [63, 45, 7]$	135	10212	$\approx 6.13\text{E}-4$	$\approx 2.45\text{E}-7$

Table 2. The implementation footprint of a practical fuzzy extractor, using the generalized IBS and von Neumann debiasing schemes respectively. The setting is identical to Table 1. A residual min-entropy $\tilde{\mathbb{H}}_{\infty}(X|P) \geq 128$ and an expected failure rate $\mathbb{E}[p_F] \leq 10^{-6}$ are imposed. All BCH codes with $n_1 \in \{7, 15, 31, 63, 127, 255\}$ are considered. For IBS and the von Neumann schemes, we consider all repetition codes with $n_2 \in \{1, 3, 5, 7, 9, 11, 13\}$ and $n_2 \in \{8, 10, 12, 14\}$ respectively. For IBS in particular, we impose the constraint $n_{\text{index}} \leq 16$.

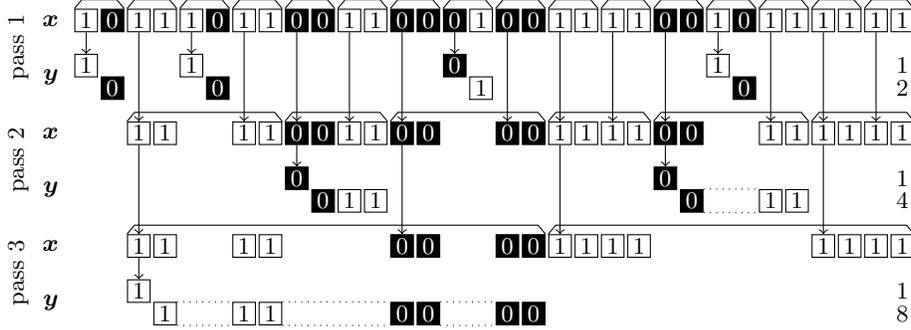


Fig. 12. Several variations of the von Neumann debiasing algorithm. In the first pass, 01 and 10 sequences are retained, while 00 and 11 sequences are discarded. Optionally, a second pass can retain previously discarded 0011 and 1100 sequences, but therefore not the 0000 and 1111 sequences. A third pass retains previously discarded 0001111 and 1111000 sequences, but therefore not the 00000000 and 11111111 sequences. Originally, only the first bit of each retained sequence contributes to y , resulting in uniformity. A joint optimization with repetition codes allows to retain sequences as a whole, i.e., 2, 4, and 8 bits are retained in the first, second and third pass respectively.

$$\begin{aligned}
\mathbb{E}_{\mathbf{v} \leftarrow V} [P_{\text{fail}, \mathcal{C}_2}] &= 1 - \sum_{i=0}^t \mathbf{f}_{\text{bino}} \left(i; \frac{n_2}{2}, \mathbb{E}_{v_i \leftarrow V_i} [P_{\text{error}, i} | X_i = 0] \right) \cdot \mathbf{F}_{\text{bino}} \left(t - i; \right. \\
&\quad \left. \frac{n_2}{2}, \mathbb{E}_{v_i \leftarrow V_i} [P_{\text{error}, i} | X_i = 1] \right) - \frac{1}{2} \sum_{i=0}^{t+1} \mathbf{f}_{\text{bino}} \left(i; \frac{n_2}{2}, \mathbb{E}_{v_i \leftarrow V_i} [P_{\text{error}, i} | X_i = 0] \right) \quad (23) \\
&\quad \cdot \mathbf{f}_{\text{bino}} \left(t + 1 - i; \frac{n_2}{2}, \mathbb{E}_{v_i \leftarrow V_i} [P_{\text{error}, i} | X_i = 1] \right).
\end{aligned}$$

Prior debiasing proposals [40, 18, 34, 35, 28] conjectured that a stand-alone sketch cannot handle biased distributions well. This corresponds to an educated guess, originating from the extrapolation of repetition code insights and/or the application of the $(n-k)$ bound. Our newly developed bounds clearly resolve this motivational uncertainty. It turns out that a stand-alone sketch is competitive in the low-bias region, e.g., $b \in [0.42, 0.58]$. Nevertheless, for high-bias situations, debiasing schemes are needed. The benefit is amplified by choosing a sketch with a k -bit output, several of which are listed in Appendix A. The uniform output is then directly usable as a key, hereby eliminating the `Hash` function and its additional min-entropy loss in case the leftover hash lemma is applied.

Finally, we highlight that the von Neumann debiasing scheme in Fig. 13 was claimed to be reusable [28]. This claim holds, despite overlooking the misuse of Boyen’s proof and stating that a stand-alone sketch is reusable. An unintended side effect of introducing placeholder pairs is that individual bit error rates cannot be estimated anymore. Helper data only allows for the estimation of pairwise

error rates. The scheme is considerably less efficient than other von Neumann variants though, showing that reusability comes at a price.

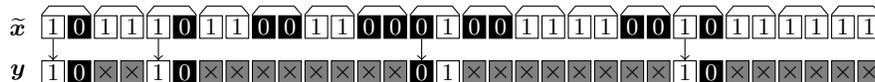


Fig. 13. A reusable von Neumann debiasing scheme that allows for the enrollment of an unlimited number of noisy PUF responses \tilde{x} . There is a single pass that retains 01 and 10 sequences as a whole. The 00 and 11 sequences merely serve as placeholders, contributing to neither the enrollment nor the reproduction, i.e., only part of $\tilde{x} \oplus w$ is released as helper data. The $[n_2, k_2 = 1, d_2 = n_2]$ repetition code with n_2 even is virtually shortened due to local placeholder pairs.

5 Conclusion

Secure sketches are the main workhorse of modern PUF-based key generators. The min-entropy loss of most sketches is upper-bounded by $(n - k)$ bits and designers typically instantiate system parameters accordingly. However, the latter bound tends to be overly pessimistic, resulting in an unfortunate implementation overhead. We showcased the proportions for a prominent category of PUFs, with bias and spatial correlations acting as the main non-uniformities. New considerably tighter bounds were derived, valid for a variety of popular but algebraically complex codes. These bounds are unified in the sense of being applicable to seven secure sketch constructions. Deriving tighter alternatives for the $(n - k)$ bound counts as unexplored territory and we established the first significant stepping stone. New techniques may have to be developed in order to tackle more advanced *second-order* distributions. Elaborating a wider range of applications would be another area of progress. We hope to have showcased the potential by debunking the main security claim of the reverse fuzzy extractor and by providing proper quantitative motivation for debiasing schemes.

Acknowledgment

The authors greatly appreciate the support received. The European Union’s Horizon 2020 research and innovation programme under grant number 644052 (HECTOR). The Research Council of KU Leuven, GOA TENSE (GOA/11/007), the Flemish Government through FWO G.0550.12N and the Hercules Foundation AKUL/11/19. The national major development program for fundamental research of China (973 Plan) under grant number 2013CB338004. Jeroen Delvaux is funded by IWT-Flanders grant number SBO 121552. Matthias Hiller is funded by the German Federal Ministry of Education and Research (BMBF) in the project SIBASE through grant number 01IS13020A.

References

1. 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.
2. A. Aysu, E. Gulcan, D. Moriyama, P. Schaumont, and M. Yung. End-To-End Design of a PUF-Based Privacy Preserving Authentication Protocol. In *Cryptographic Hardware and Embedded Systems - CHES 2015 - 17th International Workshop, September 13-16, 2015, Proceedings*, pages 556–576, 2015.
3. B. Barak, Y. Dodis, H. Krawczyk, O. Pereira, K. Pietrzak, F. Standaert, and Y. Yu. Leftover Hash Lemma, Revisited. In *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference*, pages 1–20, 2011.
4. G. T. Becker. On the Pitfalls of Using Arbiter-PUFs as Building Blocks. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 34(8):1295–1307, 2015.
5. C. H. Bennett, G. Brassard, C. Crépeau, and M. Skubiszewska. Practical Quantum Oblivious Transfer. In *Advances in Cryptology - CRYPTO 1991, 11th Annual Cryptology Conference*, pages 351–366, 1991.
6. M. Bhargava and K. Mai. An efficient reliable PUF-based cryptographic key generator in 65nm CMOS. In *Design, Automation & Test in Europe Conference & Exhibition, DATE 2014, Dresden, Germany, March 24-28, 2014*, pages 1–6, 2014.
7. C. Bösch, J. Guajardo, A. Sadeghi, J. Shokrollahi, and P. Tuyls. Efficient Helper Data Key Extractor on FPGAs. In *Cryptographic Hardware and Embedded Systems - CHES 2008, 10th International Workshop*, pages 181–197, 2008.
8. X. Boyen. Reusable cryptographic fuzzy extractors. In *Proceedings of the 11th ACM Conference on Computer and Communications Security, CCS 2004, Washington, DC, USA, October 25-29, 2004*, pages 82–91, 2004.
9. L. Carter and M. N. Wegman. Universal Classes of Hash Functions. *Journal of Computer and System Sciences*, 18(2):143–154, 1979.
10. P. Charpin, T. Helleseth, and V. A. Zinoviev. The Coset Distribution of Triple-Error-Correcting Binary Primitive BCH Codes. *IEEE Transactions on Information Theory*, 52(4):1727–1732, 2006.
11. J. Delvaux, D. Gu, D. Schellekens, and I. Verbauwhede. Helper Data Algorithms for PUF-Based Key Generation: Overview and Analysis. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 2015.
12. J. Delvaux, R. Peeters, D. Gu, and I. Verbauwhede. A Survey on Lightweight Entity Authentication with Strong PUFs. *ACM Comput. Surv.*, 48(2):26, 2015.
13. J. Delvaux and I. Verbauwhede. Fault Injection Modeling Attacks on 65 nm Arbiter and RO Sum PUFs via Environmental Changes. *IEEE Trans. on Circuits and Systems*, 61-I(6):1701–1713, 2014.
14. Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. *SIAM Journal on Computing*, 38(1):97–139, 2008.
15. W. Feller. *An Introduction to Probability Theory and Its Applications, Vol. 1, 3rd Edition*. 1968.
16. J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A Pseudorandom Generator from any One-way Function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
17. A. V. Herwege, V. van der Leest, A. Schaller, S. Katzenbeisser, and I. Verbauwhede. Secure PRNG seeding on commercial off-the-shelf microcontrollers. In *TrustED'13, Proceedings of the 2013 ACM Workshop on Trustworthy Embedded Devices*, pages 55–64, 2013.

18. M. Hiller, D. Merli, F. Stumpf, and G. Sigl. Complementary IBS: application specific error correction for PUFs. In *2012 IEEE International Symposium on Hardware-Oriented Security and Trust, HOST 2012, June 3-4, 2012*, pages 1–6, 2012.
19. M. Hiller, M. Yu, and M. Pehl. Systematic Low Leakage Coding for Physical Unclonable Functions. In *ASIA CCS 2015, 10th ACM Symposium on Information, Computer and Communications Security*, pages 155–166, 2015.
20. D. E. Holcomb, W. P. Burlison, and K. Fu. Power-Up SRAM State as an Identifying Fingerprint and Source of True Random Numbers. *IEEE Transactions on Computers*, 58(9):1198–1210, 2009.
21. A. Juels and M. Wattenberg. A Fuzzy Commitment Scheme. In *CCS 1999, 6th ACM Conference on Computer and Communications Security*, pages 28–36, 1999.
22. H. Kang, Y. Hori, T. Katashita, M. Hagiwara, and K. Iwamura. Cryptographic Key Generation from PUF Data Using Efficient Fuzzy Extractors. In *ICACT 2014, 16th International Conference on Advanced Communication Technology*, 2014.
23. P. Koeberl, J. Li, A. Rajan, and W. Wu. Entropy loss in PUF-based key generation schemes: The repetition code pitfall. In *2014 IEEE International Symposium on Hardware-Oriented Security and Trust, HOST 2014, Arlington, VA, USA, May 6-7, 2014*, pages 44–49, 2014.
24. F. J. MacWilliams and N. J. A. Sloane. *The theory of error correcting codes*. 1977.
25. R. Maes. *Physically Unclonable Functions: Constructions, Properties and Applications*. PhD thesis, KU Leuven, 2012. Ingrid Verbauwhede (promotor).
26. R. Maes. An accurate probabilistic reliability model for silicon PUFs. In *Cryptographic Hardware and Embedded Systems - CHES 2013 - 15th International Workshop, Santa Barbara, CA, USA, August 20-23, 2013. Proceedings*, pages 73–89, 2013.
27. R. Maes, P. Tuyls, and I. Verbauwhede. A Soft Decision Helper Data Algorithm for SRAM PUFs. In *ISIT 2009, IEEE International Symposium on Information Theory*, pages 2101–2105, 2009.
28. R. Maes, V. van der Leest, E. van der Sluis, and F. Willems. Secure key generation from biased PUFs: extended version. *Journal of Cryptographic Engineering*, 6(2):121–137, 2016.
29. R. Maes, A. Van Herrewege, and I. Verbauwhede. PUFKY: A Fully Functional PUF-Based Cryptographic Key Generator. In *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop*, pages 302–319, 2012.
30. L. Reyzin. Entropy Loss is Maximal for Uniform Inputs. Technical Report BUCS-TR-2007-011, Department of Computer Science, Boston University, September 2007.
31. A. Rukhin, J. Soto, J. Nechvatal, M. Smid, E. Barker, S. Leigh, M. Levenson, M. Vangel, D. Banks, A. Heckert, J. Dray, and S. Vo. A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications. National Institute for Standards and Technology (NIST).
32. P. Tuyls, A. H. M. Akkermans, T. A. M. Kevenaar, G.-J. Schrijen, A. M. Bazen, and R. N. J. Veldhuis. Practical Biometric Authentication with Template Protection. In *AVBPA 2005, Int. Conference on Audio- and Video-Based Biometric Person Authentication*, pages 436–446, 2005.
33. P. Tuyls, G.-J. Schrijen, B. Škorić, J. van Geloven, N. Verhaegh, and R. Wolters. Read-Proof Hardware from Protective Coatings. In *CHES 2006, Int. Workshop on Cryptographic Hardware and Embedded Systems*, pages 369–383, 2006.

34. V. van der Leest, G.-J. Schrijen, H. Handschuh, and P. Tuyls. Hardware Intrinsic Security from D Flip-flops. In *Proceedings of the Fifth ACM Workshop on Scalable Trusted Computing*, STC '10, pages 53–62, 2010.
35. A. Van Herrewege. *Lightweight PUF-based Key and Random Number Generation*. PhD thesis, KU Leuven, 2015. Ingrid Verbauwhede (promotor).
36. A. Van Herrewege, S. Katzenbeisser, R. Maes, R. Peeters, A. Sadeghi, I. Verbauwhede, and C. Wachsmann. Reverse Fuzzy Extractors: Enabling Lightweight Mutual Authentication for PUF-Enabled RFIDs. In *Financial Cryptography and Data Security - 16th International Conference, FC 2012, Kralendijk, Bonaire, February 27-March 2, 2012, Revised Selected Papers*, pages 374–389, 2012.
37. Y. Wang, S. Rane, S. C. Draper, and P. Ishwar. A Theoretical Analysis of Authentication, Privacy, and Reusability Across Secure Biometric Systems. *IEEE Transactions on Information Forensics and Security*, 7(6):1825–1840, 2012.
38. H. Yu, P. H. W. Leong, H. Hinkelmann, L. Möller, M. Glesner, and P. Zipf. Towards a Unique FPGA-Based Identification Circuit Using Process Variations. In *FPL 2009, Int. Conference on Field Programmable Logic and Applications*, pages 397–402, 2009.
39. M. Yu. Turn FPGAs Into “Key” Players In The Cryptographics Field, Jul 2009. Electronic Design Magazine, <http://electronicdesign.com/fpgas/turn-fpgas-key-players-cryptographics-field>.
40. M. Yu and S. Devadas. Secure and Robust Error Correction for Physical Unclonable Functions. *IEEE Design & Test of Computers*, 27(1):48–65, 2010.

A Secure Sketch Equivalency Proofs

Bounds previously derived for the code-offset method of Dodis et al. [14] apply to six other constructions equally well. For convenience, we generalize the original secure sketch so that its reconstructed output $\mathbf{y} \leftarrow \text{SSRep}(\mathbf{x}, \mathbf{p})$ is not necessarily equal to \mathbf{x} . As such, the prior notion of *fuzzy commitment* [21] can be supported as well. Hereby, we commit to a secret value \mathbf{y} by binding it to \mathbf{x} . One may decommit given an $\tilde{\mathbf{x}}$ that is sufficiently close to \mathbf{x} . Constructions that return a substring of \mathbf{x} , e.g., [22], are supported too. The fuzzy extractor definition offers intrinsic support for both cases, without any modifications from our part. The key is still computed as $\mathbf{k} \leftarrow \text{Hash}(\mathbf{y})$.

Fig. 14 specifies the seven secure sketch constructions of interest, all instantiated with a binary code \mathcal{C} . We now review additional coding theory, before transitioning to individual sketch discussions. A generator matrix is in *standard form* if $\mathbf{G} = (\mathbf{I}_k \parallel \mathbf{A})$. I.e., the first k bits of a codeword equal the message, followed by $n - k$ redundancy bits. A *parity check matrix* \mathbf{H} , with dimensions $(n - k) \times n$, determines the so-called *syndrome* $\mathbf{s} = \tilde{\mathbf{w}} \cdot \mathbf{H}^T$. The syndrome captures all the information necessary for decoding $\tilde{\mathbf{w}}$. For each codeword \mathbf{w} , the following holds: $\mathbf{0} = \mathbf{w} \cdot \mathbf{H}^T$. Therefore, the syndrome can be rewritten as $\mathbf{s} = \mathbf{e} \cdot \mathbf{H}^T$. Generator and parity check matrices can be derived from each other. E.g., for a generator matrix in standard form, $\mathbf{H} = (\mathbf{A}^T \parallel \mathbf{I}_{n-k})$. There is a one-to-one correspondence between cosets and syndromes [24].

All seven constructions exhibit an identical min-entropy loss. Or more precisely, all have the same residual min-entropy $\mathbb{H}_\infty(Y|P)$ given in (24), as long as

$\mathbf{p} \leftarrow \text{SSGen}(\mathbf{x})$	$\hat{\mathbf{y}} \leftarrow \text{SSRep}(\tilde{\mathbf{x}}, \mathbf{p})$	
Random $\mathbf{w} \in \mathcal{C}$ $\mathbf{p} \leftarrow \mathbf{x} \oplus \mathbf{w}$	$\tilde{\mathbf{w}} \leftarrow \tilde{\mathbf{x}} \oplus \mathbf{p} = \mathbf{w} \oplus \mathbf{e}$ $\hat{\mathbf{y}} = \hat{\mathbf{w}} \leftarrow \text{Correct}(\tilde{\mathbf{w}})$	(a) Code-offset method of Juels et al. [21].
	$\tilde{\mathbf{w}} \leftarrow \tilde{\mathbf{x}} \oplus \mathbf{p} = \mathbf{w} \oplus \mathbf{e}$ $\hat{\mathbf{y}} = \hat{\mathbf{x}} \leftarrow \mathbf{p} \oplus \text{Correct}(\tilde{\mathbf{w}})$	(b) Code-offset method of Dodis et al. [14].
	$\tilde{\mathbf{w}} \leftarrow \tilde{\mathbf{x}} \oplus \mathbf{p} = \mathbf{w} \oplus \mathbf{e}$ $\hat{\mathbf{y}} = \hat{\mathbf{m}} \leftarrow \text{Decode}(\tilde{\mathbf{w}})$	(c) Code-offset method of Tuyls et al. [32].
$\mathbf{p} \leftarrow \mathbf{x} \cdot \mathbf{H}^T$	$\mathbf{s} \leftarrow \tilde{\mathbf{x}} \cdot \mathbf{H}^T \oplus \mathbf{p} = \mathbf{e} \cdot \mathbf{H}^T$ Determine $\hat{\mathbf{e}}$ $\hat{\mathbf{y}} = \hat{\mathbf{x}} \leftarrow \tilde{\mathbf{x}} \oplus \hat{\mathbf{e}}$	(d) Syndrome method of Bennett et al. [5].
$\mathbf{p} \leftarrow \mathbf{x}(1:k) \cdot \mathbf{A}$ $\oplus \mathbf{x}(k+1:n)$	$\hat{\mathbf{w}} \leftarrow \text{Correct}(\tilde{\mathbf{x}} \oplus (\mathbf{0} \parallel \mathbf{p}))$ $\hat{\mathbf{y}} = \hat{\mathbf{x}} \leftarrow \hat{\mathbf{w}} \oplus (\mathbf{0} \parallel \mathbf{p})$	(e) Systematic method of Yu [39].
	$\hat{\mathbf{y}} = \hat{\mathbf{x}}(1:k) \leftarrow \text{Decode}(\tilde{\mathbf{x}} \oplus (\mathbf{0} \parallel \mathbf{p}))$	(f) Systematic method of Kang et al. [22].
$\mathbf{p} \leftarrow j$ so that $\mathbf{x} \in \mathcal{C}_j$	$\hat{\mathbf{y}} = \hat{\mathbf{m}} \leftarrow \text{Decode}_{\mathcal{C}_j}(\tilde{\mathbf{x}})$	(g) Multi-code method of Ahlswede et al. [1].

Fig. 14. Seven secure sketch constructions, all having an n -bit input \mathbf{x} . Correctness of reconstruction is guaranteed, given a noisy version $\tilde{\mathbf{x}}$ with $\text{HD}(\mathbf{x}, \tilde{\mathbf{x}}) \leq t$.

the ingoing distribution X and the code \mathcal{C} are identical. A consequence thereof is that the well-known $(n - k)$ upper bound on the min-entropy loss as well as our newly derived bounds apply to all seven sketches. Simple equivalency proofs are established in pairwise manner, as guided by Fig. 15. Several pairwise equivalencies were already established in existing literature, e.g., [37, 11], but these often impose unnecessary restrictions on the distribution. We hence make progress in terms of completeness and generality.

$$\tilde{\mathbb{H}}_{\infty}(Y|P) = -\log_2 \left(\mathbb{E}_{\mathbf{p} \leftarrow P} \left[\max_{\mathbf{y} \in \mathcal{Y}} \mathbb{P}((Y = \mathbf{y})|(P = \mathbf{p})) \right] \right). \quad (24)$$

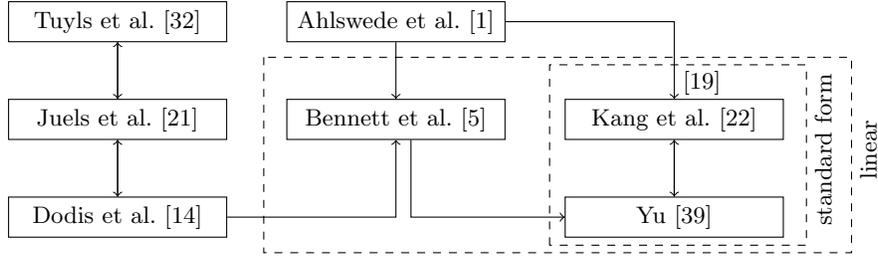


Fig. 15. Pairwise min-entropy loss equivalencies among seven sketches, as indicated by the arrows. Transitive relations apply when following the arrows. E.g., the schemes of Dodis et al. and Kang et al. are equivalent, given that both are instantiated with a linear code in standard form.

A.1 Code-Offset Methods of Juels et al., Dodis et al. and Tuyls et al.

The code-offset method of Juels et al. [21] is represented by Fig. 14(a). The code \mathcal{C} is not necessarily linear. Even more, it is not required to be a block code either. Fig. 14(b) represents a modification where Rep returns sketch input \mathbf{x} rather than codeword \mathbf{w} , as proposed by Dodis et al. [14]. For the latter, it was proven that the $(n-k)$ upper bound on the min-entropy loss $\Delta\mathbb{H}_\infty$ holds, given a block code. Fig. 14(c) represents another minor modification where Rep returns message \mathbf{m} , as suggested by Tuyls et al. [32]. This necessitates an implementation of Decode rather than Correct .

All three code-offset methods produce the same helper data \mathbf{p} but differ in their reconstructed output \mathbf{y} . Nevertheless, we argue that the residual min-entropy is identical. This follows from an underlying one-to-one correspondence, given in (25). Encode comprehends a bijection between message space \mathcal{M} and codeword space \mathcal{W} . Furthermore, for a given \mathbf{p} , there is a bijection between \mathcal{W} and a reduced response space $\mathcal{X}' = \{\mathbf{p} \oplus \mathbf{w} \mid \mathbf{w} \in \mathcal{W}\} \subseteq \mathcal{X}$. Therefore, (24) evaluates to the same value for all three methods. Note that $|\mathcal{M}| = |\mathcal{W}| = |\mathcal{X}'|$.

$$\begin{aligned} \forall(\mathbf{p}, \mathbf{m}) \in (\mathcal{P} \times \mathcal{M}), \mathbb{P}((M = \mathbf{m})|(P = \mathbf{p})) &= \mathbb{P}((W = \text{Encode}(\mathbf{m}))|(P = \mathbf{p})) \\ &= \mathbb{P}((X = \text{Encode}(\mathbf{m}) \oplus \mathbf{p})|(P = \mathbf{p})). \end{aligned} \tag{25}$$

A.2 Syndrome Method of Bennett et al.

The syndrome method of Bennett et al. [5] is represented by Fig. 14(d). Although initially proposed as part of a *quantum oblivious transfer* protocol, it maps quite easily to the secure sketch framework of Dodis et al. [14]. The method requires a linear code \mathcal{C} , given the use of a parity check matrix \mathbf{H} . The well-known $(n-k)$ upper bound on the min-entropy loss $\Delta\mathbb{H}_\infty$ holds, as proven by Dodis et al. [14].

This is a trivial consequence from the universally valid expression in (26), given that the helper data \mathbf{p} is limited to $(n - k)$ bits.

$$\tilde{\mathbb{H}}_\infty(X|P) \geq \tilde{\mathbb{H}}_\infty(X) - \log_2(|\mathcal{P}|). \quad (26)$$

The syndrome method of Bennett et al. and the code-offset method of Dodis et al. both reconstruct $\mathbf{y} = \mathbf{x}$. Furthermore, for both methods, helper data \mathbf{p} reveals in which coset \mathbf{x} resides. For the syndrome method, this is a trivial consequence from the one-to-one correspondence between cosets and syndromes. For the code-offset method, \mathbf{p} comprehends a random element in the same coset as \mathbf{x} . Note that the code-offset method is being instantiated with a linear code, given that the syndrome method is restricted to this case. The residual min-entropy of both methods can hence be written as shown in (5).

A.3 Systematic Methods of Yu and Kang et al.

The method of Yu [39] is represented by Fig. 14(e). It requires a linear code \mathcal{C} with the generator matrix in standard form, i.e., $\mathbf{G} = (\mathbf{I}_k \| \mathbf{A})$. We observe that $\Delta\mathbb{H}_\infty \leq (n - k)$ holds due to (26), given that helper data \mathbf{p} is limited to $(n - k)$ bits. Fig. 14(f) represents a slightly modified method where Rep returns $(x_1 \ x_2 \ \dots \ x_k)$ rather than \mathbf{x} . This was first proposed by Kang et al. in [22] and independently also by Hiller et al. in [19]. Nevertheless, (27) indicates that the residual min-entropy is identical. The main insight is that $(x_1 \ x_2 \ \dots \ x_k)$ and \mathbf{p} fully determine $(x_{k+1} \ x_{k+2} \ \dots \ x_n)$.

$$\begin{aligned} \forall(\mathbf{p}, \mathbf{x}) \in (\mathcal{P} \times \mathcal{X}), \mathbb{P}((X(1:k) = \mathbf{x}(1:k)) | (P = \mathbf{p})) \\ = \mathbb{P}((X = (\mathbf{x}(1:k) | (\mathbf{x}(1:k) \cdot \mathbf{A} \oplus \mathbf{p})) | (P = \mathbf{p})). \end{aligned} \quad (27)$$

The methods of Bennett et al. and Yu both reconstruct the sketch input, i.e., $\mathbf{y} = \mathbf{x}$. We are the first to observe though that the helper data is identical as well, as proven in (28). Of course, this assumes a generator matrix in standard form, i.e., $\mathbf{G} = (\mathbf{I}_k \| \mathbf{A})$, given that Yu's method is restricted to this case.

$$\mathbf{p} = \mathbf{x} \cdot \mathbf{H}^T = \mathbf{x} \cdot \begin{pmatrix} \mathbf{A} \\ \mathbf{I}_{n-k} \end{pmatrix} = \mathbf{x}(1:k) \cdot \mathbf{A} \oplus \mathbf{x}(k+1:n). \quad (28)$$

A.4 Multi-Code Method of Ahlswede et al.

The method of Ahlswede et al. [1] is represented by Fig. 14(g). Although initially proposed for secret key transport with *correlated sources*, it maps quite easily to our framework of interest, as observed by Hiller et al. [19]. A distinguishing feature is the use of multiple codes \mathcal{C}_j , covering mutually disjoint sets of codewords. We restrict our attention to $[n, k, d]$ block codes with $j \in [0, 2^{n-k} - 1]$. Every $\mathbf{x} \in \mathcal{X}$ then coincides with exactly one codeword, guaranteeing correctness. Furthermore, $\Delta\mathbb{H}_\infty \leq (n - k)$ holds due to (26), given that helper data $\mathbf{p} = j$ is limited to $(n - k)$ bits.

In [19], Hiller et al. proposed an efficient implementation where all codes are derived from a single parent code \mathcal{C}_0 . In particular, \mathcal{C}_0 is a linear code in standard form, i.e., $\mathbf{G} = (\mathbf{I}_k \parallel \mathbf{A})$, and all other codes are cosets: $\mathcal{C}_j = \{\mathbf{w} \oplus (\mathbf{0} \parallel \mathbf{p}) \mid \mathbf{w} \in \mathcal{C}_0\}$. This turns out to be fully equivalent with the method of Kang et al. in Fig. 14(f), i.e., helper data \mathbf{p} and reconstructed output \mathbf{y} are identical. We consider a slightly more general case. In particular, a linear code \mathcal{C}_0 that is not necessarily in standard form, as required by the method of Bennett et al. as well. All child codes \mathcal{C}_j are again formed as the cosets of \mathcal{C}_0 . Therefore, helper data $\mathbf{p} = j$ still reveals in which coset \mathbf{x} resides and (5) holds once again. The one-to-one correspondence of output \mathbf{y} in (29) finalizes our proof.

$$\forall(\mathbf{p}, \mathbf{x}) \in (\mathcal{P} \times \mathcal{X}), \mathbb{P}((X = \mathbf{x}) | (P = \mathbf{p})) = \mathbb{P}((M = \text{Decode}_{\mathcal{C}_p}(\mathbf{x})) | (P = \mathbf{p})). \quad (29)$$