Information Leakage of Continuous-Source Zero Secrecy Leakage Helper Data Schemes

Joep de Groot, Boris Škorić, Niels de Vreede, Jean-Paul Linnartz

Abstract

A Helper Data Scheme (HDS) is a cryptographic primitive that extracts a high-entropy noise-free string from noisy data. Helper Data Schemes are used for preserving privacy in biometric databases and for Physical Unclonable Functions. HDSs are known for the guided quantization of continuous-valued biometrics as well as for repairing errors in discrete-valued (digitized) extracted values. We refine the theory of Helper Data Schemes with the Zero Leakage (ZL) property, i.e., the mutual information between the helper data and the extracted secret is zero. We focus on quantization and prove that ZL necessitates particular properties of the helper data generating function: (i) the existence of "sibling points", enrollment values that lead to the same helper data but different secrets; (ii) quantile helper data.

We present an optimal reconstruction algorithm for our ZL scheme, that not only minimizes the reconstruction error rate but also yields a very efficient implementation of the verification. We compare the error rate to schemes that do not have the ZL property.

Keywords: Biometrics, fuzzy extractor, helper data, privacy, secrecy leakage, secure sketch.

1 Introduction

1.1 Biometric Authentication

Biometrics have become a popular solution for authentication or identification, mainly because of its convenience. Biometric features cannot be forgotten (like a password) or lost (like a token). Nowadays identity documents such as passports nearly always include biometric features extracted from fingerprints, faces or irises. Governments are storing biometric data for fighting crime and some laptops and smart phones use biometrics-based user authentication.

In general biometrics are not strictly secret. Fingerprints can be found on many objects. It is hard to prevent one's face or irises from being photographed. Nonetheless, one does not want to store biometric features in an unprotected database since this will make it easier for an adversary to misuse them.

Storage of the features introduces both security and privacy risks for the user. Security risks include the production of fake biometrics from the features, e.g., rubber fingers [14, 18]. These fake biometrics can be used to obtain unauthorized access to information or services or to leave fake evidence at crime scenes.

We mention two privacy risks. (i) Some biometrics are known to reveal diseases and disorders of the user. (ii) Unprotected storage allows for cross-matching between databases.

The security and privacy problems cannot be solved by simply encrypting the database. An important part of the standard attack model is *insider attacks*, i.e., attacks by people who are authorized to access the database. They possess the decryption keys; hence database encryption does not stop them.

All in all, the situation is very similar to the problem of password storage. The standard solution is to store *hashed* passwords. Cryptographic hash functions are one-way functions, i.e., inverting them is computationally infeasible. An attacker who gets hold of a hashed password cannot deduce the password from it.

Straightforward application of this hashing method to biometrics does not work, however. Biometric measurements are noisy, which causes (small) differences between the enrollment measurement and the later measurement during verification. Particularly if the biometric value lies near a quantization boundary, a minor noise contribution can flip the discretized value and trigger an avalanche of bit flips at the output of the hash. Despite significant progress in the signal processing of biometrics, measurements still suffer from relatively large deviations between enrollment and verification. Hence it remains essential to optimize the algorithms for extraction of reliable, repeatable yet high-entropy verification strings.

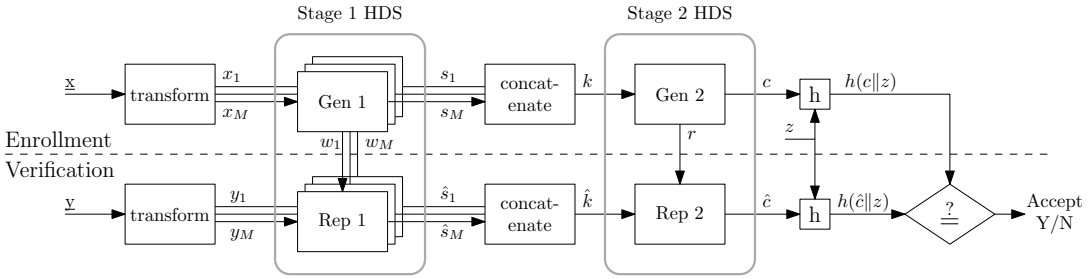

Figure 1: Common steps in a privacy-preserving biometric verification scheme.

1.2 Two Stage Approach

Fig. 1 describes a commonly adopted two-stage approach, as for instance presented in [17, Chap. 16].

A Helper Data Scheme (HDS) consists of two algorithms, `Gen` and `Rep`. The `Gen` takes a noisy value as input and generates a secret and so-called helper data. The `Rep` algorithm has two inputs: the helper data and a noisy value correlated with the first one. It outputs an estimator for the secret made by `Gen`.

The stage 1 HDS in Fig. 1 processes the signals during an optimized bit extraction [3, 6, 12]. Helper data is applied in the 'analog' domain, i.e., before 'bit-mapping': per dimension, the orthogonalized biometric x is biased towards the center of a quantization interval, e.g., by adding a helper data value w.

The stage 2 HDS employs digital error correction methods with discrete helper data, for instance the helper data is the XOR of secret k with a random codeword from an error-correcting code [9].

Such a two-stage approach is also common practice in communication systems that suffer from unreliable (wireless) channels: the signal conditioning prior to the bit mapping involves optimization of signal constellations and multidimensional transforms. In fact, the discrete mathematical operations, such as error correction decoding, are known to be particularly effective for sufficiently error-free signals. According to the asymptotic Elias bound [13, Chap. 17], at bit error probabilities above 10% one cannot achieve code rates better than 0.5. Similarly, in biometric authentication, optimization of the first stage appears essential to achieve adequate system performance.

In a preparation phase preceding all enrollments, the population's biometrics are studied and a transform is derived (using well known techniques such as PCA/LDA [21]) that splits the biometric vector $\underline{x}$ into scalar components $(x_i)_{i=1}^M$. We will refer to these components x_i as features. The transform ensures that they are mutually independent, or nearly so.

At enrollment, a person's biometric $\underline{x}$ is obtained. The transform is applied, yielding features $(x_i)_{i=1}^M$. The `Gen` algorithm of the first-stage HDS is applied to each feature independently. This gives continuous helper data $(w_i)_{i=1}^M$ and short secret strings $s_1, \ldots, s_M$ which may or may not have equal length, depending on the signal-to-noise ratio of the features and subsequent choice of allocated number of bits. All these secrets are combined into one high-entropy secret k, e.g., by concatenating them after Gray-coding. Biometric features have a within-class Probability Density Function (PDF) that, with multiple dimensions, will lead to some errors in the reproduced secret $\hat{k}$; hence a second stage of error correction is done with another HDS. The output of the second-stage `Gen` algorithm is discrete helper data r and a practically noiseless string c. The hash $h(c\|z)$ is stored in the enrollment database, along with the helper data $(w_i)_{i=1}^M$ and r. Here z is salt: a random string to prevent easy cross-matching.

In the authentication phase, a fresh biometric measurement $\underline{y}$ is obtained and split into components $(y_i)_{i=1}^M$. For each i independently, the estimator $\hat{s}_i$ is computed from y_i and w_i. The $\hat{s}_i$ are combined into an estimator $\hat{k}$, which is then input into the 2nd-stage HDS reconstruction together with r. The result is an estimator $\hat{c}$. Finally $h(\hat{c}\|z)$ is compared with the stored hash $h(c\|z)$.

1.3 Secret Extraction

Special algorithms have been developed for HDSs [3,6,7,12]: Fuzzy Extractors (FE) and Secure Sketches (SS). The FE and SS are special cases of the general concept and in our case can apply to both the stage 1 and 2 HDS. The algorithms have different requirements,

- *Fuzzy Extractor*

  The probability distribution of s given w has to be (nearly) uniform.

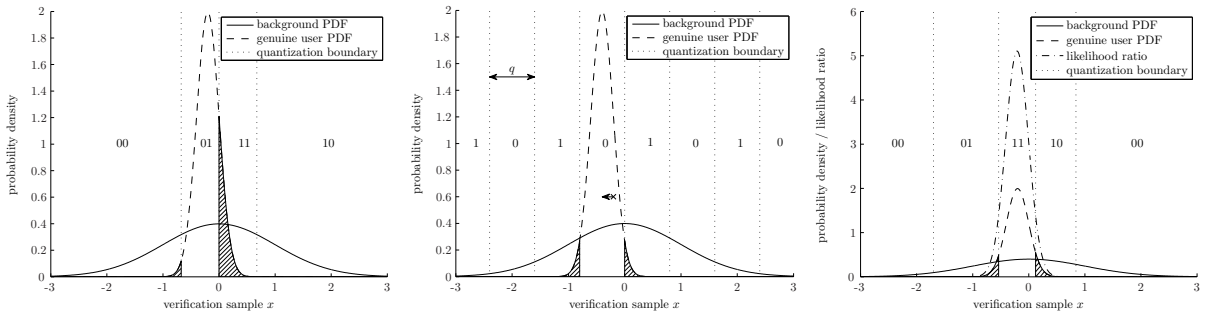

(a) Fixed equiprobable quantization *(b)* Quantization Index Modulation *(c)* Multi-bits based on likelihood ratio [3]

Figure 2: Examples (of adaptation to) the genuine user PDF during verification. Fixed equiprobable quantization does not translate the PDF, quantization index modulation centers the PDF on a quantization interval and the Multi-bits likelihood ratio based scheme of Chen et al. [3] uses the likelihood ratio to adjust the quantization regions.

- *Secure Sketch*

  s given w must have high entropy, but does not have to be uniform. Typically s is equal to (a discretized version of) x.

The FE is typically used for the extraction of cryptographic keys from noisy sources such as Physical Unclonable Functions (PUFs) [8, 15, 20]. Some fixed quantization schemes support the use of a fuzzy extractor, provided that the quantization intervals can be chosen such that each secret s is equiprobable, as in [16].

The SS is very well suited to the biometrics scenario described above.

1.4 Security and Privacy Preservation

In the HDS context, the main privacy question is how much information, and *which* information, about the biometric x is leaked by the helper data. Ideally, the helper data would contain just enough information to enable the error correction. Roughly speaking this means that the vector $\underline{\mathrm{w}} = (w_i)_{i=1}^{M}$ consists of the noisy 'least significant bits' of x, which typically do not reveal sensitive information since they are noisy anyway. In order to make this kind of intuitive statement more precise, one studies the information-theoretic properties of HDSs. In the system as sketched in Fig. 1 the mutual information $I(C; \underline{\mathrm{W}}, R)$ is of particular interest: it measures the leakage about the string c caused by the fact that the attacker observes $\underline{\mathrm{w}}$ and r. By properly separating the 'most significant digits' of x from the 'least significant digits', it is possible to achieve $I(C; \underline{\mathrm{W}}, R) = 0$. We call this Zero Secrecy Leakage or, more compactly, Zero Leakage (ZL). HDSs with the ZL property are very interesting for quantifying privacy guarantees: if a privacy-sensitive piece of a biometric is fully contained in c, and not in $(\underline{\mathrm{w}}, r)$, then a ZL HDS based database reveals *absolutely nothing* about that piece.

If in a concatenation of multiple stages of the HDS, each stage has ZL, this is a sufficient condition that also $I(C; W, R) = 0$. Therefore we will focus in particular on schemes whose first stage has the ZL property $I(S_i, W_i) = 0$. Yet, $I(S_i, W_i) = 0$ is not a necessary condition for $I(C; W, R) = 0$. In fact, a counter example is that one can simply eliminate all dimensions for which $I(S_i, W_i) > 0$ in the second stage of the HDS to achieve $I(C; W, R) = 0$. However, this would reduce the entropy in the extracted secret C and is therefore not attractive.

1.5 Biometric Quantization

Biometric quantization, or at least a translation from finely quantized features to a coarse bit-mapping, takes place in the first-stage HDS of a common biometric verification scheme as depicted in Fig. 1. Some concepts, which we will briefly discuss in this section, have already been developed for this stage.

Many biometric parameters can be converted by a Principal Component Analysis (PCA) into a vector of independent dimensions [11]. For this reason, most papers on helper data in the analog domain can restrict themselves to a one-dimensional quantization, e.g., [3, 11, 12]. Yet, the quantization strategy differs, as we will review below.

1.5.1 Fixed Quantization (FQ)

The simplest form of quantization applies a uniform, fixed quantization grid during both enrollment and verification. An example for $N = 4$ quantization regions during verification is depicted in Fig. 2a. As can be seen in the figure an unfavorably located genuine user PDF can cause a high reconstruction error (the hatched areas). The inherently large error probability can be mitigated by 'reliable component' selection [16]. Only dimensions that are likely to repeatedly deliver realizations within the same quantization interval are selected during the enrollment. Hence the verification relies on features which for that biometric prover have most of their PDF mass within the same interval. The selection of dimensions is also conveyed to the verification phase as user-dependent helper data.

In such a scheme the verification capacity is not optimally used: features that are unfavorably located w.r.t. the quantization grid, but nonetheless carry information, are eliminated. Revealing which dimensions are not near quantization boundaries may leak information about the actual result of the quantization. The outer intervals, as depicted in Fig. 2a, are wider and therefore are more likely to produce 'reliable components' than the inner intervals, which can cause leakage when no precautions are taken [10].

Implicitly, the fixed quantization case also covers systems that only implement a stage 2 HDS. Potentially such systems can be further improved by adaptive quantization, as illustrated below.

1.5.2 Quantization Index Modulation (QIM)

This method borrows principles from digital watermarking [2] and writing on dirty paper [4]. QIM uses alternating quantization intervals labeled with '0' and '1' as the values for the secret s. Based on the enrollment sample, an offset w is calculated which is added to bias the verification sample towards the center of a quantization interval as depicted in Fig. 2b. If the bias is always exactly to the center of a quantization interval of size q, the probability of an error for a Gaussian sample pair is

$$p_e^{\mathrm{QIM}} \approx \operatorname{erfc}\left(\frac{q}{2\sqrt{2}\sigma}\right) \tag{1}$$

In fact, for a fixed quantization, i.e., without the helper data bias towards the center of quantization intervals, if the biometrics were locally approximately uniform, i.e., if the biometric lies at a random position in the quantization interval, the error probability is much larger, namely [1, Ch. 7]

$$p_e^{\mathrm{FQ}} = \frac{1}{q}\int_{-q/2}^{q/2} \operatorname{erfc}\left(\frac{q/2 - x}{2\sqrt{2}\sigma}\right) \mathrm{d}x \;\rightarrow\; \sqrt{\frac{2}{\pi}}\frac{\sigma}{q} \tag{2}$$

for q sufficiently larger than the within-class standard deviation σ. Numerical evaluation of (1) and (2) shows that to ensure an error rate less than 10%, biased quantization, such as QIM can tolerate $q \approx 3.3\sigma$, while fixed quantization would require $q > 8\sigma$. This implies that in practical systems, the omission of stage 1 HDS, but relying on stage 2 HDS, jeopardizes extraction of more than one bit per dimension. Moreover, it illustrates the substantial enhancement that helper-data guided quantization can bring in a stage 1 HDS and has been a motivation for this work.

QIM allows a choice in quantization step sizes to make a trade-off between verification performance and leakage [12]. The alternating bit-mapping was adopted to reduce leakage but sacrifices entropy extraction. However, our results show quantization schemes that outperform this.

1.5.3 Likelihood based Quantization (LQ)

Chen et al. [3] improves the reconstruction rate by exploiting the likelihood ratio to set the quantization scheme. Based on some enrollment samples, the genuine user's PDF is estimated, which is used to determine a likelihood ratio. It is assumed that the background PDF is known and equal during enrollment and verification. The scheme allocates quantization regions for $|\mathcal{S}| = N$ different secret values as depicted in Fig. 2c. The first two quantization boundaries for the verification phase are then chosen such that they are at an equal likelihood ratio and at the same time enclose a probability mass of $1/N$ on the background distribution. Subsequent quantization intervals are chosen contiguous to the first and enclose again a $1/N$ probability mass. Finally the probability mass in the tails of the distribution is added up as a wrap-around interval, which also holds a probability mass of $1/N$. Since the quantization boundaries are at fixed probability mass intervals it suffices to convey a single threshold t to the verification phase.

Likelihood based quantization does not provide an equiprobable secret s and therefore does not qualify as a fuzzy extractor. Moreover, $I(T;S) \neq 0$ which implies there is some secrecy leakage from the public verification threshold t to the enrolled secret s.

The work in this paper does not a priori choose a quantization for enrollment and verification, but imposes the requirement of zero secrecy leakage and a low reconstruction error rate to come to an optimal quantization scheme.

1.6 Contributions and Organization of This Paper

In this paper we zoom in on the first-stage HDS and focus on the ZL property in particular. Our aim is to minimize reconstruction errors in ZL HDSs that have scalar input $x \in \mathbb{R}$. We treat the helper data as being real-valued, $w \in \mathbb{R}$, though of course w is in practice stored as a finite-precision value.

- We show that the ZL constraint for continuous helper data necessitates the existence of 'Sibling Points', points x that correspond to different s but give rise to the same helper data w.
- We prove that the ZL constraint for $x \in \mathbb{R}$ implies 'quantile' helper data. This holds for uniformly distributed s as well as for non-uniform s. Thus, we identify a simple quantile construction as being the generic ZL scheme for all HDS types, including the FE and SS as special cases. It turns out that the continuum limit of a FE scheme of Verbitskiy at al. [19] precisely corresponds to our quantile HDS.
- We derive a reconstruction algorithm for the quantile ZL FE that minimizes the reconstruction errors. It amounts to using a set of optimized threshold values, and is very suitable for low-footprint implementation.
- We analyze, in an all-Gaussian example, the performance (in terms of reconstruction error rate) of our ZL FE combined with the optimal reconstruction algorithm. We compare this scheme to fixed quantization and a likelihood-based classifier. It turns out that our error rate is better than that of fixed quantization, and not much worse than that of the likelihood-based classifier.

The organization of this paper is as follows. The sibling points and the quantile helper data are treated in Section 3. Section 4 discusses the optimal reconstruction thresholds. The performance analysis in the Gaussian model is presented in Section 5.

2 Preliminaries

2.1 Quantization

Random variables are denoted with capital letters and their realizations in lowercase. Sets are written in calligraphic font. We zoom in on the one-dimensional first-stage HDS in Fig. 1. For brevity of notation the index $i \in \{1, \ldots, M\}$ on x_i, w_i, s_i, y_i and $\hat{s}_i$ will be omitted.

The Probability Density Function (PDF) or Probability Mass Function (PMF) of a random variable A is denoted as f_A, and the Cumulative Distribution Function (CDF) as F_A. The helper data is considered continuous, $W \in \mathcal{W} \subset \mathbb{R}$. The secret S is an integer in the range $\mathcal{S} = \{0, \ldots, N-1\}$, where N depends on the signal-to-noise ratio of the biometric feature. The helper data is computed from X using a function g, i.e., $W = g(X)$. Similarly we define a quantization function Q such that $S = Q(X)$. The enrollment part of the HDS is given by the pair Q, g. We define quantization regions as follows,

$$A_s = \{x \in \mathbb{R} : Q(x) = s\}. \tag{3}$$

The quantization regions are non-overlapping and cover the complete feature space, hence form a partitioning:

$$A_s \cap A_t = \emptyset \quad \text{for } s \neq t\,; \qquad \bigcup_{s \in \mathcal{S}} A_s = \mathbb{R}. \tag{4}$$

We consider only quantization regions that are contiguous, i.e., for all s it holds that A_s is a simple interval. In Section 4.2 we will see that many other choices may work equally well, *but not better*; our preference for contiguous A_s regions is tantamount to choosing the simplest element Q out of a whole equivalence class of quantization functions that lead to the same HDS performance. We define

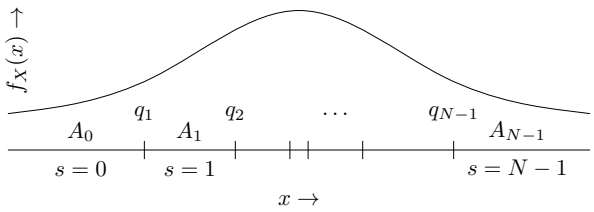

Figure 3: Quantization regions A_s and boundaries q_s. The locations of the quantization boundaries are based on the distribution of x, such that secret s occurs with probability p_s.

quantization boundaries $q_s = \inf A_s$. Without loss of generality we choose Q to be a monotonically increasing function. This gives $\sup A_s = q_{s+1}$. An overview of the quantization regions and boundaries is depicted in Fig 3.

In a generic HDS the probabilities $\mathbb{P}[S = s]$ can be different for each s. We will use shorthand notation

$$\mathbb{P}[S = s] = p_s > 0. \tag{5}$$

The quantization boundaries are given by

$$q_s = F_X^{-1}\left(\sum_{t=0}^{s-1} p_t\right), \tag{6}$$

where F_X^{-1} is the inverse CDF. For a Fuzzy Extractor one requires $p_s = 1/N$ for all s, in which case (6) simplifies to

$$q_s^{\mathrm{FE}} = F_X^{-1}\left(\frac{s}{N}\right). \tag{7}$$

2.2 Noise Model

We assume the verification sample y to be related to the enrollment sample x as follows

$$Y = \lambda X + R, \tag{8}$$

in which $\lambda \in [0, 1]$ is the attenuation parameter and R independent additive noise. Hence

$$\sigma_Y^2 = \lambda^2 \sigma_X^2 + \sigma_R^2. \tag{9}$$

In this equation σ_X^2, σ_Y^2 and σ_R^2 denote the variance of the enrollment sample, verification sample and noise respectively. One often uses the correlation $\rho \in [-1, 1]$, defined as

$$\rho = \frac{\mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]}{\sigma_X \sigma_Y}. \tag{10}$$

By using this equation we can derive an expression for λ

$$\lambda^2 = \frac{\rho^2}{1 - \rho^2} \frac{\sigma_R^2}{\sigma_X^2}. \tag{11}$$

and for the signal-to-noise ratio (SNR)

$$\mathrm{SNR} = \frac{\lambda^2 \sigma_X^2}{\sigma_R^2} = \frac{\rho^2}{1 - \rho^2} \tag{12}$$

In this model we can identify two limiting situations

1. *Perfect enrollment*
   The enrollment can be conditioned such that noise has negligible influence. However, the verification will be subject to noise. In this situation we get

$$\sigma_Y^2 = \sigma_X^2 + \sigma_R^2 \qquad \Longleftrightarrow \qquad \lambda^2 = 1. \tag{13}$$

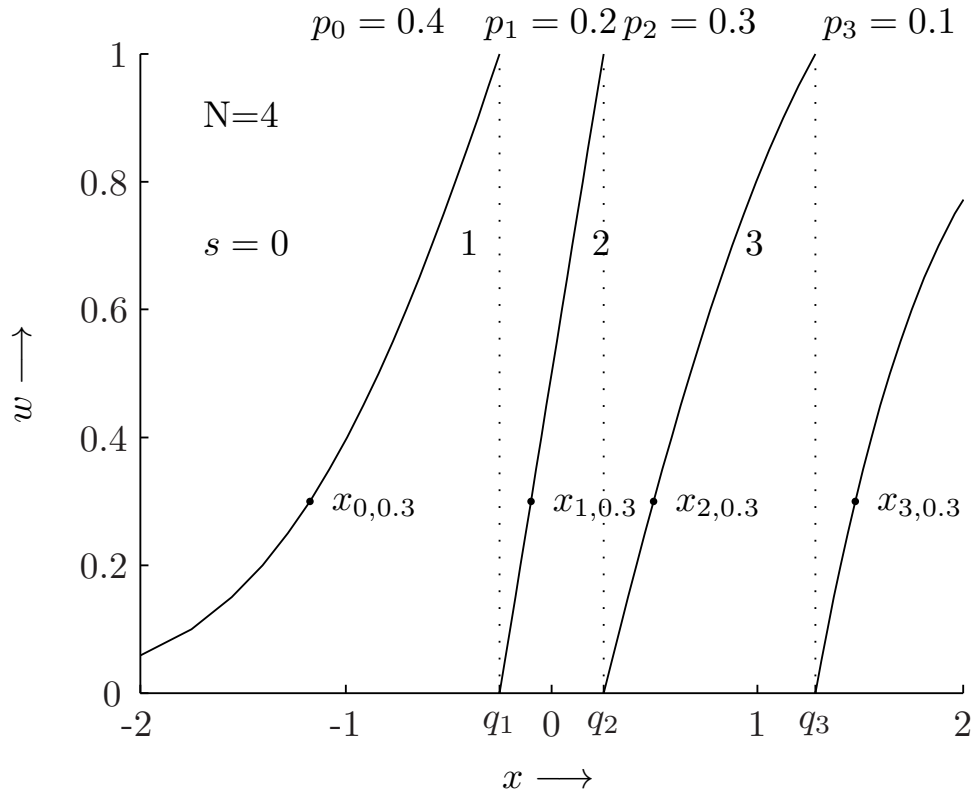

Figure 4: Example of helper data generating function g for a standard Gaussian distribution, i.e. $x \sim \mathcal{N}(0,1)$, and $N = 4$. Sibling points x_{sw} are given for $s \in \{0,\ldots,3\}$ and $w = 0.3$.

2. *Identical conditions*
   The enrollment and verification recordings are identical in terms of noise. In this situation we get

$$\sigma_Y^2 = \sigma_X^2 \qquad \Longleftrightarrow \qquad \begin{matrix} \lambda^2 = \rho^2 \\ \sigma_R^2 = (1-\rho^2)\sigma_X^2 \end{matrix} . \tag{14}$$

Definition 2.1*:* The noise is called *symmetric fading noise* if for all x, y_1, y_2 it holds that

$$|y_1 - \lambda x| > |y_2 - \lambda x| \implies f_{Y|X}(y_1|x) < f_{Y|X}(y_2|x). \tag{15}$$

An example of a distribution that satisfies the requirements, i.e., symmetric and fading, and has correlation ρ is a jointly Gaussian biometric, given by

$$X \sim \mathcal{N}(0,1), \qquad (Y|X=x) \sim \mathcal{N}(\rho x, 1-\rho^2). \tag{16}$$

This distribution corresponds to the "Identical conditions" situation and yields the quantization pattern as depicted in Fig. 6c.

3 Zero Leakage Helper Data

First we will derive properties for the helper data generating function g that are required for reproduction. These properties will be taken into account when deriving the requirements for zero leakage in the second part of this analysis.

3.1 Requirements for Reproduction

During a verification the verifier will reconstruct x_s samples based on the public helper data w and compare the measured verification sample y to them. Based on our assumptions of the noise, as explained in Section 2.2, the originator x_s scaled by λ, which is the closest to y will be selected to produce the estimate of the secret $\hat{s}$. Hence, the minimum distance between these x_s originators has to be maximized, in order to minimize reproduction errors of the secret.

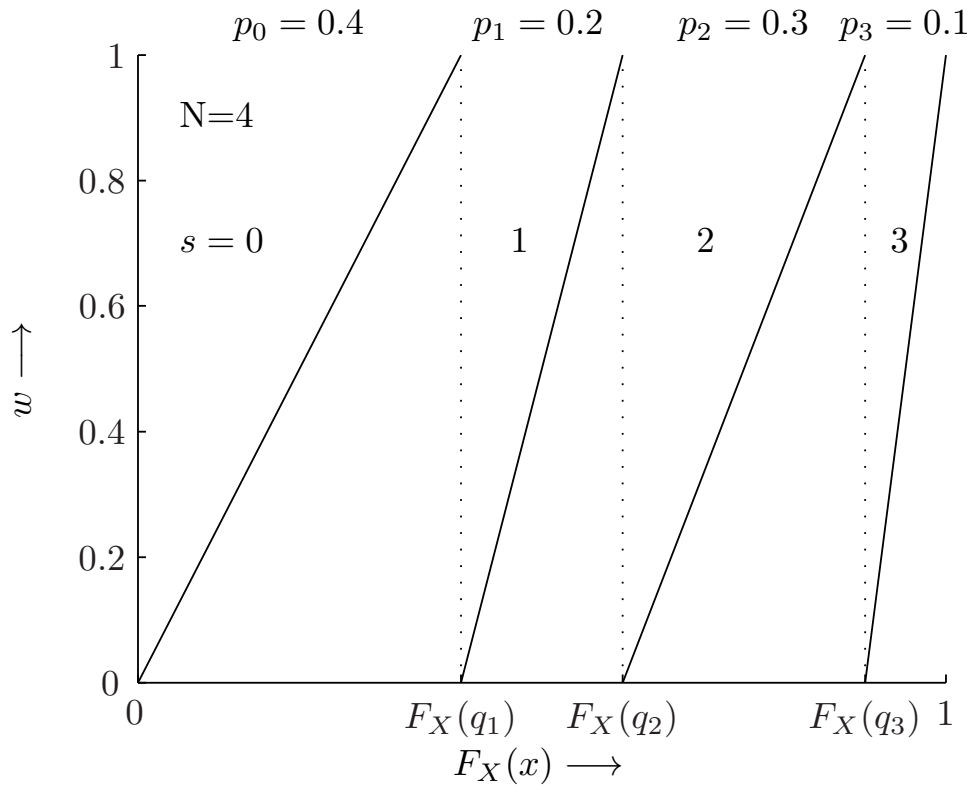

Figure 5: Example of helper data generating function g for $N = 4$ on quantile x, i.e. $F_X(x)$. Sibling points $u_{sw} = (s + w)/N$ are given for $s \in \{0, \ldots, 3\}$ and $w = 0.3$.

Lemma 3.1*:* Let $\mathcal{W} \subset \mathbb{R}$ be the set of possible helper data values. Let $w \in \mathcal{W}$. Then, to prevent excessive leakage and maximize min. distance between x originators, i.e. x values that yield $w = g(x)$, there is exactly one $x \in A_s$. In other words g is invertible on each A_s

Proof. Given a $w \in \mathcal{W}$ and the fact that $w = g(x)$, there has to be at least one $x \in \mathbb{R}$ that has generated this w.

Let $w \in \mathcal{W}$ be known. Suppose there is a region A_s that has no $x \in A_s$ that yields $w = g(x)$, then

$$\mathbb{P}(X \in A_s | W = w) = 0. \tag{17}$$

However, since the regions span the entire feature space (Eq. (4)), there is at least one other $x \in A_t, t \neq s$, therefore

$$\mathbb{P}(X \in A_t | W = w) > 0. \tag{18}$$

This implies that if regions without an originating x exist, leakage cannot be low, so we require at least one possible x per region A_s.

More than one point decreases min. distance: regions A_s contiguous and cover entire x space.

$$\mathcal{X}_{sw} = \{x \in A_s \,|\, g(x) = w\} \tag{19}$$

Show that this set can only consist of one point to have maximal min. distance. □

Definition 3.2 (sibling points)*:* Let $s, t \in \mathcal{S}$, $s \neq t$ be secrets. Given a point $x_s \in A_s$ we define the sibling point $x_t \in A_t$ as the point in A_t that has the same helper data, i.e. $g(x_s) = g(x_t)$.

By this definition we obtain a set of N sibling points for each w.

Lemma 3.3 (ordering of sibling points)*:* Let $x_1, x_2 \in A_s$ and $x_3, x_4 \in A_t$, $s \neq t$, with $g(x_1) = g(x_3)$, $g(x_2) = g(x_4)$ and $x_2 > x_1$. Then $x_4 > x_3$ leads to a higher min. distance than $x_4 \leq x_3$

Proof. Let us fix x_1, x_2, x_3 and vary $x_4 \in A_t$. We consider two cases:

- Case 1 (lower): $x_4 \leq x_3$, for which we define $x_4^L = (x \in A_t \,|\, x \leq x_3)$.

- Case 2 (higher): $x_4 > x_3$, for which we define $x_4^H = (x \in A_t \,|\, x > x_3)$.

Since the following inequalities hold

$$|x_2 - x_4^L| < |x_2 - x_4^H|, \tag{20}$$
$$|x_2 - x_4^L| < |x_1 - x_3|, \tag{21}$$

it is possible to show that

$$\max_{x_4 \in \{x_4^L, x_4^H\}} \min(|x_1 - x_3|, |x_2 - x_4|) \tag{22}$$
$$= \max(|x_2 - x_4^L|, \min(|x_1 - x_3|, |x_2 - x_4^H|)) \tag{23}$$
$$= \min(|x_1 - x_3|, |x_2 - x_4^H|). \tag{24}$$

In the second step x_4^L gets eliminated, since it never leads to a maximum min. distance. □

In fact there is an entire class of functions g that satisfy Lemma 3.1 and Lemma 3.3. For example it is possible to shift some part of the function (without causing an overlap) or apply an permutation to the outcome, e.g. encrypt the most significant bits. These alternatives will perform equally well in terms of maximal min. distance, but not any better, therefore we limit ourselves to the simples function possible. This brings us to the following conjecture:

Conjecture 3.4*:* Without loss of generality, we can choose g to be a simple function, i.e. a differentiable function on each A_s for all $s \in \mathcal{S}$.

From this point on we will assume g to be a piecewise differential function.

Theorem 3.5 (sign g' equal on each A_s)*:* Let $x_s \in A_s, x_t \in A_t$ be sibling points as defined in Def. 3.2. Let g be differentiable on each A_s. Then $\text{sign}(g'(x_s)) = \text{sign}(g'(x_t))$ leads to a higher min. distance than $\text{sign}(g'(x_s)) \neq \text{sign}(g'(x_t))$.

Proof. Based on Conjecture 3.4 we assume the derivative of g exists and can be written as

$$g'(x_s) = \lim_{\Delta \to 0} \frac{g(x_s + \Delta) - g(x_s)}{\Delta}. \tag{25}$$

Since we are only interested in the sign we can omit the limit operation. Equal sign yields

$$\text{sign}\left[g(x_s + \Delta) - g(x_s)\right] = \text{sign}\left[g(x_t + \Delta) - g(x_t)\right], \tag{26}$$

whereas unequal sign implies opposite sign, hence

$$\text{sign}\left[g(x_s + \Delta) - g(x_s)\right] = \text{sign}\left[g(x_t) - g(x_t + \Delta)\right]. \tag{27}$$

Based on Lemma 3.1 there is only one x in each interval that yields w, therefore we can add $-\Delta$ to both arguments of g on the right hand side without changing the sign, hence

$$\text{sign}\left[g(x_s + \Delta) - g(x_s)\right] = \text{sign}\left[g(x_t - \Delta) - g(x_t)\right], \tag{28}$$

which brings us to a similar situation as for Lemma 3.3 with $x_4^H = x_t + \Delta$ and $x_4^L = x_t - \Delta$. □

Based on Theorem 3.5 we restrict ourselves to piecewise monotonic g. Without loss of generality we choose g to be increasing on each interval A_s.

3.2 Requirements for Zero Leakage

The Zero Leakage requirement is formulated as

$$I(S; W) = 0 \qquad \text{or equivalently} \qquad H(S|W) = H(S), \tag{29}$$

where H stands for Shannon entropy, and I for mutual information. (See, e.g. [5, Eq. (2.35)-(2.39)].) The mutual information between the secret S and the publicly available helper data W must be zero. In other words, knowing W does not reduce the unpredictability of S.

Theorem 3.6 (ZL equivalent to quantile relationship between sibling points)*:* Let $\mathcal{W} \subset \mathbb{R}$ be the set of possible helper data values. Let g be monotonously increasing on each interval A_s, with $g(A_0) = \cdots = g(A_{N-1}) = \mathcal{W}$. Let $s, t \in \mathcal{S}$. Let $x_s \in A_s, x_t \in A_t$ be sibling points as defined in Def. 3.2. In order to satisfy Zero Leakage we have the following necessary and sufficient condition on the sibling points,

$$\frac{F_X(x_s) - F_X(q_s)}{p_s} = \frac{F_X(x_t) - F_X(q_t)}{p_t}. \tag{30}$$

Proof. The ZL requirement $I(S;W) = 0$ is equivalent to saying that S and W are independent. This is equivalent to $f_W = f_{W|S}$, which gives

$$f_W(w) = f_{W|S}(w|s) = \frac{f_{W,S}(w,s)}{p_s} \qquad \forall s \in \mathcal{S}, \tag{31}$$

where $f_{W,S}$ is the joint distribution for W and S. We work under the assumption that $w = g(x)$ is an monotonous function on each interval A_s, fully spanning $\mathcal{W}$. Then for given s and w there exists exactly one point x_{sw} that satisfies $Q(x) = s$ and $g(x) = w$. Furthermore, conservation of probability then gives

$$f_{W,S}(w,s)\,\mathrm{d}w = f_X(x_{sw})\,\mathrm{d}x_{sw}. \tag{32}$$

Since the right hand side of (31) is independent of s, we can write $f_W(w)\mathrm{d}w = p_s^{-1} f_X(x_{sw})\mathrm{d}x_{sw}$ for *any* $s \in \mathcal{S}$. Hence for any $s, t \in \mathcal{S}$, $w \in \mathcal{W}$ it holds that

$$\frac{f_X(x_{sw})\mathrm{d}x_{sw}}{p_s} = \frac{f_X(x_{tw})\mathrm{d}x_{tw}}{p_t}, \tag{33}$$

which can be rewritten as

$$\frac{\mathrm{d}F_X(x_{sw})}{p_s} = \frac{\mathrm{d}F_X(x_{tw})}{p_t}. \tag{34}$$

The result (30) follows by integration, using the fact that A_s has lower boundary q_s. □

Corollary 3.7 (ZL FE sibling point relation)*:* Let $\mathcal{W} \subset \mathbb{R}$ be the set of possible helper data values. Let g be monotonously increasing on each interval A_s, with $g(A_0) = \cdots = g(A_{N-1}) = \mathcal{W}$. Let $s, t \in \mathcal{S}$. Let $x_s \in A_s, x_t \in A_t$ be sibling points as defined in Def. 3.2. Then for a Fuzzy Extractor we have the following necessary and sufficient condition on the sibling points in order to satisfy Zero Leakage

$$F_X(x_s) - \frac{s}{N} = F_X(x_t) - \frac{t}{N}. \tag{35}$$

Proof. Immediately follows by combining Eq. (30) with the fact that $p_s = 1/N\ \forall s \in \mathcal{S}$ in a FE scheme, and the FE quantization boundaries given in Eq. (7). □

Theorem 3.6 allows us to define the enrollment steps in a ZL HDS in a very simple way,

$$\begin{aligned} s &= Q(x) \\ w &= \frac{F_X(x) - F_X(q_s)}{p_s}. \end{aligned} \tag{36}$$

Note that $w \in [0,1)$, and $F_X(q_s) = \sum_{t=0}^{s-1} p_t$. The helper data can be interpreted as a quantile distance between x and the quantization boundary q_s, normalized with respect to the probability mass p_s in the interval A_s. In the FE case, Eq. (36) simplifies to

$$F_X(x) = \frac{s+w}{N}, \qquad w \in [0,1). \tag{37}$$

Eq. (36) is the *simplest* way to implement an enrollment that satisfies the sibling point relation of Theorem 3.6. However, it is not the *only* way. For instance, by applying any invertible function to w, a new helper data scheme is obtained that also satisfies the sibling point relation (30) and hence is ZL. Another example is to store the whole set of sibling points $\{x_{tw}\}_{t \in \mathcal{S}}$; this contains exactly the same information as w. The transformed scheme can be seen as merely a different representation of the 'basic' ZL HDS (36). Such a representation may have various advantages over (36), e.g. allowing for a faster reconstruction procedure, while being completely equivalent in terms of the ZL property. We will see such a case in Section 4.2.

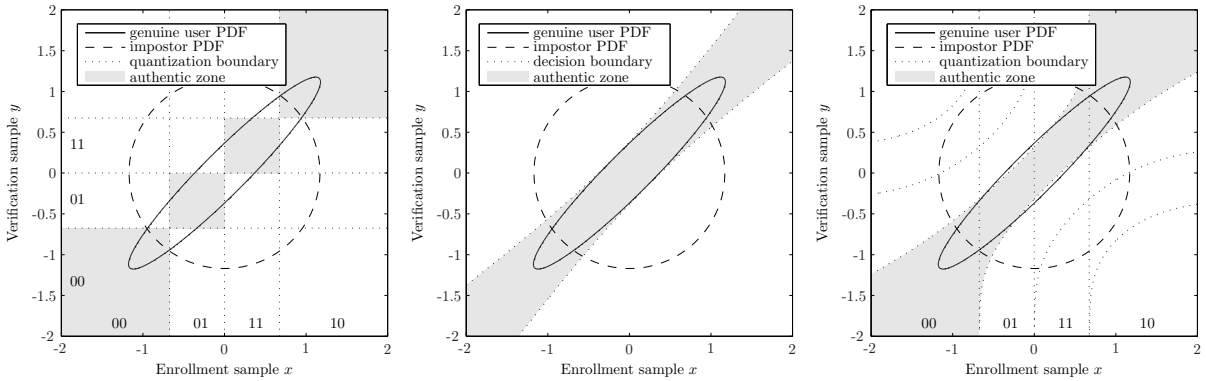

(a) Fixed equiprobable quantization (FQ)

(b) Likelihood classification (LC)

(c) Zero leakage quantization scheme with likelihood based reproduction (ZLQ-LR)

Figure 6: Quantization and decision patterns based on the genuine user and impostor PDFs. Ideally the genuine user PDF should be contained in the authentic zone and the impostor PDF should have a large mass outside the authentic zone. 50% probability mass is contained in the genuine user and impostor PDF ellipse and circle. The genuine user PDF is based on a 10 dB SNR.

4 Optimal Reconstruction

The goal of the HDS reconstruction algorithm $\mathrm{Rep}(y, w)$ is to reliably reproduce the secret s. The best way to achieve this is to choose the most probable $\hat{s}$ given y and w, i.e., a maximum likelihood algorithm.

Lemma 4.1*:* Let $\mathrm{Rep}(y, w)$ be the reproduction algorithm of a ZL FE system. Let g_s^{-1} be the inverse of the helper data generation function for a given secret s. Then optimal reconstruction is achieved by

$$\mathrm{Rep}(y, w) = \arg\max_{s \in \mathcal{S}} f_{Y|X}(y|g_s^{-1}(w)). \tag{38}$$

Proof. As noted above, optimal reconstruction can be done by selecting the most likely secret given y, w,

$$\mathrm{Rep}(y, w) = \arg\max_{s \in \mathcal{S}} f_{S|Y,W}(s|y, w) \tag{39}$$

$$= \arg\max_{s \in \mathcal{S}} \frac{f_{Y,S,W}(y, s, w)}{f_{Y,W}(y, w)}. \tag{40}$$

The denominator does not depend on s, and can hence be omitted. This gives

$$\mathrm{Rep}(y, w) = \arg\max_{s \in \mathcal{S}} f_{S,Y,W}(s, y, w) \tag{41}$$

$$= \arg\max_{s \in \mathcal{S}} f_{Y|S,W}(y|s, w) f_{W|S}(w|s) p_s. \tag{42}$$

We constructed the scheme to be ZL, and therefore $p_s = 1/N$ and $f_{W|S}(w|s) = f_W(w)$. We see that both p_s and $f_{W|S}(w|s)$ do not depend on s, which implies they can be omitted from Eq. (42), yielding

$$\mathrm{Rep}(y, w) = \arg\max_{s \in \mathcal{S}} f_{Y|S,W}(y|s, w). \tag{43}$$

Finally, knowing S and W is equivalent to knowing X. Hence $f_{Y|S,W}(y|s, w)$ can be replaced by $f_{Y|X}(y|x)$ with x satisfying $Q(x) = s$ and $g(x) = w$. The unique x value that satisfies these constraints is $g_s^{-1}(w)$. □

To simplify the verification phase we can identify thresholds τ_s that denote the lower boundary of a decision region. If $\tau_s \leq y < \tau_{s+1}$, we reconstruct $\hat{s} = s$. The $\tau_0 = -\infty$ and $\tau_N = \infty$ are fixed, which implies we have to find optimal values only for the $N - 1$ variables $\tau_1, \ldots, \tau_{N-1}$ as a function of w.

Theorem 4.2*:* Let $f_{Y|X}$ be a symmetric fading noise. Then optimal reconstruction in a FE scheme is obtained by the following choice of thresholds

$$\tau_s = \lambda \frac{g_s^{-1}(w) + g_{s-1}^{-1}(w)}{2}. \tag{44}$$

Proof. In case of symmetric fading noise we know that

$$f_{Y|X}(y|x) = \varphi(|y - \lambda x|), \tag{45}$$

with φ some monotonic decreasing function. Combining this notion with that of Eq. (38) to find a point $y = \tau_s$ that gives equal probability for s and $s-1$ yields

$$\varphi(|\tau_s - \lambda g_{s-1}^{-1}(w)|) = \varphi(|\tau_s - \lambda g_s^{-1}(w)|). \tag{46}$$

The left and right hand side of this equation can only be equal for equal arguments, and hence

$$\tau_s - \lambda g_{s-1}^{-1}(w) = \pm(\tau_s - \lambda g_s^{-1}(w)). \tag{47}$$

Since $g_s^{-1}(w) \neq g_{s-1}^{-1}(w)$ the only viable solution is Eq. (44). □

Instead of storing the ZL helper data w according to (36), one can also store the set of thresholds $\tau_1, \ldots, \tau_{N-1}$. This contains precisely the same information, and allows for quicker reconstruction of s: just a thresholding operation on y and the τ_s values, which can be implemented on computationally limited devices.

4.1 Special Case: 1-bit Secret

In the case of a one-bit secret s, i.e., $N = 2$, the above ZL FE scheme is reduced to storing a single threshold τ_1.

It is interesting and somewhat counterintuitive that this yields a threshold for verification that does not leak information about the secret. In case the average of X is zero, one might assume that a positive threshold value implies $s = 0$. However, both $s = 0$ and $s = 1$ allow positive as well as negative τ_1, dependent on the relative location of x in the quantization interval.

4.2 FE: Equivalent choices for the quantization

Let us reconsider the quantization function $Q(x)$ in the case of a Fuzzy Extractor. Let us fix N and take the $g(x)$ as specified in Eq. (37). Then it is possible to find an infinite number of different functions Q that will conserve the ZL property and lead to exactly the same error rate as the original scheme. This is seen as follows. For any $w \in [0, 1)$ there is an N-tuplet of sibling points. Without any impact on the reconstruction performance we can permute the s-values of these points; the error rate of the reconstruction procedure depends only on the x-values of the sibling points, not on the s-label they carry. It is allowed to do this permutation for every w independently, resulting in an infinite equivalence class of Q-functions. The choice we made in Section 2 yields the simplest function in an equivalence class.

5 Example: Gaussian features and BCH codes

To be able to benchmark the reproduction performance of our scheme we will give an example based on Gaussian-distributed variables. In this example we will assume all variables to be Gaussian distributed, although the scheme is capable of achieving optimal reproduction for any kind of random variable with a fading noise distribution. The results in this example are obtained by numerical integration of the genuine user PDF given the corresponding scheme's thresholds.

We will compare the reproduction performance of our Zero Leakage Quantization scheme with Likelihood based Reproduction (ZLQ-LR) to a scheme with 1) Fixed Quantization (FQ) and 2) Likelihood Classification (LC). The former is, to our knowledge, the only other scheme sharing the zero secrecy leakage property, since this scheme does not use any helper data. Instead fixed quantization intervals are defined for both enrollment and verification as explained in section 1.5.1. An example of such a quantization with $N = 4$ intervals is depicted in Fig. 6a. Likelihood classification is not an actual quantization scheme since it requires the enrollment sample to be stored in-the-clear. However, a likelihood based classifier provides an optimal trade-off between false accept and false reject according to communication theory [5] and should therefore yield the lowest error rate one can get. Instead of quantization boundaries the classifier is characterized by decision boundaries as depicted in Fig. 6b.

A comparison with QIM cannot be made since the probability for an impostor in a QIM scheme to correctly guess the enrolled secret cannot be made equal to $1/N$. This would result in an unfair

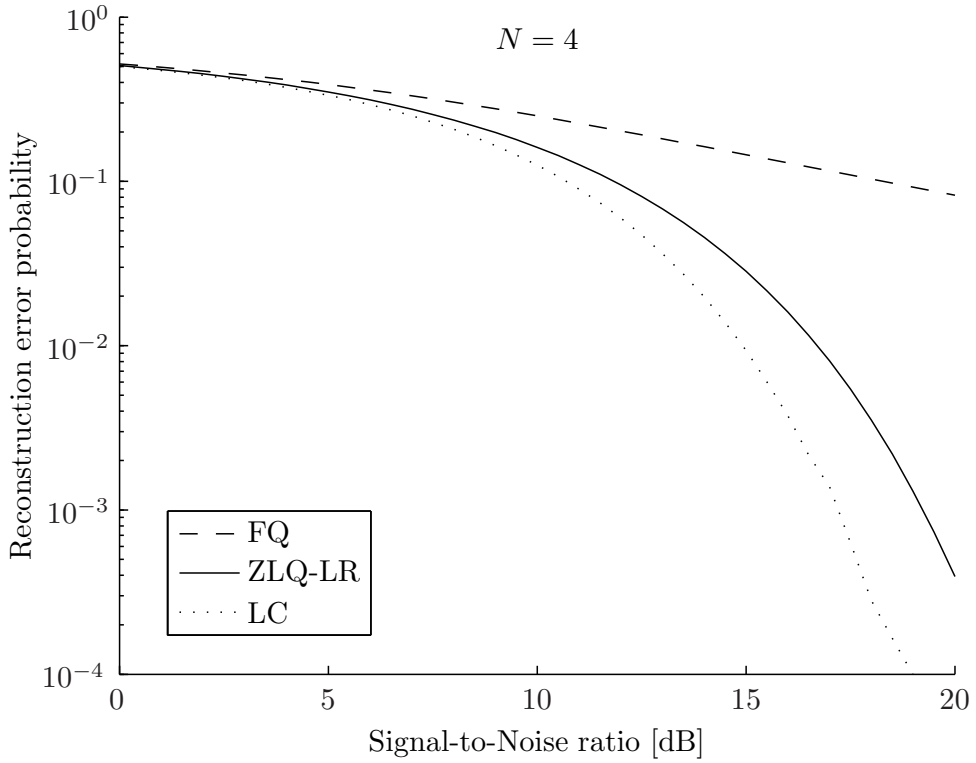

Figure 7: Reproduction performance in terms of error probability for Gaussian distributed features and Gaussian noise for $N = 4$.

comparison since the other schemes are designed to possess this property. Moreover, the QIM scheme allows the reproduction error probability to be made arbitrary small by increasing the quantization width at the cost of leakage.

Also the likelihood based classification can be tuned by setting the a decision threshold. However, for this scheme it is possible to choose a threshold such that an impostor will have a probability of $1/N$ to be accepted, which corresponds to the $1/N$ probability of correctly guessing the enrolled secret in a fuzzy extractor scheme. Note that for a likelihood classifier there is no enrolled secret since this is not a quantization scheme.

As can be seen from Fig. 7, the reproduction performance for a ZSL scheme with likelihood based reproduction is always better than that of a fixed quantization scheme. However, it is outperformed by the likelihood classifier. Differences are especially apparent for features with a higher signal-to-noise ratio. In these regions the fixed quantization struggles with a inherent high error probability, while the ZSL scheme follows the likelihood classification.

In a good quantization scheme there needs to be a small gap between the mutual information $I(X;Y)$ and $I(S;\hat{S})$. For a Gaussian channel standard expressions are known from [5, Eq. (9.16)]. The mutual information $I(S;\hat{S})$ can be calculated by using the numerical evaluation for the error probability as described above. Fig. 8 show that a fixed quantization requires a higher SNR to converge to the maximum number of bits, whereas the ZLQ-LR scheme directly reaches this value.

Finally, we will consider the vector case of the two quantization schemes discussed above. In this section we concluded that the fixed quantization will have a larger error probability during reproduction, but we will show how this relates to either false rejection of genuine users or secret length when combined with a code offset method using error correcting codes [9].

In order to derive these properties we will assume the features to be i.i.d. and therefore we can calculate false acceptance rate (FAR) and false rejection rate (FRR) based on a binomial distribution. In practice features can be made (nearly) independent, but they will in general not be identically distributed. However, results will be similar. Furthermore we assume the error correcting code can be applied such that its error correcting properties can be fully exploited. This implies we have to use a gray code to label the extracted secrets before concatenation.

In our experiment we have used 64 i.i.d. features each having a SNR of 17 dB. From these features we extract 2 bits per feature on which we apply BCH codes with a code length of 127, which implies we

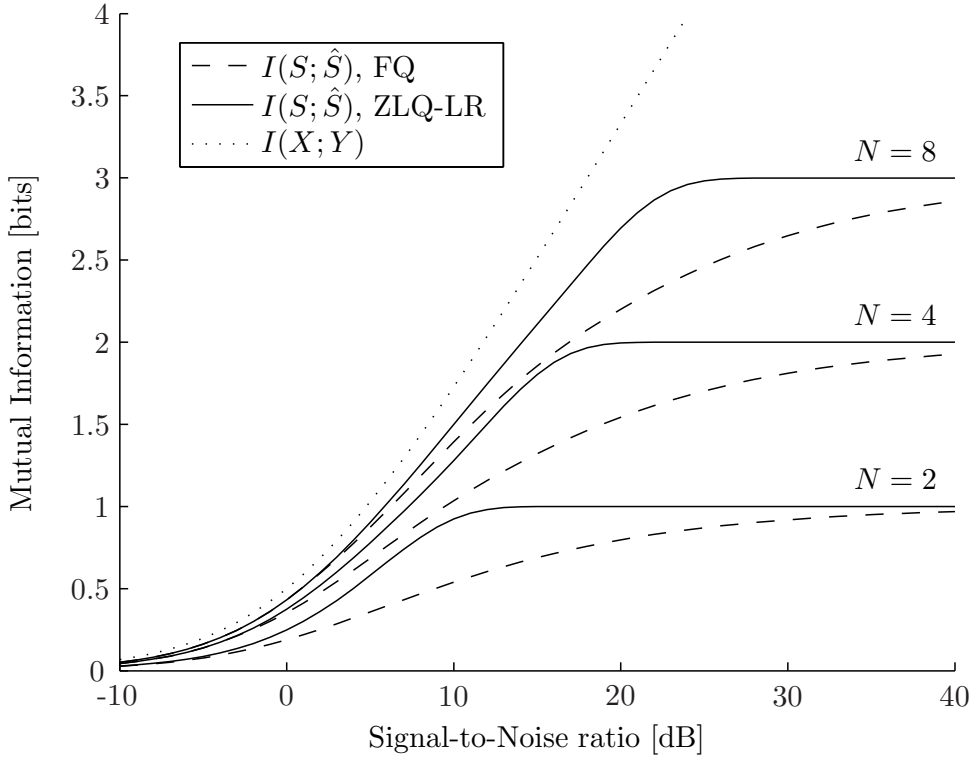

Figure 8: Mutual information between S and $\hat{S}$ for Gaussian distributed features and Gaussian noise.

omit one bit. For analysis we have also included the code $(127, 127, 0)$, which is not an actual code, but represents the case in which no error correction is applied.

Suppose we want to achieve a target FRR of $1 \cdot 10^{-3}$, the topmost dotted line in Fig. 9, then we require a BCH $(127, 92, 5)$ code for the ZSL-LR scheme, while a BCH $(127, 15, 27)$ code is required for the fixed quantization scheme. This implies we would have a secret key size of either 92 or 15 bits. Clearly the last will not be sufficient for any security application since it has a key size that can be attacked in a limited amount of time. At the same time, due to the small key size, the scheme has an increased FAR as depicted in Fig. 9.

6 Conclusion

In this paper we have studied a generic Helper Data Scheme (HDS) which comprises the Fuzzy Extractor (FE) and the Secure Sketch (SS) as special cases. In particular, we have looked at the Zero Leakage (ZL) property of HDSs in the case of a one-dimensional continuous source X and continuous helper data W.

We make minimal assumptions, justified by Conjecture 3.4: we consider only monotonic $g(x)$. We have shown that the ZL property implies the existence of sibling points $\{x_{sw}\}_{s\in\mathcal{S}}$ for every w. These are values of x that have the same helper data w. Furthermore, the ZL requirement is equivalent to a quantile relationship (Theorem 3.6) between the sibling points. This directly leads to equation (36) for computing w from x. (Applying any reversible function to this w yields a completely equivalent helper data system.) The special case of a FE ($p_s = 1/N$) yields the $m \to \infty$ limit of the Verbitskiy et al. [19] construction.

We have derived reconstruction thresholds τ_s for a ZL FE that minimize the error rate in the reconstruction of s (Theorem 4.2). This result holds under very mild assumptions on the noise: symmetric and fading. Eq. (44) contains the attenuation parameter λ, which follows from the noise model as specified in Section 2.2.

Finally we have analyzed reproduction performance in an all Gaussian example. Fixed quantization struggles with inherent high error probability, while the ZL FE with optimal reproduction follows the performance of the optimal classification algorithm. This results in a larger key size in the protected template compared to the fixed quantization scheme, since an ECC with a larger message length can be applied in the second stage HDS to achieve the same FRR.

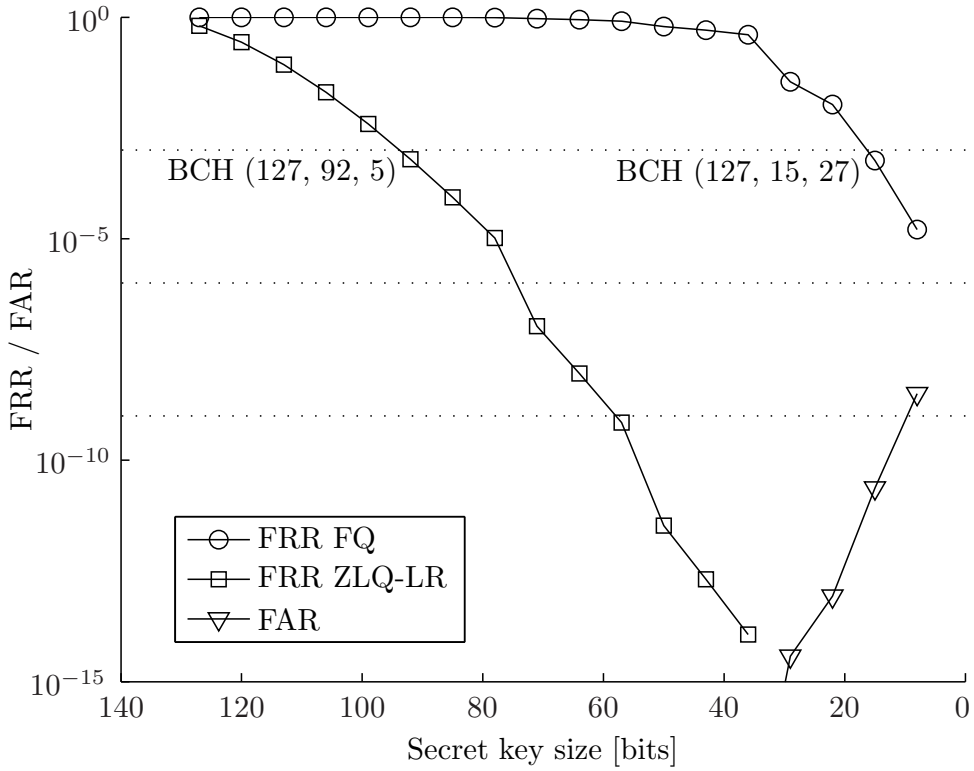

Figure 9: System performance compared to a scheme applying fixed quantization.

References

[1] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables, 9th printing.* Dover, New York, 1972.

[2] B. Chen and G. Wornell. Quantization index modulation: a class of provably good methods for digital watermarking and information embedding. *Information Theory, IEEE Transactions on*, 47(4):1423 –1443, may 2001.

[3] C. Chen, R. N. J. Veldhuis, T. A. M. Kevenaar, and A. H. M. Akkermans. Multi-bits biometric string generation based on the likelihood ratio. In *Biometrics: Theory, Applications, and Systems, 2007. BTAS 2007. First IEEE International Conference on*, pages 1–6, 2007.

[4] M. H. M. Costa. Writing on dirty paper (corresp.). *Information Theory, IEEE Transactions on*, 29(3):439–441, 1983.

[5] T. M. Cover and J. A. Thomas. *Elements of Information Theory.* John Wiley & Sons, Inc., second edition, 2005.

[6] J. A. de Groot and J.-P. M. Linnartz. Improved privacy protection in authentication by fingerprints. In *Proc of the 32st Symp on Inf Theory in the Benelux*, 2011.

[7] Y. Dodis, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *LNCS*. Springer, 2004.

[8] D. Holcomb, W. Burleson, and K. Fu. Power-Up SRAM State as an Identifying Fingerprint and Source of True Random Numbers. *Computers, IEEE Transactions on*, 58(9):1198 –1210, sept. 2009.

[9] A. Juels and M. Wattenberg. A fuzzy commitment scheme. In *CCS '99: Proceedings of the 6th ACM conf on Comp and comm security*, 1999.

[10] E. Kelkboom, K. de Groot, C. Chen, J. Breebaart, and R. Veldhuis. Pitfall of the detection rate optimized bit allocation within template protection and a remedy. In *Biometrics: Theory, Applications, and Systems, 2009. BTAS '09. IEEE 3rd International Conference on*, pages 1 –8, sept. 2009.

[11] E. Kelkboom, G. Molina, J. Breebaart, R. Veldhuis, T. Kevenaar, and W. Jonker. Binary biometrics: An analytic framework to estimate the performance curves under gaussian assumption. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 40(3):555 –571, may 2010.

[12] J.-P. Linnartz and P. Tuyls. New shielding functions to enhance privacy and prevent misuse of biometric templates. In *Audio- and Video-Based Biometric Person Authentication*. Springer, 2003.

[13] F. MacWilliams and N. Sloane. *The Theory of Error Correcting Codes*. North-Holland Mathematical Library. North-Holland, 1978.

[14] T. Matsumoto, H. Matsumoto, K. Yamada, and S. Hoshino. Impact of artificial "gummy" fingers on fingerprint systems. *Optical Security and Counterfeit Deterrence Techniques*, 4677:275–289, 2002.

[15] G. E. Suh and S. Devadas. Physical unclonable functions for device authentication and secret key generation. In *Proceedings of the 44th annual Design Automation Conference*, DAC '07, pages 9–14, New York, NY, USA, 2007. ACM.

[16] P. Tuyls, A. Akkermans, T. Kevenaar, G.-J. Schrijen, A. Bazen, and R. Veldhuis. Practical biometric authentication with template protection. In *Audio- and Video-Based Biometric Person Authentication*, pages 436–446. Springer, 2005.

[17] P. Tuyls, B. Škorić, and T. Kevenaar. *Security with Noisy Data: Private Biometrics, Secure Key Storage and Anti-Counterfeiting*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2007.

[18] T. van der Putte and J. Keuning. Biometrical fingerprint recognition: don't get your fingers burned. In *Fourth Working Conference on Smart Card Research and Advanced Applications*, pages 289–303, Norwell, MA, USA, 2001. Kluwer Academic Publishers.

[19] E. A. Verbitskiy, P. Tuyls, C. Obi, B. Schoenmakers, and B. Škorić. Key extraction from general nondiscrete signals. *Information Forensics and Security, IEEE Transactions on*, 5(2):269 –279, June 2010.

[20] B. Škorić, P. Tuyls, and W. Ophey. Robust key extraction from physical uncloneable functions. In J. Ioannidis, A. Keromytis, and M. Yung, editors, *Applied Cryptography and Network Security*, volume 3531 of *Lecture Notes in Computer Science*, pages 99–135. Springer Berlin / Heidelberg, 2005.

[21] J. L. Wayman, A. K. Jain, D. Maltoni, and D. Maio, editors. *Biometric Systems: Technology, Design and Performance Evaluation*. Springer Publishing Company, Incorporated, 1st edition, 2005.