Secure Sketch for Multiple Secrets

Chengfang Fang[1], Qiming Li[2], and Ee-Chien Chang[1,⋆]

[1] School of Computing, National University of Singapore
{c.fang,changec}@comp.nus.edu.sg
[2] Institute for Infocomm Research, Singapore
Qiming.Li@ieee.org

Abstract. Secure sketches are useful in extending cryptographic schemes to biometric data since they allow recovery of fuzzy secrets under inevitable noise. In practice, secrets derived from biometric data are seldom used alone, but typically employed in a multi-factor or a multimodality setting where multiple secrets with different roles and limitations are used together. To handle multiple secrets, we can generate a sketch for each secret independently and simply concatenate them. Alternatively, we can "mix" the secrets and individual sketches, for example, by taking the first secret as the key to encrypt the sketches of all other secrets. Hence, it is interesting to investigate how the secrets are to be mixed so as to cater for different requirements of individual secrets. We found that, by appropriate mixing, entropy loss on more important secrets (e.g., biometrics) can be "diverted" to less important ones (e.g., password or PIN), thus providing more protection to the former. On the other hand, we found that mixing may not be advisable if the amount of randomness invested in sketch construction is large, or the sketch contains high redundancy, or all secrets are of the same importance. Our analysis provides useful insights and guidelines in the applications of secure sketches in biometric systems.

Keywords: Multi-factor authentication, biometric security, sketch construction.

1 Introduction

Biometrics is potentially useful in building secure and easy-to-use security systems, since it is tightly bound to identities, cannot be easily forgotten or lost. However, these features can also make user credentials based on biometric measures hard to revoke, since once the biometric data of a user is compromised, it would be very difficult to replace it, if possible at all. A key challenge in protecting biometric data as user credentials is that they are fuzzy, in the sense that it is not possible to obtain exactly the same data in two measurements. This renders traditional cryptographic techniques used to protect passwords and keys inapplicable.

⋆ Chang is supported by Grant R-252-000-413-232/422/592 from Temasek Defence Systems Institute (TDSI).

J. Zhou and M. Yung (Eds.): ACNS 2010, LNCS 6123, pp. 367–383, 2010.

Secure sketches [6] are a recently proposed cryptographic primitive that can be used, in conjunction with other cryptographic techniques, to extend classical cryptographic techniques to fuzzy secrets, including biometric data. The key idea is that, given a secret x, we can compute some auxiliary data p, which is called a *sketch*. The sketch p will be able to correct errors from a noisy version of x and recover the original data x that was enrolled. From there, typical cryptographic schemes such as one-way hash functions can then be applied on x. In particular, an *extractor* can be further applied on the data to obtain a nearly-uniform key of certain length given that the min-entropy of the original data is known. Such a generic method of obtaining a consistent key from fuzzy data is referred to as a *fuzzy extractor*.

The work by Dodis et al. [6] on secure sketches and fuzzy extractors provides a theoretical framework that allows us to analyze the security measured by the *entropy loss* of the sketch, which gives a measure of the amount of information a sketch reveals about the underlying secret. There are also a number of schemes (e.g., [11,10,6,3]) with provable upper bounds on the entropy loss. However, no matter how small the entropy loss is, without additional protections, for most biometric representations, it is inevitable that some important information is revealed.

Biometric data is often employed together with other types of secrets as in a multi-factor setting, or in a multimodal setting where there are multiple sources of biometric data, partly due to the fact that human biometrics is usually of limited entropy. In the context of secure sketches, it is possible to treat these secrets independently: The sketches are generated independently and the final sketch is simply the concatenation of all sketches. The security analysis can also be easily carried out by investigating each secret separately. However, secrets may differ in terms of their entropies and fuzziness. More importantly, they may differ in their roles and constraints in their usage. For example, the likelihood of being lost, stolen or forgotten and the ease of revocation and replacement, would be different for different secrets. Furthermore, when exposed, biometric data, like fingerprints, can be used to infer some sensitive illness information [25] of the person. The straightforward method of combining the secrets independently treats each secret equally, thus may not be able to cater for individual security requirements.

One way to address this issue is to mix different secrets together. By mixing the secrets, we may be able to provide more protection to more important secrets at the expense of reduced protection of others. However, if mixing is not done appropriately, it could reveal more information compared to the straightforward method without mixing. Therefore, a detailed investigation is required.

Let us give a simple example here. Suppose the credential of a user consists of a fingerprint and a password, both of which are known to have relatively low entropies. Intuitively, to provide more protection to the important fingerprint, one could use the password to encrypt the fingerprint's sketch. Now, a few questions to address are: How to quantify the additional protection provided? Does

the method really provide additional protection, i.e. are there situations where more information is leaked by inappropriate mixing?

In this paper, we propose and analyze a cascaded mixing approach which is essentially the same as described above: use the less important secret to mix with the sketch of the more important secret. As the leftover entropies of password might be low, it is feasible for an adversary to carry attack by enumerating all likely passwords. Hence, we do not rely on the assumption that the mixing function is computationally one-way. Instead, we focus on information-theoretic aspect of the mixing. To address the question on how to quantify the security, note that if we treat the two secrets as a single secret and investigate the combined leftover entropy $\widetilde{\mathbf{H}}_\infty((X,K)|Q)$, where X, K and Q are random variables of the biometrics data, password and the final sketch respectively, the simple method of concatenating the two secrets could already be optimal. Hence, to capture the additional protection, we investigate the individual leftover entropy $\widetilde{\mathbf{H}}_\infty(X|Q)$ and $\widetilde{\mathbf{H}}_\infty(K|Q)$.

We show that, if the sketch construction is deterministic, cascaded mixing can divert the information leakage towards K. Such additional protection is desirable. In the above example, consider a scenario where an adversary happens to obtain some prior knowledge of the more important secret X. Without mixing, the sketch may provide additional information for the adversary to obtain the secret with high probability. By proper mixing, the adversary cannot obtain information of X from the mixed sketch Q, instead, he obtains some information of K.

Consider the second question on whether there are scenarios where mixing is not advisable. We make two observations. Firstly, we found that when there are high redundancies in the sketch, more entropy could be lost compared with the straightforward method of handling the secrets independently. More precisely, the leftover entropy $\widetilde{\mathbf{H}}_\infty(X,K|Q)$ may be less than $\widetilde{\mathbf{H}}_\infty(X,K|P)$, where P is simply the concatenation of the sketches for X and K. This observation is useful as a number of sketch constructions (e.g., [5]) would produce sketches that contain high redundancies but are difficult to compress. In the second observation, we give counter example to show that, when the randomness invested during sketch construction cannot be decoupled from the sketch, there are scenarios where the mixing is an redundant step as it does not provide more protection to the more important secret, i.e. it essentially provides the same protection as the simple concatenation method. Hence given two choices of sketch constructions where one is deterministic and the other is probabilistic, it is advisable to employ the deterministic method to achieve the protection provided by mixing.

Contributions and Organization

We observe that, in some biometric applications, different secrets have different requirements and some secrets require more protection than others. We argue that the straightforward method of constructing the sketch independently is not satisfactory as it does not address such differences.

We propose a cascaded mixing approach to mix the secrets whereby more important secrets are mixed first (Section 4.1). We analyze the approach and show

that, if the sketch construction does not involve randomness, the information leakage on the more important secrets will be "diverted" to the less important secrets (Section 5.1, Theorem 2, 3).

We provide counter-examples to demonstrate that, if the sketch construction involves randomness, there are scenarios where mixing function is unable to further protect the more important secret (Section 6.1) and in some cases it leak information of the less important secret (Section 6.2). We also give an intuitive explanation.

Based on our analysis, we provide guidelines in constructing sketches for multiple secrets (Section 7).

2 Related Work

The fuzzy commitment [11] and the fuzzy vault [10] schemes are among the first error-tolerant cryptographic techniques. More recently, Dodis et al. [6] give a general framework of secure sketches and fuzzy extractors, where the security is measured by the entropy loss of the secret given the sketch. They give specific schemes that meet theoretical bounds for Hamming distance, set difference and edit distance respectively. Another distance measure, point-set difference, motivated from a popular representation for fingerprint features, is investigated in a number of studies [5,3,4]. A different approach [14,24,23] focuses on information leakage defined using Shannon entropy on continuous data with known distributions.

There are also a number of investigations on the limitations of secure sketches under different security models. Boyen [1] studies the re-usability of sketches where the concern is whether multiple sketches of the same biometric data reveal sensitive information. This security model is further extended and studied by Boyen et al. [2] and Simoens et al. [20], where the latter work focuses more on privacy issues. Kholmatov et al. [12] and Hong et al. [9] demonstrate such limitations by giving correlation attacks on known schemes.

The idea of using a secret to protect other secrets is not new. Souter et al. [21] propose integrating biometric patterns and encryption keys by hiding the cryptographic keys in the enrollment template via a secret bit-replacement algorithm. Some other methods use password protected smartcards to store user templates [15,19]. Ho et al. [8] propose a dual-factor scheme where a user needs to read out a one-time password generated from a token, and both the password and the voice features are used for authentication. Sutcu et al. [22] study secure sketch for face features and give an example of how the sketch scheme can be used together with a smartcard to achieve better security.

Using only passwords as an additional factor is more challenging than using smartcards, since the entropy of typical user chosen passwords is relatively low [17,7,13]. Monrose [16] presents an authentication system based on Shamir's secret sharing scheme to harden keystroke patterns with passwords. Nandakuma et al. [18] propose a scheme for hardening a fingerprint minutiae-based fuzzy vault using passwords, so as to prevent cross-matching attacks.

3 Formulations and Background

Table 1 summarizes the notation we are going to use in this paper.

Table 1. Table of notations used

X:	Fuzzy secret distributed over space $\mathcal{M}$.
$\mathtt{D}$:	Distance function defined with $\mathcal{M}$.
$\mathbf{H}_\infty(A)$:	Min-entropy of random variable A.
$\widetilde{\mathbf{H}}_\infty(A\|B)$:	Average min-entropy of A given B.
Enc:	Encoder of a known sketch scheme.
P:	The sketch of X, $P = \mathsf{Enc}(X, R)$.
R:	Recoverable random string used in an encoder.
K:	A non-fuzzy secret or a key.
f:	A mixing function.
S:	Recoverable randomness used in f.
Q:	Output of a mixing function, $Q = f(P, K, S)$.
L_A:	The length of variable A, e.g. L_P is the length of sketch P

3.1 Min-Entropy and Entropy Loss

We follow Dodis et al. [6] and use the following definitions of min-entropy and entropy loss.

The *min-entropy* $\mathbf{H}_\infty(A)$ of a discrete random variable A is $\mathbf{H}_\infty(A) = -\log(\max_a \Pr[A = a])$. For two discrete random variables A and B, the *average min-entropy* of A given B is defined as $\widetilde{\mathbf{H}}_\infty(A|B) = -\log(\mathbb{E}_{b \leftarrow B}[2^{-\mathbf{H}_\infty(A|B=b)}])$

The *entropy loss* of A given B is defined as the difference between the min-entropy of A and the average min-entropy of A given B. In other words, the entropy loss $\mathcal{L}(A, B) = \mathbf{H}_\infty(A) - \widetilde{\mathbf{H}}_\infty(A|B)$. Note that for any n-bit string B, it holds that $\widetilde{\mathbf{H}}_\infty(A|B) \geq \mathbf{H}_\infty(A) - n$, which means we can bound $\mathcal{L}(A, B)$ from above by n regardless of the distributions of A and B.

3.2 Secure Sketches and Fuzzy Extractors

Assuming the original secret x is a point in a discrete domain $\mathcal{M}$ with distance function $\mathtt{D}$, a secure sketch scheme consists of two efficient algorithms: An encoder Enc, which computes a sketch p on the given x, and a decoder Dec, which computes an x' given a p and y such that $x' = \mathsf{Dec}(p, y) = x$ if $\mathtt{D}(x, y) \leq t$ for some threshold t.

More formally, let $\mathcal{M}$ be a metric space with distance function $\mathtt{D}$, we have the following definition[1].

[1] Our definition here looks slightly different from that given by Dodis et al. [6] in that we make the randomness invested during encoding more explicit.

Definition 1 ([6]). *An $(\mathcal{M}, t, \gamma)$-sketch scheme consists of two deterministic polynomial-time algorithms* $\mathsf{Enc} : \mathcal{M} \times \{0,1\}^\gamma \to \{0,1\}^*$ *and* $\mathsf{Dec} : \mathcal{M} \times \{0,1\}^* \to \mathcal{M}$ *such that for all $x, y \in \mathcal{M}$ and $r \in \{0,1\}^\gamma$, it holds that* $\mathsf{Dec}(y, \mathsf{Enc}(x, r)) = x$ *when* $\mathtt{D}(x, y) \leq t$. *We call $p = \mathsf{Enc}(x, r)$ the sketch of x. Also, we say that the randomness r is recoverable if for any x and r', if* $\mathsf{Enc}(x, r') = \mathsf{Enc}(x, r)$, *we have $r = r'$.*

A fuzzy extractor can be built on top of a secure sketch by applying an *extractor* Ext on a random secret, as shown by Dodis et al. [6]. Given a random variable X with sufficient min-entropy, an extractor[2] is able to compute a *nearly uniform* key of a length that is slightly less than the min-entropy of X. Hence, given a secret x, we can use an extractor to obtain a key k from it. When a y that is close to x with respect to $\mathtt{D}$ and t is presented, the original x can be reconstructed and hence the same key can be obtained by applying the same extractor again on the reconstructed x. In this way, fuzzy secrets can be used just the same way as consistent secrets, except that now an additional sketch p has to be stored and used to reconstruct the original secret.

A well adopted approach measures the security of such a scheme by the amount of information revealed by the sketch about the original secret. Formally, for discrete metric space $\mathcal{M}$ with distance function $\mathtt{D}$, the entropy loss of an $(\mathcal{M}, t, \gamma)$-sketch scheme with encoder Enc is defined as follows.

Definition 2. *The entropy loss of an $(\mathcal{M}, t, \gamma)$-sketch scheme is* $\mathbf{H}_\infty(X) - \widetilde{\mathbf{H}}_\infty(X \mid P)$ *for random variable X on $\mathcal{M}$ and the sketch* $P = \mathsf{Enc}(X, R)$.

Essentially, if a sketch scheme has an entropy loss bounded from above by $\mathcal{L}$, it means that $\mathbf{H}_\infty(X) - \widetilde{\mathbf{H}}_\infty(X|P) \leq \mathcal{L}$ for any distribution of X. It is possible to design an extractor Ext such that $K = \mathsf{Ext}(X)$ is nearly uniform even when P is known, and the length of K can be at least $\widetilde{\mathbf{H}}_\infty(X|P) - \delta$ for a small δ determined by how close the distribution of K is to the uniform distribution[6]. Hence, if an attacker tries to guess the extracted key, the success probability cannot be much better than $2^{-\widetilde{\mathbf{H}}_\infty(X|P)+\delta}$.

Furthermore, let R be the randomness invested by the encoder Enc during the computation of the sketch P, it is not difficult to show (as mentioned in [6]) that when R is recoverable from X and P, we have

$$\mathbf{H}_\infty(X) - \widetilde{\mathbf{H}}_\infty(X|P) \leq L_P - \mathbf{H}_\infty(R) \tag{1}$$

That is, the entropy loss is bounded from above by the difference between the length of P and $\mathbf{H}_\infty(R)$, which is just the length of R if it is uniform. Furthermore, this upper bound is independent of X, hence it holds for any distribution of X.

[2] For example, pair-wise independent hash functions.

The inequality (1) is useful in deriving a bound on the entropy loss, since typically the size of P and $\mathbf{H}_\infty(R)$ can be easily obtained regardless of the distribution of X. This approach is useful in many scenarios where it is difficult to model the distribution of X, for example, when X represents the features of a fingerprint.

4 Secure Sketch for Two Secrets

In some applications the credential of a user consists of two independent secrets. The sources of these secrets can be different. For example, they may be in different metric spaces with different distance functions and thresholds.

A straightforward extension of sketch construction to two secrets is to simply apply two sketch schemes, for the two secrets x_1 and x_2 independently. The final sketch for the two secrets is the concatenation of the sketches p_1 and p_2 computed from x_1 and x_2 respectively. That is, the sketch $p = p_1 \| p_2$, where $\|$ represents concatenation. Furthermore, the final key can be obtained by concatenating the keys k_1 and k_2 extracted from x_1 and x_2 respectively.

Suppose the entropy loss of the first secret given the sketch is at most $\mathcal{L}_1$, and that of the second secret is at most $\mathcal{L}_2$, then it is clear that the overall entropy loss is at most $\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2$, since the secrets are independent.

As we have mentioned, this straightforward approach is not able to differentiate secrets with different characteristics, and give equal protection to both secrets.

4.1 A Cascaded Mixing Approach

Instead of treating the two secrets independently, it may be desirable to combine different types of secrets to achieve additional security goals. Here we give an alternative sketch construction. Figure 1 illustrates our proposed method.

For secrets x_1 and x_2, we first compute sketches p_1 and p_2 as in the concatenating approach, and extract keys k_1 and k_2 respectively then we encrypt p_1 using k_2 as the key. That is, we compute $q_1 = f(p_1, k_2, s)$, where f is a deterministic function and s is an auxiliary random string. The final sketch q is $q = q_1 \| p_2$.

Let us call f the *mixing function* which serves as an encryption with k_2 as the key. As the leftover entropy of k_2 given p_2 could be low, we should not rely on the computational difficulty in inverting f to protect p_1. Thus, it is important to analyze how much information about the two secrets x_1 and x_2 is revealed.

Let us consider the mixing function $f : \{0,1\}^{L_P} \times \{0,1\}^{L_K} \times \{0,1\}^{L_S} \rightarrow \{0,1\}^{L_Q}$ and random variables Q, P, K and S such that $Q = f(P, K, S)$. We require f to have certain properties. First, as an encryption function, f must be invertible.

Definition 3 (Invertibility). *We say that a mixing function f is invertible if there is a function g such that for all $p \in \{0,1\}^{L_P}$, $k \in \{0,1\}^{L_K}$ and $s \in \{0,1\}^{L_S}$, $g(f(p,k,s),k) = p$.*

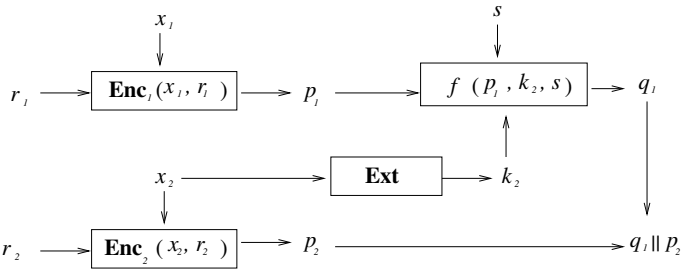

Fig. 1. Construction of cascaded mixing approach

In addition, in our analysis we consider mixing functions with the following properties on recoverability of the randomness invested.

Definition 4 (Recoverable Randomness). *For a mixing function f, the randomness S is called recoverable if for any $p \in \{0,1\}^{L_P}$, $k \in \{0,1\}^{L_K}$ and $s, s' \in \{0,1\}^{L_S}$, if $f(p,k,s) = f(p,k,s')$, we have $s = s'$.*

Definition 5 (β-Recoverable Key). *For a mixing function f, the key K is called β-recoverable if for any $p \in \{0,1\}^{L_P}$, $q \in \{0,1\}^{L_Q}$, the cardinality of the set $\mathcal{K}_{p,q} = \{k \in \{0,1\}^{L_K} | \exists s \in \{0,1\}^{\beta}, f(p,k,s) = q\}$ is at most 2^{β}.*

It is easy to construct mixing function achieving both invertability and recoverability. For example, we can obtain one from a block cipher $f(p,k,r) = r \| E_k(p \| r)$. Note that the recoverability properties are not necessary for the recovery of the secrets, but will become handy in the security analysis.

When a user presents y_1 and y_2 that are close to x_1 and x_2 respectively, x_2 is first reconstructed using y_2 and p_2, and a key k_2 is extracted from x_2, which in turn is used to retrieve p_1 if f is invertible. After that, x_1 is reconstructed using y_1 and p_1. An extractor can be further applied on $x_1 \| x_2$ to extract a key.

Intuitively, this alternative approach gives more protection to the first secret x_1, since it would require the attacker to guess x_2 using p_2 first, only when the attacker is successful can the attacker gain information on x_1 from q_1 by computing p_1 from q_1 and x_2.

5 Analysis

We now study the case of two secrets and a scheme that follows the cascaded sketch construction (Section 4.1) Let $x \in \mathcal{M}$ be a fuzzy secret (say, a fingerprint), and let $k \in \{0,1\}^{L_K}$ be an independent secret key that is not fuzzy. Consider a $(\mathcal{M}, t, L_R)$-sketch scheme with encoder Enc, and let the sketch $p = \mathsf{Enc}(x, r)$. Figure 2 illustrates the process.

It is clear that when the key K is uniform and no shorter than the sketch, we can easily hide the sketch p completely (e.g., by using a one-time pad). However,

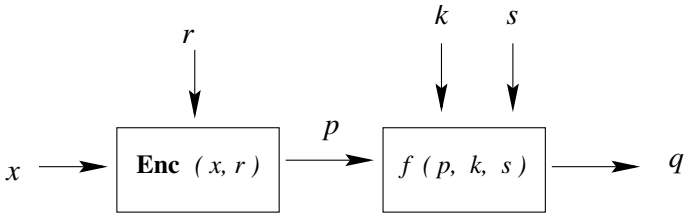

Fig. 2. Computation of mixed sketch

in practical scenarios (e.g., user chosen PIN/password as the key), K can be shorter than p, and the analysis of security may become challenging. In fact, we will show that, for shorter K, mixing is not always a better strategy than the straightforward method of treating the secrets independently. We will also show the conditions under which mixing is desirable.

5.1 Security of the Cascaded Mixing Approach

Analysis of overall remaining entropy $\widetilde{\mathbf{H}}_\infty(X, K|Q)$
First, let us investigate the remaining entropy when we treat (X, K) as a single secret, i.e. the remaining entropy $\widetilde{\mathbf{H}}_\infty(X, K|Q)$.

Lemma 1. *Given random variables X, K, R, S and mixing function f as described above, We have* $\widetilde{\mathbf{H}}_\infty(X, K|Q) \geq \mathbf{H}_\infty(X) + \mathbf{H}_\infty(K) + \mathbf{H}_\infty(R) - L_P$.

Proof: Since S is recoverable, we can consider Enc and f together as the encoding algorithm for the final sketch Q, R and S together as the recoverable randomness, and the inequality (1) in Section 3 applies. Note that $L_Q = L_P + L_S$, and we have

$$\begin{aligned}\widetilde{\mathbf{H}}_\infty(X, K|Q) &\geq \mathbf{H}_\infty(X, K) + \mathbf{H}_\infty(R) + \mathbf{H}_\infty(S) - L_Q \\ &= \mathbf{H}_\infty(X) + \mathbf{H}_\infty(K) + \mathbf{H}_\infty(R) - L_P.\end{aligned}$$

Hence the lemma holds as claimed. □

Lemma 1 gives a lower bound of the remaining entropy of X and K. In general, if both secrets are fuzzy, we can similar obtain the bound:

$$\widetilde{\mathbf{H}}_\infty(X_1, X_2|Q) \geq \mathbf{H}_\infty(X_1) + \mathbf{H}_\infty(X_2) + \mathbf{H}_\infty(R_1) + \mathbf{H}_\infty(R_2) - L_{P_1} - L_{P_2}.$$

where X_1 and X_2 are the secrets, R_1, R_2, are the randomness invested in constructing the sketch P_1, P_2 for the two respective secrets. Note that this bound is the same when we use the straightforward concatenation approach.

Analysis of individual secret $\widetilde{\mathbf{H}}_\infty(X|Q)$ and $\widetilde{\mathbf{H}}_\infty(K|Q)$
Now, let us look at the remaining entropy of individual secret, i.e. $\widetilde{\mathbf{H}}_\infty(X|Q)$ and $\widetilde{\mathbf{H}}_\infty(K|Q)$.

If the sketch is not uniformly distributed, then given the mixed q, it is possible that $(K|Q = q)$ is not uniform. That is, Q will leak some information about K. Indeed, an adversary, given q, may enumerate all possible k's and the correspond sketch p to determine the most likely k. Nevertheless, leakage of K is acceptable as long as it can provide more protection to X. Next theorem gives a lower bound on the remaining entropy of X given the mixed sketch Q.

Theorem 2. *Given three independent random variables X, K and R distributed over $\mathcal{M}$, $\{0,1\}^{L_K}$ and $\{0,1\}^{L_R}$ respectively and an $(\mathcal{M}, t, L_R)$-sketch scheme with encoder* Enc*, Let P be the sketch of X, i.e., $P = \mathsf{Enc}(X, R)$, where R is recoverable, and let $f : \{0,1\}^{L_P} \times \{0,1\}^{L_K} \to \{0,1\}^{L_Q}$ be an mixing function and $Q = f(P, K, S)$, where S is a L_S bits of recoverable randomness. If f is invertible and the key K is L_S-recoverable. Then*

$$\widetilde{\mathbf{H}}_\infty(X|Q) \geq \mathbf{H}_\infty(X) + \mathbf{H}_\infty(K) - L_Q. \tag{2}$$

We would like to refer the reader to Appendix A for the proof of the above theorem.

The theorem holds for any distributions of X and K, and for uniformly distributed K, the theorem implies that $\widetilde{\mathbf{H}}_\infty(X|Q) \geq \mathbf{H}_\infty(X) + L_K - L_Q$. Let us compare the remaining entropy if we use the simple concatenation method, which is as follows,

$$\widetilde{\mathbf{H}}_\infty(X|P) \geq \mathbf{H}_\infty(X) + L_R - L_P \tag{3}$$

Now, coming back to the question that whether it is beneficial to use a cascading function when the secret k is short compared with p. Clearly, from Theorem 2 and inequality (3), we can see that when $\mathbf{H}_\infty(K) - L_Q \geq L_R - L_P$, or equivalently, $\mathbf{H}_\infty(K) \geq L_R + L_S$, the R.H.S in (2) is larger then the R.H.S in (3), i.e. the entropy bound when using a mixing function is no worse than not using it. In particular, consider a deterministic sketch scheme (i.e. L_R =0), and a length preserving mixing function (thus $L_P = L_Q$), the difference in the right hand side of the inequality (2) and (3) is $\mathbf{H}_\infty(K)$. In other words, the bound on leftover entropy of X given Q can be increased by $\mathbf{H}_\infty(K)$. Viewing from another direction, information loss on X is "diverted" to K.

Now, we consider only the non-fuzzy secret k and analyze the entropy loss.

Theorem 3. *Given an $(\mathcal{M}, t, L_R)$-sketch scheme with encoder* Enc*, and let X, K, R, P, Q, f, S be as defined in Theorem 2, we have*

$$\widetilde{\mathbf{H}}_\infty(K|Q) \geq \mathbf{H}_\infty(K) + \mathbf{H}_\infty(R) - L_P. \tag{4}$$

Proof: Since $Q = f(P, K, S)$, we can regard Q as a sketch of K where the cascading function f is an encoder, and $P = \mathsf{Enc}(X, R)$ and S are the "randomness" invested in computing Q, which are recoverable. Clearly, we can apply the general bound (1) on K and Q, and since R is recoverable, we have

$$\mathbf{H}_\infty(X) + \mathbf{H}_\infty(P) \geq \widetilde{\mathbf{H}}_\infty(X, P) \geq \mathbf{H}_\infty(X) + \mathbf{H}_\infty(R)$$

which means that $\mathbf{H}_\infty(P) \geq \mathbf{H}_\infty(R)$, hence the inequality holds as desired. □

It is worth to note that the bound in Theorem 3 is tight in the sense that there exists random variables and functions such that the equality in (4) holds. We will see an example of such case in Section 6.2. Therefore, if L_P is large but the min-entropy $\mathbf{H}_\infty(P)$ is low, the quantity $\mathbf{H}_\infty(K)+\mathbf{H}_\infty(P)-L_P$ may be reduced to 0 or even less than 0, in which case Q may reveal all information about K.

6 Examples of Improper Mixing

In this section we give examples to illustrate the scenarios where mixing function may not be beneficial: (1) in scenarios where the sketch construction employs randomness, mixing function may not always provide protection on X. (2) when the sketch contains high redundancy from the adversary point of view, mixing function may reveal information of K.

6.1 Randomness Invested in Sketch

This section gives a simple example to illustrate the idea that mixing function may not always provide protection on X, if the sketch construction contains randomness. Hence, as a general guideline, when choosing a sketch scheme to be used in the cascaded mixing framework, it is better to select one that requires no randomness.

Consider a non-fuzzy K in $\{0,1\}^{L_K}$, and a fuzzy X in $\{1 \dots 2^{L_X}\}$ with the distance function

$$d(x_1, x_2) = \begin{cases} 0, & \text{if } x_2 = x_1 \\ 1, & \text{if } x_2 = x_1 + 1 \mod 2^{L_X} \\ \infty, & \text{otherwise} \end{cases}$$

for any $x_1, x_2 \in \{1 \dots 2^{L_X}\}$ and the noise threshold is 1. Hence, a noisy copy of an x could be either x or $(x+1) \mod 2^{L_X}$.

Consider the following two sketch constructions: a deterministic construction $\mathsf{Enc}_1(X) = X \mod 2$, and a probabilistic construction $\mathsf{Enc}_2(X, R) = X + R \mod 2^{L_X}$, where R is a uniform random even number in $\{1 \dots 2^{L_X}\}$. Without mixing, sketches output from both constructions reveal at most one bit of X.

Given a one bit secret K, let the mixing function $f(P, K, S)$ be as following: it first generates with seed S a set $\mathbf{S} = \langle \mathbf{s}_1, \mathbf{s}_2 \rangle$ of random strings of length L_P, then it output $P + \mathbf{k}_K \mod 2^{L_P}$.

Consider the case when Enc_1 is used, the mixing function is one-time pad encryption, by Theorem 2, there will be no entropy loss on X i.e. $\mathbf{H}_\infty(X) - \widetilde{\mathbf{H}}_\infty(X|Q) = 0$. However, when Enc_2 is used, there could be cases where $\mathbf{s}_i$ has same parity, for example, $\mathbf{S} = \langle 0, 2 \rangle$. In that case, the information of the sketch is not protected and $\mathbf{H}_\infty(X) - \widetilde{\mathbf{H}}_\infty(X|Q) = 1$ and there is no gain nor loss in mixing the secrets compare to the straightforward method. In other words, the secret K is unable to provide additional protection as desired.

Note that, by Lemma 1, the overall entropies $\widetilde{\mathbf{H}}_\infty(X, K|Q)$ are the same in the aforementioned two cases, as well as in the straightforward method of not mixing the secrets.

Hence, when given two choices of sketch constructions where one is deterministic and the other is probabilistic, it is advisable to employ the deterministic method to achieve the protection provided by mixing function.

6.2 Redundancy in Sketch

When the sketch has redundancy, that is, the entropy of the sketch is smaller than the length of the sketch, information on k will be leaked from the mixed sketch. There are a few known sketch constructions where the "support" of the sketch (i.e. the number of sketches which non-zero probability of occurrences) is significantly smaller than 2^{L_P} where L_P is the length of the sketch and thus their sketches contain redundancy. One example is the chaff-based method [5] proposed to protect the biometric fingerprint. Here, a fingerprint is the secret x and can be represented as a set of 2D points. The chaff-based method gives its sketch which is the original x union with a set of random 2D points, constrained by the requirement that no two points are close to each other (w.r.t Euclidean distance). It is not easy to derive a compact description of the sketch whose support has size close to 2^{L_P}. Now, suppose that the sketch is mixed with a short k. Given a mixed sketch q, it could be highly likely that among all possible K's in inverting q, only one give a point set that satisfies the constrain. Thus, immediately, the secret k and the sketch is revealed, and the remaining entropy of the combined $\widetilde{\mathbf{H}}_\infty(X, K|Q) = \widetilde{\mathbf{H}}_\infty(X|P)$. Hence, by mixing, not only there is no further protection of x, the k is revealed.

We also conducted experiment to illustrate that, even when the description of sketch is compact, i.e. its support equals 2^{L_P}, the chaff-based sketch still contains significant redundancy that leads to lost of information on k.

Consider the chaff-based method for 1D points, which is easy to derive a compact description. We simulated the chaff-based method in $\mathbb{Z}_{24}$ with a minimum distance 3. There are in total 605 possible sketches, and we randomly generated 10^5 sketches. Figure 3 shows the numbers of occurrences for all 605 sketches with x-axis descendingly sorted by the number of occurrence (and we call the position of a sketch in this descending list the *rank* of it).

Suppose the sketch is then protected by a 5 bits key k, and a mixing function f such that the inverts are always valid sketches. We then simulate an adversary who try to guess k when given $q = f(k, p, s)$, where k and s are randomly chosen from their domain and p is chosen according to the distribution approximated by Figure 3. We simulated 10^5 guesses and the adversary can succeed with probability slightly more than 0.052, instead of $1/(2^5) = 0.03125$ as in random guessing.

7 Further Discussions

7.1 The Case of Two Fuzzy Secrets

When both secrets are fuzzy and may not be uniform, we show that the bounds of Lemma 1, Theorem 2 and 3 can be obtained with slight modifications.

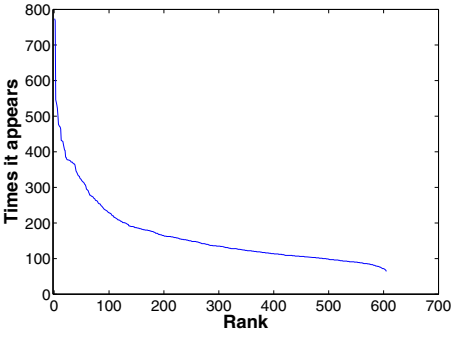

Fig. 3. Histogram of sketch occurrences

Suppose there are two independent secrets $x_1 \in \mathcal{M}_1$ and $x_2 \in \mathcal{M}_2$, and two sketch construction schemes with encoder Enc_1 and Enc_2 respectively. We assume that the first secret x_1 is more important than x_2. In this case, we can use the following steps to construct the sketch for the two secrets.

1. Compute $p_1 = \mathsf{Enc}_1(x_1, r_1)$ and $p_2 = \mathsf{Enc}_2(x_2, r_2)$.
2. Extract a key k_2 from x_2 using an extractor Ext.
3. Compute $q_1 = f(p_1, k_2, S)$ using a mixing function f.
4. Output the final sketch $q = q_1 \| p_2$.

It is possible to design Ext such that K_2 and P_2 are independent, and $\mathbf{H}_\infty(K_2)$ is only slightly smaller than $\widetilde{\mathbf{H}}_\infty(X_2|P_2)$ [6]. Let δ be a small extractor-dependent value such that $\mathbf{H}_\infty(K_2) \geq \widetilde{\mathbf{H}}_\infty(X_2|P_2) - \delta$.

The bound in Theorem 2 still applies on x_1 and k_2. Consider random variables X_1 and K_2, corresponding sketches P_1 and P_2, mixed sketch Q_1, and final sketch Q, it's not difficult to show that $\widetilde{\mathbf{H}}_\infty(X_1|Q) \geq \mathbf{H}_\infty(X_1) + \mathbf{H}_\infty(X_2) + \mathbf{H}_\infty(R_2) - L_{P_2} - \delta - L_Q$ where R_2 is the recoverable randomness used in computing P_2. In this case, the small δ can be considered as the overhead of using the extractor Ext.

As a comparison, if we treat the two secrets independently, and consider $P = P_1 \| P_2$, we have $\widetilde{\mathbf{H}}_\infty(X_1|P) = \widetilde{\mathbf{H}}_\infty(X_1|P_1) \geq \mathbf{H}_\infty(X_1) + L_{R_1} - L_{P_1}$.

Similar to the example, we can conclude that if $\mathbf{H}_\infty(K_2) \geq L_{R_1} + L_S$, we can obtain a better bound on the entropies when we choose to mix k_2 with p_1. Otherwise, doing so may reveal more information about X_1.

The entropy loss on the second secret X_2 can be obtained using the bound in Theorem 3. It's not difficult to show that $\widetilde{\mathbf{H}}_\infty(X_2|Q) \geq \mathbf{H}_\infty(X_2) + \mathbf{H}_\infty(R_2) + \mathbf{H}_\infty(R_1) - L_{P_1} - L_{P_2} - \delta$

The overall entropy loss in Lemma 1 applies to the general case. That is,

$$\widetilde{\mathbf{H}}_\infty(X_1, X_2|Q) \geq \mathbf{H}_\infty(X_1) + \mathbf{H}_\infty(X_2) + \mathbf{H}_\infty(R_1) + \mathbf{H}_\infty(R_2) - L_{P_1} - L_{P_2}.$$

7.2 Cascaded Structure for Multiple Secrets

In some systems, it may be desirable to use more than two secrets. For example, in a multi-factor system, a user credential may include a fingerprint, a smartcard and a PIN, or two fingerprints and a password. Unlike the two secret case, there are many different cascaded strategies to mix the secrets.

Given secrets $x_1, x_2, \cdots, x_s$ and the corresponding sketches $p_1, p_2, \cdots, p_s$, the following are the main strategies to mix them, assuming we have mixing functions $f_1, \cdots, f_{s-1}$.

1. (Fanning) Apply mixing functions f_i on x_1 and p_{i+1} for all $1 \leq 1 \leq s-1$.
2. (Chaining) Apply mixing function f_i on x_i and p_{i+1} for all $1 \leq 1 \leq s-1$.
3. (Hybrid) Use a combination of fanning, chaining and independent encoding. For example, we can mix x_1 with p_2 and p_3, and further mix x_2 with p_4, but x_5 is encoded independently.

With the fanning approach, the entropy loss would be mostly diverted to the first secret, which may be the most easily revocable and replaceable secret. However, this approach requires that the first secret has sufficiently high entropy, since otherwise it may be relatively easy to obtain the first secret from the mixed sketch. In practice, this approach can be used when a long revocable key is available, such as key stored in a smartcard.

On the other hand, using the chaining approach only requires that the entropy of the i-th secret is sufficient to mix with the $(i+1)$-th sketch. In this case, the secrets should be mixed in the order of their "importance", which could be, for example, the ease of revocation and replacement, or the likelihood of being lost or stolen. Note that in this approach, it is crucial to determine the exact order of importance of the secrets.

If no single secret is of sufficient entropy, and the order of importance among secrets is not always clear, a hybrid approach may become more appropriate. As a special case, when all secrets are short and no secret is more important than others, it would not be advisable to use the mixing approach and a straightforward method can be better.

7.3 Guidelines for Applying Mixing Functions on Two Secrets

To summarize, we give some guidelines for the application of cascaded mixing functions to two secrets. The same principles apply to multiple secrets.

1. If the importance of the secrets cannot be determined or is the same for both secrets, mixing is not recommended.

2. For the more important secret, if there are two secure sketch schemes that differ only in the amount of randomness used in the construction; choose the one that uses less randomness.

3. If the randomness invested cannot be decoupled from the sketch, cascaded mixing is not advisable unless the length of consistent key is longer than the length of the sketch.

8 Conclusions

In this paper, we investigate the security of secure sketches and fuzzy extractors that use more than one secret, motivated by the fact that user credentials based on biometric data are seldom used alone, but often combined with other secrets. Since the leftover entropy of each secret is not high and exhaustive search is feasible, we focus on information theoretic results and measure security using min-entropies.

In many practical applications that involve multiple secrets, the secrets may have different characteristics such as their revocability, ease of replacement, and likelihood of being lost or stolen. Hence, they often require different level of protections. However, such differentiation cannot be expressed easily in existing frameworks.

To cater for different security requirements for different secrets, we propose to analyze the security separately for different secrets, and we propose a cascaded mixing approach that combines the secrets when computing the final sketch. We show that under certain conditions, the proposed method provides more protections to more important secrets at the expense of increasing the risk of reduced security on the less important ones.

We show that there are scenarios where the cascaded mixing approach may not be advisable. These include cases where the sketch construction uses a lot of randomness, or the sketch contains a lot of redundancies, or it is difficult to determine the importance of secrets. We illustrate these subtleties with some examples.

We start with the case of two secrets and extend our discussions to the case of more secrets. We also give general guidelines as how these secrets should be mixed in practice.

References

1. Boyen, X.: Reusable cryptographic fuzzy extractors. In: Proceedings ACM Conf. on Computer and Communications Security, October 2004, pp. 82–91 (2004)
2. Boyen, X., Dodis, Y., Katz, J., Ostrovsky, R., Smith, A.: Secure remote authentication using biometric data. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 147–163. Springer, Heidelberg (2005)
3. Chang, E.-C., Li, Q.: Hiding secret points amidst chaff. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 59–72. Springer, Heidelberg (2006)
4. Chang, E.C., Shen, R., Teo, F.W.: Finding the original point set hidden among chaff. In: ACM Symposium on Information, computer and communications security, p. 188 (2006)
5. Clancy, T.C., Kiyavash, N., Lin, D.J.: Secure smartcard-based fingerprint authentication. In: ACM Workshop on Biometric Methods and Applications, pp. 45–52 (2003)
6. Dodis, Y., Reyzin, L., Smith, A.: Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 523–540. Springer, Heidelberg (2004)
7. Florencio, D., Herley, C.: A large-scale study of web password habits. In: Proceedings of the 16th international conference on World Wide Web, pp. 657–666. ACM, New York (2007)

8. Ho, P., Armington, J.: A dual-factor authentication system featuring speaker verification and token technology. In: Audio- and Video-Based Biometric Person Authentication, pp. 128–136 (2003)
9. Hong, S., Jeon, W., Kim, S., Won, D., Park, C.: The vulnerabilities analysis of fuzzy vault using password. In: Second International Conference on Future Generation Communication and Networking, FGCN'08, pp. 76–83 (2008)
10. Juels, A., Sudan, M.: A fuzzy vault scheme. In: IEEE Intl. Symp. on Information Theory, pp. 408–421 (2002)
11. Juels, A., Wattenberg, M.: A fuzzy commitment scheme. In: Proceedings ACM Conf. on Computer and Communications Security, pp. 28–36 (1999)
12. Kholmatov, A., Yanikoglu, B.: Realization of correlation attack against the fuzzy vault scheme. Security, Forensics, Steganography, and Watermarking of Multimedia Contents (January 2008)
13. Klein, D.V.: Foiling the cracker: A survey of, and improvements to, password security. In: 2nd USENIX Security Workshop, pp. 5–14 (1990)
14. Linnartz, J.-P.M.G., Tuyls, P.: New shielding functions to enhance privacy and prevent misuse of biometric templates. In: Kittler, J., Nixon, M.S. (eds.) AVBPA 2003. LNCS, vol. 2688, pp. 393–402. Springer, Heidelberg (2003)
15. Lisimaque, G.: Biometrics and smart cards. In: Proceedings of Conference of the Biometric Consortium (1999)
16. Monrose, F., Reiter, M., Wetzel, S.: Password hardening based on keystroke dynamics. In: Proceedings ACM Conf. Computer and Communications Security, pp. 73–82 (1999)
17. Morris, R., Thompson, K.: Password security: A case history. Communications of the ACM, 594–597 (1979)
18. Nandakumar, K., Nagar, A., Jain, A.K.: Hardening fingerprint fuzzy vault using password. In: Advances in Biometrics International Conference, August 2007, pp. 927–937 (2007)
19. Sanchez-Reillo, R.: Including biometric authentication in a smart card operating system. In: Bigun, J., Smeraldi, F. (eds.) AVBPA 2001. LNCS, vol. 2091, pp. 342–347. Springer, Heidelberg (2001)
20. Simoens, K., Tuyls, P., Preneel, B.: Privacy weaknesses in biometric sketches. In: IEEE Symposium on Security and Privacy, vol. 16. IEEE Computer Society, Los Alamitos (2009)
21. Soutar, C., Roberge, D., Stoianov, A., Gilroy, R., Kumar, B.V.K.V.: Biometric encryption. In: ICSA Guide to Cryptography (1999)
22. Sutcu, Y., Li, Q., Memon, N.: Protecting biometric templates with sketch: Theory and practice. IEEE Transactions on Information Forensics and Security, 503–512 (September 2007)
23. Tuyls, P., Akkermans, A.H.M., Kevenaar, T.A.M., Schrijen, G.J., Bazen, A.M., Veldhuis, R.N.J.: Practical biometric authentication with template protection. In: Kanade, T., Jain, A., Ratha, N.K. (eds.) AVBPA 2005. LNCS, vol. 3546, pp. 436–446. Springer, Heidelberg (2005)
24. Tuyls, P., Goseling, J.: Capacity and examples of template-protecting biometric authentication systems. In: Maltoni, D., Jain, A.K. (eds.) BioAW 2004. LNCS, vol. 3087, pp. 158–170. Springer, Heidelberg (2004)
25. Yousefi-Nooraie, R., Mortaz-Hedjri, S.: Dermatoglyphic asymmetry and hair whorl patterns in schizophrenic and bipolar patients. Psychiatry Research, 247–250 (2008)

Appendix A: Proof of Theorem 2

Proof: First, let $\mathcal{K}_{x,q} \subset \{0,1\}^{L_K}$ be the set of secret $k \in \{0,1\}^{L_K}$ such that there exists an $r \in \{0,1\}^{L_R}$ and $s \in \{0,1\}^{L_S}$ so that q can be computed from x, r, k and s. That is,

$$\mathcal{K}_{x,q} = \{k \in \{0,1\}^{L_K} | \exists r, s, f(\mathsf{Enc}(x,r),k,s) = q\}.$$

Since the key of the mixing function f is L_S-recoverable, it is clear that the cardinality $|\mathcal{K}_{x,q}|$ is no more than the number of all possible r's multiplied by 2^{L_S}, where $L_S = L_Q - L_P$. That is, $|\mathcal{K}_{x,q}| \leq 2^{L_R+L_S}$ for any x and q. Now, consider

$$\begin{aligned} A &= 2^{-\widetilde{\mathbf{H}}_\infty(X|Q)-L_R-L_S} \\ &= \sum_q \Pr[Q=q] \max_x \Pr[X=x|Q=q] 2^{-L_R-L_S} \\ &= \sum_q \max_x \Pr[X=x, Q=q] 2^{-L_R-L_S}. \end{aligned}$$

On the other hand, we have

$$B = 2^{-\widetilde{\mathbf{H}}_\infty(X,K|Q)} = \sum_q \max_{x,k} \Pr[X=x, K=k|Q=q].$$

For any $q_0 \in \{0,1\}^{L_Q}$, let us consider

$$\begin{aligned} &\max_x \Pr[X=x, Q=q_0] 2^{-L_R-L_S} \\ &= \max_x \sum_k \Pr[X=x, Q=q_0, K=k] 2^{-L_R-L_S} \\ &\leq \max_x \left(\max_k \Pr[X=x, Q=q_0, K=k] 2^{L_R+L_S} \right) 2^{-L_R-L_S} \\ &= \max_{x,k} \Pr[X=x, Q=q_0, K=k] \end{aligned}$$

The inequality holds because for any x, there will be at most $|\mathcal{K}_{x,q_0}| \leq 2^{L_R+L_S}$ non-zero terms in the summation, hence the sum will be at most $2^{L_R+L_S}$ times the largest term in the summation. As a result, we have

$$A \leq \sum_q \max_{x,k} \Pr[X=x, Q=q, K=k] = B.$$

This is equivalent to

$$\widetilde{\mathbf{H}}_\infty(X|Q) + L_R + L_S \geq \widetilde{\mathbf{H}}_\infty(X,K|Q).$$

By applying the bound on overall entropy loss (Lemma 1), and considering that the recoverable randomness includes the L_R bit R and L_S bit S, we have

$$\widetilde{\mathbf{H}}_\infty(X|Q) \geq \widetilde{\mathbf{H}}_\infty(X,K|Q) - L_R - L_S \geq \mathbf{H}_\infty(X) + \mathbf{H}_\infty(K) - L_Q$$

Therefore the theorem holds as claimed. □