2014 22nd International Conference on Pattern Recognition, ICPR. IEEE 2014.
4483 - 4488

DOI: http://dx.doi.org/10.1109/ICPR.2014.767

Protected Facial Biometric Templates Based on Local Gabor Patterns and Adaptive Bloom Filters

Marta Gomez-Barrero*, Christian Rathgeb†, Javier Galbally*, Julian Fierrez*, Christoph Busch†

*Biometric Recognition Group - ATVS, EPS, Universidad Autonoma de Madrid, Spain
{marta.barrero,javier.galbally,julian.fierrez}@uam.es

†da/sec - Biometrics and Internet Security Research Group, Hochschule Darmstadt, Germany
{christian.rathgeb,christoph.busch}@h-da.de

***Abstract*—Biometric data are considered sensitive personal data and any privacy leakage poses severe security risks. Biometric templates should hence be protected, obscuring the biometric signal in a non-reversible manner, while preserving the unprotected system's performance. In the present work, irreversible face templates based on adaptive Bloom filters are proposed. Experiments are carried out on the publicly available BioSecure DB utilizing the free Bob image processing toolbox, so that research is fully reproducible. The performance and security evaluations proof the irreversibility of the protected templates, while preserving the verification performance. Furthermore, template size is considerably reduced.**

I. Introduction

Biometrics are nowadays being introduced into diverse applications, representing an alternative to traditional knowledge- or token-based authentication mechanisms [1]. As any other security technology, biometric systems are exposed to external attacks that can compromise their security. In particular, it has been proven that it is feasible to recover the biometric trait (i.e., iris [2], face [3], handshape [4] or fingerprint [5]) from the information contained in the stored biometric reference data (template). This fact poses a severe security and privacy problem: a skilled impostor could have access to our biometric data. In order to prevent such privacy violations, biometric templates need to be protected. This is a challenging task due to the intra-subject variability of the acquired traits, as shown by recent studies on template protection schemes [6]. In accordance with the ISO/IEC IS 24745 [7] on biometric information protection, technologies of biometric template protection should be designed to meet the requirement of irreversibility: knowledge of the protected template should not allow to determine any information about the original biometric sample, while it should be easy to generate the protected template. In order to fulfill this security requirement, template protection technologies tend to obscure original biometric signals in an irreversible manner. As a consequence, the majority of published approaches to template protection report a significant decrease in recognition accuracy [6].

In the present work, we propose a new approach to obtain irreversible facial references (i.e., protected templates) based on adaptive Bloom filters, while maintaining the system recognition performance. A Bloom filter is a space-efficient probabilistic data structure representing a set in order to support membership queries [8]. Bloom filter-based representations of binary biometric templates enable a rapid biometric comparison, while a successive mapping of parts of a binary biometric template to a Bloom filter represents an irreversible transform. In [9] the applicability of adaptive Bloom filters in order to achieve alignment-free cancelable iris biometric templates has been demonstrated. While iris biometric templates are usually binary, face templates are represented with real values in the vast majority of facial recognition systems. In the proposed system, the Local Gabor Binary Pattern Histogram Sequence (LGBPHS) algorithm is utilized to extract facial features [10], [11] which are binarized and encoded in order to serve as input for adaptive Bloom filter based transforms. The experimental results show that the proposed approach generates irreversible facial templates maintaining at the same time the recognition performance of the unprotected biometric system.

The remainder of this paper is structured as follows: Sect. II reviews related work on biometric template protection for facial images. The proposed system which obtains irreversible facial biometric templates is summarized in Sect. III. Experimental evaluations are presented in Sect. IV. Final conclusions are drawn in Sect. V.

II. Related Work

With respect to face biometrics different techniques for biometric template protection, which are commonly categorized as biometric cryptosystems and cancelable biometrics [6], have been proposed. Biometric cryptosystems are designed to securely bind a digital key to a biometric or to generate a digital key from a biometric, offering solutions to biometric-dependent key-release and to biometric template protection. Cancelable biometrics consist of intentional, repeatable distortions of biometric signals based on transforms which enable the comparison of biometric templates in the transformed domain [12]. The inversion of such transformed biometric templates must not be feasible for potential impostors.

Focusing on biometric cryptosystems, Sutcu *et al.* [13] proposed a quantization scheme in which hash values are created from face samples. Features are distributed into intervals by mapping them on convolved Gaussian functions where correct intervals are concealed by adding noise in form of fake Gaussian functions. In [14], a key-binding scheme based on face is proposed applying quantization index modulation, which is originally targeted for watermarking applications. In [15], a facial template protection scheme based on helper data is presented, reporting a non-negligible performance degradation.

Ratha *et al.* [12] were the first to introduce the concept of cancelable biometrics applying non-invertible transforms. At enrolment, a non-invertible transform (e.g. surface folding) is applied to a facial image using application-dependent parameters. During authentication, biometric inputs are transformed and protected templates are compared. In [16], cryptographically secure biotokens are proposed and applied to existing recognition schemes for face, e.g. PCA (Principal Component Analysis). The key idea is to split biometric features into a stable part and an unstable part. For face, the authors suggest to simply split real feature values into an integer part and a fractional part. Subsequently, stable parts are encrypted and unstable parts are obscured applying non-invertible projections. In the vast majority of cancelable biometrics approaches, revocability is provided by incorporating additional secret tokens (e.g. random numbers). This way, performance evaluations have to be performed under the "stolen-secret scenario", where each impostor is in possession of valid secrets.

In [17] a technique applied to face biometrics called "BioHashing" was introduced. Basically, the BioHashing approach operates as a key-binding scheme, using secret user-specific tokens (unlike public helper data) at authentication Prior to the key-binding step, secret tokens are blended with biometric data to derive a distorted biometric template, i.e., BioHashing represents an instance of "Biometric Salting" [6]. In most biometric salting approaches [17], [18], subject-specific secrets are incorporated while experiments are performed under the non-stolen-secret scenario omitting the actual biometric performance of the system. In a more recent publication [19], a significant degradation of biometric performance is reported for the stolen-token scenario. Savvides *et al.* [18] generate cancelable face biometrics by applying so-called minimum average correlation filters which provide non-invertibility. User-specific secret PINs serve as seed for a random basis for the filters. In [20] user-specific random projections are applied to PCA-based face features followed by an error minimizing template transform. Again, the authors do not consider a stolen-token scenario.

III. System Architecture

In order to verify an identity claim, the proposed system follows three key steps (see Fig. 1):

A) **LGBPHS**: face image is preprocessed and Gabor-based features are extracted.
B) **Feature encoding and binarization**: histograms are encoded and binarized.
C) **Bloom filter computation**: Bloom filters are computed from the binarized features, and compared to the Bloom filter-based reference template in the database to obtain the final score.

A last critical step once the protected templates have been generated is to perform an:

D) **Irreversibility study**: given a specific Bloom filter, how many sequences can originate it? By answering this question, the irreversibility of the Bloom filter transform is assessed.

In the next subsections more detailed information on each step is provided.

A. LGBPHS System

The face verification system that served as baseline for the proposed approach is an implementation of the LGBPHS algorithm [10], a state-of-the-art system robust to illumination changes. In a fair benchmark among four state-of-the-art algorithms for face recognition established in [21], using the same databases and protocols, LGBPHS achieved a top performance. Furthermore, LGBPHS requires no training.

In order to extract features from a captured biometric sample, face images are convolved with a set of 40 Gabor filters where phase information is discarded, thus leading to 40 Gabor Magnitud Pictures (GMP). A Local Binary Pattern (LBP) operator is used to compute the LGBP map of each GMP. These are further divided into M non-overlapping blocks, from which histograms are computed and concatenated to form the final representation of a face image (see Fig. 1 left). For more details, the reader is referred to [10].

B. Feature Encoding and Binarization

The histograms computed by the LGBPHS system are re-arranged in the following manner: a rectangular matrix is derived from each block of the original image, taking into account the information provided by all 40 LGBP maps. Each row of the matrix is the histogram computed from the corresponding LGBP map (one such matrix is depicted in Fig. 1 center). Therefore, the re-arranged template comprises M histogram matrices.

In order to binarize these matrices, two approaches have been considered:

1) **Binarization scheme I**. A fixed threshold at 0 is used for all the bins: bins having non-zero values will be represented with 1 after the binarization process.
2) **Binarization scheme II**. Different thresholds are computed for each bin of each histogram from a pool of subjects (i.e., training set). The average value of each bin is computed across the subjects in the training set and used as threshold: values higher than the corresponding mean will be represented with a 1 after the binarization process.

C. Bloom Filter-based Transform and Comparison

A Bloom filter $\mathbf{b}$ is a bit array of length n, where initially all bits are set to 0. In order to represent a set $S = \{x_1, x_2, ..., x_m\}$, a Bloom filter traditionally utilizes k independent hash functions $h_1, h_2, ..., h_k$ with range $[0, n-1]$. For each element $x \in S$, bits $h_i(x)$ of Bloom filter $\mathbf{b}$ are set to 1, for $1 \leq i \leq k$. An index can be set to 1 multiple times, but only the first change has an effect. To test if an element y is in S, it has to be checked whether all position of $h_i(y)$ in $\mathbf{b}$ are set to 1. If this is the case, it is assumed that y is in S with a certain probability of false positive. If not, clearly y is not a member of S, hence, traditional Bloom filter are suitable for any application where a distinct probability of false positive is acceptable.

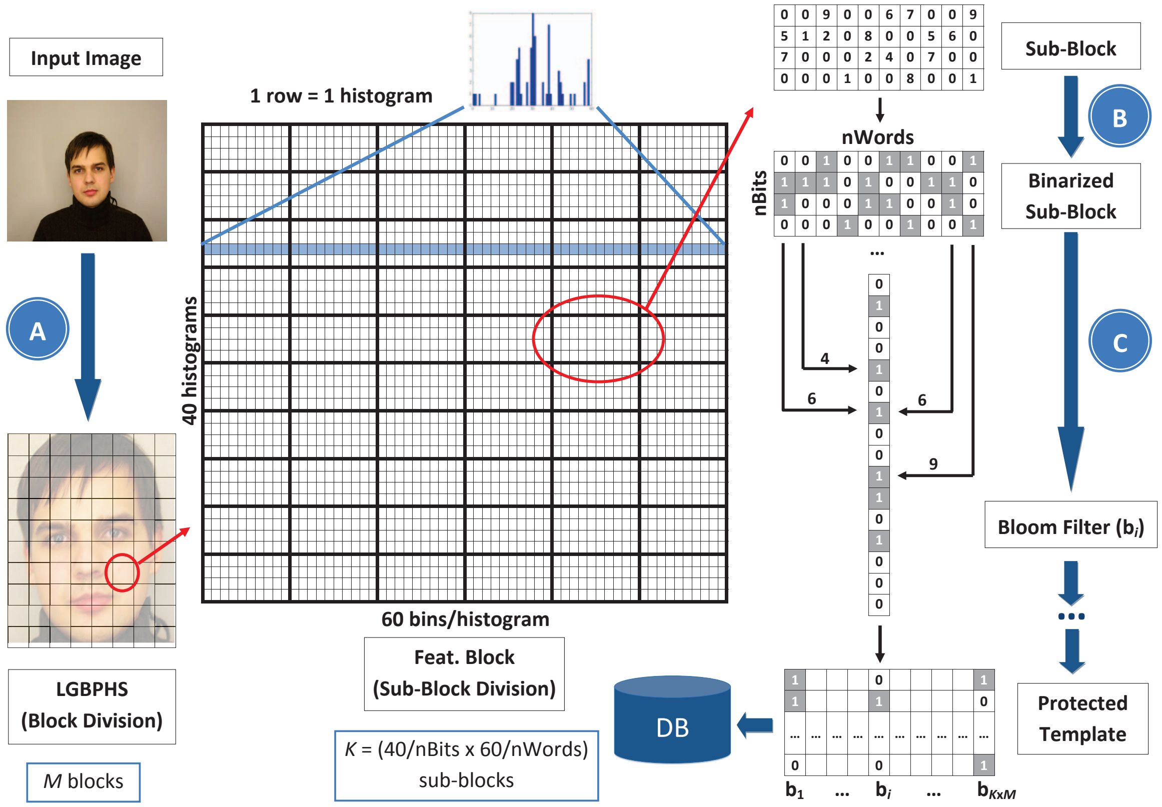

0	0	9	0	0	6	7	0	0	9
5	1	2	0	8	0	0	5	6	0
7	0	0	0	2	4	0	7	0	0
0	0	0	1	0	0	8	0	0	1

0	0	1	0	0	1	1	0	0	1
1	1	1	0	1	0	0	1	1	0
1	0	0	0	1	1	0	0	1	0
0	0	0	1	0	0	1	0	0	1

1				0					1
1				1					0
...	...	...	...	...	...	...	...	...	...
0				0					1

Fig. 1. Diagram of the processing steps undergone by the input face image: image is divided into M blocks, from which the LGBPHS are extracted (left). The 40 padded histograms are arranged in a matrix (center), further divided into sub-blocks (with K sub-blocks per block). Each sub-block in the figure is binarized and one Bloom filter is computed (right). Finally, the protected template, comprising $M \times K$ Bloom filters, to be stored in the database.

In the proposed scheme the original concept of Bloom filters is adapted in order to achieve irreversible face biometric templates. In the aforementioned encoding scheme, face biometric templates are represented as two-dimensional binary feature vectors of width W and height H. Binarized features for the whole face are then divided into M blocks and each block into K sub-blocks of equal size (see Fig. 1). Subsequently, the entire sequence of columns of each subbock, where each column consists of $nBits$ bits, is successively transformed to according locations within Bloom filters, that is, a total number of K separate Bloom filters of length $n = 2^{nBits}$ form the protected template of size $K2^{nBits}$ corresponding to one block out of M. The final template size will then be $MK2^{nBits}$. The transform is implemented by mapping columns of the 2D binary template to according indexes of their decimal value as shown for sample codewords (=columns) in Fig. 1 (right); i.e. for each column $x \in \{0,1\}^{nBits}$, the mapping is defined as,

$$\mathbf{b}[h(x)] = 1, \text{ with } h(x) = \sum_{j=0}^{nBits-1} x_j \cdot 2^j. \tag{1}$$

In the remainder of the paper these adapted versions of original Bloom filters are referred to as Bloom filters.

By applying the proposed transform the original positions of codewords are concealed; i.e. given a Bloom filter $\mathbf{b}$ it is not clear from which column a distinct 1-bit in the protected template originated. In addition, it is most likely that diverse columns are mapped to a single index and the occurrence of distinct codewords can not be established from the protected template.

Typically, the comparison between a pair of binary biometric feature vectors is carried out applying the simple XOR operator. The sum of all detected disagreements between any corresponding pairs of bits divided by the amount of compared bits yields the fractional Hamming distance (HD) as a measure of dissimilarity between pairs of binary biometric feature vectors. Let $|\mathbf{b}|$ denote the amount of bits within a Bloom filter $\mathbf{b}$ set to 1. Then the dissimilarity score DS between two Bloom filters $\mathbf{b}_i$ and $\mathbf{b}_j$ is defined as,

$$DS(\mathbf{b}_i, \mathbf{b}_j) = \frac{HD(\mathbf{b}_i, \mathbf{b}_j)}{|\mathbf{b}_i| + |\mathbf{b}_j|}. \tag{2}$$

The computation efficiency of the dissimilarity scores DS is the same as that of HD. The final verification decision is taken according to the sum of MK different DS (i.e., specific weights can be incorporated based on the local origin of extracted features).

D. Irreversibility Analysis

Within the presented scheme, irreversibility is achieved by mapping column-wise codewords to Bloom filters. Given a Bloom filter $\mathbf{b}$ of length n we restrict to inserting only $nWords$ codewords, where $nWords \leq n$ (blocks do not contain more than n columns). In case of uniformly distributed data the probability that a certain bit is set to 1 during the insertion of an element is $1/n$, i.e. the probability that a bit is still 0 is $1 - 1/n$. For inserting a total of $nWords$ elements $1 - (1 - 1/n)^{nWords}$ bits are expected to be set to 1. For $n = nWords \cdot c$ and $c \in \mathbb{N}$ (i.e. n represents a multiple of $nWords$), $\lim_{n\to\infty}(1 - 1/n)^{nWords} = 1/\mathrm{e}^{nWords/n}$. Focusing on biometric data this theoretical expectation does not apply, since bits of binary biometric feature vectors must not be expected to be mutually independent (i.e. reasonable parts of feature vectors correlate).

Consequently, a significant amount of codewords is expected to be mapped to identical positions in Bloom filters even for small values of $nWords$. Let us assume $|\mathbf{b}|$ bits are set to 1 within a Bloom filter after inserting $nWords$ codewords, i.e. $|\mathbf{b}|$ different codewords occur in a block of $nWords$. Hence, the probability of re-mapping a bit to a certain position is $1 - |\mathbf{b}|/nWords$. For a potential attacker the reconstruction of the original template part involves arranging $|\mathbf{b}|$ codewords to $nWords$ positions. For $|\mathbf{b}| \leq nWords$ the theoretical amount of possible sequences is recursively defined by $f(|\mathbf{b}|, nWords)$ where each of the $|\mathbf{b}|$ codewords have to appear at least once within $l = nWords$ columns,

$$f(|\mathbf{b}|,l) = \begin{cases} 1, & \text{if } |\mathbf{b}| = 1\ , \\ |\mathbf{b}|^l - \sum\limits_{i=1}^{|\mathbf{b}|-1} \binom{|\mathbf{b}|}{i} \cdot f(i,l) & \text{otherwise.} \end{cases} \qquad (3)$$

In other words, all sequences with less than $|\mathbf{b}|$ codewords are subtracted from the number of all possible sequences, $|\mathbf{b}|^{nWords}$. Fig. 2 illustrates the rapid increase of possible sequences even for small values of $|\mathbf{b}|$ (note the logarithmic scales of y axis). Peaks are located around $3/4 \cdot \mathrm{nWords}$, in case of $\mathrm{nWords} = |\mathbf{b}|$ we get $f(nWords, nWords) = nWords!$ and $f(1, nWords) = 1$.

For example: for $nWords = 4$ and $|\mathbf{b}| = 2$ we get $f(2,4) = 2^4 - \binom{2}{1} \cdot f(1,4) = 16 - 2 \cdot 1 = 14$ possible sequences, for $nWords = 4$ and $|b| = 3$ we get $f(3,4) = 3^4 - \binom{3}{1} \cdot f(1,4) - \binom{3}{2} \cdot f(2,4) = 81 - 3 \cdot 1 - 3 \cdot 14 = 36$ possible sequences, for $nWords = 4$ and $|\mathbf{b}| = 4$ we get $f(4,4) = 4! = 24$ possible sequences and so forth.

IV. Experimental Evaluation

The goal of the experiments is threefold: i) assess whether the described encoding approach using Bloom filters causes a loss on recognition performance, ii) analyse the template sizes of the protected and unprotected systems, and iii) study the irreversibility of the proposed scheme.

It should be noticed that a toolbox freely available online and a publicly available database are used in the experiments, thus ensuring the reproducibility of the research.

A. Experimental Setup

All the experiments are run on Bob[1] [11], a free signal and image processing toolbox available online. More specifically, the face recognition algorithm considered is implemented in the Facereclib [21], a library comprising several face verification algorithms and database interfaces, implemented over the more general Bob platform.

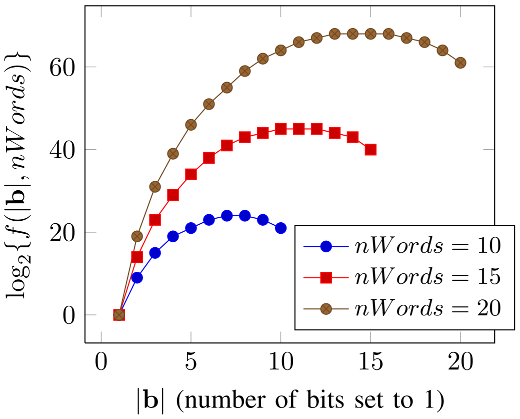

Fig. 2. Amount of possible sequences (per block) for different block sizes and proportions of re-mapped codewords.

In this particular implementation of LGBPHS, each image is divided into $M = 80$ blocks. Therefore, $80 \times 40 = 3,200$ 59-bit histograms are computed and concatenated. Prior to the binarization step, the 59-bins histograms are padded with a 0 in order to obtain 60 bins per histograms, a non-prime number that allows a further division of each block into $K = (40/nBits) \times (60/nWords)$ sub-blocks.

The experiments are carried out on the face subcorpus included in the Desktop Dataset of the Multimodal Biosecure Database[2] [22], which comprises voice, fingerprints, face, iris, signature and hand of 210 subjects, captured in two time-spaced acquisition sessions.

The face subset used in this work includes four frontal images (two per session) with an homogeneous grey background, captured with a reflex digital camera without flash ($210 \times 4 = 840$ face samples). Eyes were automatically annotated using VeriLook SDK 4.0, developed by Neurotechnology[3].

The database is divided into a training set, comprising all samples belonging to the first 10 subjects, and a test set, comprising the remaining 200 subjects. The training set is used in the second binarization scheme to estimate the bins thresholds (see Sect. III-B). Genuine and impostor scores for the performance evaluation are computed on the test set, leading to $1,200$ genuine and $238,800$ impostor comparisons.

Performance is estimated in terms of False Non-Match Rate (FNMR) at a targeted False Match Rate (FMR) and Equal Error Rate (EER). The FNMR of a biometric system defines the proportion of genuine attempt samples falsely declared not to match the template from the same subject supplying the sample. By analogy, the FMR defines the proportion of zero-effort impostor attempt samples falsely declared to match the compared non-self template. As score distributions overlap EERs are obtained (i.e., the system error rate where FNMR = FMR).

[1] Publicly available at http://idiap.github.io/bob/

[2] Publicly available at http://biosecure.it-sudparis.eu/AB

[3] http://www.neurotechnology.com/verilook.html

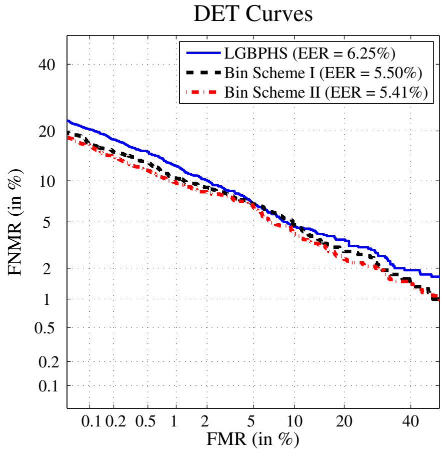

Fig. 3. DET curves for the LGBPHS system and for the best configurations found for both binarization schemes ($nBits = 4$, $nWords = 10$).

B. Performance Evaluation

First of all, the performance of the unprotected LGBPHS baseline system is evaluated, according to the protocol established in Sect. IV-A. The Detection Error Tradeoff curve (DET) is depicted in Fig. 3. The EER obtained is 6.25%. This value will be compared to the EER of the protected system in order to assess whether the template protection scheme proposed degrades the performance of the face verification system.

As described in Sect. III, the performance of the Bloom filter scheme depends on two parameters, $nBits$ and $nWords$. Exhaustive experiments determined the optimal ranges for both parameters to be $nBits \in \{2, 4, 5\}$ and $nWords \in \{10, 15, 20\}$: higher values of $nBits$ lead to template sizes that are not feasible to handle, while different values of $nWords$ result in lower verification performance.

The EER for the different configurations considered are shown in Table I. Performance is similar for both binarization schemes. In both cases, the best configuration found is $nBits = 4$ and $nWords = 10$, leading to EER = 5.50% and 5.41%, respectively. Moreover, as we may see in Fig. 3, the DET curves for the protected and unprotected systems are almost identical, showing that the proposed protection scheme maintains the performance of the original system.

It should be noted that EERs are in some cases even lower than the unprotected system EER (6.25%). This leads us to believe that the context-based comparison of bit blocks is not only a valid protection approach, but also constitutes an improved comparator with regard to the element-wise-based.

C. Template Compression

In addition to improving the performance, privacy and security of the unprotected system, Bloom filters provide smaller templates, and therefore faster verification. The original unprotected template consisted of 40 GMP × 80 blocks/GMP × 59 bins/block $= 188,800$ integer values, thus needing 184.38 KB per template. Binarizing the features divides the template size by a factor of 8 (integer values are converted to bits), resulting in templates of 23.05 KB. Finally, storage requirements are considerably reduced with the Bloom filters (BF): depending on the values of the parameters $nWords$ and $nBits$, template size is given by this formula,

		Bin Scheme I			Bin Scheme II		
		$nBits$			$nBits$		
		2	**4**	**5**	**2**	**4**	**5**
$nWords$	**10**	6.26%	**5.50%**	6.08%	5.67%	**5.41%**	5.92%
	15	6.25%	5.75%	5.75%	5.83%	5.75%	6.25%
	20	6.32%	5.75%	6.00%	5.58%	5.92%	6.07%

TABLE I. EER OF THE PROTECTED SYSTEM FOR DIFFERENT VALUES OF $nWords$ AND $nBits$ FOR BOTH BINARIZATION SCHEMES (BEST RESULTS IN BOLD).

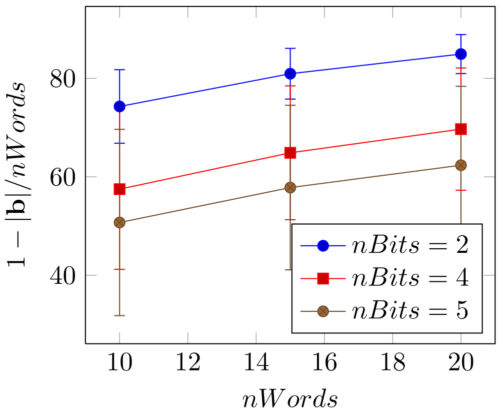

Fig. 4. Proportion of re-mapped codewords for different values of $nWords$ and $nBits$ for binarization scheme I.

$$80 \text{ blocks } \times \left(\frac{40}{nBits} \times \frac{60}{nWords} \right) \frac{\text{BF}}{\text{block}} \times 2^{nBits} \frac{\text{bits}}{\text{BF}} \quad (4)$$

Therefore, for the best performing configuration ($nWords = 10$, $nBits = 4$), templates require only 9.38 KB. Template size is thus reduced by 95% from the original integer valued features, and by 59% from its binarized form. Moreover, if the configuration with $nWords$ = 20 and $nBits = 4$ is chosen, template size is further divided by two (4.69 KB) at a small cost in terms of EER (5.75%). In that case, the compression rates rise to 97% and 80%, respectively.

D. Irreversibility Study

The security of the entire approach relies on the non-invertible mapping of codewords to Bloom filters: the transformation from binarized LGBPHS to Bloom filters obscures the number of occurrences of each codeword as well as its original position. The average percentage of re-mapped codewords (=columns) and according standard deviations for binarization scheme I is depicted in Fig. 4 (the behaviour of both binarization schemes is almost identical, as could be expected from the very similar performance shown in Table I). As expected (see Sect. III-D), two different trends may be observed, namely: i) for a fixed value of $nWords$, the bigger $nBits$, the lower the percentage of re-mapped words (i.e., less information is lost); and ii) for a fixed value of $nBits$, the bigger $nWords$, the higher the percentage of re-mapped words.

For the best performing configuration in terms of EER ($nBits = 4$, $nWords = 10$), 57.50% of the codewords are re-mapped. Therefore, on average $|\mathbf{b}| = 10 \cdot (1 - 0.575) = 4.25$. This means that an eventual brute-force attacker would have to try $\sim 2^{20}$ different sequences for each of the $K = 80 \times (10 \times 6) = 4,800$ sub-blocks (see Fig. 2, $nWords = 10$).

It was shown in Sect. III-D that the optimal re-mapping in terms of security would be $1 - |\mathbf{b}|/nWords \simeq 1 - 3/4 = 25\%$. In Fig. 4 may be observed that higher values of $nBits$ and lower values of $nWords$ will lead to percentages of re-mapped words closer to this value. However, it is shown in Fig. 2 that for higher values of $nWords$ the number of possible sequences increases. A good balance could be thus $nBits = 4$ and $nWords = 20$: security is considerably improved at a small cost in terms of performance (EER rises from 5.50% to 5.75%, see Table I), while template size is also reduced (from 9.38 KB to 4.69 KB, see Eq. 4). In this case, 69.69% of the words are re-mapped, which leads to $|\mathbf{b}| = 20 \cdot (1 - 0.697) = 6.06$ on average. As depicted in Fig. 2, an eventual attacker would thus need to try $\sim 2^{51}$ different sequences for each of the $K = 80 \times (10 \times 3) = 2,400$ sub-blocks (i.e., $\sim 2^{62}$ possible sequences).

V. Conclusion

In this work we introduced the generation of irreversible facial biometric templates based on adaptive Bloom filters. We proved that the proposed system, which builds upon a representative LGBPHS-based facial recognition system, maintains (or even improves) biometric performance, while securing the biometric information through irreversibility. Furthermore, the template size is drastically reduced and the matching process sped up, so that this scheme can be a good solution for systems with low computational capabilities such as match-on-card applications. It is also found that Bloom filter-based transforms, which have previously been applied to iris biometrics [9], represent a rather generic approach for biometric template protection which can be successfully used with different traits.

This paper is a very promising study on the feasibility of applying a Bloom filter-based protection scheme to a face verification system. Even if the binarization scheme II showed limited benefits for the configurations considered (EER slightly lower, similar irreversibility properties), it is expected to be more generic when applied to other data or configurations. As part of the future work, we will consider different face-based verification algorithms. Furthermore, in order to construct a ISO/IEC IS 24745 [7] compliant template protection system, we will additionally incorporate application-specific non-linear transformations of irreversible features to achieve unlinkability. Finally, since Bloom filters have been successfully applied to iris templates, and face and iris are commonly regarded as a user-friendly trait combination, we will investigate the feasibility of constructing a multi-biometric template protection scheme based on adaptive Bloom filters.

Acknowledgements

This work has been partially supported by projects Bio-Shield (TEC2012-34881) from Spanish MINECO, TABULA RASA (FP7-ICT-257289), FIDELITY (FP7-SEC-284862) and BEAT (FP7-SEC-284989) from EU, the Center of Applied Security Research Darmstadt (CASED) and *Cátedra UAM-Telefónica*. Marta Gomez-Barrero is supported by a FPU Fellowship from Spanish MECD. This work was conducted during a research stay at da/sec.

References

[1] A. K. Jain, A. Ross, and S. Pankanti, “Biometrics: a tool for information security,” *IEEE TIFS*, vol. 1, no. 2, pp. 125–143, 2006.

[2] J. Galbally, A. Ross *et al.*, “Iris image reconstruction from binary templates: An efficient probabilistic approach based on genetic algorithms,” *CVIU*, vol. 117, no. 10, pp. 1512–1525, 2013.

[3] A. Adler, “Sample images can be independently restored from face recognition templates,” in *Proc. CCECE*, vol. 2, 2003, pp. 1163–1166.

[4] M. Gomez-Barrero, J. Galbally *et al.*, “A novel hand reconstruction approach and its application to vulnerability assessment,” *IS*, 2013, DOI http://dx.doi.org/10.1016/j.ins.2013.06.015.

[5] R. Cappelli, D. Maio *et al.*, “Fingerprint image reconstruction from standard templates,” *IEEE TPAMI*, vol. 29, pp. 1489–1503, September 2007.

[6] C. Rathgeb and A. Uhl, “A survey on biometric cryptosystems and cancelable biometrics,” *EURASIP JIS*, vol. 3, pp. 1–25, 2011.

[7] ISO/IEC JTC1 SC27 Security Techniques, *ISO/IEC 24745:2011. Information Technology - Security Techniques - Biometric Information Protection*, International Organization for Standardization, 2011.

[8] B. Bloom, “Space/time tradeoffs in hash coding with allowable errors,” *Comm. of the ACM*, vol. 13, no. 7, pp. 422–426, 1970.

[9] C. Rathgeb, F. Breitinger, and C. Busch, “Alignment-free cancelable iris biometric templates based on adaptive bloom filters,” in *Proc. ICB*, 2013, pp. 1–8.

[10] W. Zhang, S. Shan *et al.*, “Local gabor binary pattern histogram sequence (LGBPHS): a novel non-statistical model for face representation and recognition,” in *Proc. ICCV*, vol. 1, 2005, pp. 786–791.

[11] A. Anjos, L. E. Shafey *et al.*, “Bob: a free signal processing and machine learning toolbox for researchers,” in *Proc. ACM MM*, 2012, pp. 1449–1452.

[12] N. K. Ratha, J. H. Connell, and R. M. Bolle, “Enhancing security and privacy in biometrics-based authentication systems,” *IBM SJ*, vol. 40, pp. 614–634, 2001.

[13] Y. Sutcu, H. T. Sencar, and N. Memon, “A secure biometric authentication scheme based on robust hashing,” in *Proc. WMS*. ACM, 2005, pp. 111–116.

[14] F. M. Bui, K. Martin *et al.*, “Fuzzy key binding strategies based on quantization index modulation (QIM) for biometric encryption (BE) applications,” *IEEE TIFS*, vol. 5, pp. 118–132, 2010.

[15] T. Kevenaar, G. Schrijen *et al.*, “Face recognition with renewable and privacy preserving binary templates,” in *Proc. WAIAT*, 2005.

[16] T. Boult, “Robust distance measures for face-recognition supporting revocable biometric tokens.” in *Proc. ICAFGR*, 2006, pp. 560–566.

[17] A. B. J. Teoh, D. C. L. Ngo, and A. Goh, “Personalised cryptographic key generation based on FaceHashing,” *CS*, no. 23, pp. 606–614, 2004.

[18] M. Savvides, B. Kumar, and P. Khosla, “Cancelable biometric filters for face recognition,” in *Proc. ICPR*, vol. 3, 2004, pp. 922–925.

[19] A. B. J. Teoh, Y. W. Kuan, and S. Lee, “Cancellable biometrics and annotations on biohash,” *PR*, vol. 41, no. 6, pp. 2034–2044, 2008.

[20] Y. Kim and K. Toh, “A method to enhance face biometric security,” in *Proc. BTAS*, 2007, pp. 1–6.

[21] M. Günther, R. Wallace, and S. Marcel, “An open source framework for standardized comparisons of face recognition algorithms,” in *Proc. ECCV*, ser. LNCS, vol. 7585, 2012, pp. 547–556.

[22] J. Ortega-Garcia, J. Fierrez *et al.*, “The multi-scenario multi-environment BioSecure multimodal database (BMDB),” *IEEE TPAMI*, vol. 32, pp. 1097–1111, 2010.