When a Bloom Filter is a Doom Filter: Security Assessment of a Novel Iris Biometric Template Protection System

Jens Hermans, Bart Mennink, and Roel Peeters
KU Leuven, ESAT/COSIC and iMinds
firstname.lastname@esat.kuleuven.be

Abstract: Biometric template protection systems are expected to meet two major security requirements: irreversibility and unlinkability. We analyze the Bloom filter based iris biometric template protection system recently introduced by Rathgeb et al. at ICB 2013 and IET Biometrics 2014. We demonstrate that the scheme does not achieve unlinkability, presenting a simple attack that in the worst case succeeds with probability at least 96%. We also present a security analysis on generating false positives or recovering the key, both leading to undesirably low attack complexities: 2^{25} for generating false positives for the smaller versions of the scheme, and a complexity between 2^2 and 2^8 for recovering the secret key.

1 Introduction

Security systems based on iris biometric recognition [BHF08,Dau93,RUW13] find myriad practical applications, including border control, forensics, access control, and cryptosystems [Dau09, Ros10]. At a high level, iris biometric template protection identification systems follow Daugman's principle [Dau04], which covers the entire process from the enrollment of the image of an eye via the feature extraction to the authentication phase.

Along with their broad usage in practical applications with high security, ISO/IEC IS 24745 [ISO11] prescribes two major security requirements: irreversibility and unlinkability. Irreversibility covers the case that the original iris data cannot be recovered from the transformed features, and unlinkability means that different extracts from the same iris cannot be linked, hence that they appear like mutually independent extracts. Needless to say, iris biometric template protection systems need to comply with a wide range of other security and efficiency properties, including privacy (see also Cimato et al. [CGP+09, CGP+08]), security against key recovery, a low probability of false positives (see also [ISO06]), speed, and so on. We refer to Jain et al. [JNN08] for a broad discussion of the security requirements of various biometric template protection systems.

Rathgeb et al.'s Iris Biometric Template Protection System

Biometric template protection schemes are conventionally divided [RU11] into biometric cryptosystems (such as fuzzy commitment and vault schemes [JS06, JW99] and shielding functions [LT03]) and cancelable biometrics as introduced by Ratha et al. [RCB01]. The idea of cancelable biometrics is to introduce an intentional, repeatable distortion of the biometric input using a fixed transform. Upon authentication, the input signal is transformed the same and verification is done in the transformed domain. Zuo et al. [ZRC08] introduced various techniques for cancelable iris biometrics, and further noteworthy improvements have been presented by Hämmerle-Uhl et al. [HPU09], Pillai et al. [PPCR11], and Chong et al. [CJL06a, CJL06b].

The focus in this work lies on a recently proposed cancelable iris biometric template protection scheme from Rathgeb et al. [RBB13, RBBB14], which we call RTPS throughout this work. At the heart of RTPS are Bloom filters [Blo70]. Bloom filters are randomized data structures that concisely represent a set in order to support membership queries, and nowadays find a wide range of applications [BM03]. We introduce Bloom filters and basic mathematical preliminaries in Section 2. RTPS is relatively simple, allows for high biometric data compression, and is highly efficient as it mostly relies on binary operations. The usage of Bloom filters enables irreversibility of RTPS, as can be demonstrated by basic mathematics. Rathgeb et al. [RBB13, RBBB14] claim that unlinkability follows from the usage of application-specific secrets. A formal proof of the latter claim is, unfortunately, lacking. The RTPS scheme is explained in more detail in Sect. 3.

Our Contributions

While the irreversibility argument on RTPS for uniform random data is correct (it is paraphrased in Sect. 4.1 for completeness), we observe that the scheme does not provide unlinkability. To the contrary, in Sect. 4.2 we derive a simple and highly efficient attack that, in the worst possible setting and most tolerant security model, succeeds with probability at least 96%. Most importantly, two different templates coming from the same iris always have the same Hamming weight, yet ideally they would have unrelated Hamming weights. This observation is already enough to break the unlinkability. Yet, also beyond this undesirable property we present efficient combinatorial tricks to verify whether or not two templates come from the same biometrics. While the main attack is described for the case of two protected templates derived from the same feature vector, we also show how it generalizes to the case of templates derived from different but related feature vectors in Sect. 5.

We further analyze the scheme with respect to additional properties: false positives in Sect. 4.3 and key recovery in Sect. 4.4. We derive attacks in reasonably efficient complexities, such as generating false positives for the smaller versions of RTPS in a complexity of at most 2^{25}, and recovering the key with complexity between 2^2 and 2^8. We nevertheless remark that Rathgeb et al. [RBB13, RBBB14] state that the application-specific secret is only used to provide unlinkability security of RTPS, and no other security properties are deduced of this secret.

The work is concluded in Sect. 5. In this section, we also identify the main pitfall of the system that causes itself to be insecure, and discuss attempts of salvation. Unfortunately, we remark that the attacks, and particularly the unlinkability attack, generalize to the most straightforward fix, which consists of using multiple application-specific secret values per transformation. We advocate for the usage of non-linear and non-invertible functions to derive the Bloom filters, rather than the currently used linear mapping, a strengthening which has also already been suggested in [RBBB14]. This fix will, however, naturally degrade the efficiency of the scheme.

2 Bloom Filters

We start with a brief introduction on Bloom filters, a principle dating back to 1970 [Blo70]. We refer to Broder and Mitzenmacher [BM03] for a detailed discussion on Bloom filters and their applications. Let $k, n \geq 1$, and let $h_1, \ldots, h_k$ be hash functions with range $[0, n-1]$. A Bloom filter b is a binary array of length $n \geq 1$, initialized $(0, \ldots, 0)$. To add an element v to the Bloom filter, the bits in the Bloom filter at positions $h_1(v), \ldots, h_k(v) \in [0, n-1]$ are set to 1. Likewise, to verify that an element w is in the Bloom filter, one checks if b is 1 at positions $h_1(w), \ldots, h_k(w)$.

Bloom filters allow for false positives, incorrectly suggesting an element is in the Bloom filter, but these are rather rare, if we assume the hash functions are random. In more detail, if ℓ elements are added to a Bloom filter, the probability that a certain position of b is still 0 equals:

$$\mathbb{P}\left(b \text{ is 0 at certain position}\right) = (1 - 1/n)^{k\ell},$$

and a false positive is thus triggered with probability $\left(1 - (1 - 1/n)^{k\ell}\right)^k$. Additionally, the expected number of 1's in b is

$$\mathbb{E}\left(|b|\right) = n\left(1 - (1 - 1/n)^{k\ell}\right) \approx n\left(1 - e^{-k\ell/n}\right).$$

3 Bloom Filter Based Iris Biometric Template Protection System

The iris biometric template protection system recently proposed by Rathgeb et al. [RBB13, RBBB14], which we call RTPS throughout, is a mapping that takes as input a binary matrix M of width W and height H, which is derived from an iris in some way. We disregard the generation of this matrix M from iris biometric, and refer to [RBB13, RBBB14, BHF08] for a detailed discussion on this topic. Throughout, we consider uniformly randomly generated M, unless specified otherwise. RTPS then transforms M into K Bloom filters of length 2^H, for some security parameter $K \geq 1$.[1] In more detail, the mapping RTPS operates as follows (see also Fig. 1). Firstly, the input matrix M is parsed into K submatrices of width W/K and height H (silently assuming that W/K is integral):

$$M \longrightarrow [M_1 \cdots M_K].$$

[1] The original scheme allows for Bloom filters of length $2^w \leq 2^H$, but we focus on $w = H$.

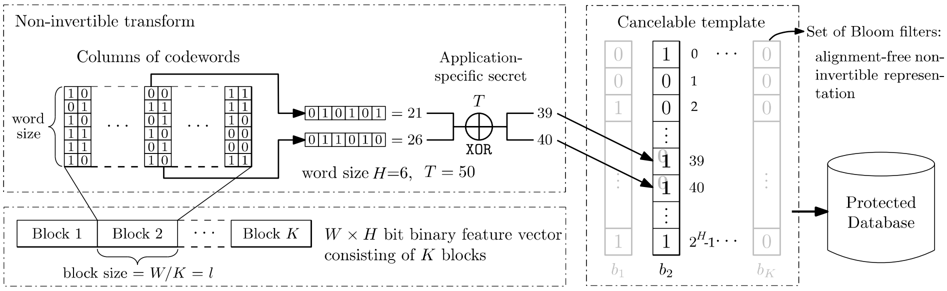

Figure 1: The iris biometric template protection system RTPS of Rathgeb et al. [RBB13, RBBB14]

For $i = 1, \ldots, K$, the submatrices M_i are then transformed to Bloom filters b_i as follows: every column $M_i[j] \in \{0,1\}^H$ (for $j = 1, \ldots, W/K$) is XORed with some predetermined application-specific secret value T, the obtained value $M_i[j] \oplus T$ is transformed to an integer in $[0, 2^H - 1]$, and the Bloom filter b_i is set to 1 at this position.

For further analysis, we briefly introduce two definitions. We define by bin2int the function that transforms an H-bit binary string to an integer in $[0, 2^H - 1]$ and by int2bin the inverse of bin2int.

Formally, RTPS employs hash function $h(v) = \mathsf{bin2int}(v \oplus T) \in [0, 2^H - 1]$, where $T \in \{0,1\}^H$ is an application-specific secret value, and applies it to all columns $M_i[j]$ of M_i. As remarked by Rathgeb et al. [RBB13, RBBB14], this secret T is used in order to provide unlinkability between multiple different templates of a single subject and it does not serve any security properties.

Verification is done the obvious way, by verifying if b_i is set to 1 at position $h(M_i[j])$ for $j = 1, \ldots, W/K$ and $i = 1, \ldots, K$.

Typical parameter sets are $W = 1024$, $H \in \{8, 9, 10\}$, and $\ell = W/K \in \{2^5, 2^6, 2^7, 2^8\}$. That is, most analysis in [RBB13,RBBB14] is done for these parameters, with best claimed performance for $H = 10$ and $\ell = K = 32$. We will stick to these parameters choices.

4 Security Assessment

We present a security analysis of RTPS. In Sect. 4.1, we elaborate on the irreversibility analysis of RTPS as presented in [RBB13, RBBB14]. Then, we present attacks on unlinkability, false positives, and key recoveries in Sects. 4.2-4.4.

4.1 Irreversibility

An irreversibility argument for RTPS for the case of uniformly random data is given in [RBB13, RBBB14], but we summarize the findings in our own terminology. Suppose that

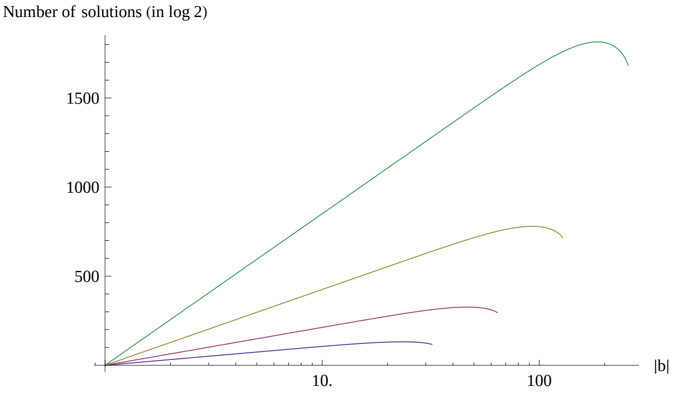

Figure 2: Number of possible matrices of width $\ell \in \{2^5, 2^6, 2^7, 2^8\}$ for given Bloom filter b, as function of $|b|$. Here, $\ell = 2^5$ is the bottom line and $\ell = 2^8$ the top line. The graphs are in $\log_2$ scale.

a Bloom filter b, after $\ell = W/K$ elements added to it, contains $|b|$ ones. Then the number of possible matrices M of width ℓ that could have lead to b is given by

$$f(|b|, \ell) = \sum_{i=1}^{|b|} (-1)^{|b|-i} \binom{|b|}{i} i^{\ell},$$

which follows from a simple application of the inclusion-exclusion principle.[2] A recursive variant of this function is given in [RBB13, RBBB14]. For various choices of $|b|$ and ℓ, this function $f(|b|, \ell)$ is plotted in Fig. 2. An irreversibility attack would consist of guessing a matrix M^* that could have lead to b, and is typically successful with probability $1/f(|b|, \ell)$. It is clear that RTPS provides irreversibility. Indeed, for $\mathbb{E}(|b|) \approx 2^H \left(1 - e^{-\ell/2^H}\right)$, this success probability equals

	$\ell = 2^5$	$\ell = 2^6$	$\ell = 2^7$	$\ell = 2^8$
$H = 8$	$2^{-124.493}$	$2^{-317.698}$	$2^{-777.341}$	$2^{-1805.82}$
$H = 9$	$2^{-121.617}$	$2^{-311.085}$	$2^{-763.644}$	$2^{-1809.62}$
$H = 10$	$2^{-117.663}$	$2^{-304.911}$	$2^{-748.228}$	$2^{-1779.99}$

We remark that these computations only hold for uniformly randomly generated data M. This is usually not the case as correlation may occur among the data, and this may make the Bloom filters reversible.

[2] In more detail, given the $|b|$ positions, there are $|b|^{\ell}$ possible matrices whose columns set one only at these $|b|$ positions, but this includes matrices that set one at only $|b| - 1$ of the positions. By the inclusion-exclusion principle, we have to subtract these $\binom{|b|}{|b|-1}(|b| - 1)^{\ell}$, and proceed similarly for $i = |b| - 2, \ldots, 1$.

4.2 Unlinkability

Unlinkability means that templates from a single object look indistinguishable from each other. Rathgeb et al. [RBB13, RBBB14] claim that unlinkability is provided by incorporating an application specific bit vector $T \in \{0,1\}^H$, although no proof is given. We demonstrate that RTPS does *not* provide unlinkability, by presenting a practical attack that matches two templates derived from a single subject using different secret bit vectors. We first consider the general attack for the case of one matrix block. Then, we elaborate on the case of $K > 1$ matrix blocks.

The idea of unlinkability is that two Bloom filters b, b' derived from the same data (using a different secret) should appear mutually independent. In other words, we consider the case where an adversary is given either these two filters b, b', or two completely random and mutually independent filters c, c', and its goal is to distinguish the two settings. Before proceeding, we remark that our attacks do not imply a distinguishability attack in case we consider two protected templates created out of two different biometric feature vectors. Nevertheless, the attacks generalize to the case of two related (yet not entirely the same) biometric feature vectors. We elaborate on this setting in Sect. 5.

One Block

Let M be an arbitrary matrix of width $\ell = W/K$ and height H. Let $T, T' \in \{0,1\}^H$ be two independently and uniformly randomly generated secret values. Denote by b the Bloom filter of M under secret T and similarly by b' the Bloom filter under T'.

The attack relies on basic mathematics and combinatorial tricks, and we start with a high-level intuition of it. As first observation, we point out that b and b' always have the same Hamming weight, while this would ideally not be the case. In more detail, two completely random Bloom filters of size 2^H for $H \in \{8, 9, 10\}$ have the same Hamming weight with probability at most 0.04 (formal computation below). We additionally demonstrate that even if we are comparing (b, b') with two random Bloom filters (c, c') of the same Hamming weight, there exist efficient combinatorial tricks to verify whether or not two templates come from the same biometrics. These tricks in essence use that every index i for which b is non-zero corresponds to a column $M[j]$ of M such that $\mathsf{bin2int}(M[j] \oplus T) = i$. The same holds for b'. Then, if $|b| = |b'|$ is even, an XOR of all indices for which b is set equals an XOR of all indices for which b' is set (as an XOR of an even number of T's and an even number of T''s cancel out). For the case of odd $|b| = |b'|$ a slightly more involved but similar approach is taken. In the remainder of this section, we present the formal mathematics behind these attacks.

More formally, the attack relies on the following simple but important observation. For any $j, j' \in [1, \ell]$:

$$M[j] \oplus T = M[j'] \oplus T \iff M[j] = M[j'] \iff M[j] \oplus T' = M[j'] \oplus T',$$

which means that collisions in b occur if and only if they occur in b' if and only if they occur in M in the first place. Define by $I, I' \subseteq [0, 2^H - 1]$ the index sets of all positions at which

b (resp. b') is 1. Let $J \subseteq [1, \ell]$ be the maximal set of indices such that $\{M[j] \mid j \in J\}$ contains no duplicate elements. By above observation, $|I| = |I'| = |J| =: \alpha$. This already allows for a linkability attack: two truly random Bloom filters c, c' would not satisfy this condition in the first place, except with probability at most

$$\begin{aligned} \mathbb{P}\left(|c| = |c'|\right) &= \sum_{i=0}^{2^H} \mathbb{P}\left(|c| = |c'| = i\right) = \sum_{i=0}^{2^H} \mathbb{P}\left(|c| = i\right) \cdot \mathbb{P}\left(|c| = i\right) \\ &= \sum_{i=0}^{2^H} \left(\binom{2^H}{i} \left(\frac{1}{2}\right)^{2^H} \right)^2 = \binom{2^{H+1}}{2^H} \left(\frac{1}{2}\right)^{2^{H+1}} . \end{aligned} \tag{1}$$

This term is at most 0.04 for $H \in \{8, 9, 10\}$. We nevertheless proceed, assuming b, b' are compared with two filters c, c' of the same weight.

First assume α is even. Then,

$$\bigoplus_{i \in I} \mathsf{int2bin}(i) = \bigoplus_{j \in J} M[j] \oplus T = \bigoplus_{j \in J} M[j] \oplus T' = \bigoplus_{i \in I'} \mathsf{int2bin}(i), \tag{2}$$

where the first and third step are by construction and the middle step as $\alpha = |J|$ is even and thus $\bigoplus_{j \in J} T = 0 = \bigoplus_{j \in J} T'$. If α is odd, (2) does not exactly hold and a slightly more elaborate analysis is needed. Let $i^* \in I$ be an arbitrarily chosen index. By construction, there exist $i'^* \in I'$ and $j^* \in J$ such that

$$\bigoplus_{i \in I \setminus \{i^*\}} \mathsf{int2bin}(i) = \bigoplus_{j \in J \setminus \{j^*\}} M[j] \oplus T = \bigoplus_{j \in J \setminus \{j^*\}} M[j] \oplus T' = \bigoplus_{i \in I' \setminus \{i'^*\}} \mathsf{int2bin}(i).$$

Unfortunately, these values i'^*, j^* are unknown. Yet, there are only α possible values i'^* and we have:

$$\bigoplus_{i \in I \setminus \{i^*\}} \mathsf{int2bin}(i) \in \left\{ \bigoplus_{i \in I' \setminus \{i'^*\}} \mathsf{int2bin}(i) \;\middle|\; i'^* \in I' \right\}. \tag{3}$$

In other words, for two Bloom filters b, b' derived from the same M, either (2) or (3) holds. Two truly random Bloom filters c, c' would set (2) with probability $1/2^H$ and (3) with probability $\alpha/2^H \leq \ell/2^H$. (We remark that this bound is rather loose, as many collisions may occur once $\ell \geq 2^{H/2}$. In more detail, we have $\mathbb{E}\left(\alpha/2^H\right) \approx 1 - e^{-\ell/2^H}$ (cf. Sect. 2).) Recall that random Bloom filters c, c' would not satisfy $|c| = |c'|$ in the first place, except with the probability computed in (1). Combining these observations, the linkability attack is successful with probability at least

$$1 - \frac{\ell}{2^H} \mathbb{P}\left(|c| = |c'|\right) = 1 - \frac{\ell}{2^H} \binom{2^{H+1}}{2^H} \left(\frac{1}{2}\right)^{2^{H+1}} ,$$

which achieves its minimum for $H = 8$ and $\ell = 2^8$ at 0.964755. The advantage only increases for higher values of H and smaller values of ℓ. This means that our attack succeeds with a probability of more than 96%.

K Blocks

In RTPS, the matrix M is first parsed into K submatrices $[M_1 \cdots M_K]$. The above-mentioned unlinkability attack can be applied to all of these blocks, leading to a success with probability of at least

$$1 - \left(\frac{\ell}{2^H} \binom{2^{H+1}}{2^H} \left(\frac{1}{2} \right)^{2^{H+1}} \right)^K ,$$

which achieves its minimum for $H = 8$ and $\ell = W/K = 2^8$ at 0.999998.

We remark that in the general case of $K > 1$ blocks the protected templates show another undesirable feature [Rat14]. Denote by $b = [b_1 \cdots b_K]$ the Bloom filters of M under secret T and similarly by $b' = [b'_1 \cdots b'_K]$ the Bloom filters under T'. Then, every row of b appears as a row of b' and vice versa. Formally, $b' = P \cdot b$ for some permutation matrix P of size 2^H. This is, in fact, a generalization of the above-mentioned observation for $K = 1$, and also allows an adversary to easily link b and b' with probability close to 1.

4.3 False Positives

We consider the probability of an adversary to generate a false positive for the scheme, i.e., to generate an input that is incorrectly viewed as a legitimate input. We remark, as we will also elaborate on later, that the probability of generating a false positive equals the expected false accept ratio.

In Sect. 2 we computed the probability of a false positive, provided the Bloom filters are generated based on uniformly randomly generated data M. However, if v gives a false positive for Bloom filter b_i, then the matrix M_i^* that consists of W/K repetitions of v gives a false positive. Admittedly, M_i^* does not look like a legitimate block from an iris, but is does not need to be: an adversary may spoof the system in any way. Based on this observation, we note that a matrix M^* consisting of W repetitions of a randomly chosen vector v results in a false positive for RTPS with probability

$$\mathbb{P}_{\text{fp}} := \left(1 - \left(1 - 1/2^H\right)^{W/K}\right)^K .$$

We remark that $\mathbb{P}_{\text{fp}}$ is equal to the expected false accept rate. In more detail, the false accept rate is the number of successful attempts divided by the number of attempts. If the adversary makes X random attempts, the expected number of successful attempts equals $X\mathbb{P}_{\text{fp}}$, and hence

$$\mathbb{E}\left(\text{false accept rate}\right) = \frac{X\mathbb{P}_{\text{fp}}}{X} = \mathbb{P}_{\text{fp}} .$$

For $W = 1024$ and the various choices of H and K, this value equals:

	$K = 2^2$	$K = 2^3$	$K = 2^4$	$K = 2^5$
$H = 8$	$2^{-2.64035}$	$2^{-10.748}$	$2^{-34.7856}$	$2^{-98.7706}$
$H = 9$	$2^{-5.37836}$	$2^{-17.4027}$	$2^{-49.4065}$	$2^{-129.391}$
$H = 10$	$2^{-8.70385}$	$2^{-24.7085}$	$2^{-64.7067}$	$2^{-160.697}$

While the values for $K = 2^5$ are adequate and meet current standards, the remaining configurations yield questionable (in case $H = 10, K = 2^4$) to non-sufficient (the remaining cases) security levels.

We stress that this computation holds for the case of uniformly random input data. In case the matrix M is not entirely random, or more detailed if columns of M can be guessed with a probability higher than usual, the success probability of a false positive increases drastically. As an example, suppose one single vector v_i of a submatrix M_i is leaked. Clearly, a false positive for b_i is generated with probability one (just input a matrix consisting of ℓ repetitions of v_i). What is more, correlations among the submatrices render a significant increase in the construction of false positives for the remaining submatrices. This is particularly perilous as data obtained from an iris shows high correlations between neighboring columns [Dau03,Dau04,Dau06,HBF09]. We refer to [VS11] for a more detailed treatment on how to generate false successful iris textures from an original iris texture.

4.4 Key Recovery

As a bonus, we consider the possibility to recover the secret value T, given the input data M and output Bloom filters b. We stress that Rathgeb et al. [RBB13, RBBB14] state that the application-specific secret is only used to provide unlinkability security of RTPS, and no other security properties are deduced of this secret.

Without loss of generality, we discuss the case of one block only. We will present a naive guessing attack and a more sophisticated attack, and both rely on basic probability theory. Let M denote a uniformly randomly generated matrix of width ℓ and height H, and denote its corresponding Bloom filter by b. Our goal is to recover secret key $T \in \{0,1\}^H$, given M and b. Inheriting notation of Sect. 4.2, denote by $I \subseteq [0, 2^H - 1]$ the index set of all positions at which b is 1, and let $J \subseteq [1, \ell]$ be the maximal set of indices such that $\{M[j] \mid j \in J\}$ contains no duplicate. Again, we have $|I| = |J| =: \alpha$. By construction, for every $i \in I$ there is a unique $j \in J$ such that

$$\mathsf{int2bin}(i) = M[j] \oplus T. \qquad (4)$$

Hence, any choice $(i^*, j^*) \in I \times J$ satisfies (4) with probability $1/\alpha \geq 1/\ell$, in which case it leads to a key recovery. For the proposed parameter choices of $\ell \in \{2^5, 2^6, 2^7, 2^8\}$, this implies that the secret value T can be recovered with probability ranging between $1/2^5$ and $1/2^8$.

The more sophisticated attack consists of smartly verifying links between I and J. Fix arbitrary distinct i, i', and write $Z = \mathsf{int2bin}(i) \oplus \mathsf{int2bin}(i')$. Denote by JJ the set of all pairs $(j, j') \subseteq J$ such that $M[j] \oplus M[j'] = Z$. By basic probability we have $\mathbb{E}(|JJ|) = \binom{\alpha}{2}/2^H$, and Markov's inequality states that, for any $A \in \{1, \dots, \ell/2\}$,

$$\mathbb{P}(|JJ| \leq A) \geq 1 - \mathbb{E}(|JJ|)/A.$$

We proceed with the key recovery attack. The trick we will use is that the couple (i, i') corresponds to exactly one $(j, j') \in JJ$, in which case i corresponds to either $M[j]$ or $M[j']$ (and i' to the other one). Formally:

$$\mathsf{int2bin}(i) \oplus T \in \Big\{ M[j], M[j'] \;\Big|\; (j, j') \in JJ \Big\}.$$

A key recovery consists of selecting any j^* among the $2|JJ|$ possibilities, and guessing $T^* = \mathsf{int2bin}(i) \oplus M[j^*]$. We find:

$$\begin{aligned} \mathbb{P}(T^* = T) &\geq \mathbb{P}\big(T^* = T \;\big|\; |JJ| \leq 2\mathbb{E}(|JJ|)\big) \cdot \mathbb{P}(|JJ| \leq 2\mathbb{E}(|JJ|)) \\ &\geq \frac{1}{4\mathbb{E}(|JJ|)} \cdot \frac{1}{2} \geq \frac{2^H}{4\ell^2}. \end{aligned}$$

This attack improves over the naive one as long as $\ell \leq 2^{H-2}$.

5 Conclusions

We presented a security analysis of the recently proposed iris biometric template protection system of Rathgeb et al. [RBB13, RBBB14]. While on the one hand we reconfirm Rathgeb et al.'s irreversibility security analysis for uniformly random data, we debunk the unlinkability claim by presenting a practical attack that distinguishes two Bloom filters b, b' generated from the same data from two independent ones c, c' with a probability of at least 96%. We additionally analyzed adversarial success probabilities in generating false positives and in key recoveries, leading to undesirably low attack complexities: 2^{25} for generating false positives for the smaller versions of the scheme, and a complexity between 2^2 and 2^8 for recovering the secret key.

The weaknesses are mainly caused by the fact that RTPS uses only one hash function and that it is a very simple one. At first sight, a possible solution lies in employing two hash functions based on different secret values T_1 and T_2. While this would, indeed, be a countermeasure against the attacks of Sect. 4, we remark that the linkability attack would still persist, be it as a slightly more elaborate combinatorial exercise. Here, the trick is to observe that although $|I|$ and $|I'|$ increase and are not necessarily the same, $|J|$ remains unchanged. The procedure of Sect. 4.2 should then be applied on all subsets of I of size $|J|$. A similar reasoning applies to the case two Bloom filters b, b' are derived from two different but related feature vectors M, M'. For this, assume M, M' are the same at $\ell' < \ell$ columns. Identify the set J as before, but then for the ℓ' columns only. Then, the procedure of Sect. 4.2 should similarly be applied on all subsets of I of size $|J|$.

A possible fix to salvage RTPS, which has also been suggested in [RBBB14], is the usage of non-linear and non-invertible functions to derive the Bloom filters, instead of the linear mapping currently employed, but this will degrade the efficiency of the scheme.

ACKNOWLEDGMENTS. This work was supported by the European Commission through the FIDELITY EU-FP7 project (Grant No. SEC-2011-284862) and the Research Council KU Leuven: GOA TENSE (GOA/11/007).

References

[BHF08] Kevin W. Bowyer, Karen Hollingsworth, and Patrick J. Flynn. Image Understanding for Iris Biometrics: A Survey. *Comput. Vis. Image Underst.*, 110(2):281–307, May 2008.

[Blo70] Burton H. Bloom. Space/Time Trade-offs in Hash Coding with Allowable Errors. *Commun. ACM*, 13(7):422–426, July 1970.

[BM03] Andrei Z. Broder and Michael Mitzenmacher. Network Applications of Bloom Filters: A Survey. *Internet Mathematics*, 1(4):485–509, 2003.

[CGP$^+$08] Stelvio Cimato, Marco Gamassi, Vincenzo Piuri, Roberto Sassi, and Fabio Scotti. Privacy-Aware Biometrics: Design and Implementation of a Multimodal Verification System. In *ACSAC*, pages 130–139. IEEE Computer Society, 2008.

[CGP$^+$09] Stelvio Cimato, Marco Gamassi, Vincenzo Piuri, Roberto Sassi, and Fabio Scotti. *Privacy in Biometrics*, pages 633–654. John Wiley & Sons, Inc., 2009.

[CJL06a] Siew Chin Chong, Andrew Teoh Beng Jin, and David Ngo Chek Ling. High security Iris verification system based on random secret integration. *Computer Vision and Image Understanding*, 102(2):169–177, 2006.

[CJL06b] Siew Chin Chong, Andrew Teoh Beng Jin, and David Ngo Chek Ling. Iris Authentication Using Privatized Advanced Correlation Filter. In David Zhang and Anil K. Jain, editors, *International Conference on Biometrics – ICB*, volume 3832 of *Lecture Notes in Computer Science*, pages 382–388. Springer, 2006.

[Dau93] John Daugman. High confidence visual recognition of persons by a test of statistical independence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11):1148–1161, Nov 1993.

[Dau03] John Daugman. The importance of being random: statistical principles of iris recognition. *Pattern Recognition*, 36(2):279–291, 2003.

[Dau04] John Daugman. How iris recognition works. *IEEE Trans. Circuits Syst. Video Techn.*, 14(1):21–30, 2004.

[Dau06] John Daugman. Probing the Uniqueness and Randomness of IrisCodes: Results From 200 Billion Iris Pair Comparisons. *Proceedings of the IEEE*, 94(11):1927–1935, Nov 2006.

[Dau09] John Daugman. Iris Recognition at Airports and Border-Crossings. In StanZ. Li and Anil Jain, editors, *Encyclopedia of Biometrics*, pages 819–825. Springer US, 2009.

[HBF09] Karen P. Hollingsworth, Kevin W. Bowyer, and Patrick J. Flynn. The Best Bits in an Iris Code. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(6):964–973, June 2009.

[HPU09] Jutta Hämmerle-Uhl, Elias Pschernig, and Andreas Uhl. Cancelable Iris Biometrics Using Block Re-mapping and Image Warping. In Pierangela Samarati, Moti Yung, Fabio Martinelli, and Claudio Agostino Ardagna, editors, *Information Security – ISC*, volume 5735 of *Lecture Notes in Computer Science*, pages 135–142. Springer, 2009.

[ISO06] ISO/IEC 19795-1:2006. Information technology – Biometric performance testing and reporting – Part 1: Principles and framework, 2006.

[ISO11] ISO/IEC 24745:2011. Information technology – Security techniques – Biometric information protection, 2011.

[JNN08] Anil Jain, Karthik Nandakumar, and Abhishek Nagar. Biometric Template Security. *EURASIP Journal on Advances in Signal Processing*, 2008(1):579416, 2008.

[JS06] Ari Juels and Madhu Sudan. A Fuzzy Vault Scheme. *Des. Codes Cryptography*, 38(2):237–257, 2006.

[JW99] Ari Juels and Martin Wattenberg. A Fuzzy Commitment Scheme. In Juzar Motiwalla and Gene Tsudik, editors, *ACM Conference on Computer and Communications Security*, pages 28–36. ACM, 1999.

[LT03] Jean-Paul M. G. Linnartz and Pim Tuyls. New Shielding Functions to Enhance Privacy and Prevent Misuse of Biometric Templates. In Josef Kittler and Mark S. Nixon, editors, *Audio-and Video-Based Biometrie Person Authentication – AVBPA*, volume 2688 of *Lecture Notes in Computer Science*, pages 393–402. Springer, 2003.

[PPCR11] Jaishanker K. Pillai, Vishal M. Patel, Rama Chellappa, and Nalini K. Ratha. Secure and Robust Iris Recognition Using Random Projections and Sparse Representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(9):1877–1893, Sept 2011.

[Rat14] Christian Rathgeb. Personal communication, 2014.

[RBB13] Christian Rathgeb, Frank Breitinger, and Christoph Busch. Alignment-free cancelable iris biometric templates based on adaptive bloom filters. In Julian Fiérrez, Ajay Kumar, Mayank Vatsa, Raymond N. J. Veldhuis, and Javier Ortega-Garcia, editors, *International Conference on Biometrics – ICB*, pages 1–8. IEEE, 2013.

[RBBB14] Christian Rathgeb, Frank Breitinger, Christoph Busch, and Harald Baier. On application of bloom filters to iris biometrics. *IET Biometrics*, 2014. To appear.

[RCB01] Nalini K. Ratha, Jonathan H. Connell, and Ruud M. Bolle. Enhancing Security and Privacy in Biometrics-based Authentication Systems. *IBM Syst. J.*, 40(3):614–634, March 2001.

[Ros10] Arun Ross. Iris Recognition: The Path Forward. *Computer*, 43(2):30–35, Feb 2010.

[RU11] Christian Rathgeb and Andreas Uhl. A survey on biometric cryptosystems and cancelable biometrics. *EURASIP Journal on Information Security*, 2011(1):3, 2011.

[RUW13] Christian Rathgeb, Andreas Uhl, and Peter Wild. *Iris Biometrics - From Segmentation to Template Security*, volume 59 of *Advances in Information Security*. Springer, 2013.

[VS11] Shreyas Venugopalan and Marios Savvides. How to Generate Spoofed Irises From an Iris Code Template. *IEEE Transactions on Information Forensics and Security*, 6(2):385–395, June 2011.

[ZRC08] Jinyu Zuo, Nalini K. Ratha, and Jonathan H. Connell. Cancelable iris biometric. In *International Conference on Pattern Recognition – ICPR*, pages 1–4. IEEE, 2008.