Privacy by Design Solutions for Biometric One-to-Many Identification Systems

IPC Technical Report

June 2014

Ann Cavoukian, Ph.D.
Information and Privacy Commissioner
Ontario, Canada

Alex Stoianov, Ph.D.
Senior Policy Specialist –
Surveillance, Biometrics, IT Security

Acknowledgements

Our deepest thanks go to Dr. George Tomko, whose invention of Biometric Encryption and early work inspired this paper and formed the basis of the hybrid OLG System.

Information and Privacy Commissioner
Ontario, Canada

2 Bloor Street East
Suite 1400
Toronto, Ontario
M4W 1A8
Canada

416-326-3333
1-800-387-0073
Fax: 416-325-9195
TTY (Teletypewriter): 416-325-7539
Website: www.ipc.on.ca
Privacy by Design: www.privacybydesign.ca

Privacy by Design Solutions for Biometric One-to-Many Identification Systems

TABLE OF CONTENTS

1. Introduction

Over the last two decades there has been an increase in the interest in, and uptake of, automated biometric systems. Biometrics are now commonly being integrated into a range of large and complex information communication technology systems and processes, and access to this data is becoming virtual rather than physical. We see the use of contemporary biometric systems being implemented throughout the world in areas such as national ID systems, border security control, crime prevention, fraud detection, forensics, war zone applications, attendance recording, access control and financial transactions.

These advanced automated systems use a scanner to take a biometric sample or what is known as a digital image from an individual during enrolment. Data are then extracted from the sample image to create a unique biometric template. The biometric data, either in image form or the template or both, can then be stored centrally in a database or in a distributed environment, such as a smart card. The biometric data can now serve to either verify or identify an enrolled individual.

Verification involves a "one-to-one" match where an individual presents a "live" biometric which is matched to a stored image or biometric template. The matching of the live biometric to the stored biometric data is all that is necessary to authenticate or validate the claimed identity of the individual. Since the individual is present, there is greater control over the use of the biometric data and increased security, since there is no need for additional identifying information to be attached to the stored biometric data. Moreover, the biometric data can be stored or even processed on a smart card or other device that is in the user's possession.

Identification, on the other hand, refers to the ability of a system to uniquely distinguish an individual from a larger set of centrally stored biometric data or what is often referred to as a one-to-many match. Identification systems require storing large amounts of data and, in general, are more prone to errors.

There are several main types of biometric one-to-many systems:

- *De-duplication*, or *multiple enrolment ("double dipping") prevention*: the most notable example is a national ID system. Such a system is usually mandatory for everyone. After the registration, the acquired biometric samples are run, either in real time or not, against the entire database. At the end, very few true hits are expected. The system is characterized by high throughput, and the database size varies from large to extra-large.

- *Watch list*: A relatively small percentage of population (e.g., terror suspects or soccer hooligans) is registered. However, a biometric sample of virtually every visitor is run against the database in real time. Very few true hits are expected. The database size is small, but the throughput rate can be very high. The system must be able to work in an unsupervised and non-cooperative environment.

- *Access control*: Each user is registered. During authentication, the user does not claim any identity. Instead, a live biometric sample is run in real time against the entire database. At the end, one true hit is always expected for

everyone. The database size varies from small to medium, and the throughput is from low to medium. The users are usually cooperative.

- *Suspects/strangers identification*: It has been a subject of forensic analysis for a long time when a suspect's fingerprint was run against a very large database of Automated Fingerprint Identification System (AFIS). These days, the same has become feasible with facial images obtained by surveillance cameras or even with mobile devices. The databases of facial images are not necessarily under the law enforcement control; they may be, for example, photos from the drivers' licence databases or social networks. This makes identification of a complete stranger in a crowd possible.

De-duplication and especially watch list and suspects/strangers identification systems are often referred to as "open sets", while access control belongs to "closed set" systems. It is obvious that all those types of one-to-many systems are very different in their objectives, design principles, privacy and security requirements, etc. For example, the main objective of a one-to-many access control system is convenience (the users do not have to carry ID cards or tokens) and overall high speed of authentication. The privacy concerns are less serious than for other types of one-to-many systems. On the other hand, the privacy implications of new developments using facial recognition in the crowd are far-reaching and not the subject of the present paper. In the following sections, we will focus primarily on state or private sector-led efforts to create a one-to-many system for a specific purpose.

As biometric uses and databases grow, so do concerns that the personal data collected will not be used in reasonable and accountable ways, especially given the revelations about the state surveillance programs [1]. The threat to privacy arises not from the positive identification that biometrics provide best, but from the issues related to informational privacy rights that include potential data misuse, function creep, linkage of databases via biometric templates, that make surveillance, profiling and discrimination without the knowledge of the individual, all possible. Biometric data transmitted across networks and stored in various databases by others can also be stolen, copied, or otherwise misused in ways that can materially affect the individual involved, such as identity theft or fraud. Moreover, unlike passwords, biometric data are unique, permanent and therefore irrevocable.

Unlike one-to-one systems, biometric one-to-many systems always imply a central database. The privacy risks here are, in general, more serious than for distributed one-to-one systems. Consequently, the global privacy and data protection community have consistently argued against the use of biometrics for most 1:many identification purposes and against the creation of large centralized or interoperable databases of biometric data. The examples include the opinion of the Dutch Data Protection Authority on the passing of the new Dutch Passport Act of 2009 that would regulate the inclusion of biometric identifiers in the passport chip, also the establishment of a central fingerprint database and the opposition to a proposed Biometric Documents Identification Law of 2009 in Israel that would enable the government to set up a database that would include biometric identification on every Israeli citizen [2, 3, 4]. In 2011, the largest biometric project in the world – India's Unique ID – suffered a serious, albeit temporary, backlash [5]. Also, there is an opposition in Argentina against the President's decision to integrate the biometric (fingerprints and face) national ID database with the law

enforcement biometric database to create a new centralized, nationwide biometric ID service (SIBIOS) [6].

However, it would be short-sighted, at best, for the data protection community to reject all biometrics, including 1:many, across the board, as being privacy-invasive. We are witnessing a world wide growth of government, law enforcement and business applications. Even though a 1:1 biometric verification is indeed preferred from a privacy perspective, there are important applications where 1:many searches are inevitable. The examples include a watch list scenario or a de-duplication system.

We need to change the zero-sum paradigm by realizing that the same technology that serves to threaten or erode our privacy may also be enlisted to assure its protection. This entails the use of Privacy by Design (PbD) – embedding privacy directly into technologies and business practices, resulting in privacy and functionality through a "positive-sum" paradigm [7]. This paper describes the ways in which a potentially privacy-invasive biometric 1:many system can be transformed to enhance privacy without diminishing functionality.

The objective of this paper is not to provide a comprehensive review of literature, but rather to explore and discuss new ideas and solutions that, in our view, can lead to deployment of privacy-protective and secure biometric one-to-many systems. In such systems, privacy is designed as a core functionality; that is, privacy protection relies not only on regulatory measures but is embedded into the system on a technological level.

The paper is organized as follows. We begin with a brief overview of the *Privacy by Design* approach as a viable privacy-protective alternative to conventional biometrics (Section 2). The next three sections deal with the following groups of privacy-protective technologies: *Biometric Encryption* (Section 3), *Cryptographically Secure Biometric Architectures* (Section 4), and *Biometric Setbase/Weak Links* (Section 5). We discuss both advances and technological challenges for each group of technologies. As a real-life example, we present a case study of the first Biometric Encryption application using facial recognition in a watch list scenario known to be the largest successful deployment in a casino and gaming context (Subsection 3.2). We also propose a cryptographically secure architecture for one-to-many system using Blum-Goldwasser cryptosystem (Subsection 4.3). Finally (Section 6), we show that most of these solutions are complementary to each other with a high degree of synergy and, by combining them in application-specific context, it is possible to create a one-to-many system which is fully compliant with *Privacy by Design* principles.

2. Biometric one-to-many systems and *Privacy by Design*

In today's world, privacy is often "trumped" by more pressing social, legal, and economic imperatives because adding privacy to the system means subtracting something else. This is the classic "zero-sum" thinking [8].

The right to control the collection, use and disclosure of information about oneself is an essential foundation upon which free societies are built [9, 10, 11, 12]. Personal information, be it biographical, biological, genealogical, historical, transactional, locational, relational, computational, vocational or reputational, is what comprises our modern identity.

Privacy by Design was developed to address the ever-growing and systemic effects of information systems, and of large-scale networked infrastructure, in a comprehensive manner. It advances the view that the future of privacy cannot be assured solely by compliance with legislation and regulatory frameworks [13, 14]; rather, privacy assurance must ideally become an organization's default mode of operation. It means building in privacy up front – right into the design specifications and architecture of new systems and processes. This approach to privacy has gained widespread international recognition, and was recognized in 2010 by data commissioners as the focus for a consolidated global approach to privacy in each of their jurisdictions [15].

Before applying *PbD* principles to a biometric 1:many system, we assume that the system has passed a *reasonableness and proportionality* test, i.e. its suitability, necessity, and appropriateness in a given context was established, and a decision to deploy was made in a transparent, responsible and accountable way.

To protect the privacy in biometric 1:many systems, a database separation was proposed in one of the versions of the Biometric Documents Identification Law of 2009 in Israel and also in the new ISO/IEC standard [16] (see also [17]). In these scenarios, the anonymous database of biometric templates is stored separately from the database containing personal information (PI) of the users. Both databases are administered by separate government entities. The databases are linked only by digital identifiers. Upon a positive biometric identification, a digital identifier is released and corresponding PI is retrieved from the second database.

While the idea of database separation is a step in right direction, this is not enough. It is obvious that, in fact, only legal measures provide protection of privacy. Both biometric and personal information are still fully under the government control. Moreover, the anonymous biometric database can be linked with other biometric databases because of the permanent nature of biometrics.

The database separation appears to be promising in conjunction with other technologies that provide a viable alternative to centrally storing biometric images/templates: *Biometric Encryption, Cryptographically Secure Biometric Architectures,* and *Biometric Setbase/Weak Links*. These technologies will be reviewed in the following sections.

3. Biometric Encryption: Making Biometrics "Untraceable" by Design

"Untraceable Biometrics" (UB) is the term that defines privacy-enhancing biometric technologies [18, 19]. At present, Untraceable Biometrics includes two major groups of emerging technologies: Biometric Encryption (BE) and Cancelable Biometrics (CB). The features of this technology approach embody standard fair information principles, providing user control, data minimization, and data security and are as follows:

- there is no storage of a biometric image or a conventional biometric template;
- it is computationally difficult to recreate the original biometric image/template from the stored information;
- a large number of untraceable templates for the same biometric can be created for different applications;
- the untraceable templates from different applications cannot be linked; and
- the untraceable templates can be renewed or revoked.

Biometric Encryption (a.k.a biometric template protection, biometric cryptosystems, etc.) was proposed [20] as a viable approach to meeting the intent of conventional biometric systems while at the same time addressing the privacy issues. BE is a group of emerging technologies that securely bind a digital key to a biometric or generate a digital key from the biometric, so that no biometric image or template is stored. It must be computationally difficult to retrieve either the key or the biometric from the stored BE template, which is also called "helper data." The key will be recreated only if the genuine biometric sample is presented on verification. The output of the BE authentication is either a key (correct or incorrect) or a failure message. See surveys [21, 22, 18, 23, 24] for more details on BE.

Cancelable Biometrics [25], which is also known as feature transformation techniques [22], does the feature transformation and stores the transformed template. The transform is stored separately, e.g., on a token, or generated from a user's password. On verification, the transformed templates are compared. There is a large number of transforms available, so that the templates are revocable. The difficulty with this approach is that the transform is in most cases fully or partially invertible, meaning that it should be kept secret. The feature transformation usually degrades the system accuracy. The system remains vulnerable to a substitution attack and to overriding the Yes/No response. Also, if the attacker knows the transform, a "masquerade" biometric sample can be created (i.e. it is not necessarily the exact copy of the original but, nevertheless, can defeat the system).

Note that it would be challenging to apply CB to 1:many systems, since the transforms would have to be centrally stored alongside with the templates rather than on tokens. This is because the users are unknown to the system on authentication and are not required to present any token or a password. The only

way around would be to store the transforms in a database separated from the CB templates. We are not aware of any such system.

There are two BE approaches: *key binding*, when a key is generated in random and then is bound to the biometric, and *key generation*, when a key is directly derived from the biometric. Both approaches usually store biometric dependent helper data and are often interchangeable for most BE schemes. After the digital key is recreated on BE verification, it can be used as the basis for any physical or logical application.

BE can greatly enhance both the privacy and security of a biometric system. In general, it is believed that BE is less susceptible to high level security attacks on a biometric system, such as substitution, Trojan horse, and masquerade attacks, tampering, overriding Yes/No response, etc. BE can work in a non-trusted or, at least, in less trusted environment, and is less dependent on hardware, procedures, and policies. The random keys are usually longer than conventional passwords and do not require user memorization. The BE helper data are renewable and revocable. At present, the most popular are the following BE schemes: Fuzzy Commitment, QIM, and Fuzzy Vault [23].

It is important to distinguish BE from the traditional encryption of the stored biometric data. This traditional encryption approach does not fully protect the users' privacy since the encryption keys are usually possessed by the data custodian. Moreover, the templates must be decrypted before each authentication attempt, so that they will be eventually exposed. Also, it is not possible to use an approach common for the password-based systems when only the hashed versions of the passwords are stored and compared: because of the natural variability of biometric samples, the hash will be completely different for any fresh biometric sample. This is a fundamental problem in bridging biometrics and cryptography, since the latter usually does not tolerate a single bit error.

Over the last few years there has been substantial progress made in developing BE technologies [23]. This includes improving accuracy and speed; multimodal solutions; fusion with conventional and cancellable biometrics; developing commercial products; pilots and deployments; new applications; addressing security issues; integration with conventional cryptography; standardization. In particular, BE has been brought to the next level, the 1:many search.

3.1. BE for 1: many systems

Even though a 1:1 biometric verification is preferred from a privacy perspective, there are important applications where 1:many searches are inevitable. The examples include a watch list scenario or de-duplication. Originally BE was designed as a pure 1:1 system, however, in recent years the 1:many search using BE has become possible.

The idea of using BE for 1:many system ascends to a 1998 patent by Tomko [26]. The BE key bound to the person's biometric encrypts, or points to, the personal information associated with the biometric identity. If a biometric sample presented during the identification matches one of the stored BE templates, the key is released and the pointer and/or the decryption key is regenerated. Thus,

the system prevents double dipping while maintaining the privacy of the personal information.

However, there are technological challenges for running BE in a 1:many mode:

- speed: BE authentication algorithms are typically slower than those of conventional biometric systems;
- accuracy: FRR results in missed hits; on the other hand, it is difficult to sort out the false hits (resulting from FAR) since no image or template is stored;
- filtering is difficult;
- security issues, especially offline FAR attack, linkage attack;
- non-repudiation and remediation: no forensic proof of misuse (no image is stored);
- all biometrics, including BE, suffer from the "glass slipper effect", i.e., a given individual's biometric will be capable of revealing a key [27]; and
- database FAR attack (a.k.a doppelganger attack) – for a sufficiently large database, the attacker can always find a false match with a legitimate user – is common for all large-scale (BE or non-BE) biometric systems [28].

In the following two subsections we will discuss the practical solutions for 1:many BE systems.

3.2. Hybrid system: conventional biometrics followed by BE (OLG Self-Exclusion Program)

To the best of our knowledge, the first real-life application of BE for one-to-many system is the Ontario Lottery and Gaming Corporation (OLG) Self-Exclusion Program.

In order to foster an environment of responsible gaming, OLG offers voluntary self-exclusion that helps individuals to take a break from gambling at all gaming sites in Ontario. If a self-excluded person is found by OLG staff at a gaming site, he/she is escorted from the premises.

However, designing a practical system is really challenging. With about 18,000 (and growing) self-excluded people, it is virtually impossible for the security personnel to memorize them all.

An automated facial recognition system seems to be the only alternative to manual inspection. Such a system captures facial images at a distance, with no requirement for direct physical interaction or even cooperation from a user (in Ontario, the casino visitors are not registered and in general, do not provide any identities). It also addresses the legacy issue, as the previously taken photographs of self-excluded individuals must be incorporated into the new system.

The self-exclusion context at OLG is a watch list where the system must identify self-excluded persons from a crowd of other patrons - the general public. Once a watch list has been created, and the system is installed, an image of each individual in range of a camera is temporarily captured. Biometric features are extracted from these new images, which are then compared to each of the templates collected during the enrollment phase – the watch list.

Since a 1:many comparison requires much more processing power than a one-to-one biometrics system, a hybrid system has been proposed [29] to use Biometric Encryption in a watch list scenario. A facial recognition module is used to reduce the normal 1:many comparison to near-1:1 comparison by filtering out all but the top few matches. The number of top matches sent to the BE module is configurable and can be as low as one (typically, below 5). An important distinction of this solution is that it offers some of the privacy-protective properties of BE to the OLG watch list application scenario by combining the conventional biometric identification (i.e. vendor-supplied facial recognition system) and BE in sequence.

In the OLG watch list application, BE is used to conceal the relationship between a person's facial recognition (FR) template and their other personal information (PI). To reveal the corresponding PI record for a matched FR record, BE will attempt to release the pointer key (PI-ID) from the biometrically encrypted data in the BE helper database.

The input during enrollment is the self-excluded person's facial image as well as a unique enrollee ID. This unique identifier does not directly relay personal information about the subject (e.g., the subject's name is not used); it simply connects the extracted facial feature record stored in the vendor-supplied identification database with a particular secure BE template (called "helper data") in the BE helper database.

A commercial facial recognition system extracts biometric features (v1) and generates a template that is stored in the FR database under the enrollee ID. Another set of biometric features, v2, is sent to the BE key binding algorithm, which creates BE helper data from v2 and a pointer key (normally, generated randomly). The BE helper data is stored in another database under the same enrollee ID. It is important that the feature sets v1 and v2 be different and not interoperable. The pointer key is used to point to the subject's facial image and other personal information (PI) that is stored in a PI database.

During identification (Fig. 1), the vendor-supplied system will attempt to do a one-to-many match of subjects entering the facilities to those in the FR database. The "Preliminary Identification" stage is typical of a watch list scenario. The enrolled IDs of the top matches are output to "BE Verification" which is a BE key retrieval algorithm. If a key is retrieved from one of the candidate BE helper data, it enters "Final Verification." The pointer to the PI record is regenerated, and the record is retrieved. An operator manually compares the retrieved facial image with the candidate image. Thus, the final decision (i.e. whether to approach the person and ask to confirm his/her identity) is left to the human operator.

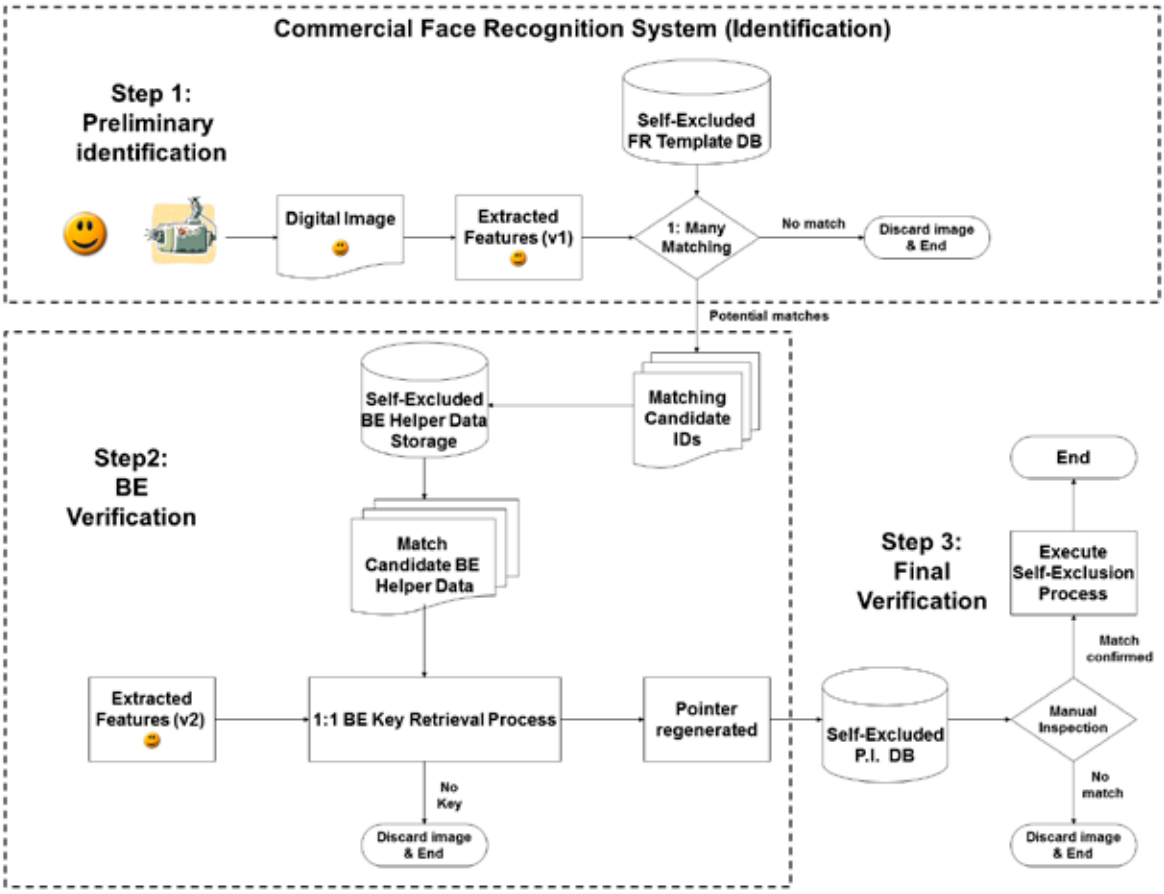

Figure 1 - Hybrid commercial facial recognition system and Biometric Encryption for identification and key retrieval.

The above-described system architecture also uses several advanced IT techniques, such as conventional cryptography, access rights, secure communications, etc., to increase the overall privacy and security of the system.

One of the principal privacy protections of this system, though, is that the link to the photo and other personal information of a self-excluded person can only be determined by accessing a key that is bound to the person's biometric. In order to reveal the stored information, the BE key retrieval algorithm must be able to regenerate a biometrically encrypted pointer key. To achieve this, the person's live facial image is required – control, thus, rests with the individual. This control prevents any lower rank security or other personnel who has access to the system from "browsing" through the personal information of self-excluded individuals. It also makes it much more difficult for the information to be linked with other, third-party databases without the user's consent.

In the OLG application, a connection from a PI record to the FR database must also be accessible for the authorized personnel to allow for updates to records in the FR database (for example, to de-enroll someone from the program). To ensure that BE cannot be circumvented to reveal the link from FR record to the PI record,

a one-way hashing algorithm is used to reveal the link between a PI record and a corresponding FR record (not shown in Fig. 1).

By providing an extra security layer, the system with BE offers better safeguards over information either stored or transmitted to a third party. To access and decrypt personal data, a would-be attacker would face the additional task of cracking the BE helper data. Even if successful, such an attack would be limited in scope, as though a "crack" would cause a breach of a single person's information, with the rest of the system remaining uncompromised (as there is no single key or template to decrypt *all* records).

Finally, BE may also enhance privacy protection for the casino patrons who are not self-excluded – the general public. BE works as a second classifier that follows the vendor's biometric engine. As shown in [30], this can substantially reduce the system FAR without a significant impact on FRR. In other words, fewer legitimate users will be inconvenienced by identity checks, as compared to the use of a facial recognition system without BE.

The system described above has moved beyond the conceptual phase. In fact, proof of concept testing has allowed OLG, in collaboration with iView Systems, University of Toronto and IPC, to show that a facial recognition application with BE is viable for development and deployment in a casino environment [31]. As this was a first-of-its-kind effort, expectations were unknown. The proof of concept testing had the following main steps and results:

- Facial recognition technology was proven effective in a casino environment through several tests in real field conditions (up to 20,000 visitors on Saturday) involving OLG control group participants (simulating the self-excluded persons) and the general public. The control group data samples were used to estimate the Correct Identification Rate (CIR). The general public data sample was used to measure the False Acceptance Rate (FAR). When there was a match from the facial recognition (FR) system, the test operator manually compared the image of the person in question with the corresponding enrolled image of a self-excluded person and decided whether the person should be approached. At the end, it was known if the system detected a real self-excluded person (correct match) or if a false match was produced by the FR system.

- The CIR was increased to a maximum of approximately 91per cent from a starting point of approximately 30 per cent. Those advances in CIR were achieved using a measured approach of several field tests and were mainly due to corrections in lighting, camera position and subject pose through the use of "attention-getting" devices like marketing screens. This compares positively to, for instance, a 2007 German Federal Criminal Police Office study which achieved a 30-60 per cent recognition rate for a facial recognition watch list field tested at a railway station [32].The False Acceptance Rate at OLG was of the order of 1 per cent (it depends on the self-excluded database size). In all the tests, the FAR remained at manageable levels.

- Biometric Encryption did not decrease the efficiency of facial recognition in a pipeline test. Positive matches from facial recognition were fed into the BE system; the BE system marginally affected the original CIR (by less than 1 per cent) while reducing the FAR by 30 per cent to 50 per cent. This result

was an unexpected benefit, which occurs due to the status of the BE module as a secondary classifier. The system used a variant of QIM BE scheme developed in [30] which allows tuning the FAR/FRR rates by varying the size of the quantization interval.

- Facial recognition was field tested using the actual OLG self-excluded images to determine by how much the detection rates for real self-excluded persons would improve. Preliminary results showed that the proposed system was a valuable tool in the overall goal of assisting self-excluded patrons from staying out of gaming facilities.
- The system architecture was successfully created to integrate BE into a commercial face recognition product from iView Systems while maintaining OLG's core requirements. This architecture treated BE as one part in a multi-layered approach to privacy and security of the overall system.
- The system has been successfully deployed and is fully operational in most of Ontario's 27 gaming sites. To the best of our knowledge, this is by far the largest installation of a BE system and the first ever application of BE in a watch list scenario. The overall identification accuracy of self-excluded people has already improved by more than one order of magnitude compared to the previous manual system.

More information about this project can be found in [33, 23, 31].

3.3. Priv-ID/GenKey one-to-many BE solutions

The Netherlands-based priv-ID, which was a spinoff of Philips Research and now is a part of GenKey, offers a variety of BE products. The technology (also called Helper Data Systems) is based on the Fuzzy Commitment scheme with a selection of most reliable components [34]. Recent advances in the technology resulted in creating a one-to-many BE system. The most important features of the priv-ID solution are the following:

3.3.1. Privacy bootstrapping [35]

This approach that fuses BE with a conventional biometric system is somewhat opposite to the hybrid solution proposed for the OLG project [31]: the initial search is done through the BE templates followed by conventional biometrics templates, as shown in Fig. 2.

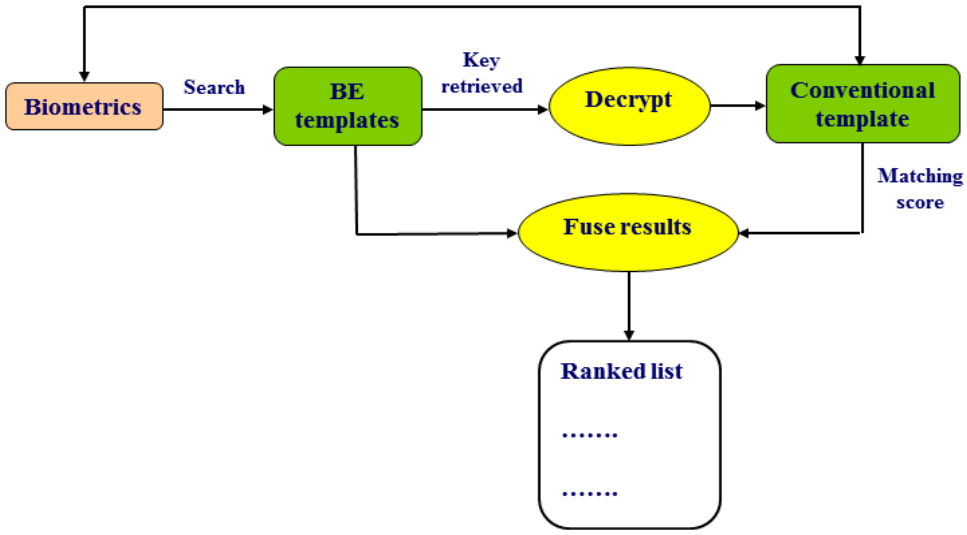

Figure 2 - Privacy bootstrapping

If a key is retrieved from a BE template, it creates a decryption key and possibly a pointer to the corresponding conventional template. This template, which had been encrypted using a strong cryptographic algorithm (e.g., AES), is then decrypted. In the next stage, the live biometrics is matched with the decrypted conventional template. This matching can use any thirdparty biometric product. The matching score is fused with the BE results (either on a decision or a score level[1] to generate the ranked list, i.e. an output of the one-to-many system. For the top ranked templates, a new decryption key and/or a pointer to the personal information can be generated, similar to the OLG system. This information should be stored in a separate database.

The priv-ID bootstrapping solution has clear benefits:

- The priv-ID system protects conventional templates using strong encryption. The encryption key, however, is not stored anywhere and can only be retrieved from a BE template by a legitimate user who, therefore, exercises control over his biometric information. The protection of the conventional template is "bootstrapped" on the protection of the BE template.

- The system is very flexible: it can use either multimodal biometrics, e.g., fingerprint BE templates followed by finger vein conventional templates, or the same biometrics but with different features, e.g., fingerprint non-minutiae BE templates followed by fingerprint conventional minutiae templates. The only requirement for a multimodal system is that at least one modality should allow conversion to BE.

- Any third party conventional biometric product can be used without integrating it closely with a BE system, thus simplifying the integration process. In fact, any existing multimodal biometric database can be protected if for at least one of the modalities a BE implementation is available.

1 BE "score" can be related, for example, to a number of errors returned by the ECC decoder.

– The proprietary techniques of the third party are protected since the knowledge of these techniques is not required for the integration with BE.

– The solution can improve the overall recognition performance, so that the resulting Receiver Operating Characteristic (ROC) curve can be even better than the best of two.

3.3.2. GPU acceleration [36]

One of the most critical issues of a BE one-to-many system is the search speed. It is known that BE, which was initially designed for one-to-one systems, is slower than many high-speed conventional biometric systems. The slowest part of the Fuzzy Commitment BE scheme is the Error Correcting Code (ECC) decoder. However, decoding algorithms for some well-known ECCs (e.g., BCH) can be ported into specialized hardware which could speed up the processing.

Priv-ID used off-the-shelf hardware acceleration – Graphics Processing Unit (GPU), which is present in most modern PCs and laptops. This is a specialized microprocessor that accelerates computer graphics. For the last several years, there has been a tendency to use GPU for general purpose computing. The GPU manufacturers developed new programming tools to allow applications to access the resources of a GPU for these purposes. In certain instances, GPU acceleration can yield several orders of magnitude higher performance than a conventional CPU.

The priv-ID implementation of their BE algorithm allowed achieving 100,000 searches/s on a PC with a single GPU [36]. The processing for a one-to-many system can easily be parallelized, meaning that searching a database of millions of records becomes feasible.

3.3.3. Fast AFIS [37]

Another Priv-ID product, Fast AFIS, claims the search speed of 120,000,000 matches per second for a fingerprint database. It can greatly increase the throughput of AFIS systems that use traditional minutiae matchers. The details of this new technology are unknown. It is clear that some sort of template filtering is used, such as the database indexing described in [38]. The latter bears some resemblance to the Beacon Guided Search for iris biometrics [56] that we will discuss later. Note that any filtering would inevitably degrade the system accuracy, so that thorough experiments are required to keep this degradation at an acceptable level.

The Fast AFIS technology, together with privacy bootstrapping and GPU acceleration, brings BE for one-to-many to the level that is suited for large-scale identification systems.

***

Both OLG and priv-ID one-to-many BE solutions are suitable for large-scale database applications. However, one of the challenges to BE adoption is the potential vulnerability to attacks specifically designed against BE [18]. For example, if an attacker gets access to a BE template, he or she can run offline a biometric database of a sufficient size to obtain a false acceptance. In this case the correct key will be recreated and the encrypted PI, including biometric templates

or images, will be decrypted. There are many ways to protect the system against the attacks, such as the advanced IT techniques used for the OLG application. Using multimodal biometrics (e.g., multiple fingers, face, iris) would definitely make the attacks much less feasible. BE data separation [16] would also help thwart the attacks. In our opinion, however, the most promising way would be integrating BE with conventional cryptography and performing the search in the encrypted domain. These techniques will be described in the following section.

4. Cryptographically secure biometric architectures

Cryptographically Secure Biometrics, also called Biometric Authentication Protocols [39, 40] or Biometric Cryptography, is a group of emerging technologies that have several parts of a biometric system (e.g., sensor, database, and matcher) communicate via secure cryptographic protocols, so that each part learns only minimal information to ensure that the users' privacy is protected. Such technologies include Authentication in Encrypted Domain, Secure Multiparty Computation [41, 42], *Cryptographically Secure Filtering, and One-to-many Access Control-type System in Match-on-Card Architecture* [43]. In the following subsections, we will discuss some of those solutions.

4.1. Ideal cryptographically secure biometric system

We note that any 1:many biometric system consists of at least three components that communicate with each other:

- *Client*, which includes a biometric sensor and also may include a processor that extracts the biometric features and generates a biometric template;
- *Database*, which stores the biometric templates;
- *Service Provider (SP)*, which does the authentication of a user. For a one-to-many system, SP is also called Identity Provider. In some publications, SP is preceded by Matcher but, for the sake of clarity, we will not make this distinction.

The SP is always separated from two other components. The secret encryption keys are usually stored on SP. (Here the word "encryption" refers to a strong cryptographic algorithm.)

An ideal biometric system that provides strong privacy protection and is cryptographically secure should satisfy the following requirements [44]:

- biometric data are stored and transmitted in encrypted form only;
- on authentication, the encrypted biometric data are not decrypted;
- the encrypted templates from different applications cannot be linked;
- the Service Provider never obtains unencrypted biometric data;
- the Client never obtains secret keys from the Service Provider;
- the encrypted template can be re-encrypted or updated without decryption or re-enrollment;
- the system is resilient to a template substitution attack;
- the encrypted template is resilient to all low level attacks (such as FAR attack);

- the system must be computationally feasible and preserve the acceptable accuracy.

In most cases it is assumed that all parts of the biometric system follow "honest-but-curious" behavior. However, as explained in [39], a possibility of malicious internal adversaries should also be taken into account. The authors of [39] suggested several black box attacks and also included collusion attacks (i.e. when two or more parts of the system, such as the Sensor and the Database, collude) in their analysis. In our opinion, including collusion attacks would be somewhat far fetched, given many advanced additional IT techniques that are used in any system to prevent those attacks. Indeed, in most systems a possibility of collusion is assumed only when several malicious users are colluding in order to attack, e.g., protected content. Collusion between a user and a database would be extraordinary.

4.2. Authentication in Encrypted Domain using Homomorphic Encryption

Most proposed solutions to the problem of authentication in the encrypted domain deal with homomorphic encryption [45, 46, 47, 48]. Homomorphism is a property of some cryptographic algorithms that allows performing some computations (e.g., additions or multiplications) in the encrypted domain, i.e. without decryption and re-encryption. These algorithms are most suitable for biometric systems that use simple operations (e.g., binary XOR) for matching. In this regard, a good candidate is an iris-based system with the standard binary templates of 2048 bits.

USING HOMOMORPHIC ENCRYPTION FOR SECURE BIOMETRIC AUTHENTICATION

A consistent protocol was developed by Bringer and Chabanne [47] using a combination of BE Fuzzy Commitment scheme with the Goldwasser-Micali [49] homomorphic encryption. In this asymmetric encryption scheme, a pair of public, pk, and secret (private), sk, keys is generated. One bit at a time is encrypted, so that, in order to encrypt a binary string, m, every bit must be individually encrypted. The Goldwasser-Micali scheme possesses a homomorphic property $\mathrm{Enc}(m) \times \mathrm{Enc}(m') = \mathrm{Enc}(m \oplus m')$, where $\oplus$ denotes the bitwise XOR operation.

For the Client – SP architecture, a simplified version of the protocol is the following:

On enrollment, the SP generates a Goldwasser-Micali (pk, sk) key pair and sends pk to the Client. The Client captures the user's iris biometric and creates the binary biometric template, b. A random Error Correcting Code (ECC) codeword, c, is generated and XOR-ed with the template to obtain BE helper data, $c \oplus b$. The result is encrypted with pk to obtain $\mathrm{Enc}(c \oplus b)$ and is put into the Database. Also, a hashed codeword, $H(c)$, is stored separately by the SP.

On verification, a fresh binary template, b', is obtained by the Client. The enrolled encrypted helper data, $\mathrm{Enc}(c \oplus b)$ is retrieved from the Database. Then, using the homomorphic property of the Goldwasser-Micali encryption, the product is computed: $\mathrm{Enc}(c \oplus b) \times \mathrm{Enc}(b') = \mathrm{Enc}(c \oplus b \oplus b')$. The result is sent to the SP, where it is decrypted with the private key sk to obtain $c \oplus b \oplus b'$. Then the ECC decoder obtains a codeword c'. Finally, the service provider checks if $H(c) = H(c')$.

The Service Provider never obtains the biometric data, which stay encrypted during the whole process. The BE helper data, $c \oplus b$, is stored in encrypted form. Since the codeword, c, is not stored anywhere, the BE helper data cannot be substituted or tampered with. Overall, this system would solve most BE security problems. The authors [47] also proposed using yet another homomorphic encryption, Paillier, on top of Goldwasser-Micali to further enhance the privacy and security protection for the database application.

It should be noted that the XOR-based Fuzzy Commitment scheme is the only BE scheme suitable for this system.

A similar biometric system that uses a homomorphic encryption, but without BE, was earlier proposed by Bringer et al [45] and Schoenmakers and Tuyls [46]: it simply encrypts the binary template, $\mathrm{Enc}(b)$, without the key binding (i.e. there is no ECC codeword, c). On verification, the result $\mathrm{Enc}(b) \times \mathrm{Enc}(b') = \mathrm{Enc}(b \oplus b')$ is sent to SP, where it is decrypted to obtain $b \oplus b'$. The Hamming distance is computed and verified against a threshold.

It is interesting to note that this approach can be viewed as a method for Cancelable Biometrics (CB), where the biometric features are transformed by the homomorphic encryption. However, like any other CB scheme, it is vulnerable to a substitution attack: if the attacker knows the public key, pk, he or she can substitute the encrypted template, $\mathrm{Enc}(b)$. Also, it does not allow re-encryption without the template decryption or re-enrollment. Finally, by having a codeword c XOR-ed with the template, the BE version [47] offers higher security as it does not completely rely on the secrecy of the homomorphic private key.

It should be noted that Simoens et al [39] described an attack on the scheme of Bringer et al [45]. The attack takes place on the side of the Authentication Server (or Database, in our notations) and uses the fact that all bits are encrypted separately. By exploiting the malleability of the homomorphic encryption, the attacker creates an encryption of a zero-bit and then sends a fake encrypted template, $\mathrm{Enc}(b'')$, for authentication. This allows the attacker to recover the stored binary template, b, bit by bit. To counteract this threat, an additional signature scheme could be employed.

Despite apparent progress achieved in the implementation of homomorphic encryption [50, 51, 52], it is still impractical in terms of processing speed, especially for one-to-many systems.

4.3. Authentication in Encrypted Domain: Secure BE using Blum-Goldwasser cryptosystem

A new approach that combines Biometric Encryption with classical Blum-Goldwasser cryptosystem was proposed in [44] for one-to-one applications. In this section, we will outline how to extend this solution to one-to-many systems.

APPLYING BLUM-GOLDWASSER CRYPTOSYSTEM TO BE IN ONE-TO-ONE ARCHITECTURE

The Blum-Goldwasser (BG) cryptosystem is a public key (i.e. asymmetric) encryption algorithm [49]. It generates a keystream using the Blum Blum Shub (BBS) pseudo-random number generator. The keystream is XOR-ed with the plaintext. It is a probabilistic, semantically secure, and quite efficient scheme both in terms of storage and speed.

In brief, Alice generates a pair of public, pk, and private (secret) keys, sk, and sends the public key to Bob. To encrypt a plaintext message, m, Bob select a *random seed*, x_0 . Using the public key and the seed, Bob computes a BBS keystream, S, which is XOR-ed with the message, m, and the encrypted value of the random seed, x_{t+1} (t is the number derived from the private key). The ciphertext, $(S \oplus m, x_{t+1})$ is sent to Alice. Using the private key and x_{t+1}, Alice recovers the seed, x_0, and reproduces the keystream, S, which decrypts the message, m.

The BG scheme is known to be vulnerable to a *chosen ciphertext attack* [49]. In our opinion, this limitation can be thwarted by the very design of a biometric system (the SP should never release a decrypted string to an attacker). Note that any homomorphic encryption scheme is also vulnerable to this kind of attack.

The BG cryptosystem can be applied to BE in one-to-one architecture in the following way [44]:

On enrollment, the SP generates the Blum-Goldwasser public, pk and private, sk, key pair and sends pk to the Client. The Client captures the user's biometric (e.g., iris) and creates the binary biometric template, b. A random ECC codeword, c, is generated and XOR-ed with the template to obtain BE helper data, $c \oplus b$. The Client chooses a random seed, x_0 , and generates the Blum Blum Shub keystream, S (of the same length as b), and the encrypted value of the random seed, x_{t+1} . The BE helper data are secured by XOR-ing them with S to obtain $S \oplus c \oplus b$. The encrypted BE template that comprises $(S \oplus c \oplus b;\ x_{t+1})$ is put into the Database. The hashed codeword, $H(c)$, is sent to the SP, where it is stored. The public key, pk, and the seed, x_0 , are discarded.

On verification, a fresh binary template, b', is obtained by the Client. The enrolled BG-encrypted helper data, $S \oplus c \oplus b$, is retrieved from the Database and XOR-ed with b'. The result and the encrypted seed are sent to the SP: $(S \oplus c \oplus b \oplus b';\ x_{t+1})$. The SP, using the private key sk and x_{t+1}, obtains S and XORs it with the data received to obtain $c \oplus b \oplus b'$. Then the ECC decoder is run and a codeword, c', is obtained. If b and b' are close enough, it will be the same codeword, c, as was on enrollment. This is verified by comparing two hashes, $H(c)$ and $H(c')$.

BG cryptosystem can be applied to two of the most popular BE schemes, Fuzzy Commitment and Quantization Index Modulation (QIM), and to the conventional XOR-based biometrics. For simplicity, we briefly describe the Fuzzy Commitment scheme in the Client – SP one-to-one architecture (see the Sidebar). Further details for one-to-one system can be found in [44].

In one-to-many system (see APPENDIX A), no hashes of the codewords are stored on SP. On verification, the system output is the ranked list of IDs with the number of corrected errors below a certain threshold. The corresponding codewords may generate pointers and/or decryption keys to personal information, such as in the OLG system or for privacy bootstrapping.

Both for one-to-one and one-to-many systems the biometric data are stored and stay encrypted during the whole process. The Service Provider never obtains biometrics b or b', or even the BE helper data, $c \oplus b$ (it is stored in the database in the BG-encrypted form). These measures assure strong privacy protection. For the one-to-one system, the accuracy (i.e. FAR/FRR numbers) remains at least the same as for the unencrypted Fuzzy Commitment or QIM scheme. Note that for the QIM scheme the accuracy can be even higher since the size of the quantization interval can increase (i.e. the scheme becomes more error-tolerant). The accuracy of the one-to-many system is also affected by the things like filtering, etc., but at its core, the BG encryption does not lower the accuracy.

The proposed approach allows updating the template without decryption or re-enrollment. Most attacks against BE, including FAR-attack, are thwarted. The BG cryptosystem can be more efficient than RSA for longer messages and, of course, is much more efficient than homomorphic encryption schemes that have to do the encryption/decryption component-wise (e.g., 2048 times for an iris 2048-bit template).

4.4. Cryptographically secure filtering

A one-to-many system often has a filtering step before the main search algorithm. This is usually the case for large databases when the computational efficiency is a crucial issue. The filtering is even more important for a BE system which is generally slower than a conventional biometric system.

The filtering algorithm uses biometric features that are somewhat different from the ones used by the main search algorithm. The main requirement here is that the filtering FRR should be low, since it adds up to the overall system FRR. On the other hand, the filtering FAR can be relatively high, since it only affects the system efficiency (for example, FAR = 5 per cent just means that the search space would be reduced by a factor of 20). For a BE or other privacy-protective one-to-many system there is also a security requirement: the filtering should not reveal much information about the biometrics.

Bringer et al [53, 54] and Adjedj et al [55] proposed new cryptographic constructions called Error-Tolerant Searchable Encryption. The authors of [53-55] (see also [42]) generalized the previous work by Hao et al [56] and made the construction cryptographically secure. We will first describe the Hao et al scheme.

Beacon Guided Search (Hao et al [56])

Hao et al [56] proposed a new filtering technique called Beacon Guided Search (BGS) for iris biometrics. The beacons are dispersed in the search space. A 2048-bit iris code is first permuted to get rid of correlations of the adjacent bits. Then it is divided into 128 two-byte chunks. From each chunk, 10 most robust bits are extracted to form a beacon, thus producing $128 \times 2^{10} = 131{,}072$ beacons. Each beacon stores $N/2^{10}$ IDs on average, or 617 IDs for the United Arab Emirates (UAE) iris database with N = 632,500. During identification, the algorithm counts the number of times (so-called collisions) when there is an exact match with one of the beacons containing the same ID. If there are at least 3 collisions, this ID passes the filtering step, i.e. goes to the full search algorithm. Hao et al obtained very good results for the UAE database: while the exhaustive search would require ~ 317262 (i.e. N/2) matches on average, the beacon search reduces this number to just 41, i.e. by four orders of magnitude. The accuracy degradation (i.e. the FRR increase) was very small.

The BGS technique is especially valuable for large size databases with N ~ 10 million. If such a database contains BE data, the filtering technique may solve the slowdown problem. However, if the filtering data are stored and processed unprotected, it will reveal 1280 bits (i.e. 62.5 per cent) for each enrolled iris code.

The authors of [53-55] showed that Hao et all construction is a special case of Locality-Sensitive Hashing (LSH) functions. They apply either Private Information Retrieval protocol [53,54] or Symmetric Searchable Encryption [55]. In both cases the secret keys are stored on the sensor client side but not on the database side, which ensures the privacy of the stored data.

Even though the security of the proposed solutions [53-55] was proven, they are not sufficient for the entire one-to-many system, as they are suitable for filtering only. After the filtering step, the full search algorithm must be run for the remaining IDs. Assuming that filtering reduces the search space by a factor of about 10^4, as in Hao et al [56], one can see that there will be about 1000 full searches for a database of about 10 million records. The authors of [53 -55] leave this issue behind the scene, just briefly suggesting [54] the use of a homomorphic encryption (or, in more general terms, the Extended Private Information Retrieval [57]).

As we already discussed, those solutions are currently infeasible for one-to-many purposes. A more practical solution using BE in a cryptographic domain was proposed in [44]. Overall, the cryptographically secure filtering [53 -55] will undoubtedly help solve the efficiency problem for large databases.

5. Biometric Setbase or Weak Links

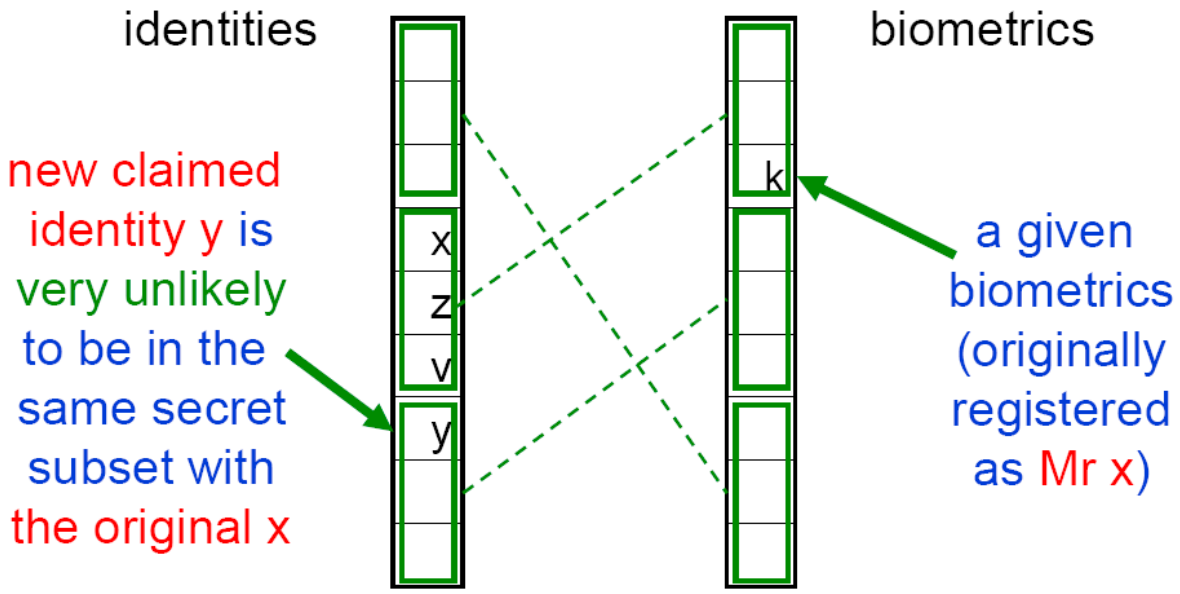

Figure 3 - Biometric setbase: disproving a wrong claim that Mr. X is Mr. Y. Adopted from [58]

To address the privacy issues of the Israel national ID program, Adi Shamir, the renowned cryptographer, proposed a new structure called "setbase" [58, 59]. Unlike traditional biometric databases where each biometric is linked to its identity through a one-to-one correspondence, here the ID setbase consists of drawers filled with about 1000 identities (this number, as Shamir mentions, should be about $\sqrt{N}$, where N is the total number of identities). Another setbase, biometric, is partitioned into drawers of the same size but filled with biometric data. Only the drawers of both setbases, but not the files themselves, are linked (Fig. 3). Those links are kept secret. The entire design is based on the cryptographic random graphs theory [59].

Since each biometric is linked to about 1000 IDs, it is not possible to identify a person solely by his/her biometrics. For that, both ID and biometrics are needed. As Shamir mentions, this separation should remain true even if the law changes after the database is set up, or even if anyone colludes with the government.

Now suppose that a cheater, originally registered as Mr. X, claims the identity of Mr. Y (Fig. 3). It is very unlikely that identities X and Y will be in the same secret drawer of the identities setbase (in the foregoing example, the probability of coincidence is about 0.001). The system will perform the following searches:

- if identity Y is in the setbase, all registered biometrics from the corresponding drawer will be run against the person's live biometrics, but no match will be found;
- after that, the person's live biometrics will be run against the entire biometrics setbase and a correct biometric sample, k, will be found. This hit corresponds

to another drawer of the identities setbase, namely, containing X's identity (Fig. 3);

– if identity Y is not in the setbase, the person's live biometrics is run against the entire biometrics setbase. Again, a correct biometric sample, k, will be found.

As a result, the cheater will be caught. Thus, the system prevents issuing multiple ID cards, yet, at the same time, makes it somewhat difficult to use the biometric setbase for other purposes or to link biometrics with other databases.

A similar idea called "weak links" was earlier proposed by Bernard Didier and François Rieul in 2004 – 2007 [60, 61].

The database is separated into two parts: an identity (ID) database containing personal information, and a biometric database. Normally, each record in the ID database would have a bidirectional link to the biometric database and vice versa. In the "weak links" construction, the links between two databases are deliberately weakened. In a simplified example, if the database contains 10 million individuals, each link would normally be a 7-digit number, from 0 to 10^7 -1. If we deliberately reduce it from 7 to 3 digits, there will be 10^4 links corresponding to each biometric record, which becomes unusable on its own. The same weakening is done for the links to each ID. Thus, the identification of one individual on the basis of biometric data only is no longer possible. For that, two entries are required: both ID and the biometric. Each will generate 10^4 weak links. Their crossover will bring a correct result in 99.9 per cent cases. The remaining 0.1 per cent will have two or more possible records, one of which is correct.

The difference between Shamir and Didier is that in Shamir there is a drawer-to-drawer link (i.e. all IDs in one drawer are linked to all biometrics in another), while in Didier the links are not necessarily (although likely) grouped into drawers. Also, Shamir uses the random graphs theory and keeps the links secret, while Didier does not specify any encryption.

The Shamir setbase proposal (or a similar Didier and Rieul weak links proposal), which is in line with the concept of the database separation [16], is a significant step towards achieving the goal of privacy protection for biometric databases.

However, we see the following potential shortcomings of the proposed system:

– Neither Shamir nor Didier describe a mechanism of handling the system's errors, such as FRR and FAR. It is assumed that the biometric system is perfect, i.e. FAR = FRR = 0. Besides affecting the system accuracy, those errors may lead to difficulties in dealing with the non-repudiation and remediation issues: both a cheater and an innocent person, Y, may claim that a false rejection occurred when all registered biometrics from the drawer corresponding to Y identity were run against the Y's live biometrics, but no match was found; and a false acceptance occurred when Y's live biometrics was run against the entire biometrics setbase and a hit k was found (Fig. 3). Usually those issues are addressed by manual inspection of the stored biometric images and other information, e.g., a photo. However, since the setbase system does not actually identify person Y but just approves or disproves his or her claim, this additional information is not available. If, alternatively, the photo is

stored and linked to the person's biometric template to facilitate this kind of investigation, this will create a backdoor and defeat the privacy protection, since photos are widely available (e.g., in social networks).

– The system stores biometric templates openly, albeit anonymously. Therefore, a template-to-template linkage with another biometric setbase is still possible: if, for example, a malicious third party wants to find out the identity of a biometric template X in two biometric setbases, they will generate $\sqrt{N}$ weak links for both setbases. The crossover of two ID lists, which requires up to N checks, will likely identify the individual in both setbases. Of course, this will even be easier if the other database is conventional, i.e. links biometrics and IDs.

– The links between drawers must be kept secret.

– As mentioned in [59], it is difficult to remove a person (e.g., who died or emigrated) from the setbase.

– Obviously, the setbase system makes sense only in case of a relatively large database (10^7 or more records), which is a limitation of its applicability to privacy protection.

In our opinion, the setbase/weak links system would significantly benefit if it stored BE helper data instead of conventional biometric templates. The links between biometric and ID drawers would not be stored at all – they would be regenerated from a BE key when a live biometric is run against the database and unlocks the key. Therefore, there is no need to keep secret links between the biometric and ID drawers (but the links from ID to biometric drawers still have to be kept secret). As with any BE system, the template-to-template authentication is not possible, so that there will be no linkage with another BE database or setbase. As in the OLG BE system, the biometric images and/or photos can be stored encrypted with the BE key. This key can only be retrieved for manual investigation purposes by applying correct biometrics to the BE helper data. This addresses the non-repudiation and remediation issue.

It should be noted that the BE setbase would have some benefits for BE as well: if an attacker cracks the BE helper data (e.g. by running certain offline attacks) and retrieves the key, he/she won't learn the exact identity of the person, only the fact that this person's ID is in the drawer with at least 1,000 others. The same will be true if the attacker obtains a person's biometric image and uses it to retrieve a key from the BE template. In other words, the setbase is an effective tool against the "glass slipper effect." Also, the setbase thwarts to some extent the doppelganger attack: if the attacker finds the template that generates a false match with his biometrics, he won't be able to identify the rightful owner of the template.

Therefore, we strongly advocate the setbase construction that would store BE helper data instead of conventional biometrics.

6. Everything is coming together: Next steps

As shown in the previous sections, several technological solutions that try to address privacy and security issues of a biometric one-to-many system have recently been proposed. An important observation is that all those solutions are not in conflict, but are rather complementary to each other with high degree of synergy, as seen from Table 1:

Table 1: Challenges and possible solutions.

Challenges	Possible solutions
Speed	BE hardware acceleration (3.3.2) Filtering (3.3.3), (4.4) OLG hybrid solution: conventional templates followed by BE (3.2)
Accuracy	BE privacy bootstrapping (3.3.1) Two classifiers – OLG hybrid solution (3.2) False hits are manually sorted out after BE regenerates a pointer and a decryption key to other personal information (3.2)
Filtering is difficult	Fast AFIS (3.3.3) Cryptographically secure Beacon Guided Search (4.4)
Offline FAR and other attacks against BE	Authentication in encrypted domain (4.2, 4.3) Multimodal BE (3.3) Separation of the databases [16] Setbase with BE (5)
"Glass slipper" effect and doppelganger attack	Setbase (5)
Non-repudiation and remediation	Regenerating (using BE) a pointer and/or a decryption key to other personal information (3.2) – can be combined with Setbase (5)

The examples of how these solutions can complement each other include:

Hardware acceleration [36] and secure filtering [37, 38, 53-55] will be beneficial in terms of speed for any one-to-many system;

Separation of the databases [16], which helps to mitigate both privacy concerns and vulnerability against attacks, is present in most systems described in the previous sections;

Using BE templates instead of conventional ones within a setbase [58] would have clear advantages for privacy and security;

Privacy bootstrapping [35] and the OLG hybrid solution [31] could work within a setbase, thus improving the system accuracy;

A cryptosystem doing authentication in encrypted domain [47, 44] could be integrated with a setbase [58]. This would solve most privacy and security problems of a "double-dipping" prevention one-to-many system.

In general, there is no single solution that fits all; rather, in each case several approaches can be combined to address the issues of a specific application. We hope that the efforts of all interested stakeholders will lead to creating working prototypes suitable for pilot projects.

7. Conclusions

Biometric one-to-many systems have been used for a variety of purposes, such as multiple enrollment prevention, watch list, access control, forensics, stranger identification, etc. While these systems often serve legitimate purposes, such as combatting fraud or catching a villain, the rise of the ubiquitous use of biometrics can be viewed as an integral part of the emerging surveillance society.

Biometrics, especially one-to-many systems, can pose some serious threats to privacy due to uniqueness, permanent nature and irrevocability of biometric data. However, it does not have to be that way. The same technology that serves to threaten or erode our privacy may also be enlisted to strengthen its protection. In particular, Biometric Encryption (BE) technologies, or, in more general terms, "Untraceable Biometrics" was proposed as a privacy-protective alternative to conventional one-to-one biometrics.

In this paper, we apply a *Privacy by Design* approach to exploring new ideas and solutions that can lead to deployment of privacy-protective and secure biometric one-to-many systems. We showed that new advances in BE can be complemented with other innovative solutions, such as Cryptographically Secure Biometric Architectures and Biometric Setbase/Weak Links. We presented a case study of the first BE application using facial recognition in a watch list scenario at the Ontario Lottery and Gaming Corporation (OLG). We proposed a cryptographically secure architecture for a one-to-many system using Blum-Goldwasser cryptosystem.

In our opinion, these solutions can be combined with each other in the application-specific context in order to create a one-to-many system that addresses privacy, security and functionality issues, all the hallmarks of a *Privacy by Design* approach.

Appendix A

Technical details of the proposed solution for Blum-Goldwasser cryptosystem for one-to-many

The proposed system is shown in Fig. 4.

On enrollment, the SP generates the Blum-Goldwasser public, *pk*, and private, *sk*, key pair and sends *pk* to the Client. The Client generates the BG-encrypted BE template ($S \oplus c \oplus b$; x_{t+1}) in the same way as for the token-based one-to-one architecture. However, only the first part, $S \oplus c \oplus b$, is sent to the Database for storage. The encrypted seed, x_{t+1} is sent to the SP. The latter regenerates the keystream, *S*, and stores it under the user's identity ID. Note that the hash, *H*(c), is not stored in this version.

On verification, a fresh binary template, *b*′, is obtained by the Client. The new seed, x_0' , and the new keystream, *S*′, are generated by the Client. The XOR-ed data $S' \oplus b'$ are sent to the Database where they are XOR-ed with all the stored templates to obtain $S \oplus S' \oplus c \oplus b \oplus b'$. The results for each ID are sent from the Database to SP through a high speed communication line, such as a 100 Gigabit Ethernet fiber optic network. At the same time, the new encrypted seed, x'_{t+1} , is sent by the Client to SP. The SP, using the private key *sk* and x'_{t+1}, obtains S' and XORs it with all the stored keystreams, *S*. Then $S \oplus S'$ are XOR-ed with $S \oplus S' \oplus c \oplus b \oplus b'$ that have been received from the Database to obtain $c \oplus b \oplus b'$ for each ID. Then the ECC decoder is run and all the codewords, *c*', are obtained. The IDs with the number of corrected errors below certain threshold are put on the ranked list. The corresponding codewords may generate pointers and/or decryption keys to personal information, such as in the OLG system or for privacy bootstrapping.

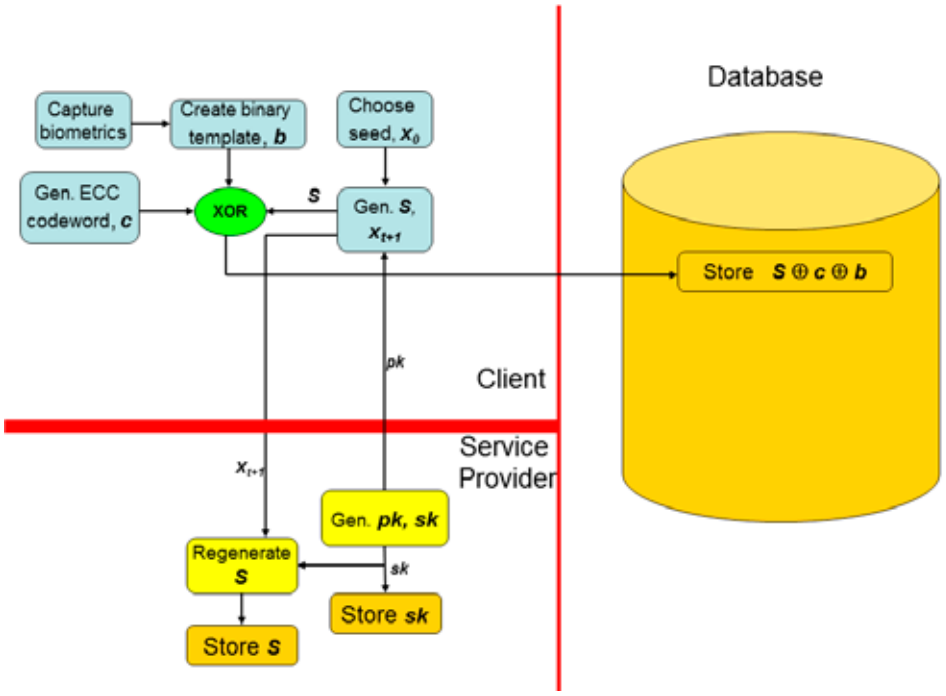

Figure 4 - Secure Fuzzy Commitment scheme using Blum-Goldwasser cryptosystem in the Client – Database – Service Provider architecture (enrollment).

The biometric data are stored and stay encrypted all the time, and the SP or a database never obtains b or b' . Due to the probabilistic nature of the BG cryptosystem, the SP can have the same pair of (pk, sk) for all the users (only the seeds are different), which reduces the requirements to the key management. Since the encrypted seeds, x_{t+1} , are not stored in the Database, the template size is the same as for the unencrypted version (e.g., 2048 bits for iris).

Note that the BG-encrypted templates in the Database can be additionally encrypted (by XOR-ing with the new keystream, ΔS) using a stream cipher (which is independent from the BG cryptosystem) to generate a new keystream, ΔS. The latter is XOR-ed with the stored templates and is re-created during verification on both the Database and the SP. This is to ensure better system security, so that by the time of verification the keystream will be different, $\check{S}$, and the stored template will be $\check{S} \oplus c \oplus b$.

The proposed solution is well suited both for the priv-ID system that is based on Fuzzy Commitment scheme and for the OLG system that uses QIM.

It should be mentioned that Simoens et al [39] proposed a Hill Climbing attack on this scheme. The attack originates on the Authentication Server (or Database, in our notations) side. First, it is assumed that the Database knows $S' \oplus b'$ that leads to a positive decision. (Note that this is already a far-fetched assumption since the SP does not communicate its decision back to the Database.) Second, the Database starts from this $S' \oplus b'$ and adds progressively some errors until a negative result is obtained and then backtracks back by one error, and so on. Following this strategy, it is possible to recover all bits in $b \oplus b'$ (see further details in [39]). (Note that the attacker still cannot recover either b or b' !) However, the attack, as described in [39], will not work at all because of the very design of the

system: each verification query must directly send a new encrypted seed, x'_{t+1} , from the Sensor to the SP (i.e. bypassing the Database), while the Ref. [39] wrongly assumes that x'_{t+1} would remain the same during the entire Hill Climbing process.

References

1 Laura Poitras, Glenn Greenwald. *NSA whistleblower Edward Snowden: 'I don't want to live in a society that does these sorts of things' – video*. The Guardian, June 9 (2013). Retrieved from http://www.theguardian.com/world/video/2013/jun/09/nsa-whistleblower-edward-snowden-interview-video

2 The Association for Civil Rights in Israel. *ACRI applauds cancellation of biometric database bill—For now* (2009). Retrieved from http://www.acri.org.il/en/?p=694

3 Pinchuk, A. *The campaign against the biometric database act. The Public Voice Civil Society Meeting*, Jerusalem. (2010).

4 Lebovic, N. and Pinchuk, A. *The state of Israel and the biometric database law: Political centrism and the post-democratic state.* (2010). Retrieved from The Israeli Democracy Institute website http://www.idi.org.il/sites/english/BreakingTheNews/Pages/IsraelAndTheBiometricDatabase.aspx

5 Standing Committee On Finance (2011-12). *The National Identification Authority Of India Bill, 2010. Forty-Second Report.* Lok Sabha Secretariat, New Delhi, December 2011.

6 K. Rodriguez. *Biometrics in Argentina: Mass Surveillance as a State Policy. A joint campaign with Fundacion Via Libre*, January 10, 2012. https://www.eff.org/deeplinks/2012/01/biometrics-argentina-mass-surveillance-state-policy

7 Cavoukian, A. *Privacy by Design: The 7 foundational principals, Implementation and Mapping of Fair Information Practices*. Toronto: Office of the Information and Privacy Commissioner of Ontario (2010).

8 European Parliament, Directorate-General for Internal Policies, Policy Department, Citizenship and Constitutional Affairs. *Developing biometrics in the EU.* (2010).

9 Harris, H. A. *Privacy rights according to the Supreme Court of Canada*. Office of the Privacy Commissioner of Canada to CAPA Conference (1997). Retrieved from http://www.priv.gc.ca/speech/archive/02_05_a_971030_e.cfm

10 The Australian Privacy Charter Council. *Australian privacy charter*. (1994). Retrieved from http://www.privacy.org.au/About/PrivacyCharter.html

11 California Office of Privacy Protection. *Constitutional right to privacy. Article 1, Section 1*. Retrieved from http://www.leginfo.ca.gov/.const/.article_1

12 United Nations General Assembly. *The universal declaration of human rights.* (1948). Retrieved from http://www.un.org/en/documents/udhr/

13 Inskeep, T., & Claypoole, T. F. *Unintended consequences of biometrics*. In W. Sloan Coats, A. Bagdasarian, T. J. Helou, & T. Lam (Eds.), The practitioner's guide to biometrics (pp. 175–213). Chicago, IL: American Bar Association (2007).

14 Lodge, J., & Snijder, M. *Developing biometrics in the EU study*. Brussels: European Parliament Directorate General for Internal Policies (2010).

15 *Privacy by Design Resolution*. 32nd international conference of data protection and privacy commissioners, Jerusalem, Israel, 27–29 October (2010).

16 ISO/IEC 24745. *Information technology—Security techniques—Biometric information protection* (2011).

17 Tom Kevenaar, Ulrike Korte, Johannes Merkle, Matthias Niesing, Heinrich Ihmor, Christoph Busch, and Xuebing Zhou, *A Reference Framework for the Privacy Assessment of Keyless Biometric Template Protection Systems*. BIOSIG, volume 164 of LNI, pp. 45-56. GI (2010)

18 A. Cavoukian and A. Stoianov, *Biometric Encryption: The New Breed of Untraceable Biometrics*, in Biometrics: fundamentals, theory, and systems, Ch. 26, N. V. Boulgouris, K. N. Plataniotis, and E. Micheli-Tzanakou, Eds. London: Wiley-IEEE Press, pp. 655-718 (2009).

19 Cavoukian, A. and Snijder, M. *The relevance of untraceable biometrics and biometric encryption: A discussion of biometrics for authentication purposes*. Toronto: Office of the Information and Privacy Commissioner of Ontario (2009).

20 Tomko, G. J., Soutar, C., & Schmidt, G. J. *Fingerprint controlled public key cryptographic system*. U.S. Patent 5541994 (1996).

21 P. Tuyls, B. Škorić, and T. Kevenaar, Eds., *Security with Noisy Data: Private Biometrics, Secure Key Storage and Anti-Counterfeiting*. London : Springer-Verlag (2007).

22 A. K. Jain, K. Nandakumar, and A. Nagar, *Biometric Template Security*, EURASIP Journal on Advances in Signal Processing, v. 2008, Article ID 579416, pp. 1-17 (2008).

23 Ann Cavoukian, Michelle Chibba, and Alex Stoianov, *Advances in Biometric Encryption: Taking Privacy by Design from Academic Research to Deployment.* Review of Policy Research, V. 29, Issue 1, pp. 37-61 (2012).

24 Rathgeb C, Uhl A (2011). *A survey on biometric cryptosystems and cancelable biometrics*. EURASIP Journal on Information Security 2011:3–25. http://jis.eurasipjournals.com/content/2011/1/3

25 Ratha NK, Connell JH, Bolle RM. *Enhancing security and privacy in biometrics-based authentication systems*. IBM Systems Journal 40(3):614–634 (2001).

26 G. J. Tomko, *Method and apparatus for securely handling data in a database of biometrics and associated data*. U.S. Patent 5790668, August 4 (1998).

27 Kim Cameron, *Identity Weblog*, 1 April (2007). http://www.identityblog.com/?p=735

28 T.E. Boult and R. Woodworth, *Privacy and security enhancements in biometrics*. Advances in Biometrics: Sensors, Algorithms and Systems, N.K. Ratha & V. Govindaraju, eds., New York, N.Y.: Springer (2007).

29 A. Stoianov. *Private communication*, IPC, Aug. 2007.

30 Martin, K., Lu, H., Bui, F., Plataniotis, K. N. and Hatzinakos, D., *A biometric encryption system for the self-exclusion scenario of face recognition. IEEE Systems Journal: Special Issue on Biometrics Systems*, vol. 3, no. 4, pp. 440-450 (2009).

31 Ann Cavoukian, Tom Marinelli, Alex Stoianov, Karl Martin, Konstantinos N. Plataniotis, Michelle Chibba, Les DeSouza, Soren Frederiksen. *Biometric Encryption: Creating a Privacy-Preserving 'Watch-List' Facial Recognition System*. In: Security and Privacy in Biometrics, Patrizio Campisi (ed.), Ch. 9, pp. 215-238. Springer-Verlag London, 2013.

32 Bundeskriminalamt. *Face recognition as a search tool*. Final Report. http://www.eucpn.org/download/?file=GER%20FAce%20Recognition.pdf&type=14

33 A. Cavoukian and T. Marinelli. *Privacy-Protective Facial Recognition: Proof-of-Concept for Biometric Encryption: A Research Report on Using Biometric Encryption to Limit 'Self-Excluded' Problem Gambler Access to Gaming Venues* (2010). http://www.ipc.on.ca/images/Resources/pbd-olg-facial-recog.pdf

34 M. van der Veen, T. Kevenaar, G.- J. Schrijen, T. H. Akkermans, and Fei Zuo, *Face biometrics with renewable templates*, in Proceedings of SPIE, Vol. 6072: Security, Steganography, andWatermarking of Multimedia Contents VIII (2006).

35 T. A. M. Kevenaar and A. M. Lemma. *Method and system for verifying the identity of an individual by employing biometric data features associated with the individual*, WIPO Patent Application WO/2010/080020 (2010).

36 Retrieved from http://www.genkey.com/technology/technology (April 17, 2014)

37 Retrieved from http://www.genkey.com/technology/fastafisr-sdk (April 17, 2014)

38 T. A. M. Kevenaar and A. H. M. Akkermans, *Method and apparatus for calculating an index key*, WIPO Patent Application WO/2008/035251 (2008).

39 Koen Simoens, Julien Bringer, Hervé Chabanne, Stefaan Seys, *Analysis of Biometric Authentication Protocols in the Blackbox Model*. arXiv:1101.2569 [cs.CR] (13 Jan 2011). http://arxiv.org/pdf/1101.2569v1.pdf

40 K. Simoens, J. Bringer, H. Chabanne, and S. Seys. *A Framework for Analyzing Template Security and Privacy in Biometric Authentication Systems*. IEEE Transactions on Information Forensics and Security, Vol. 7, No. 2, pp. 833 - 841 (2012).

41 M. D. Raimondo, M. Barni, D. Catalano, R. D. Labati, P. Failla, T. Bianchi, D. Fiore, R. Lazzeretti, V. Piuri, F. Scotti, and A. Piva. *Privacy-preserving fingercode authentication*, in Proc. of the 12th ACM workshop on Multimedia and security (MMSec'10). ACM, pp. 231–240 (2010).

42 Bringer, J. , Favre, M. , Chabanne, H., and Patey, A. *Faster Secure Computation for Biometric Identification Using Filtering*. The 5th IAPR International Conference on Biometrics, March 29 - April 1, 2012, New Delhi, India (ICB 2012), pp.257-264 (2012).

43 Julien Bringer and Herve Chabanne, *Two efficient architectures for handling biometric data while taking care of their privacy*. In: Security and Privacy in Biometrics, Patrizio Campisi (ed.), Ch. 11, pp. 275-295. Springer-Verlag London (2013).

44 A. Stoianov, *Cryptographically secure biometrics*. Proc. SPIE, Vol. 7667, pp. 76670C-1 - 76670C-12 (2010).

45 J. Bringer, H. Chabanne, M. Izabachène, D. Pointcheval, Q. Tang, and S. Zimmer, *An Application of the Goldwasser-Micali Cryptosystem to Biometric Authentication*. LNCS, v. 4586, pp. 96-106 (2007).

46 B. Schoenmakers and P. Tuyls, *Computationally secure authentication with noisy data*. In Security with Noisy Data: Private Biometrics, Secure Key Storage and Anti-Counterfeiting. Eds. P. Tuyls, B. Škorić, and T. Kevenaar. Springer-Verlag, London, pp. 141 -149 (2007).

47 J. Bringer and H. Chabanne, *An authentication protocol with encrypted biometric data*. LNCS, v. 5023, pp. 109-124, 2008.

48 M. Upmanyu, A. M. Namboodiri, K. Srinathan, and C. V. Jawahar, *Efficient Biometric Verification in Encrypted Domain*, ICB 2009, LNCS 5558, pp. 899-908 (2009); also: M. Upmanyu, A. M. Namboodiri, K. Srinathan, and C. V. Jawahar, *Blind Authentication: A Secure Crypto-Biometric Verification Protocol*, IEEE Trans. on Info. Forensics and Security, Vol. 5, No. 2 (2010).

49 Menezes, A.J., Van Oorschot, P.C. and Vanstone, S.A. *Handbook of Applied Cryptography*, CRC Press (2001).

50 Z. Erkin, M. Franz, J. Guajardo, S. Katzenbeisser, I. Lagendijk, and T. Toft, *Privacy-preserving face recognition, in PETS '09: Proceedings of the 9th International Symposium on Privacy Enhancing Technologies.* Berlin, Heidelberg: Springer-Verlag, pp. 235–253 (2009).

51 A. Sadeghi, T. Schneider, and I. Wehrenberg, *Efficient privacy preserving face recognition, in ICISC '09: Proceedings of the 12th Annual International Conference on Information Security and Cryptology*, ser. LNCS, vol. 5984. Springer-Verlag, December 2-4, pp. 235–253 (2009). Available at http://eprint.iacr.org/2009/507.

52 M. Barni, T. Bianchi, D. Catalano, M. D. Raimondo, R. D. Labati, P. Failla, D. Fiore, R. Lazzeretti, V. Piuri, F. Scotti, and A. Piva, *Privacy-preserving fingercode authentication, in 12th ACM Multimedia and Security Workshop* (2010). Available at http://www.dmi.unict.it/diraimondo/uploads/papers/fingercodeprotocol-unpublished.pdf

53 Bringer, J., Chabanne, H., Kindarji, B. *Error-tolerant searchable encryption*. In: IEEE International Conference on Communications, 2009. ICC 2009, June 2009, pp. 1–6 (2009)

54 Bringer, J., Chabanne, H., Kindarji, B. *Identification with encrypted biometric data*. CoRR abs/0901.1062 (2009). Full version of previous publication.

55 Adjedj, M., Bringer, J., Chabanne, H. and Kindarji, B., *Biometric Identification over Encrypted Data Made Feasible*, LNCS, Springer 5905, 86-100 (2009).

56 Hao, F., Daugman, J., Zielinski, P. A *Fast Search Algorithm for a Large Fuzzy Database*. IEEE Transactions on Information Forensics and Security 3(2), 203–212 (2008).

57 J. Bringer, H. Chabanne, D. Pointcheval, and Q. Tang. *Extended private information retrieval and its application in biometrics authentications*. In Feng Bao, San Ling, Tatsuaki Okamoto, Huaxiong Wang, and Chaoping Xing, editors, CANS, volume 4856 of Lecture Notes in Computer Science, pages 175–193. Springer (2007).

58 Adi Shamir, *Adding Privacy To Biometric Databases: The Setbase Approach*. Presentation at 31st Annual Conference of Data Protection and Privacy Commissioners, Madrid, November 4-6 (2009). http://www.privacyconference2009.org/program/Presentaciones/common/pdfs/adhi_shamir_madrid.pdf

59 Adi Shamir. *Random graphs in security and privacy*, ICITS, December 5-th 2009, The Weizmann Institute, Israel (2009). https://www.rcis.aist.go.jp/ICITS2009/presentations/Shamir.pdf

60 B. Didier, *Design and use of a central biometric base of the issue of identity documents*. Sagem, France, March 2007.

61 Bernard Didier and Francois Rieul. *Person identification control method and system for implementing same*. United States Patent 7724924. Priority Date: March 17, 2004 (2010).


PbD
www.privacybydesign.ca