Biometric Encryption:
A Positive-Sum Technology That Achieves Strong Authentication, Security, *and* Privacy

March 2007

Biometric Encryption: A Positive-Sum Technology That Achieves Strong Authentication, Security, *and* Privacy

Abstract

This paper discusses privacy-enhanced uses of biometrics, with a particular focus on the privacy and security advantages of Biometric Encryption (BE) over other uses of biometrics. The paper is intended to engage a broad audience to consider the merits of the Biometric Encryption approach to verifying identity, protecting privacy, and ensuring security. Our central message is that BE technology can help to overcome the prevailing "zero-sum" mentality, namely, that adding privacy to identification and information systems will necessarily weaken security and functionality. This paper explains how and why BE technology promises a "positive-sum," win-win scenario for all stakeholders involved.

Background/Context

Identification and authentication requirements are steadily increasing in both the on-line and off-line worlds. There is a great need on the part of both public and private sector entities to "know" who they are dealing with. The current security model for the verification of identity, protection of information, and authorization to access premises or services is based on using a token, tied to and thereby representing an individual, to either authenticate identity or allow access to information, premises or services. This token may be a password or shared secret (something you know), an identity card (something you have), or a biometric (something you are). In all of these cases, the details of the token are held by a third party whose function is to authorize and, at times, allow the transaction to proceed if the details of an individual's token match those stored in a database. The biometric is increasingly viewed as the ultimate form of authentication or identification, supplying the third and final element of proof of identity. Accordingly, it is being rolled out in many security applications.

Privacy-related areas involving the protection of personal information, however, are not as strong – biometrics have not yet been able to fill this need. When an individual provides his or her personal information (financial or medical) to a second party, this party often stipulates that it will only use the personal information for the agreed-upon function, and will thereafter protect the information from access by unauthorized parties. The relationship between the individual who provides the information and the second party is largely based on a model of trust.

The trust model is becoming far less effective as current technological and geo-political situations evolve. The selling or sharing of personal information is now a lucrative business model practised by many companies. Similarly, with increased threats of terrorism, governments and law enforcement agencies can now demand access to more and more personal information. With the growing powers of the Internet, extensive electronic dossiers may now be developed about an individual, without his or her knowledge or consent. Of even greater concern, perhaps, are the errors that can easily arise, which may then adversely affect that individual's life.

These dossiers may also include the details of token-based transactions such as biometrics, resulting in surprisingly complete dossiers about individuals and their transactional histories, again without their knowledge or consent. In turn, this precludes one's ability to ever correct any errors which may be contained in such databases, presenting an ever-growing problem. In short, unauthorized access to one's personal information can result in a host of negative consequences, ranging from identity theft and harassment to the perpetuation of mistakenly used personal information.

We acknowledge that government and law enforcement agencies require personal information to protect public safety and national security, while businesses require personal information to improve business practices and customer service. However, within these scenarios, the existing model of protecting privacy and safeguarding information invariably leads to a zero-sum game – protecting privacy often leads to less security and more costly business practices. This need not be the case.

Protecting public safety and a nation's security is a necessary and important function of a civilized society; developing more efficient business practices that are more cost effective and lead to better customer service are also highly desirable. Social and economic well-being are served by both of these functions.

However, liberty and freedom of choice are also essential to the functioning of prosperous and free societies. Technological advances in the collection and processing of information over the last few decades have positioned this resource as vital to the health, well-being and freedom of individuals. More specifically, abuses of personal information can cause untold harm, wasted resources, and generally lead to the detriment of society. For example, a society of individuals perpetually anxious about identity theft, misuses of their information, or unwarranted search and seizures cannot function at optimum levels.

It is our belief that the security model in current use must change from a zero-sum to a positive-sum paradigm, where both the need for privacy/protection of personal information and the need for security can be satisfied. Accordingly, in this paper, we present what we believe to be the first step in the achievement of that goal through a new positive-sum model for both protecting information and providing security, based on "Biometric Encryption."

Growing Public Awareness and Interest

Biometrics are expected to add a new level of security to applications, as a person attempting access must prove who he or she really is by presenting a biometric to the system. Such systems may also have the convenience, from the user's perspective, of not requiring the user to remember a password.

There is evidence of growing public awareness and interest in the use of biometrics.

Border Security Control: Perhaps the most visible (and controversial) use of biometrics is taking place in the transportation sector. Identification requirements at airports and border crossings may now involve the collection and processing of travellers' fingerprints, facial images, and iris patterns. Increasingly, machine-readable travel documents such as passports, driver's licences, and other identity or travel cards may also contain biometric data or images. Frequent travellers who apply for and pass extensive background checks may use their biometrics for speedy passage through customs and immigration.

Crime and Fraud Prevention, Detection, and Forensics: The use of fingerprints by law enforcement has taken place for many years, but now that fingerprints can be digitized, stored, retrieved, and matched instantaneously, many new uses have emerged, such as for populating watch lists and carrying out private sector background checks. In some parts of the United States, cashing a cheque can require a biometric imprint to be placed on the obverse side. Not a day goes by where the public is not apprised of some new "revolutionary" biometric technology that promises to solve crimes, catch villains, and generally make the world a better place to live.

Attendance Recording: Employees and students are being required, in growing numbers, to present a biometric (such as a finger or hand) in order to "check in" to premises, much like a punchclock, or to claim some entitlement such as a luncheon meal or to check out a library book.

Payment Systems: We are seeing increasing uses of biometrics by the private sector for enhanced convenience services, such as "pay 'n' go" systems that allow enrolled customers to pay for groceries or gasoline using only their finger – at times, an enormous convenience.

Access Control: One of the most widespread uses of biometrics has been for physical and logical access to secure areas or resources (e.g., to a database of medical records, or accessing a laptop). In such circumstances, biometrics can enhance security by helping to ensure that access to sensitive resources is strictly restricted to authorized individuals.

A Biometrics Primer

"Biometrics" refers to automatic systems that use measurable, physical or physiological characteristics or behavioural traits to recognize the identity, or verify/authenticate the claimed identity of an individual. The examples of biometric characteristics that have been used for automated recognition include fingerprints, iris, face, hand or finger geometry, retina, voice, signature, and keystroke dynamics.

These systems are based on the following steps: a biometric sample is taken from an individual, for instance a fingerprint or iris scan. This physical characteristic may be presented by an image. Often data are extracted from that sample. These extracted data constitute a *biometric template*. The biometric data, either the image or the template or both, are then stored on a storage medium. The medium could be a database or a distributed environment, such as smart cards. These preparatory phases together constitute the process of *enrolment*. The person whose data are thus stored is called the enrolee.

The actual purpose of the biometric system is only achieved at a later stage. If a person presents herself to the system, the system will ask her to submit her biometric characteristic(s). The system will then compare the image of the submitted sample (or the template extracted from it) with the biometric data of the enrolee. If the match succeeds, the person is then recognized and the system will "accept" her. If the match does not succeed, she is not recognized and she will be "rejected."

Traditional Biometrics: Privacy vs. Security – A Zero-Sum Game

We thought it might be useful to begin with a table (see next page) that summarizes the essential differences between the traditional zero-sum approach to biometrics vs. the positive-sum, Biometric Encryption approach. Such a comparison facilitates ease of reference and differentiates one from the other; this is also followed by the page number where a full discussion of the issue takes place.

Applicable law and regulation will vary, but biometric data, being derived from human bodies (and especially when used to identify or verify those bodies) is considered **personally identifiable information (PII)**. The collection, use, and disclosure of biometric data – image or template – invokes rights on the part of an individual and obligations on the part of an organization.

Difficult ethical and operational questions surround the collection and use of video images used for facial recognition (which may be collected without the knowledge or consent of the individual), and of fingerprints and DNA samples, which may also reveal far more than identity.

As biometric uses and databases grow, so do concerns that the personal data collected will not be used in reasonable and accountable ways. Privacy concerns arise when biometric data are used for secondary purposes, invoking "function

	Traditional Biometrics: Privacy or Security – A Zero-Sum Game	**Biometric Encryption: Privacy and Security – A Positive-Sum Game**
1	The biometric template stored is an identifier unique to the individual.	There is no conventional biometric template; therefore no unique biometric identifier may be tied to the individual. (pp. 127, 128)
2	Secondary uses of the template (unique identifier) can be used to log transactions if biometrics become widespread.	Without a unique identifier, transactions cannot be collected or tied to an individual. (pp. 128, 129, 138)
3	A compromised database of individual biometrics or their templates affects the privacy of all individuals.	No large databases of biometrics are created, only biometrically encrypted keys. Any compromise would have to take place one key at a time. (p. 135)
4	Privacy and security not possible.	Privacy and security easily achieved. (pp. 128-132, 139-142)
5	Biometric cannot achieve a high level of challenge-response security.	Challenge-response security is an easily available option. (pp. 139-141)
6	Biometrics can only indirectly protect privacy of personal information in large private or public databases.	BE can enable the creation of a private and highly secure anonymous database structure for personal information in large private or public databases. (pp. 130-131, 140-141)
7	*1:many* identification systems suffer from serious privacy concerns if the database is compromised.	*1:many* identification systems are both private and secure. (pp. 128, 131)
8	Users' biometric images or templates cannot easily be replaced in the event of a breach, theft, or account compromise.	Biometrically encrypted account identifiers can be revoked and a new identifier generated in the event of breach or database compromise. (p. 129)
9	Biometric system is vulnerable to potential attacks.	BE is resilient to many known attacks. (p. 129)
10	Data aggregation.	Data minimization. (p. 128)

creep," data matching, aggregation, surveillance and profiling. 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.

A broad discussion of the various privacy implications of biometrics is available on the website of the Information and Privacy Commissioner of Ontario, **www.ipc.on.ca**[1].

Biometric Identification vs. Verification

Regardless of specific uses and deployment scenarios, most biometric systems will serve one of two foundational purposes: **identification** or **verification/ authentication**.

Identification refers to the ability of a computer system to uniquely distinguish an individual from a larger set of individual biometric records on file (using only the biometric data). So, theoretically, a national biometric identification system could allow a citizen to prove who he or she is without recourse to any document – assuming the citizen was already registered in the system. The presented biometric data would simply be compared with all other entries in the national database for a match, and upon a successful match the associated citizen's identity data would be released from the database. This is often referred to as a *"one-to-many"* match, and is used by police to identify criminals on watchlists, as well as by governments to identify qualified recipients for benefit-entitlement programs and registration systems such as voting, driver's licence, and other applications. So, for example, the facial images supplied in support of passport or driver's licence applications could be routinely compared against large databases to ensure that multiple documents had not been issued to the same applicant (i.e., fraud detection).

Biometric **verification** or authentication involves a *"one-to-one"* search whereby a live biometric sample presented by a person is compared with a stored sample (on a smart card or contained in a database) previously given by that individual, and the match confirmed. The eligibility of the person for the service or benefit has already been previously established. The matching of the live biometric to the sample is all that is necessary to authenticate the individual as an eligible user. There need not be any search or matching to a central database, although a central database can still be used, provided that some other identification data is used. For example, an identity card's serial number could be used to "look up" an individual in a biometric database, and the live biometric sample could then be matched against the sample stored on record to verify the individual as the rightful bearer of the card. Even simpler, the person could just type in his username, so that his biometric template could be called up from the database for verification.

1 e.g., "Privacy and Biometrics," "Biometrics and Policing: Comments from a Privacy Perspective," and "Biometrics and Consumer Applications." All documents are freely available at **www.ipc.on.ca**

Identification templates are always stored in a database that is controlled by a custodian. *One-to-one* templates can be stored either in a database or in a distributed medium carried by a user (e.g., a passport, a smart card, or token). In the latter case, the user retains control over his biometric template.

Some current deployments require both identification and verification. For example, if a person applies for a passport/ID card, his biometric samples enter a *one-to-many* search first. This is done to check his background, i.e., to make sure that the person has not been listed in a criminal/ terrorist database before, usually under different identity. If the person is cleared, he is issued the passport/ID card to be used in a *one-to-one* system later on.

Somewhere between *"one-to-many"* identification and *"one-to-one"* authentication lies *"one-to-few"* biometric data uses, where "few" is of an order of 2–10,000. For example, a biometric lock may store the templates from all the members of a household or a firm. Some tokenless access control systems operate on this basis: the employee or user simply presents a biometric sample to the system, which then compares the sample against a small database of authorized users. If a match occurs, access is granted. The individual is both "identified" and "verified" as an authorized user – no other form of identification takes place.

Problems with Using Biometrics for Identification Purposes

In the futuristic film *Minority Report,* starring Tom Cruise, individuals are automatically and instantaneously identified via a millisecond remote scan of their irises. To escape detection, individuals must literally change their eyeballs. Thankfully, this scenario isn't likely to happen for some time because, for various reasons, biometric technologies are not well suited for large-scale *one-to-many* real-time identification purposes.

It is important to bear in mind that the collection of biometric samples and their processing into biometric templates for matching is subject to great variability. Simply put, biometrics are "fuzzy" – no two samples will be perfectly identical. Facial recognition technologies, for example, are notoriously prone to variability due to different lighting conditions, angle, subject movement, and so forth. This is the reason, for example, that we are asked not to smile in our passport photos. Similarly, numerous factors affect the ability to obtain reliable and consistent fingerprint samples. Among the various biometric types, irises seem to be the most accurate and consistent.

As a consequence, live biometric samples can be at some variance with stored reference samples, making comparison, matching, and identification an inexact process. In other words, biometric systems do not have 100 per cent accuracy. When the biometric system cannot perform a proper match and (incorrectly) rejects a legitimate user, this is called a *false reject,* and the user must typically resubmit one or more biometric samples for further comparison by the system.

Biometric system designers can and do take measures to lower the *false rejection rate* (FRR) of their systems so this variability is smoothed out and the system can function properly. Apart from controlling the conditions under which fresh samples are taken, and improving the mathematical algorithms, one way to do this is to lower the threshold for matches to occur. However, the difficulty with this approach is that this often increases the *false acceptance rate* (FAR) of the system, that is, the system will incorrectly match a biometric to the wrong stored reference sample, resulting in misidentification. Usually there is a tradeoff between FRR and FAR, i.e., one error rate may only be reduced at the expense of the other (for example, some applications require lower FRR but can tolerate higher FAR, and vice versa).

The FRR/FAR numbers quoted by biometric vendors are often unreliable. The reader is advised to consult reputable independent sources of information, such as, for example, biometric competitions organized by the U.S. National Institute of Standard (NIST)[2], or International Fingerprint Verification Competitions (FVC2000/2002/2004)[3]. For most biometric systems, FRR ranges from 0.1% to 20%, meaning that a legitimate user will be rejected from one out of 1,000 times to one out of five times on average. FAR ranges from one in 100 (low security applications) to one in 10,000,000 (very high security applications).

Other challenges for a biometric system are speed (the system must make an accurate decision in real time), and security (the system must be resilient against attacks).

So far, we have presented a straightforward technical discussion of the critical concepts of FAR and FRR. Now we will consider the operational consequences and impacts of these rates for *one-to-many* identification purposes.

Assume, for example, a biometric identification system with a 0.01% FRR and 0.0001% FAR (an unlikely high accuracy, we acknowledge). That is, the system is able to consistently match a genuine biometric sample 9,999 times out of 10,000 attempts on average. As remarkably efficient as this system sounds, a single biometric sample, when compared against a database of 1,000,000 samples, will generate on average one false accept in addition to one exact match (if the user was actually enrolled in the database).

2 **http://www.frvt.org/; http://fpvte.nist.gov/; http://fingerprint.nist.gov/minex04/**

3 **http://bias.csr.unibo.it/fvc2004/**

Now assume a database of 30,000,000 entries; each biometric sample would generate about 30 false accepts, each and every time! Clearly, this would be unacceptable for any real-time automatic identification system and would require significant human intervention in order to function.

Consequently, biometric system designers have resorted to other techniques to overcome the inherent technological problems of *one-to-many* identification. One way to significantly improve accuracy is to collect and compare *multiple* biometric samples. Multi-modal biometrics, for example, can involve collecting and using two (or more) fingerprints instead of one. If one fingerprint generates dozens or hundreds of false accepts, then the likelihood that two fingerprints will falsely match others in the database diminishes considerably. This is the primary reason behind emerging international requirements for including two separate biometrics (face and finger, for example), in machine-readable travel documents such as passports.

The privacy issue here, of course, involves the fact that more and more biometric samples of personal information need to be collected, transmitted, stored, and processed in order for the system to function properly. The FBI Integrated Automated Fingerprint Identification System (AFIS), containing hundreds of millions of records, for example, uses all 10 fingerprints for increased accuracy and speed. The US-VISIT program also plans to migrate from two fingerprints to 10 fingerprints and to develop the interoperability between US-VISIT and IAFIS.[4]

Significant privacy (and operational) concerns arise with unrestricted collection and use of more and more biometric data for identification purposes. To begin with, the creation of large centralized databases, accessible over networks in real time, presents significant operational and security concerns.

If networks fail or become unavailable, the entire identification system collapses. Recognizing this, system designers often build in high redundancy in parallel systems and mirrors (as well as failure and exception management processes) to ensure availability. However, this can have the effect of increasing the security risks and vulnerabilities of the biometric data.

Large centralized databases of biometric PII, hooked up to networks and made searchable in a distributed manner, represent significant targets for hackers and other malicious entities to exploit. It is also a regrettable reality that large centralized databases are also more prone to function creep (secondary uses) and insider abuse. There are also significant risks associated with transmitting biometric data over networks where they may be intercepted, copied, and actually tampered with, often without any detection.

4 **http://www.gao.gov/new.items/d07278.pdf**

Some large-scale biometric identification databases (such as the IAFIS, cited above) not only collect and file multiple biometric samples but, in an effort to preserve maximum compatibility with other fingerprint identification systems, store the full and complete *images* of the biometrics involved in addition to the templates! Proposed international standards for biometric-enabled machine-readable travel documents, for example, call for storage of the biometric *images* in the document rather than a structured reduction of the biometric into a unique template, in order to facilitate cross comparison and identification with other databases.

Storing, transmitting and using biometric *images* only exacerbates the privacy concerns with large-scale identification systems, since a very important privacy protection afforded by templates is removed, namely, the inability to *exactly* reconstruct the original biometric image from the template.

The image, conversely, can be converted into hundreds of templates for matching and identification (or other unknown or illegal) purposes, such as creating personal profiles and, let us not forget, committing identity theft. **At this point, the privacy implications explode.**

It should be evident that the loss or theft of one's biometric image opens the door to massive identity theft if the thief can use the biometric for his or her own purposes. For example, the ability to create low-cost duplicate fake fingerprints from "gummy bears," which are capable of fooling nine out of 10 biometric systems, has been well documented.[5] Others have even documented how easy it is to fool a biometric system by presenting it with a photograph! Of course, the biometric industry has come up with countermeasures, such as "liveness detection" of a finger, or capturing 3D face images, but so will the attackers in this perpetual game. Moreover, in the digital realm, there may be no need to even present a "fake finger" if all that is required is the digital equivalent, which can be supplied to the network instead.

Even worse, in all of these identification scenarios, the biometric effectively serves as an index or key to the database involved, much like login usernames serve to identify registered users of a computer network.

But, because people usually only have two thumbs, two eyes, and one head, it is nearly impossible to change these if and when the related biometric data become compromised. In this sense biometrics operate like shared secrets or passwords – learn the secret and you're in! But there are some very important differences between biometrics and passwords: you cannot change them and have no choice but

5 T. Matsumoto, H. Matsumoto, K. Yamada, S. Hoshino, "Impact of Artificial Gummy Fingers on Fingerprint Systems," Proceedings of SPIE Vol. #4677, *Optical Security and Counterfeit Deterrence Techniques IV*, 2002.

to keep them for life. Lose control of your lifetime password and you will have some explaining to do! This, regardless of the fact that security experts roundly condemn using unchangeable passwords as shared secrets (e.g., birthdates and SINs).

Views of the Privacy Community

The global privacy and data protection community have consistently argued *against* the use of biometrics for most *one-to-many* identification purposes, and *against* the creation of large, centralized or interoperable databases of biometric data:

- Resolution of International Data Protection Authorities;[6]
- Opinions of the European EDPS and Article 29 Working Party;[7] and
- Publications and testimony of Ontario Information and Privacy Commissioner.

The global privacy community has insisted on building privacy-enhancing technologies (PETs) directly into biometrics systems wherever possible, to ensure that they reflect the requirements of Fair Information Principles and Practices and applicable privacy laws regarding the collection, use and disclosure of PII. Privacy, consumer, and civil rights advocates around the world have strongly favoured limiting the use of biometrics for verification/authentication purposes, especially in distributed environments (where the biometric sample is retained by the user on a token, say, a smart card[8]).

Deployment Experience to Date

The reality is that the highly lauded use of privacy-enhanced *one-to-one* biometric authentication technologies has simply not been widespread. Perhaps the best-known example has been its deployment in laptop computers, where users must match their biometric (fingerprint) in order to gain access to the laptop.

Public sector government bodies, on the other hand, have tended to insist on building large-scale interoperable biometric databases. The reasons for this preference are complex and worthy of exploration in a separate research paper. Briefly, however, some possible explanations are as follows:

- The claim of overriding public interests or (secondary) purposes that override individual privacy interests. It is here that the "zero-sum" game mentality prevails, i.e., more individual privacy equals less public security, and vice versa;

6 International Data Protection Commissioners, "Resolution on the Use of Biometrics in Passports, Identity Cards and Travel Documents," Montreux (September 2005) available at: **www.edps.europa.eu/legislation/05-09-16_resolution_biometrics_EN.pdf**

7 See Appendix 1 for documents and sources.

8 In the "real" world, the template or biometric image would be stored in a database as a backup in case the user lost his or her card. Otherwise, users would have to re-enroll every time they misplaced or lost their token. However, these databases would be limited and not networked, and encrypted.

- Unwillingness of system designers and operators to relinquish control over biometrics to individual users. Here, too, adding privacy is often viewed as compromising system functionality, control, and effectiveness;
- Requirements to carry out more and more background checks (e.g., against criminal records, terrorist watch lists, etc.) or to prevent multiple identity registrations and benefits fraud (welfare, medicare, driver's licences, immigration applications, etc.);
- Need to retain evidence and to make a criminal case when necessary (only biometric images verified by a human expert are accepted by courts, not just templates);
- Backup needs and escrow requirements – copies of biometric data need to be retained on file and made available to system operators and other authorities "just in case" the system fails;
- Unavailability of suitable, reliable, and cost-efficient privacy-enhanced biometric technologies and systems;
- Unreliable biometric enrolment/verification procedures and practices, which undermine ALL biometric systems if attackers can fraudulently impersonate others;
- Strong pressure from technology vendors and/or advice from independent consultants and integrators who may lack incentives to pursue privacy-enhanced biometric system options;
- The simplistic conflation of privacy and security, i.e., the misguided (and erroneous) belief that all biometric privacy interests can be satisfied by building system controls that seek to ensure confidentiality and integrity of the biometric data. This is a very common problem among security professionals, who tend to undervalue privacy as a separate and unique set of design principles; and
- Weak public demand and guidance from the privacy and data protection communities.

The reader will note that most of these explanations are predicated on zero-sum game thinking, i.e., more individual privacy and user control equals less of virtually everything else! Taken from this view, building true biometric privacy into an information system is invariably seen as a cost, rarely as an enhancement.

A more common deployment scenario is to carry out *one-to-one* biometric authentication *against a single stored sample in a database*. For example, a biometric-enabled identity card may have a serial number that acts as an index or lookup

key to the database, calling up the biometric "password" for *one-to-one* comparison and authentication against a live sample.

Security Vulnerabilities of a Biometric System

Biometric systems, especially *one-to-one*, may become vulnerable to potential attacks.[9]

Some of those security vulnerabilities include the following:

- **Spoofing.** It has been demonstrated that a biometric system can sometimes be fooled by applying fake fingerprints, face or iris image, etc.
- **Replay attacks,** e.g., circumventing the sensor by injecting a recorded image in the system input – much easier than attacking the sensor.
- **Substitution attack.** The biometric template must be stored to allow user verification. If an attacker gets an access to the storage, either local or remote, he can overwrite the legitimate user's template with his/her own – in essence, stealing their identity.
- **Tampering.** Feature sets on verification or in the templates can be modified in order to obtain a high verification score, no matter which image is presented to the system.
- **Masquerade attack.** It was demonstrated[10] that a digital "artefact" image can be created from a fingerprint template, so that this artefact, if submitted to the system, will produce a match. The artefact may not even resemble the original image. This attack poses a real threat to the remote authentication systems (e.g., via the Web), since an attacker does not even have to bother to acquire a genuine biometric sample. All he needs is just to gain an access to the templates stored on a remote server (this perfectly fits a description of a typical hacker operating from a rat hole).
- **Trojan horse attacks.** Some parts of the system, e.g., a matcher, can be replaced by a Trojan horse program that always outputs high verification scores.
- **Overriding Yes/No response.** An inherent flaw of existing biometric systems is due to the fact that the output of the system is always a binary Yes/No (i.e., match/no match) response. In other words, there is a fundamental disconnect between the biometric and applications, which makes the system open to potential attacks. For example, if an attacker were able to interject a false Yes response at a proper point of the communication between the biometrics and the

9 N.K. Ratha, J.H. Connell, R.M. Bolle. "Enhancing Security and Privacy in Biometrics-Based Authentication Systems". *IBM Systems Journal*, vol. 40, no. 3, pp. 614-634, 2001.

10 C.J. Hill, "Risk of masquerade arising from the storage of biometrics," B.S. Thesis, Australian National University, 2001 (supervisor, Dr. Roger Clarke). **http://chris.fornax.net/biometrics.html**

Figure 1: Privacy and security issues involving a biometric system

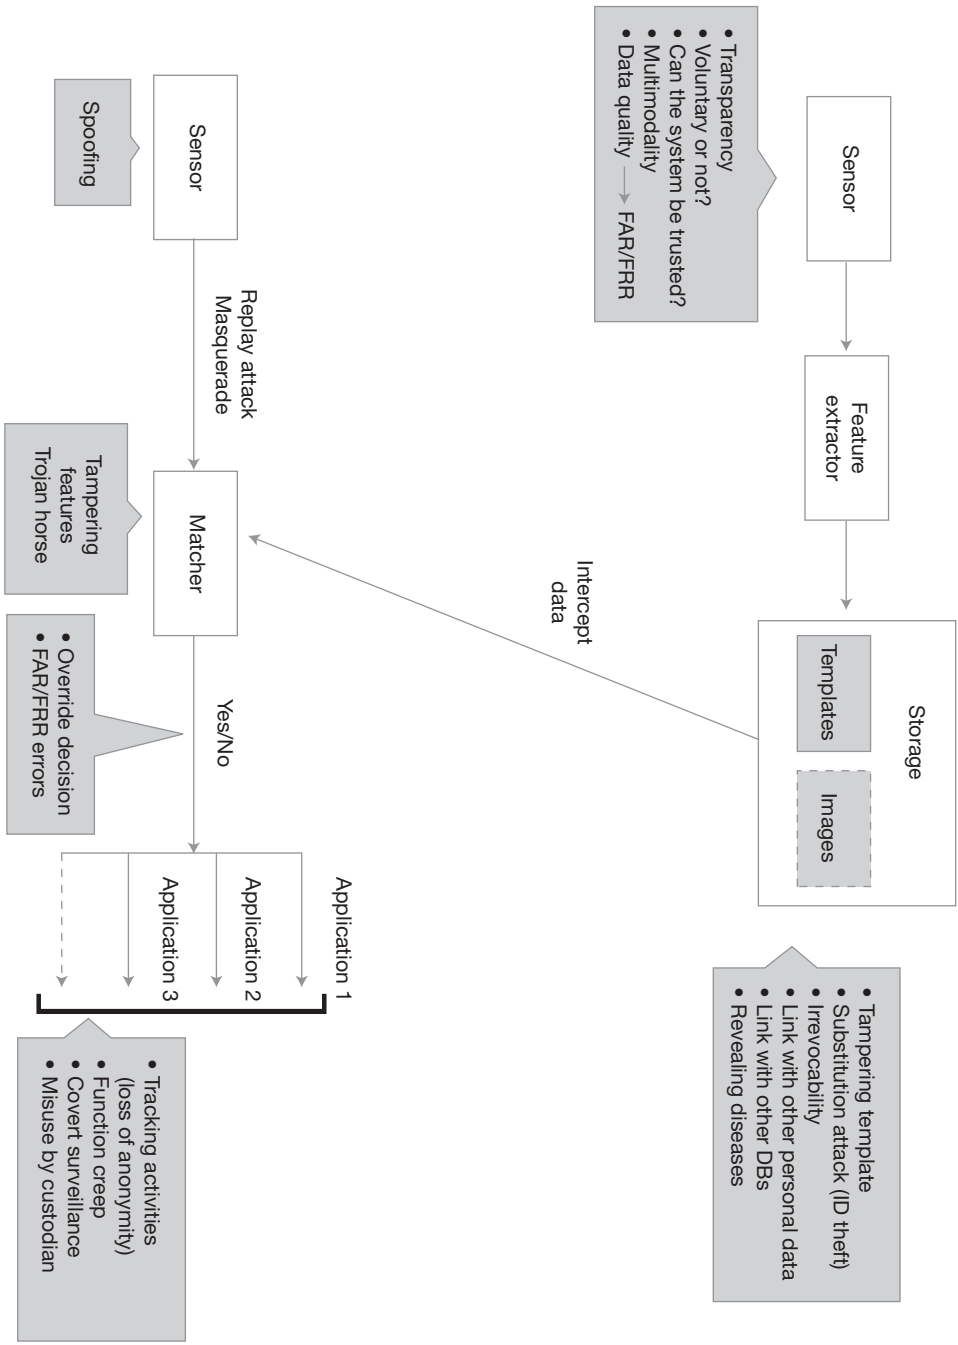

application, he could pose as a legitimate user to any of the applications, thus bypassing the biometric part.

- **Insufficient accuracy of many commercial biometric systems**, both in terms of FRR and FAR. High FRR causes inconvenience for legitimate users and prompts the system administrator to lower a verification threshold. This inevitably gives rise to FAR, which, in turn, lowers the security level of the system.

The privacy and security issues of a biometric system outlined in this section are illustrated in Fig. 1.

An enrolment part of any conventional biometric system consists of at least three blocks: a biometric sensor that acquires an image, a feature extractor that creates a biometric template, and a storage for the templates or images, or both. The storage can be either a database or a distributed medium.

A verification or identification part contains (at a minimum) a sensor to acquire a new image sample, and a matcher, which compares the image with the previously enrolled template(s) received from the storage. The output of the matcher is a Yes/No (i.e., match/no match) response that may go to the variety of applications.

A user of the system faces several privacy issues immediately at enrolment:

- Transparency, i.e., if the purpose of the system is clear to the user;
- If the enrolment is voluntary, and what the consequences are of not getting enrolled (for a variety of reasons);
- If the system can be trusted, i.e., if the personal data are adequately protected;
- Quality of biometric data: poor quality may lead to higher FRR and FAR. While FAR increases security risks for the system, a false rejection often causes some follow-up procedures, which can be privacy-invasive to the individual.

Other privacy/security issues were explained in the foregoing sections.

Biometric Encryption

Biometrics and Cryptography

Conventional cryptography uses encryption keys, which are just bit strings long enough, usually 128 bit or more. These keys, either "symmetric," "public," or "private," are an essential part of any cryptosystem, for example, Public Key Infrastructure (PKI). A person cannot memorize such a long random key, so that the key is generated, after several steps, from a password or a PIN that can be memorized. The password management is the weakest point of any cryptosystem, as the password can be guessed, found with a brute force search, or stolen by an attacker.

On the other hand, biometrics provide a person with unique characteristics which are always there. Can they be used as a cryptographic key? Unfortunately, the answer is negative: biometric images or templates are variable by nature, i.e., each new biometric sample is always different. Conventional cryptography does not tolerate a single bit error.

As noted in the previous section, a biometric system always produces a Yes/No response, which is essentially one bit of information. Therefore, an obvious role of biometrics in the conventional cryptosystem is just password management, as mentioned by Bruce Schneier.[11] Upon receiving a Yes response, the system unlocks a password or a key. The key must be stored in a secure location (so-called "trusted" device). This scheme is still prone to the security vulnerabilities noted in Fig. 1, since the biometric system and the application are connected via one bit only.

Biometric templates or images stored in a database can be encrypted by conventional cryptographic means. This would improve the level of system security, since an attacker must gain the access to the encryption keys first. However, most privacy issues associated with a large database remain, since the keys and, therefore, the biometric data are controlled by a custodian.[12]

A comprehensive review of the issues involving biometrics and cryptography can be found elsewhere.[13]

11 B. Schneier, "The Uses and Abuses of Biometrics," Comm. ACM, vol. 42, no. 8, p. 136, Aug. 1999.

12 There has been recent activity of the International Organization for Standardization in order to support the confidentiality and integrity of the biometric template by using cryptographic means (ISO/IEC WD 24745, "Biometric Template Protection"): **www.nia.din.de/sixcms/media.php/1377/SC27N4997rev1_SD7_Catalog_Proj&Stand_May2006.htm?backend_call=true#24745; www.incits.org/tc_home/CS1/2007docs/cs1070006.pdf**

13 "Future of Identity in the Information Society" (FIDIS) report, "D3.2: A study on PKI and biometrics," 2005. **www.fidis.net/fileadmin/fidis/deliverables/fidiswp3del3.2.study_on_PKI_and_biometrics.pdf**

What Is Biometric Encryption?

Because of its variability, the biometric image or template itself cannot serve as a cryptographic key. However, the amount of information contained in a biometric image is quite large: for example, a typical image of 300x400 pixel size, encoded with eight bits per pixel has 300x400x8 = 960,000 bits of information. Of course, this information is highly redundant. One can ask a question: Is it possible to consistently extract a relatively small number of bits, say 128, out of these 960,000 bits? Or, is it possible to bind a 128-bit key to the biometric information so that the key could be consistently regenerated? While the answer to the first question is problematic, the second question has given rise to the new area of research, called Biometric Encryption (BE).[14]

Biometric Encryption is a process that securely binds a PIN or a cryptographic key to a biometric so that neither the key nor the biometric can be retrieved from the stored template. The key is recreated only if the correct live biometric sample is presented on verification.

> **"In Biometric Encryption, you can use the biometric to encrypt a PIN, a password, or an alphanumeric string for numerous applications – to gain access to computers, bank machines, to enter buildings, etc. The PINs can be 100s of digits in length; the length doesn't matter because you don't need to remember it. And most importantly, all one has to store in a database is the biometrically encrypted PIN or password, not the biometric template."**
>
> Dr. George Tomko, OECD *Report on Biometric-Based Technologies* (2004)[15]

The digital key (password, PIN, etc.) is randomly generated on enrolment so that the user (or anybody else) does not even know it. The key itself is completely independent of biometrics and, therefore, can always be changed or updated. After a biometric sample is acquired, the BE algorithm securely and consistently binds the key to the biometric to create a protected BE template, also called "private template." In essence, the key *is encrypted* with the biometric. The BE template provides excellent privacy protection and can be stored either in a database or locally (smart card, token, laptop, cellphone, etc.). At the end of the enrolment, both the key and the biometric are discarded.

14 Other terms used for this technology: biometric cryptosystem, private template, fuzzy commitment scheme, fuzzy vault, fuzzy extractor, secure sketch, biometric locking, biometric key binding, biometric key generation, virtual PIN, biometrically hardened passwords, biometric signature, bioHashing. We use the term "Biometric Encryption" in a broad sense.

15 OECD *Report on Biometric-Based Technologies* (June 2004). Directorate for Science, Technology and Industry, Committee for Information, Computer and Communications Policy, DSTI/ICCP/REG(2003)2/FINAL, p. 64.

On verification, the user presents her fresh biometric sample, which, when applied to the legitimate BE template, will let the BE algorithm retrieve the same key/password. In other words, the biometric serves as a *decryption key*. At the end of verification, the biometric sample is discarded once again. The BE algorithm is designed to account for acceptable variations in the input biometric. On the other hand, an attacker, whose biometric sample is different enough, will not be able to retrieve the password. This encryption/decryption scheme is *fuzzy,* as the biometric sample is different each time, unlike an encryption key in conventional cryptography. Of course, it is a big technological challenge to make the system work.

After the digital key, password, PIN, etc., is retrieved, it can be used as the basis for any physical or logical application. The most obvious way lies in the conventional cryptosystem, such as a PKI, where the password will generate a pair of Public and Private keys.

Thus, Biometric Encryption is an effective, secure, and privacy-friendly tool for biometric password management, since the biometric and the password are bound on a fundamental level.

Advantages of Biometric Encryption (over Other Biometric Systems)

Biometric Encryption technologies have enormous potential to enhance privacy and security. Some of the key benefits and advantages of this technology include:

1 No retention of the biometric image or template

From a privacy perspective, the best practice is not to collect any personally identifiable information (PII) at all in the first place, to the fullest extent possible. This is referred to as "data minimization" – minimizing the amount of personal data collected and retained, thus eliminating the possibility of subsequent abuse.

Most privacy and security concerns derive from storage and misuse of the biometric data.

A common concern is that "if you build it [the database], they will come [for the data]." The topline privacy and security concerns include fears of potential data matching, surveillance, profiling, interception, data security breaches, and identity theft by others. Misuse and mismanagement of biometric data by others invokes "negative externalities" and costs that fall primarily upon individuals rather than the collecting organization, but also at stake is the accountability and credibility of the collecting organization, and with them, the viability of the entire program.

Biometric Encryption directly addresses these risks, threats and concerns.

Users retain complete (local) control and use of their own biometrics.

Local control enhances confidence and trust in the system, which ultimately promotes greater enrolment and use.

2 Multiple/cancellable/revocable identifiers

Biometric Encryption allows individuals to use a single biometric for multiple accounts and purposes without fear that these separate identifiers or uses will be linked together by a single biometric image or template.

Thus, if a single account identifier becomes compromised, there is far less risk that all the other accounts will also be compromised.

Even better, Biometric Encryption technologies make possible the ability to change or recompute account identifiers. That is, identifiers may be revoked or cancelled, and substituted for newly generated ones calculated from the same biometric!

Traditional biometric systems simply cannot do this.

3 Improved authentication security: stronger binding of user biometric and identifier

Account identifiers are bound with the biometric and recomputed directly from it on verification. This results in much stronger account identifiers (passwords) because:

- they are longer and more complex;
- there is no need for user memorization; and
- they are less susceptible to security attacks.

Many security vulnerabilities of a biometric system listed in Fig. 1 are addressed:

No substitution attack: An attacker cannot create his own template since he, or anybody else, does not know the digital key and other transitory data that had been used to create the legitimate template;

No tampering: Since the extracted features are not stored, the attacker has no way to modify them;

No masquerade attack: Again, the system does not store the biometric template, so the attacker cannot create a digital artefact to submit to the system. Biometric Encryption provides an effective protection for remote authentication systems;

No Trojan horse attacks: BE algorithm does not use any score, either final or intermediate, to make a decision; it just retrieves (or does not retrieve) a key. Therefore, the attacker has no means to fool the system by outputting a high score;

No overriding Yes/No response: The output of BE algorithm is a 128-bit (or longer) digital key, as opposed to the binary Yes/No response. The attacker cannot obtain the key from a private template.

The security of Biometric Encryption technology can be augmented by the use of tokens (e.g., smart cards, PDA) and additional PINs, if needed.

4 Improved security of personal data and communications

As an added bonus, users can take advantage of the convenience and ease of Biometric Encryption technologies to encrypt their own personal or sensitive data. See Case Study #1 for an example.

Since the key is one's own biometric, used locally, this technology could place a powerful tool directly in the hands of individuals.

Biometric Encryption could be viewed as encryption for the masses, made easy!

5 Greater public confidence, acceptance, and use; greater compliance with privacy laws

Public confidence and trust are necessary ingredients for the success of any biometric system deployment. One major data breach or horror story involving a large centralized database of biometric templates could set back the entire industry for years.

Data governance policies and procedures can only go so far to foster public trust. However, if privacy, security, and trust can be built directly into the biometric system, then the public and data protection authorities are far more likely to accept the privacy claims being made.

Putting biometric data firmly under the exclusive control of the individual, in a way that benefits that individual and minimizes risk of surveillance and identity theft, will go a long way toward satisfying the requirements of privacy and data protection laws, and will promote broader acceptance and use of biometrics.

6 Suitable for large-scale applications

Biometric Encryption technologies speak directly to the clear preference and recommendations of the privacy and data protection authorities for using biometrics to authenticate or verify identity, rather than for identification purposes alone.

Therefore, we prefer seeing biometrics used to positively link the bearer to a card or token, and to avoid creating systems that rely upon centralized storage and remote access/lookup of biometric data.

A prevailing reason for this view is that it is not known if biometric technology is sufficiently accurate and reliable to permit real-time identification in large *n* samples, where *n* is of an order of several million or higher. Despite these views, many large-scale *one-to-many* public biometric projects are being proposed and are well underway.

Often the biometric data in these systems are actually used for authentication purposes and not identification, but the lines between these two concepts can be blurred when multiple data items are collected and transmitted to a database for comparison. What becomes the identifier and what becomes the authenticator is somewhat arbitrary.

From a privacy point of view, transmitting biometric image or template data to a central database to be authenticated is risky enough without compounding the risks by sending more and more personal identifiers with it. "Multimodal" biometric solutions depend on collecting and comparing more than one biometric. It should be noted that the main reason for using "multimodal" solutions, besides providing a fallback for problem users, is insufficient accuracy/speed/security of existing biometrics. So the technical "solution" to using biometrics for authentication seems to be to collect more and more biometric and other personal data.

In 2006, the European Data Protection Supervisor (EDPS), Peter Hustinx, warned, in a formal opinion, of the privacy dangers of using biometric images or templates as an index or key to interoperable databases.[16]

Fortunately, Biometric Encryption technologies make possible database applications (see Case Study #3 as an example), minimizing the risks of traditional biometric systems (although we still prefer *one-to-one* applications with local template storage). It is possible to create secure and local biometric-enabled bindings of users to some other token identifiers without the need to reveal the actual biometric image or data.

It is further possible to create a so-called "anonymous database," where a link between an anonymous identifier and encrypted (by conventional cryptographic means) user's record is controlled by a Biometric Encryption process. This is very useful for a database containing sensitive information, such as medical records (see Case Study #2 for more details).

Another promising application of BE is a privacy-protected *one-to-many* database for "double dipping" prevention. The database is multimodal: it contains conventional but anonymous templates for one biometric (e.g., fingerprints) and private templates (e.g., for iris) that control a link with the user's encrypted records. A user's record would only be decrypted and displayed if there was a positive match on both conventional and private templates. Otherwise, all the information is inaccessible even to the system administrator.

16 See Appendix 1 for references and URLs.

With Biometric Encryption, users would be empowered by the ability to securely prove who they are to anyone, for any purpose, using their own biometrics, but without having to disclose the biometric data itself!

A high-level diagram of a Biometric Encryption process is shown in Fig. 2 (next page).

An enrolment part of a Biometric Encryption system consists of at least four blocks: a biometric sensor, a key generator that normally outputs a random key, a binding algorithm that creates a BE (private) template, and a storage for the BE template. Neither the key nor the image can be recovered from the BE template. The key, the image, and some transitory data are discarded at the end of the enrolment process.

A verification part contains at least a sensor to acquire a new image sample, and a key retrieval algorithm, which applies the image to the previously enrolled BE template received from the storage. The algorithm either retrieves the key, if the image on verification is close enough to the one enrolled, or fails to do so, in which case the user is rejected. The key enters an application, such as a PKI. Each application has its unique key. The biometric image is discarded at the end of the verification process.

Figure 2: High-level diagram of a Biometric Encryption process

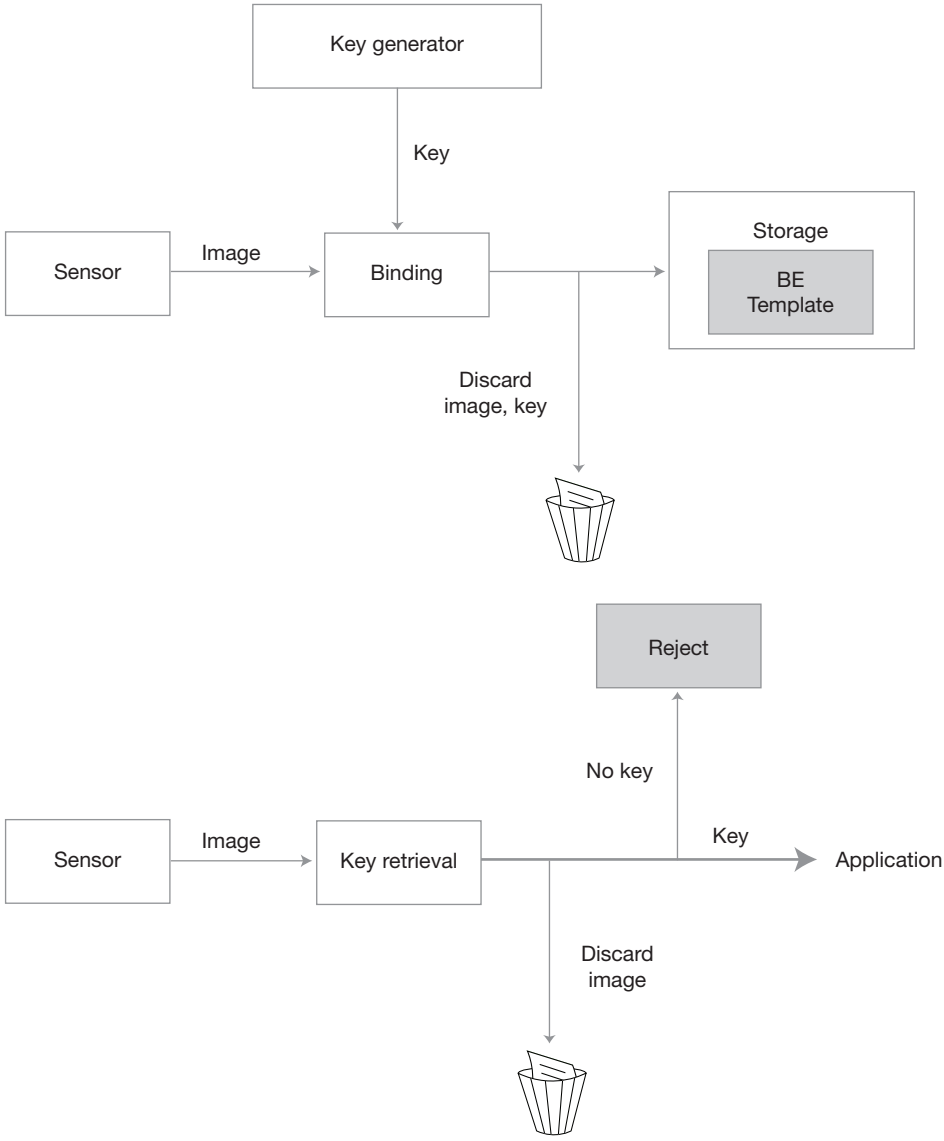

Current State of Biometric Encryption

The original concept of Biometric Encryption for fingerprints was pioneered in 1994 by Dr. George Tomko, founder of Mytec Technologies (Toronto, Canada). Since then, many research groups have taken part in the development of BE and related technologies. There are about 50 articles and patents published to date, most of which have appeared since 2002.The list of publications, with a brief review, is presented in Appendix 2.

Besides Biometric Encryption (BE), other terms have been used for this technology, such as: biometric cryptosystem, private template, fuzzy commitment scheme, fuzzy vault, fuzzy extractor, secure sketch, biometric locking, biometric key binding, biometric key generation, virtual PIN, biometrically hardened passwords, biometric signature, and bioHashing.

BE and related technologies have drawn attention from major academic research centres specializing in biometrics, such as Michigan State University, West Virginia University, Carnegie Mellon University, University of Cambridge (U.K.), and University of Bologna (Italy). Among current industry leaders, those worth noting include IBM T.J. Watson Research Center, RSA Laboratories, Lucent Technologies, Sandia National Laboratories, and Philips Research.

Virtually all types of biometrics have been tested to bind (or to generate) a digital key: fingerprints, iris, face, keystroke dynamics, voice, handwritten signatures, palm prints, acoustic ear recognition. The most promising results have been achieved with an iris: FRR = 0.47%, FAR = 0 (or at least less than one in 200,000) to generate a 140-bit key. These error rates are only marginally larger than for a conventional iris-based biometric system with the same input images[17]. The use of fingerprints is also feasible in terms of accuracy for BE, with FRR greater than 10% at present. Unlike an iris, there is a noticeable degradation in accuracy from a conventional fingerprint system. This is understandable since fingerprints are more prone to distortions and other factors that degrade accuracy. It is more difficult to compensate those factors in the case of Biometric Encryption, since BE works in a "blind" mode (the enrolled fingerprint or its minutiae template are not seen). There are several ways to overcome this problem, for example, by using a free air (i.e., contactless) fingerprint sensor, or by using more than one finger from the same person, or by combining several biometrics.[18]

Face recognition, which is usually considered third (after irises and fingerprints) in terms of accuracy in conventional biometrics, has shown a significant

17 The iris images were acquired in close to ideal conditions of a laboratory environment. In real life systems, some degradation of performance is expected, which is always the case with biometrics.

18 Note that even a 10% to 20% false rejection rate still may be acceptable for some applications with relatively low traffic and cooperative users: it simply means that a person would be rejected each fifth or tenth time on average and asked by the system to place the finger on the reader again.

improvement of performance over the last few years. This allowed Philips Research to create a working BE system using a face biometric. The published results range from FRR = 3.5% for a face database with low to medium variability of images to FRR = 35% for a database with high variability; FAR = 0 (or at least less than 1 in 100,000) in both cases. The key size used is 58 bits, which may be sufficient as a password replacement. According to communication from Dr. Michiel van der Veen of Philips Research, their technology, called privID™, is now operational and ready for deployment; in particular, it will be a part of a EU 3D Face project (WP2.5)[19]. To the best of our knowledge, the Philips system will be the first real-life application of BE technology.

It is not clear if other biometrics have enough entropy (i.e., the amount of non-redundant information) in order to bind a sufficiently long key (e.g., 128 bit). This is an area of future research.

Some works published since 2002 provide a general theoretical foundation for BE technologies from a cryptographic point of view. They prove that the system can be made secure against "brute force" search attacks. In other words, an attacker checks at random all possible combinations in order to retrieve a key (or a biometric). Like conventional cryptography, it is assumed that the attacker is fully familiar with the algorithm, and may have a template in hand, but does not have a proper biometric to unlock the secret (i.e., the key bound to the biometric).

However, the attacker may try more sophisticated attacks, exploiting inherent weaknesses (if any) of the BE system and biometrics in general. This area of research has been largely overlooked. If such an attack is successful, the effective security of the system would be reduced from 128 bits to, perhaps, 69, 44, or an even lower number of bits. "This may seem an alarmingly small number to the crypto purist" (Hao, Anderson, and Daugman, 2005). On the other hand, BE is not just another cryptographic algorithm; it is rather a key/password management scheme. Key management has always been the weakest part of any cryptosystem, as it relies on passwords that may be forgotten, stolen, guessed, shared, etc. Biometric Encryption binds the key/password with the biometric and, thus, makes the system more secure. By comparison, a conventional biometric has only 1-bit security – a Yes/No response!

It is interesting to note that code-breaking becomes reduced to a security problem, not a privacy issue with BE, e.g., with an encrypted database of templates, breaking the encryption key exposes all the templates, and one has both a security and a privacy issue. Breaking a biometrically encrypted key, however, only exposes that key, but not necessarily the biometric, let alone the entire database, making it a far more secure system.

19 **www.3Dface.org**

With the notable exception of Philips privID™, to the best of our knowledge, there is no other commercially available BE system being used to date. The reason for this lies in both the technological challenges and existing market conditions. Not only the general public, but most high-tech developers are unaware of this emerging technology. Consequently, resources and funding in this area have, to date, been quite poor. We believe that the technological challenges have largely been overcome using an iris or face, and partially for fingerprints, bringing BE technology very close to the prototype development stage, and could soon be ready for testing in pilot projects.

Related Technologies

1. Storing a key in a trusted system

There have been some products[20] that store a cryptographic key or a PIN in a so-called trusted system (e.g., a computer or a Digital Signal Processor (DSP)). The key is released upon successful biometric verification and then enters a conventional cryptosystem, e.g., Public Key Infrastructure (PKI). The biometric template (or image) is also stored somewhere, often in encrypted (by conventional means) form.

If properly implemented, such systems may offer some security benefits. However, most problems outlined in the foregoing sections remain. For example, a binary Yes/No response is still required to release the key – this part of the algorithm is just hidden better. Most privacy issues associated with the template storage are also there.

Note that these systems often use the same terminology and/or claim the same benefits as BE, while in fact they do not provide a true binding between a key and a biometric.

2. Cancellable biometrics

A new area of research, closely related to BE, is called cancellable biometrics. It has been developed by IBM T.J. Watson Research Center, and by some academic groups. In this privacy-protecting technology, a distortion transform (preferably, irreversible) is applied to a biometric template. Only those distorted templates are stored, and they are matched also in the distorted form. If a distorted template is compromised, it can be "cancelled" by choosing just another distortion transform (i.e., the biometric is not lost). The transforms are application dependent, meaning that the templates cannot be reused by another applications (function creep is prevented).

20 See, for example:
www.ceelox.com; www.sequiam.com; www.lacie.com/products/product.htm?id=10166; and
www. axistech.com/Biomeric_Time_attandance_Axis_Technology_Encryption.asp

Cancellable biometrics shares some other similarities with BE; for example, a technique called bioHashing can be used for both technologies. Unlike BE, a key is not generated or released in cancellable biometrics, so the system still produces a binary Yes/No response and is more vulnerable to attacks. The distortion transform should be truly irreversible (i.e., one way only) and kept secret. Otherwise, an attacker can either reconstruct the original biometric or create his own impostor template for a substitution attack, or even create an "artefact" image for a masquerade attack. Since the key is not generated, the variety of potential applications is narrower than for BE; for example, an anonymous database cannot be created. On the other hand, BE possesses all the functionality of cancellable biometrics, and, therefore, is a *method* for cancellable biometrics. Both technologies face similar accuracy/security challenges.

3. Fuzzy Identity Based Encryption

Another related technology, called Fuzzy Identity Based Encryption (FIBE), was proposed by A. Sahai and B. Waters in 2005. This technology also combines biometrics and cryptography on a fundamental level. Unlike BE, the user's biometric is made somewhat public. In an example provided by D. Nali, C. Adams and A. Miri (see also a webcast presentation by B. Waters)[21], a user (*A*) could go to a Driver Licensing Agency (*D*) and identify herself via an iris scan, under the ongoing surveillance of a trained agent. *D* could then use this scan to encrypt *A*'s information (e.g., an annual driver's licence), when this information needs to be securely sent to *A* (e.g., via the Web). In order to obtain her biometric *private keys*, *A* would have to go in person to a trusted third party (e.g., a state agency), which would deliver keys via the same authenticating procedure as that used by *D. A* could then decrypt the message addressed to her using FIBE. She does not need a biometric reading at that point. In other words, *A* leaves her biometrics in at least two places, *D* and the trusted third party (often called Trusted Authority (TA)).

This scheme prevents impersonation of A by surreptitiously capturing her biometric sample, such as an iris photograph or latent fingerprints. "FIBE allows biometric measurements to be public" (Nali, Adams, and Miri) and, therefore, those surreptitious samples would become useless. While interesting from a scientific point of view, this technology is not privacy protecting, at least in the sense adopted by the privacy community (biometric data are considered personal information). There are also problems in handling a false rejection: user *A* may not have a chance to present another biometric sample if the false rejection occurs during decryption.

21 **http://www.researchchannel.org/prog/displayevent.aspx?rID=3913**

Scientific, Technological, and Privacy-Related Merits

Encryption with a fuzzy key (such as a biometric) was only recently introduced in conventional cryptography. Beyond such trivial things like accepting a few spelling errors in a password, or letting Alice partially share a list of her favourite movies with Bob, Biometric Encryption technologies are by far the most important application of those theoretical works. Market demand for such a technology would provide a great incentive to this promising area of modern mathematics and cryptography.

BE results in tougher requirements for distortion tolerance, discrimination, and the security of a biometric system. Solving these problems would be a significant scientific breakthrough both in the area of biometrics and cryptography. This would accelerate research and development of better biometric sensors and other hardware, as well as new, more accurate algorithms and software. No doubt this would bring technological benefits for the entire biometrics.

BE overcomes many security vulnerabilities of a biometric system, especially in a distributed environment. This could facilitate deployment of biometric systems on portable and handheld devices (laptops, cellphones, PDAs, etc.).

It would not be an overstatement to say that biometrics is perceived, in general, as a privacy-invasive technology. As we have shown, this perception is not baseless. Biometric Encryption, on the other hand, is a *privacy-enhancing technology*. It allows a user to retain full control over her biometric and, at the same time, to stay anonymous in many applications, i.e., to be represented only by a randomly generated (and cancellable) identifier linked to her biometric. No other personal data, e.g., address, telephone, date of birth, have to be revealed.

BE can render databases privacy protected, as they will comprise "private templates." While such databases cannot be used for a background check, they are perfectly suitable for *one-to-one* access control systems or even for systems to prevent multiple registrations and related fraud. The user regains control over his or her sensitive information, such as medical or financial records, stored in the database.

Proliferation of BE technology may ultimately change the public's perception of biometrics. This would raise the benchmark for biometric technologies, such that the industry would be prompted to develop and adopt new privacy-friendly solutions. If the "private templates" generated by BE make a significant presence in the market, this could reshape the entire biometric industry. Increased user acceptance and confidence would be extremely beneficial for the industry.

Case Study #1: Small-scale use of Biometric Encryption

To demonstrate the power of BE, we will briefly present a biometric authentication protocol (remote or local) with third party certification. We use a simplified and re-worded description from Boyen's paper on Fuzzy Extractors.[22]

Suppose that Alice wishes to authenticate herself to Bob using biometrics. Due to privacy concerns, she does not wish to reveal any biometric information to Bob. Conversely, for the authentication to be meaningful, Bob wants some assurance that Alice is in fact in possession of her purported biometrics at the time the authentication is taking place (i.e., that no one is impersonating her). We assume that there is a third party (often called the Trusted Authority), Trent, whom Bob trusts to honestly certify Alice's biometrics, and to whom Alice will temporarily grant access to her biometrics for the purpose of generating such a certificate. Alice will want to be able to obtain as many or as few of those certificates as she wants, and to reuse as many of them with multiple Bobs, some of whom may be even dishonest, without fear of privacy leaks or risk of impersonation. The protocol is as follows:

Enrolment and certification: Under Trent's supervision, and using Alice's own biometric:

1 Alice creates a Biometric Encryption template from her biometric and a randomly selected PIN. Neither the biometric nor the PIN can be recovered from the template;

2 The PIN is used to generate a pair of keys, called *public* and *private keys*;

3 The biometric, the PIN, and the *private key* are discarded;

4 If Trent is satisfied that Alice has executed the steps honestly, he certifies the binding between Alice's name and the *public key*, i.e., he digitally signs the pair ["Alice," *public key*]. At this point, Alice may send the *public key* to Bob, or even publish it for all to see.

Verification: A challenge/response scheme is used to verify Alice:

1 At any time when appropriate (e.g., whenever Alice desires to authenticate herself to Bob), Bob sends Alice a fresh random challenge;

2 By obtaining her new biometric sample and applying it to her Biometric Encryption template, Alice recovers on-the-fly her PIN, which, in turn, regenerates her *private key*;

3 Alice signs the challenge with her *private key* and gives Bob the signature;

4 Bob authenticates Alice by checking the validity of the signature under her authentic *public key*.

The protocol does not require Alice to remember or store her PIN or her *private key*.

22 X. Boyen, "Reusable Cryptographic Fuzzy Extractors," CCS 2004, pp. 82–91, ACM Press.

The Biometric Encryption template may be stored on a smart card or in Alice's laptop, which also has a biometric sensor. For different applications ("multiple Bobs"), a new pair of *public* and *private keys* is generated from the PIN. Those keys are periodically updated. Some applications may require different PINs, in which case several Biometric Encryption templates can be stored. A proper template can be automatically recognized by the application.

The system based on digital signatures may be adopted both for a remote and local access. The important point is that the most critical part of any cryptosystem, the PIN (or a password), is securely bound to the biometrics.

In summary, Alice has in her possession and under her control as many BE templates as necessary. She can use them to digitally sign in, either for remote authentication or for logical or physical access. The authentication is done simply by checking the validity of her digital signature using standard cryptographic means. Neither Alice's biometric nor her PIN is stored or revealed. As a result, the system is both secure and highly privacy protective.

Case Study #2: Anonymous database; large or medium-scale applications

Suppose that a clinic, a hospital, or a network of hospitals maintains a database of medical records. Alice does not want her record to be accessed by unauthorized personnel or third parties, even for statistical purposes. For that the latter, her record is made anonymous and encrypted (by conventional means). The only public entry in the database is her personal identifier, which may be her real name or, in certain cases (e.g., drug addiction clinic), an alias ("Jane Doe"). The link between Alice's identifier and her medical record is controlled by Biometric Encryption:

On enrolment, a BE template is created from Alice's biometric and a randomly generated PIN (Alice does not even know the PIN). The PIN is used to generate a pointer to Alice's medical record and a crypto-key that encrypts the record, and also a pair of keys called *public* and *private keys* (similar to Case Study #1). The BE template and the *public key* are associated with Alice's ID and stored in the database (they can also be stored on Alice's smart card); other temporary data, such as Alice's biometric, the PIN, the *private key*, the pointer, and the crypto-key, are discarded.

Suppose that Alice visits a doctor, to whom she wants to grant remote access to her medical record, or part of it, if the record is structured. From the doctor's office, Alice makes a request to the database administrator, Bob. The authentication procedure using challenge/response scheme is similar to that in Case Study #1:

1 If Alice does not have her smart card with her (e.g., in the case of an emergency), Bob sends Alice's BE template to the doctor's office;

2 Alice applies her new biometric sample to the BE template and recovers on-the-fly her PIN;

3 The PIN is used to regenerate her *private key*, the pointer to her medical record, and the crypto-key;

4 Bob sends Alice a fresh random challenge;

5 Alice signs the challenge with her *private key* and gives Bob the signature;

6 Bob authenticates Alice by checking the validity of the signature under her *public key*;

7 Alice securely sends Bob the pointer to her medical record;

8 Bob recovers Alice's encrypted medical record (or a part of it, also encrypted) and sends it to Alice;

9 Using her crypto-key, which was regenerated from her PIN, Alice decrypts her medical record for the doctor;

10 Alice's biometric, the PIN, the *private key*, the pointer, and the crypto-key, are discarded.

In summary, Bob (the database administrator) has an assurance that Alice is, in fact, who she claims to be (she was able to unlock her BE template in the doctor's office); he is also assured that her medical record was sent to the right person. On the other hand, Alice retains full control over her medical record, so that even Bob has no access to it, since he does not have the crypto-key to decrypt it. **The privacy protection is embedded into the system at a very basic technological level.**

Case Study #3: Travel documents; large-scale database applications

Using biometrics for travel documents has been a hot topic of discussion. To illustrate how BE can protect the user's privacy and, at the same time, improve the level of security, we will consider a reworded description of a system proposed by Dr. van der Veen et al. (Ref. [40] in Appendix 2).

The International Civil Aviation Organization (ICAO) dictates international standards for Machine Readable Travel Documents (MRTD), including those for ePassports. Among the recommendations is the "three-way-check" for secure verification at a border crossing. It involves comparing data originating from (i) the biometric sensor, (ii) the biometric image stored on the ePassport, and (iii) biometric data stored in external (centralized) databases.

BE technology provides the opportunity to do this in a privacy-preserving manner: in addition to the biometric templates stored on the ePassport, their secure versions, namely, the BE templates, are also stored in a third-party database. The biometric images or conventional templates are not stored in the database. A "three-way check" is then performed by matching the BE template from the database to that appearing on the ePassport, and the live biometric measurement scanned at the kiosk. Border passage now involves the following steps:

1 At a kiosk, a user claims his identity (ID), and presents his biometric (e.g., facial image, fingerprint or iris) for measurements;

2 The ID is sent to the third-party database to extract the corresponding BE template;

3 The BE template is transmitted to the kiosk;

4 The BE template and the biometric measurement are combined to derive a cryptographic key, or rather a hashed version of it;

5 The image of the iris, face or fingerprint is extracted from the ePassport and used together with the BE template to derive another hashed version of the cryptographic key. This will validate the biometric stored on the ePassport;

6 Both hashed versions of the key derived in Steps 4 and 5 are transmitted to the border-control authority and verified against the database version. A positive authentication is achieved when all three versions are exactly the same.

In summary, the user's privacy is protected since the biometric image or template is not stored in a central database; instead, a secure BE template is stored. The database is inherently secure, meaning there is no need for complicated encryption and key management protocols. The ePassport is protected against tampering, since a potential attacker or any unauthorized user will not know the cryptographic key that was used to create the BE template.

Next Steps to Bringing Biometric Encryption to the Prototype Stage

Biometric Encryption has been researched since the mid-'90s. Technologically, this area is much more challenging than conventional biometrics. But now BE is fast approaching the next phase, i.e., the creation and testing of a prototype. The following issues still need to be addressed:

Select a Proper Biometric

The most promising results in terms of accuracy have been obtained for irises. Low variability of image samples and the presence of a natural alignment feature (eye pupil) make this biometric the number one candidate for BE.

Face recognition is the most publicly acceptable type of biometric. Recent advances in the technology allowed Philips Research to create the first operational BE system. At the present time, one of the drawbacks of the face-based BE system, however, is the relatively small size (~ 58 bits) of the encryption key that may be securely bound to the biometric.

Fingerprints, for which the BE was originally pioneered, are also a prime choice. The fingerprint biometric is used more widely than the iris or face, and most privacy concerns relate to fingerprints. On the other hand, using fingerprints for BE turns out to be much more challenging. The reasons are: high skin distortions can be introduced when the finger presses upon the sensor; and the difficulty of aligning a fingerprint on verification with the one enrolled. As mentioned before, the situation is more difficult for BE than for a conventional fingerprint verification, since BE works in a "blind" mode (the enrolled fingerprint or its minutiae template are not seen). Some of these issues can be overcome with a free-air image. Although this would present other optical issues, we believe they could be resolved by current technology. In general, face and especially iris are less vulnerable to distortion and alignment problems.[23]

Other biometrics, e.g., voice, signature, palm prints, etc., may not have enough entropy (i.e., the amount of non-redundant information to support a long enough cryptographic key). They could possibly be put on the list of "auxiliary" biometrics, i.e., used for BE in combination with irises, faces, or fingerprints or, perhaps, with conventional passwords (called "hardening").

Improve the Image Acquisition Process

For fingerprints, this means choosing a proper fingerprint sensor that is less susceptible to skin distortions (e.g., a free air sensor), or changing the existing sensor ergonomics to keep the distortions under control. Image quality can also be improved at the algorithm level (i.e., through software).

Make BE Resilient Against Attacks

This area of research, i.e., the analysis of potential vulnerability of BE against attacks, has been largely overlooked. By that we mean that a sophisticated attacker could gain access to both the BE templates and the algorithm. The only thing he cannot obtain is a user's biometric. Such an attacker, fully familiar with the algorithm and exploiting its weaknesses, will not be doing just a brute force search (i.e., about 2^{128} computations for a 128-bit key) in order to break the BE template. Instead, he will devise various attacks that can be run in a realistic time frame. The

23 There have been independent tests, such as BioPII in Germany, that reported unusually high error rates for iris recognition: **www.bsi.de/literat/studien/biop/biopabschluss2.pdf**; **www.europeanbiometrics.info/images/resources/90_264_file.pdf**. Those results were questioned by Prof. John Daugman ("BioPII Controversy to Be Tackled," *Biometric Technology Today*, vol. 13, no. 10, pp. 1-2, 2005).

BE algorithm must be resilient against those off-line attacks. The same approach (i.e., resilience against attacks) is adopted in conventional cryptography.

Improve Accuracy and Security of BE Algorithm
There have been substantial advances in algorithm development in conventional biometrics in the past few years, as demonstrated by a series of international competitions. Many of those advances are applicable to BE.

Exploit Multimodal Approaches
This has been a hot area of research and development in conventional biometrics. The performance of a biometric system is significantly improved when different algorithms, or different fingers, or different biometrics (e.g., fingerprints and face) are combined. The modes that are combined should be "orthogonal," i.e., statistically independent. It is reasonable to expect that the multimodal approach would also work for BE.

Develop BE Applications
The applications, such as those described in the case studies, should clearly demonstrate the benefits for privacy and security brought about by the use of BE.

Summary and Conclusions
Biometric Encryption technology is a fruitful area for research and has become sufficiently mature for broader public policy consideration, prototype development, and consideration of applications.

This paper has explored the possibilities and privacy-enhancing benefits of Biometric Encryption technologies for meeting the needs of businesses and government agencies.

We believe that BE technology exemplifies fundamental privacy and data protection principles that are endorsed around the world, such as data minimization, user empowerment, and security, better than any other biometric technology solution in existence.

We hope that our paper will form a valuable contribution to current national and international discussions regarding the most appropriate methods to achieve, in a privacy-enhanced manner, strong identification and authentication protocols.

While introducing biometrics into information systems may result in considerable benefits, it can also introduce many new security and privacy vulnerabilities, risks, and concerns, as discussed above. However, novel Biometric Encryption techniques have been developed that can overcome many, if not most, of those risks and vulnerabilities, resulting in a win-win, positive-sum scenario.

One can only hope that the biometric portion of such systems is done well, and preferably not modelled on a zero-sum paradigm, where there must always be a winner and a loser. A positive-sum model, in the form of Biometric Encryption, presents distinct advantages to both security AND privacy.

Appendices

For the two extensive appendices, please see the online version of *Biometric Encryption: A Positive-Sum Technology That Achieves Strong Authentication, Security AND Privacy* at **www.ipc.on.ca**.

About the Authors

Ann Cavoukian, Ph.D.
Information and Privacy Commissioner of Ontario Dr. Ann Cavoukian is recognized as one of the leading privacy experts in the world and is the author of two groundbreaking books on privacy – *Who Knows: Safeguarding Your Privacy in a Networked World* (1997), written with Don Tapscott, and *The Privacy Payoff: How Successful Businesses Build Customer Trust* (2002), written with Tyler Hamilton. Overseeing the operations of the access and privacy laws in Canada's most populous province, Commissioner Cavoukian serves as an Officer of the Legislature, independent of the government of the day.

Alex Stoianov, Ph.D.
Dr. Alex Stoianov began working in the field of biometrics after joining Mytec Technologies Inc. (Toronto, Canada) in 1994, where he was one of the originators of the privacy-enhancing technology Biometric Encryption. Working for Bioscrypt Inc., the successor of Mytec, as a Principal Scientist from 2001 to 2006, he developed numerous technological breakthroughs and improvements for fingerprint verification algorithms. He also won the Third International Fingerprint Verification Competition (FVC2004), viewed by many as the "Fingerprint Olympics," on the company's behalf. Dr. Stoianov has co-authored over 30 scientific papers and seven patents.

The authors gratefully acknowledge the work of Fred Carter, IPC Senior Policy and Technology Advisor, in the preparation of this paper.

The authors would also like to thank Prof. Dr. Christoph Busch of Fraunhofer IGD, Germany, and Mr. Bernard Didier and Mme. Alexandra Michy, both of Sagem Défense Sécurité, France, for their review and contributions to the pre-publication draft.

In addition, we would like to thank Dr. Michiel van der Veen, Senior Manager, Business Development Biometrics, Philips Research, of the Netherlands, for bringing to our attention their recent white paper, *privID™: Privacy Protection in Biometric Security Applications,* as well as the fact that Philips now has biometric encryption applications that are operational and ready for deployment.