Gene Itkis, Venkat Chandar, Benjamin Fuller, Joseph P. Campbell, and Robert K. Cunningham

Iris Biometric Security Challenges and Possible Solutions

For your eyes only—Using the iris as a key

©ISTOCKPHOTO.COM/GREYFEBRUARY

Biometrics Security and Privacy Protection

Biometrics were originally developed for identification, such as for criminal investigations. More recently, biometrics have been also utilized for authentication. Most biometric authentication systems today match a user's biometric reading against a stored reference template generated during enrollment. If the reading and the template are sufficiently close, the authentication is considered successful and the user is authorized to access protected resources. This binary matching approach has major inherent vulnerabilities.

An alternative approach to biometric authentication proposes to use fuzzy extractors (also known as *biometric cryptosystems*), which derive cryptographic keys from noisy sources, such as biometrics. In theory, this approach is much more robust and can enable cryptographic authorization. Unfortunately, for many biometrics that provide high-quality identification, fuzzy extractors provide no security guarantees.

Digital Object Identifier 10.1109/MSP.2015.2439717
Date of publication: 13 August 2015

1053-5888/15©2015IEEE

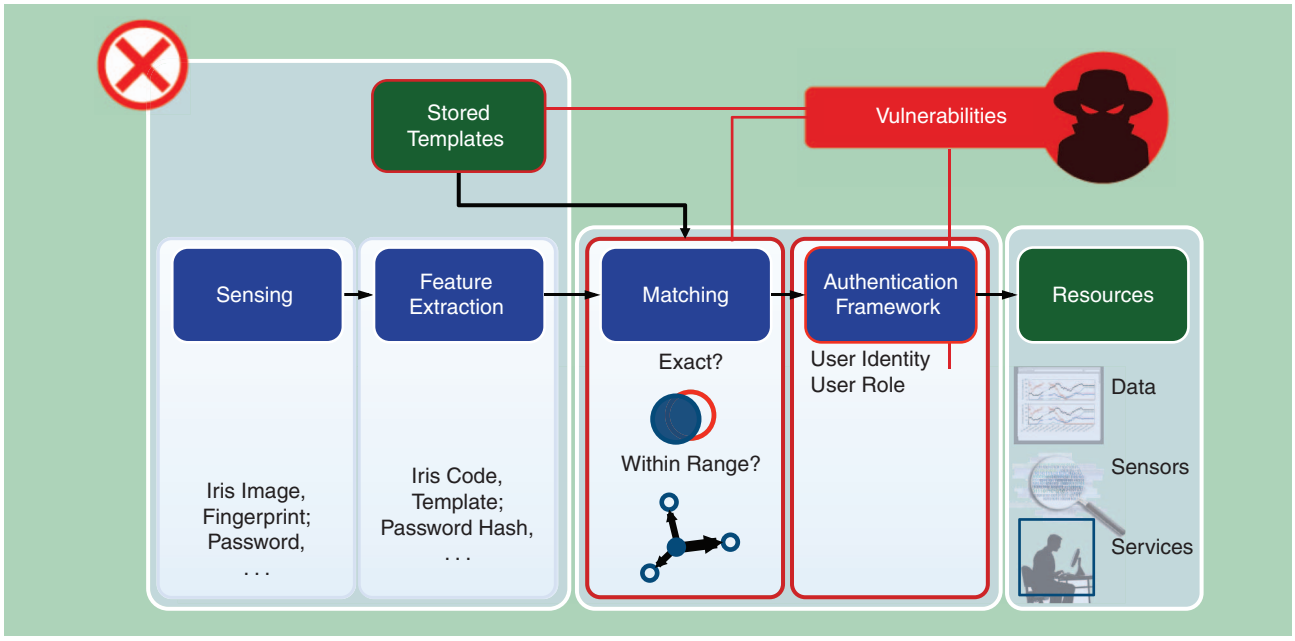

[FIG1] Binary authentication and authorization. First, authentication data, such as a password and/or biometric, is collected from the user and transformed into the appropriate canonical form, or template. Next, the acquired template is matched against the stored reference template. If the match is successful, the authorization framework grants the user access to the appropriate resources (e.g., via appropriate cryptographic keys). The three major vulnerable components of this approach are highlighted: matching (with its fragile binary decision), stored templates, and authorization framework (and its cryptographic keys).

This gap arises in part because of an objective mismatch. The quality of a biometric identification is typically measured using false match rate (FMR) versus false nonmatch rate (FNMR). As a result, biometrics have been extensively optimized for this metric. However, this metric says little about the suitability of a biometric for key derivation.

In this article, we illustrate a metric that can be used to optimize biometrics for authentication. Using iris biometrics as an example, we explore possible directions for improving processing and representation according to this metric. Finally, we discuss why strong biometric authentication remains a challenging problem and propose some possible future directions for addressing these challenges.

INTRODUCTION

AUTHENTICATION AND AUTHORIZATION

Security systems commonly include two components: 1) an authentication framework that validates a user's identity and 2) an authorization framework (sometimes called a *reference monitor*) that controls access to resources for the validated identity [1], [2]. Biometrics represent one important way to verify a user's identity [3]–[7]. An adversary able to impersonate you can do surprisingly bad things to you and in your name. Authentication is a crucial aspect of security, but it is surprisingly difficult to do in a way that is easy to deploy and use, as well as provide proper protection [8].

BINARY MATCHING AUTHENTICATION PARADIGM

In typical modern authentication systems, during the enrollment stage—when a user is added to the system—an original reading is collected from the user; transformed into a canonical form; and stored as a reference value, often referred to as a *stored template*.

Later, when the user authenticates to the system, a new reading is collected, transformed into the same canonical form, and matched against the stored reference template. The authentication is successful if the two values match. Then an authorization framework may grant the user access to protected resources, e.g., by providing cryptographic keys (called *content keys*) to the appropriate resources (see Figure 1).

For example, for passwords, a user enrolls by selecting a password, the canonical form is a cryptographic hash, and an authenticated match is required to be exact [9]. Biometrics are typically noisy: readings vary even for the same subject—hence, the match is approximate. We call this paradigm *binary matching* (sometimes using only one of these words) since the authentication results in just a single bit: match or not match.

INHERENT WEAKNESSES OF BINARY MATCHING PARADIGM

This paradigm has a number of crucial inherent weaknesses (see [10] for their manifestations in biometric systems):

- Binary authentication decisions in the matching step are fragile and can be skipped or flipped even by accidental errors.
- Matching requires the reference templates to be readily available during authentication, creating opportunities for an attacker to steal the templates, which in turn enables further attacks (see the section "Weaknesses of Binary Matching Paradigm: Details").

The binary nature of the decision also implies that the system must have access to all the resources that a properly authenticated

FUZZY EXTRACTORS

Fuzzy extractors are a pair of algorithms for deriving keys from a noisy source of entropy. The first algorithm, *generate* or Gen, is used at enrollment time. It takes an initial reading w, producing a key as well as public information P. The second algorithm, *reproduce* or Rep, is used at authentication time, taking w' (a nearby reading of an iris) and the public value P. The correctness guarantee is that Gen and Rep should give the same key if the distance between w and w' is at most some bounded parameter denoted t. To protect against the attacks described in the Introduction, the key should be strong even in the presence of P. The problem is trivial if P is private. A private P can store a key and the original reading. Then Rep outputs the key if and only if the new reading w' is close enough. This essentially reduces the problem to having a good biometric source.

Bennett, Brassard, and Robert identified two crucial tasks for deriving keys from noisy data [47]. The first, information-reconciliation, removes errors from w'. The second, privacy amplification, converts w to a uniform value. Traditionally, a fuzzy extractor uses two separate algorithms to accomplish these tasks. A secure sketch [48] performs information-reconciliation and a randomness extractor [49] performs information-reconciliation. A fuzzy extractor that separates information-reconciliation and privacy amplication is called the *sketch-and-extract construction*. See the work of Dodis et al. for formal definitions of the requirements of fuzzy extractors and secure sketches [48, Sec. 2.5–4.1]. Here we provide a brief review of standard constructions and recent advances. The goal of secure sketch is to map nearby w' back to the original w without revealing unnecessary information about w.

The simplest construction of a secure sketch uses the syndrome of an error correcting code. The public information P consists of applying a parity check matrix to the original reading w. This allows decoding of the original w from a nearby w' and P. The entropy of w conditioned on this secure sketch is at least the starting entropy minus the length of the syndrome. The length of a syndrome must increase as the error tolerance increases. This means the lower bound on the remaining entropy of w decreases as the error tolerance increases.

There are many coding-based constructions of secure sketches. The security analysis of these constructions usually considers the difference between the size of the metric space and the deficiency of the best code correcting enough errors in the metric space. (The measure can be relaxed considerably by allowing the secure sketch to occasionally output the wrong value.) This imposes a tradeoff between the remaining entropy of w and the noise that can be tolerated.

Correcting more errors decreases FNMR and decreases the length of the derived key. Standard constructions of fuzzy extractors work well when the source has a high entropy rate (nearly uniform). Recent work builds fuzzy extractors from the face biometric when the entropy rate is almost full, even for an error rate of nearly 30% [50]. However, standard fuzzy extractors are not known to be secure on sources with low entropy rates. This is not a limitation of a particular construction; there are probability distributions with the same entropy and error rate as irises where key derivation is impossible (see [40] and [48, Appendix C]).

Recent works have also built fuzzy extractors using properties of a distribution other than entropy and desired error tolerance [40], [51], [52]. Unfortunately, these constructions are not known to work for the iris distribution. These constructions assume properties of the physical source that irises do not appear to satisfy. Thus, authentication from the iris remains challenging.

user might need, thus forcing a violation of the principle of least privilege [11]. As a result, by compromising the system, an attacker can also obtain this access; this is what, in particular, makes privilege escalation attacks so attractive. These inherent weaknesses become especially apparent in settings where it is not clear who can be trusted to perform the matching and grant access, e.g., in the cloud.

THE BIOMETRIC CRYPTOSYSTEM PARADIGM

An alternative approach is to derive cryptographic keys from biometrics. These keys can then be used to access the resources. The challenge here is that the biometrics are inherently noisy. This approach originated with the work on fuzzy commitment and fuzzy vault [12]–[14], building on ideas from [15] and [16]. These ideas were formalized in a cryptographic object known as fuzzy extractors that reliably produce a uniform key from a noisy source of entropy (see "Fuzzy Extractors"). In particular, this assures that even if the entropy was distributed unevenly among the bits of the source, the output will have all bits random. This approach has also been investigated under the names of *biometric cryptosystems* and *biometric key generation* [17], [18]. We show the main stages in this approach in Figure 2.

We have implemented a full authentication system where authorization is implicit based on the knowledge of the proper cryptographic keys. Compared to the single bit produced by binary matching, this approach enables leveraging the full entropy (to the extent possible) that the user provides as part of the authentication. We call such an approach *cryptographic authorization* or *cryptographic access control*. The advantages of this approach according to the metrics of [8] are discussed in "Benefits of the General Full-Entropy Approach." During the development of this system, we found that fuzzy extractors often do not provide meaningful guarantees on the key strength (KS) of keys derived from biometrics. For example, the iris code [3], which is a representation of one of the best biometrics [19], produces a key with no provable security using standard fuzzy extractors [20, Sec. 5]. The goal of this work is to examine why this is the case and how to derive stronger keys from biometrics.

THE GAP BETWEEN BIOMETRICS AND FUZZY EXTRACTORS

Extensive work on fuzzy extractors and similar techniques may lead one to believe that deriving cryptographic keys from

biometrics is a solved problem. Unfortunately, authentication from many noisy sources remains challenging.

Biometric techniques have been developed and optimized for identification (these optimizations naturally carry over to binary matching). The metrics used for evaluating biometrics are typically variants of FMR versus false FNMR plots. But these characteristics say little about the cryptographic security of the keys derived from the biometric. A different metric must be identified for the cryptographic authentication task, and biometric techniques need to be optimized according to this metric.

In this work, we propose that biometric quality for authentication should be measured as FNMR versus the strength of the key (see the section "Metric for Cryptographic Authentication"). This will help focus the development of biometrics for authentication as their quality will be measured for that task.

Using the iris as an example, we provide initial optimizations of biometrics to this authentication metric. However, key derivation from the iris remains a challenging problem.

For the purposes of simplicity, we assume the goal is to derive a 128-bit key. This is the key length for the standard symmetric cipher AES-128 approved by the National Security Agency to protect information classified up to secret level, so this should be useful for the most commonly used systems and accounts [21].

THE CASE FOR KEY DERIVATION FROM BIOMETRICS

In this article, we consider the suitability of biometrics for strong authentication. We discuss authentication modalities in "Other Authentication Modalities."

The central problem that complicates the use of biometrics is noise: readings of the same biometric of the same user can differ significantly, even under the most controlled (and thus least flexible and convenient for the user) environments. For binary matching, this forces use of approximate comparisons (using an appropriate metric). When biometrics are used for cryptographic authorization, noise represents a larger challenge.

WEAKNESSES OF BINARY MATCHING PARADIGM: DETAILS

As discussed in the section "Inherent Weaknesses of Binary Matching Paradigm," attackers can exploit fragility of the matching step of the binary matching paradigm, obtaining access to a single session. Similarly, privilege escalation attacks are extremely powerful.

In the case of biometrics, the transformation must preserve locality (i.e., similar readings, such as those taken from the same user, must remain close according to an appropriate metric, even after the transformation into the canonical form). This makes it difficult to design truly one-way transformations. Indeed, for commonly used transformations, it turns out that the stored reference templates can be reverse-engineered to produce realistic biometrics that would match the corresponding templates. Adversaries can manufacture a natural looking iris biometric that will pass the identification test [22], [23]; the same is true for fingerprints [24].

CRYPTOMATCHING: GOING HALFWAY

It is, in principle, possible to use cryptographic key derivation to obtain a cryptographic key, but then revert back to a binary match paradigm—comparing the key to the stored reference. Since, in this case, a cryptographic key is matched against a stored reference, we call this approach a *cryptomatch* authentication.

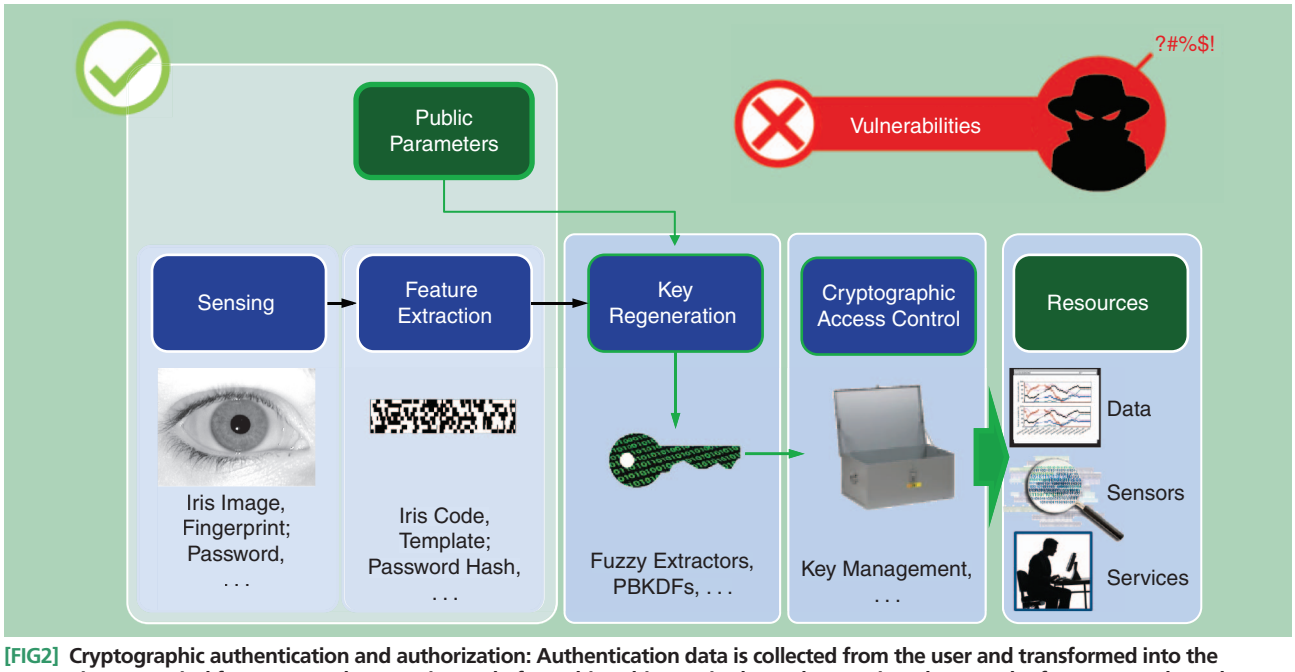

[FIG2] Cryptographic authentication and authorization: Authentication data is collected from the user and transformed into the appropriate canonical form or template. But instead of matching this acquired template against the stored reference template, the acquired template is used (with some nonsecret public parameters) to regenerate the cryptographic key. This key can then be used to obtain access to the appropriate resources via cryptographic access control and key management. The vulnerabilities highlighted in Figure 1 are no longer present.

BENEFITS OF THE GENERAL FULL-ENTROPY APPROACH

In a comprehensive study of two decades of general-purpose user authentication on the web, [8] proposes a broad set of usability, deployability, and security properties (constituting benefits, when satisfied), and uses these to compare different authentication techniques.

However, these properties had been formulated for the binary match approach. We believe that they should be revised in the context of the cryptographic authorization and cryptographic access control paradigms. Whenever it makes sense, we compare binary match to cryptomatch, rather than cryptographic authorization.

The change from binary match to cryptographic authorization has little or no effect on many properties listed in [8]: e.g., U1-U6 and D1-D2. For other properties, effect might be nonobvious. For example, Property U7: Infrequent-Errors aims to minimize the number of failed authentications. For the exact match (such as for passwords), this property would not be affected by the transition from binary matching to cryptomatching and cryptographic authorization. But for biometrics, the matching and key regeneration may involve somewhat different techniques, resulting in somewhat different error rates.

But furthermore, if proper cryptographic access control authorization is used—for any type of authentication modality, including passwords and biometrics—the failed authentication may result in a failure during access. Such access failures can be made easily detectable, eliminating the difference between binary match and cryptographic authorization approaches. But this failure detection can be exploited by an attacker as well—e.g., in an exhaustive search for the correct password. So, this represents a potential tradeoff between the usability and security (specifically, Property S4: Resilience-to-Unthrottled-Guessing). This example also illustrates the aforementioned need for revision of the properties in [8]: for example, inserting Property U7′: Easy-Recovery-from-Authentication-Failure.

Property U8: Easy-Recovery-from-Loss provides an example of a different tradeoff. A simple implementation, dogmatically following the least privilege principle, will result in a system where if a user loses the authentication data (whether public parameters or, say, the password), the data cannot be accessed by anyone at all and thus can be considered lost. However, a wiser implementation will build in various recovery mechanisms, which would trade off security and usability once again: making the data easier to recover but also easier to steal, or making the data more secure but also more difficult to recover. While passwords are typically considered to have the property U8 satisfied, this is typically because the choice in the security-usability trade off is made for the users, favoring usability at the cost of security (any administrator able to reset the password is also able to impersonate the user).

Property D3: Server-Compatibility requires the (authenticating) server to be compatible with the passwords. In our case, the servers only need to provide the right public information, and thus our approach completely eliminates the need for the server to do anything special for the authentication. D3 does not quite capture this benefit of cryptographic authorization. Cryptomatch, on the other hand, can be made server-compatible by using the key as the password.

The client in the cryptographic authorization and cryptomatch approaches needs to compute the cryptographic key. Currently,

For passwords, the traditional binary matching implementations can be seen as a specific implementation of exactly this cryptomatch principle. For biometrics, even this halfway approach can yield significant advantages compared to the binary matching the way it is typically practiced today.

As discussed previously, binary match biometrics suffer from an additional vulnerability: stealing a reference template enables an attacker to generate realistic biometrics. Some research on so-called cancelable biometrics and similar techniques helped to address this issue (e.g., see [18] for a survey). However, many of the considered methods do not result in strong security, say, compatible to provable security of common cryptographic tools (there are many notable exceptions, such as [25], [26], and others). In contrast, an ability of users to consistently regenerate cryptographic keys immediately results in a very strong version of cancelable biometrics. Once a key is derived, cryptomatch authentication can be easily implemented by deriving different (and easily replaceable) keys for each verifier, using the regenerated key as a master secret and unique random value (salt)—this is essentially similar to the way different keys are derived from a master secret in many secure protocols, e.g., such as transport-layer security [27].

BENEFITS OF CRYPTOGRAPHIC AUTHORIZATION

However, a cryptographic matching approach is still a binary decision. For high-security applications, authorization decisions should be based on knowledge of cryptographic keys derived from authentication material. Because there is no single-bit match/no match decision (one instead recovers a key), it is more difficult for an adversary to cause hardware to fail or software to branch in a way that will enable access.

Even in the most primitive case when the key regenerated in authentication stage is used as content key to access protected data directly, there is no information on the system that the attack can use to break security. The templates are replaced with public parameters and the content key is no longer stored, but rather recreated when needed from the biometric. This approach allows for strong authentication even against adversaries with physical access to the system when it is not actively used by a legitimate user. It is crucial for this approach that the public parameters produced by the fuzzy extractor do not compromise security of the biometric.

This approach also allows the creation of multiple keys from the same biometric in such a way that compromising one or a few of these keys (by leaking them to an adversary) in no way compromises the others. More sophisticated cryptographic access

our implementations run in Java, but not from a browser. It may be possible to implement similar tools as applets or run them from a browser in some other way, but until that happens, we need to consider our approach as having a negative impact on Property D4: Browser-Compatibility. While the change to cryptographic authorization has mixed effects on deployment and usability, it has positive impacts in the security area.

S1: Resilience-to-Physical-Observation, S5: Resilience-to-Internal-Observation, and S10: Requiring-Explicit-Consent are not directly affected by the shift from binary match to cryptographic authorization. These properties are typically provided only by special hard tokens (for which providing, or even better—encapsulating, so it can be used securely—a cryptographic key would typically be rather trivial). Similarly, S2: Resilience-to-Targeted-Impersonation (when personal knowledge of the target user person can be utilized in the attack) should not be affected by the shift, but it does eliminate some of the easiest pitfalls: while a user might use her mother's maiden name as a password, in cryptographic authorization, the user would not be prompted to use such information for security.

As discussed above, resilience to guessing (properties S3 and S4) can be strengthened at some usability cost. However in the throttled variant (S3), the verifier can limit the rate of attacker's guessing, but in our approach, there is no verifier. So, implementing throttling requires a different approach: e.g., a PBKDF2 function [53] can be used to increase the time to derive the key. Such an approach can also benefit even S4 (the unthrottled version, since the above mechanism requires no help from any servers). More sophisticated throttling mechanisms, however, can be designed for our approach, making it at least as resilient to guessing as the binary matching, and probably even more so.

Property S6: Resilience-to-Leaks-from-Other-Verifiers can be provided even by the cryptomatch version. For the cryptographic authorization, this property can be strengthened significantly (to be resilient) to leaks from any verifiers, since no verifiers are even present. The only other approach that assures this property is the zero-knowledge proofs of identity [54]. Such proofs are cryptographic protocols requiring significant computational power from both the prover and the verifier. Furthermore, just as in binary matching authentication, this protocol results in a single bit (pass or fail), and hence is vulnerable to the corresponding weaknesses.

Since cryptographic authorization requires no verifiers, it also has the other properties that rely on the verifiers not failing: S7: Resilient-to-Phishing and S11: Unlinkability. In fact, even the crypto-match approach can achieve unlinkability.

Cryptographic authorization also satisfies S9: No-Trusted-Third-Party. By necessity, the binary matching approach must exclude the verifier from the consideration. This is a direct consequence of the binary nature of the traditional approach. In contrast, cryptographic authorization has no verifiers and no (implicit or explicit) reference monitor acting upon the result of the binary authentication, it only relies on the client security not to steal the authentication data provided by the user (if hardware tokens or devices are an option, then even this vulnerability can be reduced or eliminated, depending on the specifics of the token). Cryptographic authorization improves security in a number of ways, with the improvements to S9 and S6 being the most significant.

control structures can be built as well, for example, traditional operating system permissions.

THE GAP BETWEEN BIOMETRICS AND FUZZY EXTRACTORS

Biometrics have been extensively used for the task of identification: discriminating effectively between two individuals. Rightfully, biometric systems were evaluated according to metrics for identification. The standard metric is a function between how reliably a single person's biometric can be recognized as such (FNMR) and how often two individuals are confused (FMR). Obviously, for a system always reporting a match, FNMR = 0 (since nonmatch, false or not, is never reported) and FMR = 1 (since different subjects are always falsely reported as matched). Conversely, never reporting a match makes FMR = 0 but FNMR = 1.

In practice, biometric systems allow adjusting their parameters to achieve some tradeoff between these characteristics depending on needs of specific applications. This tradeoff can be depicted as a function comparing FNMR versus FMR and is often used as a measure of quality of the biometric systems. We call this the *identification* metric. Current iris biometrics techniques produce a very strong biometric according to the identification metric [28].

There can be different ways to improve biometric systems according to this metric. For example, we can fuse iris codes from three readings by taking a majority for each bit in the iris code, as proposed in [15] and [29]. Then we use the result in matching, improving its quality according to the match metric (see Figure 3).

According to the identification metric, fusing provides an impressive improvement—see Figure 3, which depicts our experiments using the multispectral iris data set, licensed through the Scitor Corporation. But what this metric says is that we can get the FMR rate down to between 0.001–0.1%. In other words, by trying between 100,000 and 1,000 irises would give an attacker a good chance of impersonating a targeted user.

MEASURING BIOMETRIC SUITABILITY FOR CRYPTOGRAPHIC AUTHORIZATION

In this section, we show that the identification metric may be inappropriate for the cryptographic authorization task. Recall that the goal of this task is to create a strong cryptographic key from the same user, and note that the strength of the cryptographic key implies that different users rarely map to the same key.

Intuitively, FNMR corresponds to an authorized user failing an authentication attempt. Thus, FNMR can be viewed as,

OTHER AUTHENTICATION MODALITIES

It is common to organize all the authentication methodologies into three categories, according to the nature of the input provided by the user—we call them modalities: 1) what you have (e.g., hardware tokens, such as smart cards, etc.), 2) what you know (e.g., a password), and 3) what you are (biometrics).

Each of these categories comes with its own typical characteristics. Often multifactor authentication (using multiple modalities) is recommended to achieve higher security. But it is convenient to consider each modality, with its pros and cons, separately.

The strength of hardware tokens is that they can store plenty of entropy and perform complex computations far beyond human capacity. On the other hand, such tokens impose deployability limitations and can also be forgotten, lost or destroyed, or even stolen. Furthermore, in some cases (e.g., theft), the token might be used by an attacker. To defend against this, the user should authenticate to the token using another authentication modality [e.g., a short password, or personal identification number (PIN)]. Since token integrity can often be assumed, this authentication to the token is typically easier. Tokens are also subject to various hardware attacks and might also be used without explicit consent (e.g., if a PIN is cached after the first use of a smartcard—as done by many drivers—then malware can request smartcard to decrypt or authenticate data without the user knowing it).

Passwords—what you know—suffer from limitations of the human mind: since our memory is relatively weak, passwords have notoriously little entropy and are, hence, open to exhaustive search attack, such as password cracking. On the other hand, it is ultimately deployable, imposing the least amount of restrictions and only moderate inconvenience.

Finally, the biometric approach can be easier for the user, since there is nothing to remember or carry. Biometric reading can also be made very easy, requiring minimal effort from the user. But the easier the reading, the harder it is to enforce explicit consent. There may also be a separate tradeoff between the ease of reading and the reliability and the amount of entropy collected. This approach also requires some special equipment.

Biometric modality has another important feature: it is almost like a password that we wear literally on our face (or hands, in that case leaving its copies on everything we touch). In other words, biometrics can be easy to steal when the subject is present. So, the best use of biometrics is in remote authentication, where an attacker may not have physical access to the target user. Sometimes biometrics is used more as a test of physical presence, rather than the authentication. In that arena there is a constant arms race between biometric device manufacturing and the attackers, who use anything from jelly beans to cameras and special contact lenses. The most expert attackers tend to be leading in that race most of the time.

When a theft does occur, unlike passwords, biometrics cannot be easily replaced, no matter which authentication approach is used.

loosely speaking, a nuisance factor for authentication systems: e.g., it reflects how many attempts you would have to make before successfully logging into your own account. While this can have a serious impact on the system usability, it has no effect on security.

In contrast, FMR reflects probability that a user can be impersonated by someone else: e.g., if my iris is accepted when authenticating to log into your account. This is extremely important for security, and so FMR reflects insecurity of the system—the higher it is, the easier it is to trick the system into letting unauthorized users in.

FNMR versus FMR assumes an attacker that attempts authentication using random biometrics from a suitable population. Crucially, it says nothing about what is revealed by the authentication system. In a binary matching system, the authentication system writes down a template, revealing the original reading to the attacker. In a cryptographic authorization system, using fuzzy extractors, the system writes down some public information necessary to map nearby readings to the same cryptographic key. A dedicated attacker will use all information available; a metric for the authorization task must include this information.

As noted in [20, Sec. 5], there is no known fuzzy extractor for the iris. The core of the problem is that irises a relatively low entropy rate compared to their noise. For provable security, the known constructions of fuzzy constructors need the entropy rate to be significantly higher than the noise.

METRIC FOR CRYPTOGRAPHIC AUTHENTICATION

In an authorization system, the two important parameters are how often a user completes authentication (FNMR) and the strength of the resulting cryptographic key (in the presence of public parameters). Thus, when a biometric is used for authentication, the relevant metric is FNMR versus KS. It may be possible to generate a key whose strength and length are unequal. Since in our key derivations our goal is to produce a key indistinguishable from random, we can assume that key has strength proportional to $2^{-|\text{key}|}$.

As described in the section "Weaknesses of Binary Matching Paradigm: Details," any information written by the authentication system should be assumed to be available to the attacker. For this reason, we call any authentication information public parameters or P. Fuzzy extractors and other biometric cryptosystems write down error correcting information to ensure the same user reliably generates the same key. Therefore, the strength of a key should be measured relative to an adversary with access to public parameters P (such as the public information produced by fuzzy extractors in "Fuzzy Extractors").

We suggest measuring the quality of the biometric when used for cryptographic authorization as a function of FNMR and

the strength, KS, remaining after the key is derived (with public information known). We call this the *authorization metric*. The metric specifies how often for a given KS a legitimate user will be rejected. There are two ways of viewing this metric:

1) measuring the authentication strength possible from the source with the best known fuzzy extractor.

2) measuring the information theoretic capacity of the source for key derivation.

Both the identification and authentication metrics measure how frequently the legitimate user is granted access. The identification metric measures how frequently other users from the distribution are also granted access. When considering a determined attacker, this is not a sufficient threat model. For the authentication task, it is prudent to assume that the attacker has access to public parameters, P, and then measures the probability of recovering the key, given this information.

OPTIMIZING FOR THE ENTROPY METRIC: IRIS EXPERIENCE

THE IRIS CODE

The human iris is believed to be a strong biometric attribute [19]. The iris pattern develops in utero and is fairly stable throughout the lifetime of an individual [30]. Irises are diverse in small homogeneous populations (even right and left eyes of the same subjects appear to be independent) and are believed to be largely epigenetic (not dependent on genetic information), although some correlations may be observed [31], [32]. Typically, the near-infrared (NIR) images are used, although some research into using multispectral images has been undertaken [33].

For biometric applications, an iris is typically transformed into an iris code as follows (see Figure 4). First the image of the iris is segmented, locating the iris and the pupil. Then the iris image is unwrapped using polar coordinates, translating the iris from a two-dimensional (2-D) tor (doughnut) into a rectangle. Finally, the rectangle is subjected to various special processing (typically, filtering and quantization), producing a bit-vector called the *iris code*.

Modern transforms derive from the work of Daugman based on 2-D Gabor wavelets [3].The filtering phase uses 2-D Gabor wavelets [34] at various angular and radial coordinates. The image is divided into a polar grid of some angular and radial resolution and a wavelet is computed at each coordinate. We denote these values

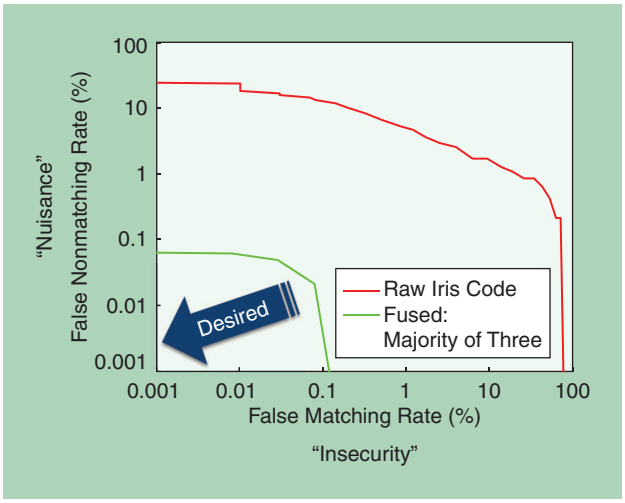

[FIG3] The match metric: FNMR versus FMR. The red plot is obtained using IrisCodes computed on iris images from the multispectral iris data set, licensed through the Scitor Corporation, using open source IrisCode software [35]. The green plot is obtained by faking the bit-wise majority of three repeated IrisCodes from the same user. As we discuss in the section "The Gap Between Biometrics and Fuzzy Extractors," while it is desired to reduce both FNMR and FMR, the latter corresponds to insecurity of the system, while the former reflects inconvenience for the users. For FNMR, the rates below 0.1% (eligible user failing to authenticate in less than one out of 1,000 attempts) can be considered quite acceptable. However, FMR range needs to be tens of orders of magnitude lower to be compatible with cryptographic security. This preference of FMR over FNMR is reflected by the angle of the "desired" arrow.

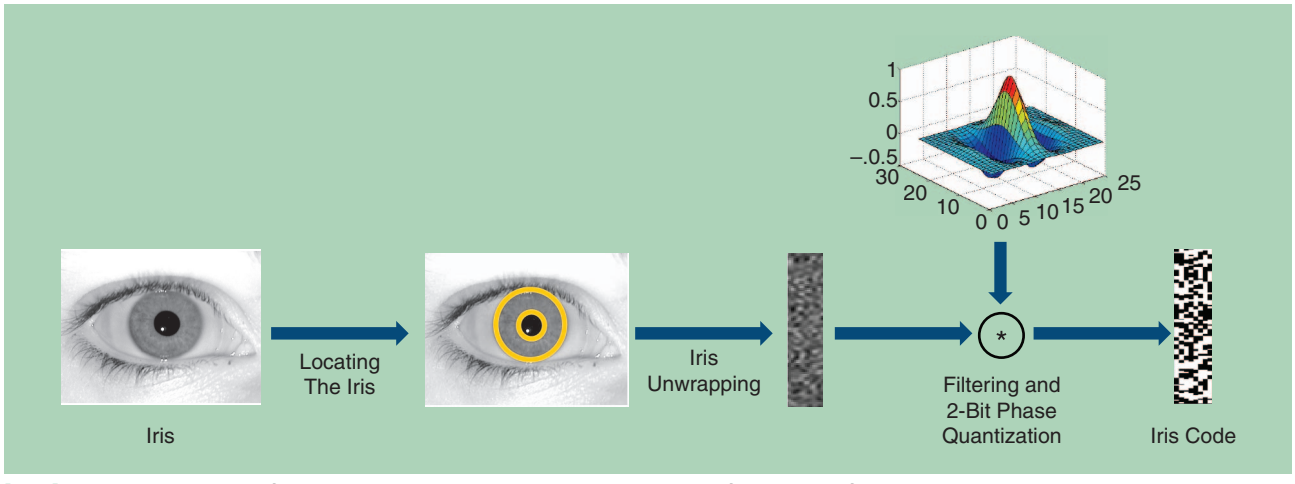

[FIG4] Deriving an iris code from the iris. The process starts with an image of an iris. The first step is segmentation: the iris is located in the image. Then the iris is unwrapped into a rectangle. The rectangular image is filtered using 2-D Gabor filters and quantized. The result is a bit array, called an iris code.

by ang_{res} and rad_{res}, respectively. The sign of the real and imaginary components yield 2-bit values for each location. The total length of the transform is $2 * \text{ang}_{\text{res}} * \text{rad}_{\text{res}}$. We then utilize the transform of Masek [35], which uses an angular resolution of $\text{ang}_{\text{res}} = 240$ and $\text{rad}_{\text{res}} = 20$ for an overall length of 9,600 bits. The transform of Daugman [3] provides superior performance but it is not publicly available. Performance is improved by applying simple automatic tests detecting some unsuccessful segmentations, to eliminate these from the statistical analysis, when working with various iris corpora.

We denote as w the iris code from an original reading collected from a user and stored as a reference template. During authentication, a new reading is then collected producing IrisCode w'. Typically, fractional Hamming distance (FHD)—the fraction of bits that differ between w and w'—is used as the distance between the two readings. Typical FHD between two readings of the same iris of the same person is between 10–30% (more careful tools can get it much lower). This is what is referred to as *in-class* FHD. For readings from different people, or different irises, FHD is within 40–60%. This is called *out-of-class* FHD.

In-class FHD can be increased by rotational distortions; e.g., when the image is taken at slightly different angles. Binary matching can compensate for this distortion by comparing the reference and authentication templates w, w' multiple times, applying a serious of small relative rotations, and picking the best (smallest) FHD. For the out-of-class FHD, this has a relatively negligible effect. Also, when reflections and occlusions occur, it is possible to simply ignore (mask) some portions of the image when the matching is performed. Both of these optimizations present additional challenges for cryptographic authorization. Next we consider the suitability of the iris for the cryptographic authorization task.

IRIS BIOMETRIC SUITABILITY FOR CRYPTOGRAPHIC AUTHORIZATION

In cryptographic authentication, we view users' inputs—in this case, irises—as sources of entropy.

IRIS CODE ENTROPY

Random strings have entropy essentially equal to their length. Unfortunately, most biometrics are not fully random. Estimating entropy of nontrivial distributions is a difficult problem [36], [37].

Assuming the existence of pseudorandom generators [38] (implied by one-way functions [39]), it is possible for distributions to appear to have significantly more entropy than they actually possess. This means it may be fundamentally impossible to estimate the entropy of distributions occurring through complex unknown processes.

Pseudorandom distributions are sophisticated and one may hope that they do not often appear in nature, or at least that biometric distributions appear to have the same entropy to all parties. There are several heuristics used to estimate entropy by comparison to well-known probability distributions. Daugman notes that the FHD between different individuals in an iris corpus fits a binomial distribution with mean $p = .5$ and $N = 249$ [3]. This yields an estimate of 249 bits of entropy.

As described in the section "Metric for Cryptographic Authentication," when measuring KS, we must include public parameters of the authentication system. This means we also need to consider the noise between repeated readings w and w'.

Rederiving the key essentially requires correcting w' to w. If we assume an error rate of ε, then the number of possible errors that might take place is $(n!/(n\varepsilon)!(n(1-\varepsilon))!)$. Hence, the error-correction requires $n(\varepsilon \lg \varepsilon + (1-\varepsilon)\lg(1-\varepsilon))$ bits of information, where n is the number of bits in w. For ε around 10%, this is approximately $n/2$ bits. For the transforms above, $n = 9,600$, this means that 4,800 bits of error correction information is necessary. Standard fuzzy extractor constructions may lose security proportional to this information. Since irises are estimated to have 249 bits of entropy, fuzzy extractors provide no guarantee on the resulting cryptographic key. The key challenge in deriving keys from irises is that the entropy rate and error rate are approximately the same. Fuzzy extractors provide good performance when the entropy rate is significantly higher than the error rate (see [40] for more details).

OPTIMIZING THE IRIS FOR THE AUTHENTICATION METRIC

As described in the previous section, a major obstacle to deriving keys from irises is the low entropy rate of current transforms. In this section, we discuss techniques for trying to improve the entropy rate of irises. We emphasize that none of these techniques currently seem sufficient to derive strong keys from irises. However, similar ideas will be necessary to derive strong keys from the iris.

SUBSAMPLING WAVELETS

Irises have a low entropy rate but each bit on its own behaves like a Bernoulli coin with $p = .5$. This indicates that each bit has full entropy and that the correlations in the iris exist between multiple bits. Thus, one approach to improve the iris transform is to try and find sets of bits that are uncorrelated with a similar error rate as the overall transform. Random subsampling preserves both entropy [41] and error rate. If iris bits are uncorrelated on large sets, then subsampling should produce an entropy rate higher than 10%. There has been work in the iris community on producing better transforms using structured subsampling. The entropy rate and error rate can be maintained while reducing length using random subsampling. If each bit of an iris is entropic on its own, random subsampling may be helpful. In particular, if any 249 bits can be used to reconstruct the iris, we can randomly subsample to 249 bits while maintaining all entropy and keeping error rate constant. The goal of structured subsampling techniques is to find better strategies.

The work of Gentile, Ratha, and Connell [42] introduced short-length iris codes, which were designed to improve processing speed of iris codes by reducing their length. Their work contains several observations: 1) the inside and outside of an iris tend

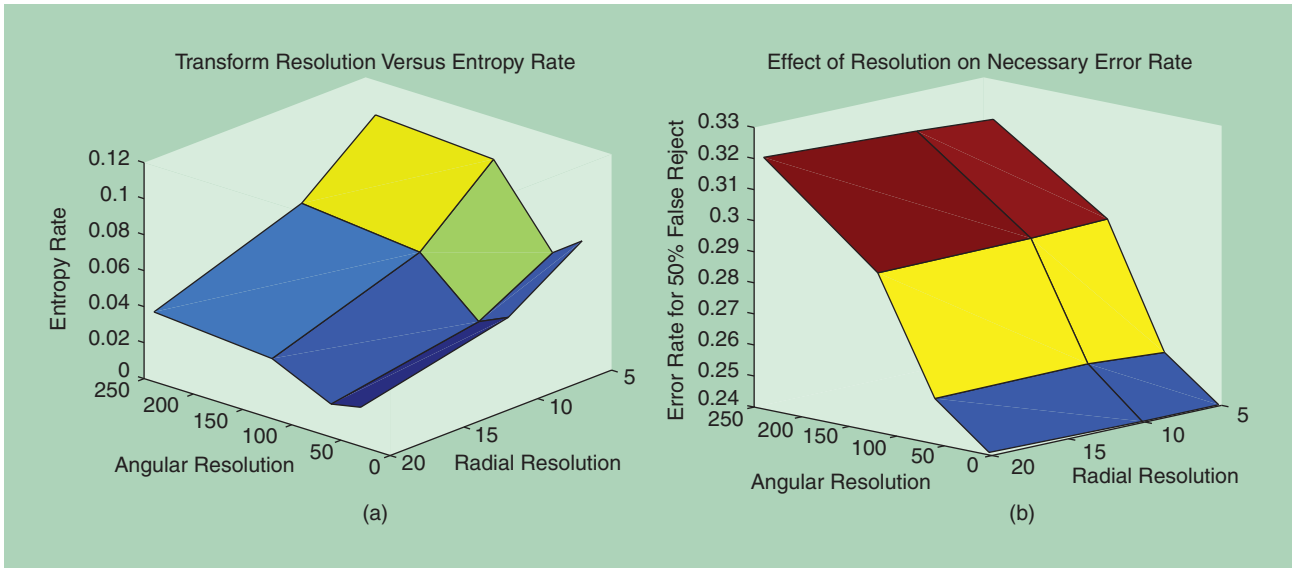

[FIG5] The effect of varying resolution on entropy and error rates. (a) Entropy rates and at various angular and radial resolutions. (b) Error rates at various angular and radial resolutions

to be less reliable due to increased deformations and occlusions, respectively, and 2) there is significant correlation between radially adjacent bits and little correlation between angularly adjacent bits. This leads them to create a transform that subsamples every tenth row of the iris code starting from the fifth row.

In Figure 5(a) and (b), we show the entropy rate and error rate at various resolutions. Observe the following:

- The entropy rate increases as the radial dimension is subsampled.
- The error rate remains constant as the radial dimension is subsampled.
- The entropy rate decreases slightly as the angular dimension is subsampled.
- The error rate decreases as the angular dimension is subsampled.

Figure 5 confirms the observations of Gentile et al. that there is significant redundancy in the radial dimension and this can safely be subsampled. We do note that although subsampling in the radial dimension improves the entropy rate, it does reduce the overall entropy. Careful analysis is needed to determine where the maximum strength key is possible.

FINDING THE BEST POSTTRANSFORM BITS

We will now look at subsampling in the bit domain (postwavelet transform). We start from a 9,600-bit transform. Each bit of the distribution follows a Bernoulli distribution with $p = 0.5$. However, it may be that some bits are more likely to contain errors and contribute more to the overall error rate. We now will try and find the bits that have the lowest error rate. This idea stems from the work of Hollingsworth et al. [43]. The results are shown in Figure 6; unfortunately, this graph is very flat, meaning there are not a large number of bits with lower error rate. Each bit has roughly the same error probability. This means subsampling at the bit level is unlikely to be helpful. We note that for a particular individual, there are bits that consistently have a lower error rate. Writing down such bits for a particular individual may reveal information about the original reading w. Thus, we only consider consistency of bits across the population.

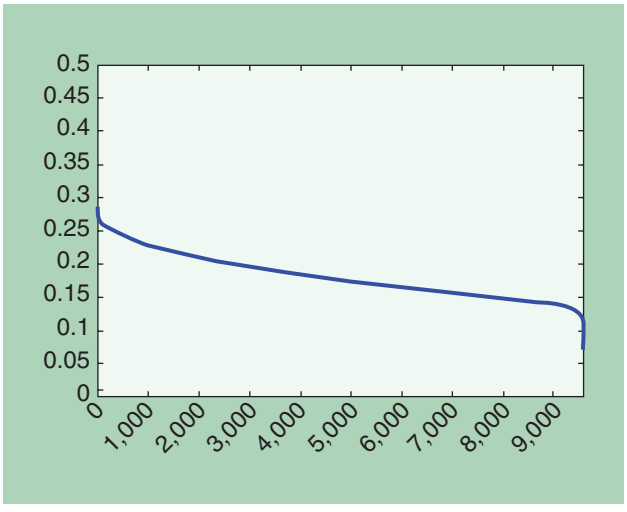

[FIG6] The best bits of an iris code.

DISCUSSION AND SUMMARY

For decades, authentication has relied on matching an original reading from a user against a previously captured and stored reference template. The binary outcome of the matching required a corresponding authorization mechanism to control access to the resources, granting the access based on the result of the matching. Implementations of this paradigm suffer from inherent weaknesses: fragility of the binary decisions, vulnerability of the stored reference templates, and the high value target of the authorization mechanisms.

Biometrics represent an important authentication source but with significant new challenges. Biometrics can be stolen and are not replaceable. Research on cancelable biometrics has tried to mitigate this issue (e.g., see [18] for a survey). These methods do not aim to replace the actual biometrics, of course, but rather they aim to replace a compromised (stolen) stored reference template with a different one. Most of these methods add certain distortions to the transformation to the canonical template, so that if a reference template is compromised for one distortion, a different one can be generated and used. However, only some (e.g., [26]) of these considered methods result in strong security, compatible to provable security of common cryptographic tools. This provides more motivation to move away from models where a template is stored, precisely because if that template is successfully attacked, then an attacker will be able to leverage that information remotely and at scale.

The main challenge to authentication using biometrics is their noisy nature: repeated readings can differ significantly. Current techniques for eliminating noise, such as fuzzy extractors, come at a significant entropy cost. However, we believe this approach has promise and that key derivation from noisy sources can be improved significantly. For noisy sources such as biometrics, existing processing algorithms have been optimized for identification, not authentication. Revisiting feature extraction for such sources with authentication in mind should reduce the entropy loss.

In this article, we show that the traditional FMR versus FNMR identification metric does not properly optimize for the authentication task. We instead propose using an authentication-specific metric, such as KS, rather than FMR versus FNMR. To illustrate the difference between these two approaches, we discuss attempts to optimize the iris biometric according to the authentication metric. Unfortunately, key derivation from the iris still remains a challenge. This article, and the works of our predecessors, lay the foundation for future progress in optimized key derivation from biometrics and their application to authentication systems.

ACKNOWLEDGMENTS

This research was enabled by access to the multispectral iris data set licensed through the Scitor Corporation. The multispectral iris data set was collected by Southern Methodist University (SMU) under the oversight of the SMU Institutional Review Board [44], [45] that enabled our research and experiments. We thank Jonas Borgstrom and Jason Thornton for their help in the early stages of our project, and Salil Prabhakar of Delta ID for providing us with an iris camera that we used for our prototypes. We are grateful to Shawn Campbell and Carolyn Greenberg for their assistance in preparing the manuscript and to Jessica Barragué for her assistance in the production of the article. This work is sponsored by Assistant Secretary of Defense for Research and Engineering under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are our own and are not necessarily endorsed by the United States Government.

AUTHORS

Gene Itkis (itkis@ll.mit.edu) received his Ph.D. degree in computer science from Boston University (BU) in 1996. From 1999 to 2009, he was on the faculty of the Computer Science Department at BU, where he founded the Applied Cryptography and e-Security group, the Industry Talks series, and served as the associate director for research at the Center for Reliable Information Systems and Cyber Security, National Security Agency National Center of Academic Excellence in Information Assurance Education. In 2009, he joined the Massachusetts Institute of Technology Lincoln Laboratory, where he is working on applied security projects including cloud security, key management, biometrics, and antitamper technologies, among others.

Venkat Chandar (chandarvenkat@verizon.net) received S.B. degrees in electrical engineering and computer sciences and mathematics in 2006, an M.Eng. degree in electrical engineering and computer sciences in 2006, and a Ph.D. degree in electrical engineering and computer sciences in 2010, all from the Massachusetts Institute of Technology (MIT). His current research interests include coding theory and algorithms, optical communications and quantum information theory, and more recently, finance theory. He was with MIT Lincoln Laboratory from 2010 until 2014, and is now a quantitative researcher at D.E. Shaw and Co.

Benjamin Fuller (bfuller@ll.mit.edu) joined the Massachusetts Institute of Technology (MIT) Lincoln Laboratory in 2007. His research focuses on cryptography and practical solutions to secure communication with a past focus on cryptographic key management and key derivation. He received the B.S. degree from Rensselaer Polytechnic Institute in 2006 and the M.A. and Ph.D. degrees from Boston University in 2011 and 2015, respectively. He completed his Ph.D. degree at Boston University under the direction of Prof. Leonid Reyzin, focusing on cryptography with imperfect and noisy randomness. His Ph.D. research focused on new approaches for the construction of fuzzy extractors.

Joseph P. Campbell (j.campbell@ieee.org) received the Ph.D. degree in electrical engineering from Oklahoma State University in 1992. Since 2001, he has been with the Massachusetts Institute of Technology's Lincoln Laboratory. He is currently the associate group leader of the Human Language Technology Group. He is a member of the IEEE Awards Planning and Policy Committee and chaired the Biometric Consortium and the IEEE Jack S. Kilby Signal Processing Medal Committee. He was a member of the IEEE Information Forensics Security Committee, the IEEE Signal Processing Technical Committee, and the IEEE Signal Processing Society's Board of Governors. He was an associate editor of *IEEE Transactions on Speech and Audio Processing*, the vice president of Technical Activities of the IEEE Biometrics Council, and a member of the IEEE Signal Processing Society Fellow Reference Committee. He was an IEEE Distinguished Lecturer and is an IEEE Fellow.

Robert K. Cunningham (rkc@ll.mit.edu) received the Ph.D. degree in computer engineering and cognitive and neural systems from Boston University in 1998. He is currently the leader of the Secure, Resilient Technology Group at the Massachusetts Institute

of Technology (MIT) Lincoln Laboratory. He won the 2015 MIT Excellence Award for Bringing Out the Best. He has published works on computer intrusion detection, computer worms, system protection, software development best practices, and on signal and image processing. He has served the IEEE as a program member for multiple conferences and workshops, and as a program and general chair for the IEEE Symposium on Technologies for Homeland Security and the IEEE Security and Privacy Symposium. He is a Senior Member of the IEEE.

REFERENCES

[1] J. H. Saltzer and M. D. Schroeder, "The protection of information in computer systems," *Proc. IEEE*, vol. 63, no. 9, pp. 1278–1308, 1975.

[2] H. M. Levy, *Capability-Based Computer Systems*. Newton, MA: Butterworth-Heinemann, 1984.

[3] J. Daugman, "How iris recognition works," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 14, no. 1, pp. 21–30, Jan. 2004.

[4] M. Zviran and W. J. Haga, "A comparison of password techniques for multilevel authentication mechanisms," *Comput. J.*, vol. 36, no. 3, pp. 227–237, 1993.

[5] S. Brostoff and M. Sasse, "Are passfaces more usable than passwords? A field trial investigation," *People Comput.*, pp. 405–424, 2000.

[6] C. Ellison, C. Hall, R. Milbert, and B. Schneier, "Protecting secret keys with personal entropy," *Future Generation Comput. Syst.*, vol. 16, no. 4, pp. 311–318, 2000.

[7] F. Monrose, M. K. Reiter, and S. Wetzel, "Password hardening based on keystroke dynamics," *Int. J. Inform. Security*, vol. 1, no. 2, pp. 69–83, 2002.

[8] J. Bonneau, C. Herley, P. C. v. Oorschot, and F. Stajano. (2012, Mar.). The quest to replace passwords: a framework for comparative evaluation of Web authentication schemes. Univ. Cambridge, Computer Lab., Tech. Rep. UCAM-CL-TR-817. [Online]. Available: http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-817.pdf

[9] R. Morris, R. Morris, K. Thompson, and K. Thompson, "Password security: A case history," *Commun. ACM*, vol. 22, pp. 594–597, Nov. 1979.

[10] N. K. Ratha, J. H. Connell, and R. M. Bolle, "Biometrics break-ins and band-aids," *Pattern Recognit. Lett.*, vol. 24, no. 13, pp. 2105–2113, 2003.

[11] J. Saltzer and M. Schroeder, "The protection of information in computer systems," *Proc. IEEE*, vol. 63, no. 9, pp. 1278–1308, Sept. 1975.

[12] A. Juels and M. Wattenberg, "A fuzzy commitment scheme," in *Proc. 6th ACM Conf. Computer and Communication Security.*, Nov. 1999, pp. 28–36.

[13] A. Juels and M. Sudan, "A fuzzy vault scheme," in *Proc. IEEE Int. Symp. Information Theory*, 2002, p. 408.

[14] A. Juels and M. Sudan. (2006). A fuzzy vault scheme. [Online]. *Designs, Codes Cryptogr.*, 38, 237–257. Available: http://dx.doi.org/10.1007/s10623-005-6343-z

[15] G. I. Davida, Y. Frankel, and B. J. Matt, "On enabling secure applications through off-line biometric identification," in *Proc. IEEE Symp. Security and Privacy.*, 1998, pp. 148–157.

[16] G. Davida, B. Matt, Y. Frankel, and R. Peralta, "On the relation of error correction and cryptography to an off line biometric based identification scheme," in Proc. *Workshop on Coding and Cryptography*, 1999, pp. 129–138.

[17] U. Uludag, S. Pankanti, S. Prabhakar, and A. K. Jain, "Biometric cryptosystems: issues and challenges," *Proc. IEEE*, vol. 92, no. 6, pp. 948–960, 2004.

[18] C. Rathgeb and A. Uhl, "A survey on biometric cryptosystems and cancelable biometrics," *EURASIP J. Inform. Security*, vol. 2011, no. 1, pp. 1–25, 2011.

[19] S. Prabhakar, S. Pankanti, and A. K. Jain, "Biometric recognition: Security and privacy concerns," *IEEE Security Privacy*, vol. 1, no. 2, pp. 33–42, 2003.

[20] M. Blanton and W. M. Hudelson, "Biometric-based non-transferable anonymous credentials," in *Information and Communications Security*. New York: Springer, 2009, pp. 165–180.

[21] E. Grosse and M. Upadhyay, "Authentication at scale," *IEEE Security Privacy*, vol. 11, no. 1, pp. 15–22, Jan./Feb. 2013.

[22] J. Galbally, A. Ross, M. Gomez-Barrero, J. Fierrez, and J. Ortega-Garcia, "From the iriscode to the iris: A new vulnerability of iris recognition systems," in *Black Hat Briefings USA*, 2012. [Online]. Available: https://www.blackhat.com/html/bh-us-12/bh-us-12-briefings.html#Galbally

[23] J. Galbally, A. Ross, M. Gomez-Barrero, J. Fierrez, and J. Ortega-Garcia, "Iris image reconstruction from binary templates: An efficient probabilistic approach based on genetic algorithms," *Comput. Vis. Image Understand.*, vol. 117, no. 10, pp. 1512–1525, Oct. 2013.

[24] A. Ross, J. Shah, and A. Jain, "From template to image: Reconstructing fingerprints from minutiae points," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 4, pp. 544–560, Apr. 2007.

[25] T. Ignatenko and F. M. Willems, "Biometric security from an information-theoretical perspective," in *Foundations and Trends in Communications and Information Theory*. Now Publishers, pp. 135-316, 2012.

[26] J. Bringer and H. Chabanne, "An application of the naccache-stern knapsack cryptosystem to biometric authentication," in *2007 IEEE Workshop on Automatic Identification Advanced Technologies*, June 2007, pp. 180–185.

[27] T. Dierks and E. Rescorla, "The transport layer security (TLS) protocol," *IETF RFC 5246*, 2008. [Online]. Available: https://tools.ietf.org/html/rfc5246

[28] J. Daugman, "Probing the uniqueness and randomness of iriscodes: Results from 200 billion iris pair comparisons," *Proc. IEEE*, vol. 94, no. 11, pp. 1927–1935, 2006.

[29] S. Ziauddin and M. N. Dailey, "Iris recognition performance enhancement using weighted majority voting," in *Proc. 15th IEEE Int. Conf. Image Processing (ICIP),* 2008, pp. 277–280.

[30] S. P. Fenker and K. W. Bowyer, "Experimental evidence of a template aging effect in iris biometrics," in *2011 IEEE Workshop on Applications of Computer Vision (WACV)*, 2011, pp. 232–239.

[31] J. Howard and D. Etter, "The effect of ethnicity, gender, eye color and wavelength on the biometric menagerie," in *Proc. IEEE Int. Conf. Technologies for Homeland Security (HST)*, 2013, pp. 627–632.

[32] S. Lagree and K. W. Bowyer, "Predicting ethnicity and gender from iris texture," in *Proc. 2011 IEEE Int. Conf. Technologies for Homeland Security (HST),* 2011, pp. 440–445.

[33] C. K. Boyce, "Multispectral iris recognition analysis: techniques and evaluation," M.S.E.E. thesis, West Virginia Univ., Dept. Comput. Sci. Electric. Eng., Citeseer, 2006.

[34] I. Daubechies, "The wavelet transform, time-frequency localization and signal analysis," *IEEE Trans. Inform. Theory*, vol. 36, no. 5, pp. 961–1005, 1990.

[35] L. Masek, "Recognition of human iris patterns for biometric identification," Bachelor's thesis, School Comp. Sci. Software Eng., Univ. Western Australia, 2003.

[36] J. Beirlant, E. J. Dudewicz, L. Györfi, and E. C. Van der Meulen, "Nonparametric entropy estimation: An overview," *Int. J. Math. Stat. Sci.*, vol. 6, no. 1, pp. 17–39, 1997.

[37] T. Schürmann, "Letter to the editor: Bias analysis in entropy estimation," *J. Phys. A,* vol. 27, pp. L295–L301, July 2004.

[38] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby, "A pseudorandom generator from any one-way function," *SIAM J. Comput.*, vol. 28, no. 4, pp. 1364–1396, 1999.

[39] L. A. Levin. (2000). The tale of one-way functions. [Online]. CoRR, vol. cs.CR/0012023. Available: http://arxiv.org/abs/cs.CR/0012023

[40] R. Canetti, B. Fuller, O. Paneth, L. Reyzin, and A. Smith. (2014). Key derivation from noisy sources with more errors than entropy. [Online]. Cryptology ePrint Archive, Report 2014/243. Available: http://eprint.iacr.org/

[41] S. P. Vadhan, "Constructing locally computable extractors and cryptosystems in the bounded-storage model," *J. Cryptol.*, vol. 17, no. 1, pp. 43–77, 2004.

[42] J. Gentile, N. Ratha, and J. Connell, "SLIC: Short-length iris codes," in *IEEE 3rd Int. Conf. Biometrics: Theory, Applications, and Systems (BTAS'09),* 2009, pp. 1–5.

[43] K. P. Hollingsworth, K. W. Bowyer, and P. J. Flynn, "The best bits in an iris code," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 6, pp. 964–973, 2009.

[44] D. Etter, J. Webb, and J. Howard, "Collecting large biometric datasets: A case study in applying software best practices," *Immutable Laws Softw. Develop., CrossTalk: J. Defense Softw. Eng.*, pp. 4–8, May/June 2014.

[45] J. J. Howard, "Large scale pattern recognition models for identifying subject specific match probability across datasets with controlled variability," Ph.D. dissertation, Southern Methodist Univ., 2014.

[46] C. H. Bennett, G. Brassard, and J.-M. Robert, "Privacy amplification by public discussion," *SIAM J. Comput.*, vol. 17, no. 2, pp. 210–229, 1988.

[47] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith, "Fuzzy extractors: How to generate strong keys from biometrics and other noisy data," *SIAM J. Comput.*, vol. 38, no. 1, pp. 97–139, 2008.

[48] N. Nisan and D. Zuckerman, "Randomness is linear in space," *J. Comput. Syst. Sci.*, vol. 52, no. 1, pp. 43–52, Feb. 1996.

[49] O. Dunkelman, M. Osadchy, and M. Sharif, "Secure authentication from facial attributes with no privacy loss," in *Proc. ACM SIGSAC Conf. Computer & Communications Security.*, 2013, pp. 1403–1406.

[50] B. Fuller, X. Meng, and L. Reyzin, "Computational fuzzy extractors," in *Proc. Advances in Cryptology (ASIACRYPT),* 2013, pp. 174–193.

[51] C. Herder, L. Ren, M. van Dijk, M.-D. M. Yu, and S. Devadas. (2014). Trapdoor computational fuzzy extractors. [Online]. Cryptology ePrint Archive, Rep. 2014/938. Available: http://eprint.iacr.org/

[52] B. Kaliski, "Pkcs# 5: Password-based cryptography specification version 2.0," in *IETF RFC* 2898. RFC Editor, 2000. [Online]. Available: http://dx.doi.org/10.17487/RFC2898

[53] U. Feige, A. Fiat, and A. Shamir, "Zero-knowledge proofs of identity," *J. Cryptol.*, vol. 1, no. 2, pp. 77–94, 1988.

[SP]