

Secure Smartcard-Based Fingerprint Authentication ^{*}

[full version]

T. Charles Clancy[†]

Computer Science
University of Maryland, College Park

tcc@umd.edu

Negar Kiyavash, Dennis J. Lin

Electrical and Computer Engineering
University of Illinois, Urbana-Champaign

{kiyavash, djlin}@uiuc.edu

ABSTRACT

In this paper, the fundamental insecurities hampering a scalable, wide-spread deployment of biometric authentication are examined, and a cryptosystem capable of using fingerprint data as its key is presented. For our application, we focus on situations where a private key stored on a smartcard is used for authentication in a networked environment, and we assume an attacker can launch off-line attacks against a stolen card.

Juels and Sudan's *fuzzy vault* is used as a starting point for building and analyzing a secure authentication scheme using fingerprints and smartcards called a *fingerprint vault*. Fingerprint minutiae coordinates m_i are encoded as elements in a finite field F and the secret key is encoded in a polynomial $f(x)$ over $F[x]$. The polynomial is evaluated at the minutiae locations, and the pairs $(m_i, f(m_i))$ are stored along with random (c_i, d_i) chaff points such that $d_i \neq f(c_i)$. Given a matching fingerprint, a valid user can separate out enough true points from the chaff points to reconstruct $f(x)$, and hence the original secret key.

The parameters of the vault are selected such that the attacker's vault unlocking complexity is maximized, subject to zero unlocking complexity with a matching fingerprint and a reasonable amount of error. For a feature location measurement variance of 9 pixels, the optimal vault is 2^{69} times more difficult to unlock for an attacker compared to a user possessing a matching fingerprint, along with approximately a 30% chance of unlocking failure.

^{*}Supported by The Boeing Company under Illinois Technology Challenge Grant 03-115. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of The Boeing Company or the State of Illinois.

[†]This work was completed while T. Clancy was with the Electrical and Computer Engineering department at the University of Illinois, Urbana-Champaign.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WBMA'03, November 8, 2003, Berkeley, California, USA.
Copyright 2003 ACM 1-58113-779-6/03/00011 ...\$5.00.

Categories and Subject Descriptors

E.3 [Data]: Data Encryption

Keywords

biometrics, fingerprint, smartcard, authentication

1. INTRODUCTION

In our increasingly electronic environment, everything is becoming a network, from the computer terminals we sit at, to the supermarket checkout, and even the locks on our doors. In these abstract networks, all forms of authorization and access control require networks to have a secure method of authenticating users.

Smartcards offer a new paradigm for authentication. Now, users' private keys are stored on smartcards. These users can prove their identity by using the card to provide a correctly signed message to an authentication server. Relying on the security of public-key authentication, the new task is to protect the private key on the smartcard itself. We believe that biometric authentication, in particular fingerprints, is a practical method of providing this protection.

There are two main approaches for using biometric information. The first is fingerprint matching, where the smartcard stores a template of the user's fingerprint and requires the user to present a matching template before it will sign messages on the user's behalf. The second method is fingerprint *mapping*, where a fingerprint is used to obscure the private key, without storing a template. The private key can only be recovered and consequently used to sign an authentication message if a valid fingerprint is provided.

From a fundamental security standpoint, the key distinction between the two techniques is whether a physical attack on the smartcard could yield any useful information. If the private key and biometric template are both stored unencrypted on the smartcard, they would both be susceptible to a physical attack. The goal here is to provide guaranteed security, rather than the illusion of security.

This paper focuses on using a modified version of Juels and Sudan's fuzzy vault [10] termed the *fingerprint vault* to encrypt the private key on the smartcard using fingerprint information, and derives theoretic bounds on its security. By analyzing the techniques used to perform the authentication and the entropy of the biometric data used, the complexity of various attacks can be quantified. In all cases, there is a trade-off between reproducibility and security.

The overall goals of this work are as follows:




Figure 1: Authentication Diagram

- define the fingerprint vault scheme and its associated algorithms;
- present probabilistic and systematic bounds on some of the vault parameters;
- define a class of unlocking techniques that can be used for both attacks and by legitimate users; and
- find the optimal fuzzy vault parameter choice to maximize attackers' complexity and minimize users' complexity, while ensuring reasonable reproducibility

The remaining sections are organized as follows. Section two discusses background information and motivates the research. Section three outlines prior similar research efforts. Section four presents the fingerprint vault and analyzes its applicability from a probabilistic sense. Section five analyzes the security of the fingerprint vault from a complexity-theoretic sense. Section six describes empirical results stemming from sample fingerprint data. Section seven concludes. Proofs are provided in the appendix of the extended version found on the authors' website.

2. BACKGROUND AND MOTIVATION

An airport is a motivating example where a fingerprint and smartcard solution is ideal. Passengers each have smartcards storing their personal information, including a private key. In order to access the information on their smartcard at airline terminals, they must provide the smartcard with valid fingerprint data.

First, we present a simplified, generic protocol for use with smartcard-based biometric authentication. It is similar to those found in many popular public-key authentication schemes, such as SESAME [24, 16] and SSH [26].

In general, a server should believe a user is who they say they are if they can provide a signed message containing their identity and a nonce or challenge (i.e. random number) selected by the server. If the public key associated with the specified identity correctly verifies the signature, only the valid user could have sent it. If a certificate is not provided by the user, the server will need a database of public keys.

From the client standpoint, signing ability is required to authenticate. Figure 1 demonstrates how this can be done with biometric data. First, a fingerprint image is captured from a scanner. This data is sent to a terminal which translates it into some smaller numeric representation, or template. Both the message digest we wish to sign and the biometric template are sent to the smartcard. Provided the biometric data is valid, the smartcard will use its internally stored private key to generate and return a signature for the message digest.

Our analysis shall focus on the methods by which smartcards use biometric data to sign a message. This assumes

the worst case scenario: a smartcard has been stolen, and an attacker with complete physical access is attempting to retrieve the private key.

Given physical access, there are two main classes of physical attacks against smartcards: noninvasive or side-channel attacks, and invasive attacks.

Three of the most popular noninvasive attacks are power analysis [13], timing analysis [12], and electromagnetic (EM) analysis [1]. All of these attacks can be used to determine sensitive information on a smartcard; however, noninvasive attacks can generally be thwarted by clever algorithms, data obfuscation, and shielding techniques.

Invasive attacks [14, 15] generally involve dissolving the chip packaging and reverse engineering the processor itself. Given complete physical access, it is impossible to prevent an attacker from retrieving data stored in memory. The only way to protect against such an attack is to encrypt the contents of memory using a key not stored on the card itself. Then, an attacker may retrieve the data from the card, but it will be of no use.

Protecting the smartcard now requires efforts on two fronts. First, the smartcard processor itself must be designed such that it is immune to the various on-line side-channel attacks. Secondly, without a matching fingerprint, an attacker should not be able to obtain any sensitive information from the card. In particular, users' private keys must be stored in an encrypted format. This second front is the focus of the work presented here.

3. PAST WORK

In the past few years, there have been several research efforts aimed at addressing the intersections between cryptography and biometrics. Here we address biometric cryptosystems in general, delve more deeply in to the fuzzy vault, and then briefly examine various polynomial reconstruction techniques.

3.1 Biometric Cryptosystems

In 1998, Davida, et al. [6], were among the first to suggest off-line biometric authentication. It moved biometric data from a central server into a signed form on a portable storage device, such as a smartcard. Their system was essentially a PKI-like environment that did local fingerprint matching. Its main flaw is that it required some local authentication authority to have a key capable of decrypting the template stored on the storage device. While they address the key management issues, the basic premise is still that of local fingerprint matching, and is therefore inherently insecure.

The next year there were three innovative, yet similar methods that did not perform biometric matching. The first is the fuzzy commitment scheme [11]. Here, a secret (presumably a private key used for later authentication) is encoded using a standard error correcting code such as Hamming or Reed-Solomon, and then XOR-ed it with a biometric template. To retrieve the secret, a slightly different biometric template can again be XOR-ed, and the result put through an error correcting decoder. Some small number of bit errors introduced in the key can be corrected through the decoding process. The major flaw of this system is that biometric data is often subject to reordering and erasures, which cannot be handled using this simple scheme.

In [19], a technique was proposed using the phase information of a Fourier transform of the fingerprint image. The fin-

gerprint information and a randomly chosen key are mixed together to make it impossible to recover one without the other. In order to tolerate errors, the system used a filter that minimizes the output variance corresponding to the input images. To provide further redundancy, an encoding process stores each bit multiple times. The work does not address how much these steps reduce the entropy of the original image; thus, it is not clear that there exists a set of parameters which will allow the system to reliably recognize legitimate users while providing a reasonable amount of security.

A third paper [18] has a similar theoretical foundation to this work, but aims toward a completely different application. Here, Monrose, et al., attempt to add entropy to users' passwords on a computer system by incorporating data from the way in which they type their password. Since the biometric being used here so so radically different from fingerprints, their results are not applicable to this work.

Recently, Juels and Sudan [10] proposed the *fuzzy vault*, a new architecture with applications similar to Juels and Wattenberg's *fuzzy commitment scheme*, but is more compatible with partial and reordered data. The fuzzy vault is used here as a starting point for the biometric scheme presented in this paper.

3.2 Fuzzy Vault

Here, we describe the original fuzzy vault, with some slight notational differences. As with any cryptosystem, there is some message m that needs to be encrypted, or in this case *locked*. Some symmetric fuzzy key can be used to accomplish this task, and then used again later to decrypt, or *unlock* the original message. Here, our message m is first encoded as the coefficients of some degree k polynomial in x over a finite field \mathbb{F}_q .

This polynomial $f(x)$ is now the secret to protect. The locking set \mathcal{L} is a set of t values $l_i \in \mathbb{F}_q$ making up the fuzzy encryption key, where $t > k$. The locked vault contains all the pairs $(l_i, f(l_i))$ and some large number of chaff points (α_j, β_j) , where $f(\alpha_j) \neq \beta_j$. After adding the chaff points, the total number of items in the vault is r .

In order to crack this system, an attacker must be able to separate the chaff points from the legitimate points in the vault. The difficulty of this operation is a function of the number of chaff points, among other things. A legitimate user should be able to unlock the vault if they can narrow the search space. In general, to successfully interpolate the polynomial they have an unlocking set \mathcal{U} of t elements such that $\mathcal{L} \cap \mathcal{U}$ contains at least $k + 1$ elements

To summarize the vault parameters:

- $f(x)$ is a degree k polynomial in $\mathbb{F}_q[x]$
- $t \geq k$ points in \mathcal{L} interpolate through $f(x)$
- $r \geq t$ is total number of points in the vault

This vault shall be referred to as $\mathcal{V}(\mathbb{F}_q, r, t, k)$.

Spurious polynomials of degree k interpolated by t points may show up in the randomly selected chaff points. In [10], the authors present a lemma describing the security of their scheme based on the number of these polynomials that exist in a vault with general parameters.

LEMMA 1. *For every $\mu > 0$, with probability $1 - \mu$, a vault of size t contains at least $\frac{\mu}{3} q^{k-t} (r/t)^t$ polynomials $f'(x)$ of degree less than k such that the vault contains exactly t points of the form $(x, f'(x))$.*

3.3 Polynomial Interpolation

In order to actually reconstruct the secret locked within the fuzzy vault, the points in the unlocking set must be used to interpolate a polynomial. The unlocking set will contain both real points and chaff points.

The simplest mechanism for recovering the polynomial is a brute-force search, where various $k + 1$ element subsets of the unlocking set are used to interpolate a degree k polynomial, using Newtonian Interpolation [8].

A second method is to use a Reed-Solomon decoder [3], as suggested by Juels and Sudan. While RS codes are traditionally used to correct errors in messages transmitted over noisy channels, they are essentially a generalization of the polynomial reconstruction problem. Using a (t, k) code, t points can be fed into the decoder, and the degree k polynomial will be returned.

There are two main RS decoding algorithms: the Berlekamp-Massey algorithm [17], and the Guruswami-Sudan algorithm [7]. Berlekamp-Massey requires $\frac{k+t}{2}$ real points in the unlocking set while Guruswami-Sudan only requires \sqrt{kt} . Unfortunately, this extra error correcting capability requires significantly more computation. Interestingly enough, the two algorithms provide nearly identical unlocking complexities for a wide range of vault parameters, so we shall focus on the Berlekamp-Massey method.

Another field called *noisy polynomial interpolation* has had some recent advances, notably by Arora and Khot [2] and Bleichenbacher and Nguyen [5]. However, the results of this work are not applicable to the fuzzy vault. In [2], they examine the problem of finding all polynomials that interpolate through the points $(x_i, [y_i - \delta, y_i + \delta])$. The *noise* in our points has been removed by the existence of the fuzzy vault, so this is not useful to us. In [5], they look at the problem of interpolating the points $(x_i, y_{i,1}), (x_i, y_{i,2}), \dots$, but since we do not have chaff points overlapping with real points, their new algorithm is also not applicable.

4. FINGERPRINT VAULT

In this section, we describe our modified *fingerprint vault*, and the algorithms used to lock and unlock data. Additionally, theoretic bounds on our ability to successfully unlock the vault are determined.

4.1 Generalized Fuzzy Vault

First, the fuzzy vault specifies that the size of the locking set, unlocking set, and RS codewords are all the same size. Here, we loosen this restriction.

- t is the number of points in \mathcal{L}
- τ is the number of points in \mathcal{U}
- n is the RS codeword size

In [10], the authors consider the case where \mathcal{L} and \mathcal{U} are taken from discrete sets, however for biometric purposes these values are not discrete (or are taken from a sufficiently high-resolution discrete set). Hence, for all the elements in \mathcal{U} , we need to find the closest elements in vault, and then try unlocking with those values. Consequently, a quantization problem is introduced. How closely can we pack chaff points and still maintain a reasonable probability of quantization error?

The key being used to lock our fuzzy vault is pixel coordinate locations, (x_i, y_i) , of features on a fingerprint image.

Consequently, an ideal field for the vault is \mathbb{F}_{p^2} , such that p is prime.

Lemma 1 probabilistically gives us the number of spurious polynomials of a particular degree in our vault. The presence of many such polynomials is the key to proof of security in [10]. Unfortunately, for reasonable vault parameters (see Section 6), there exists a δ with $k \leq \delta \leq t$ such that the expected number of spurious degree k polynomials interpolated by more than δ points is less than one. The consequence is that the brute-force search for a degree k polynomial interpolating δ points will yield a unique result, namely, the vault secret.

The precise value of δ such that the results of an unlocking attempt can be verified will be required later:

COROLLARY 1. *A value of δ satisfying the above requirement is*

$$\delta \geq \left\lceil \frac{\log \frac{1}{3} p^{2k}}{\log \frac{kp^2}{r}} \right\rceil \quad (1)$$

However, an attacker will still have to search through the space of possible polynomials. We will show in Section 5 that this complexity can still be formidable. First, however, we will define our algorithms.

4.2 Feature Extraction

Feature extraction has been the focus of much research in past years, and this paper does not address it in detail. For the tests performed in this paper, the VeriFinger toolkit from Neurotechnologija Ltd. [25] was used to extract fingerprint features. The feature extraction process is visually represented in Figure 2.

In 2(a), one can see an image scan of a fingerprint. This is received directly from the fingerprint capture device. Various edge detection algorithms are then used to convert that information into 2(b). This figure features a cleaned-up version of the original scan. From there, features can be readily identified, as in 2(c). Each identified feature in 2(c) represents a fingerprint *minutiae*, which is a location where a fingerprint ridge either splits or ends.

Here, we shall consider the feature extraction and alignment as a black box, yielding normalized (x, y) pixel coordinates of fingerprint minutiae. The outputs from the black box for several scans of the same finger are generally close to one another, unless the fingerprint image was severely clipped. The intra-scan variance imposes limitations on our system.

4.3 Feature Noise

Each person’s fingerprint consists of a fixed set of minutiae locations $m_i = (x_i, y_i) \in \mathcal{M}$. However, due to systematic errors in image capture, processing, and alignment, noise is added to each point, such that our final points are

$$m'_i = (x_i + n_{xi}, y_i + n_{yi})$$

Additionally, the processing noise may discard features or add additional features that are not present on the actual fingerprint.

The next step in the analysis is to derive a model for the noise introduced by the capture device and extraction algorithms. From this information, we can derive matching error probabilities. To simplify analysis, an additive Gaussian noise model will be assumed.




Figure 3: Plot of minutiae from 5 scans of the same person, with reliable regions marked

To estimate the distribution on the minutiae locations, statistical data is needed. To accomplish this, features were extracted from N sample fingerprint images of the same person, and then aligned. The data used here was associated using the bounded nearest neighbor averaging technique described in the next section. The result is a set of expected values (\bar{x}_i, \bar{y}_i) for each detected minutiae, the number $n \leq N$ of samples having a minutiae in that neighborhood, and the variance and covariance of those n minutiae, $(\sigma_{x,i}^2, \sigma_{y,i}^2, \rho_i)$. Figure 3 shows these regions where features reliably appear.

These values will be used later to compute a bound on the possible density of chaff points for a given quantization error probability.

4.4 Locking Set

The locking set \mathcal{L} is computed in a method similar to the method for finding minutiae variances. A user’s finger is scanned and processed N times, resulting in N sets of minutiae, b_1, \dots, b_N . These are correlated using the following algorithm with distance threshold T and multiplicity threshold S :

1. let A be the set of average points with multiplicity
2. for each minutiae set b_i
3. for each minutiae $m_j \in b_i$
4. find element in $n_k \in A$ such that $|n_k - m_j| < T$
5. select closest n_k that has not already been used
6. if no matches, add m_j to A with multiplicity one
7. else add m_j to average and increase multiplicity
8. $\mathcal{A} = \{a \in A : \text{multiplicity}(a) > S\}$

Features appearing in S or fewer scans are discarded as noise. These points generally occur in the edge regions of the image, where feature extraction is less reliable. This locking set can then be used to create the fingerprint vault.

4.5 Chaff Points

The number and location of chaff points is limited by the variance in the fingerprint capturing and feature extraction algorithm. Chaff points cannot be placed too close to real points, or they will cause quantization problems. Given some acceptable distance d is found, chaff points can be



Figure 2: Feature Extraction Process: (a) original image, (b) after edge detection, (c) including feature points

placed anywhere as long as they are at least distance d from any real points. Additionally, there is no reason to place chaff points next to each other at any distance less than d , because an attacker can immediately ignore them as unlikely candidates, as they are so close together.

LEMMA 2. *Given elements of \mathbb{F}_{p^2} have pairwise Euclidean distance no less than d , the total number of elements r with packing density ρ is less than $\frac{4\rho p^2}{d^2\pi}$.*

The optimal packing technique for circles is using a hexagonal lattice, and has a packing density of $\rho = \frac{\pi}{2\sqrt{3}} \approx 0.91$ [23]. Unfortunately, this density could never be achieved, as we require the locations of chaff points to look random. If they all existed on a lattice, any discontinuities in the lattice pattern would be the real points.

Points can be randomly packed by repeatedly selecting random field elements and putting a chaff point there if it is at least distance d from all other points. This yields a packing density of $\rho \approx 0.45$, and is guaranteed to be random. Another technique, *random close packing* [9], could yield densities closer to $\rho \approx 0.75$, but these techniques have not been sufficiently studied in the two-dimensional case, and their randomness has never been quantified.

For this results to be useful, a minimum distance d between points needs to be computed in terms of the vault parameters. Here, we make a simplifying assumptions which will slightly loosen our bound, but yield simpler results: the noise distribution is spherically Gaussian, or the two axes are independent and identically distributed.

LEMMA 3. *The probability of successfully decoding a single point at distance at least d from all others using the maximum likelihood rule is*

$$P_s \leq 1 - \exp\left(\frac{-d^2}{8\sigma^2}\right) \quad (2)$$

THEOREM 1. *The probability of error for decoding points at least δ out of t points in the fuzzy vault $\mathcal{V}(\mathbb{F}_{2^p}, r, t, k)$ is bounded below by*

$$P_e \geq \sum_{i=\delta}^t \binom{t}{i} \exp\left(-\frac{\rho p^2}{2r\pi\sigma^2}\right)^i \left(1 - \exp\left(-\frac{\rho p^2}{2r\pi\sigma^2}\right)\right)^{t-i} \quad (3)$$

for a given point variance σ^2 .

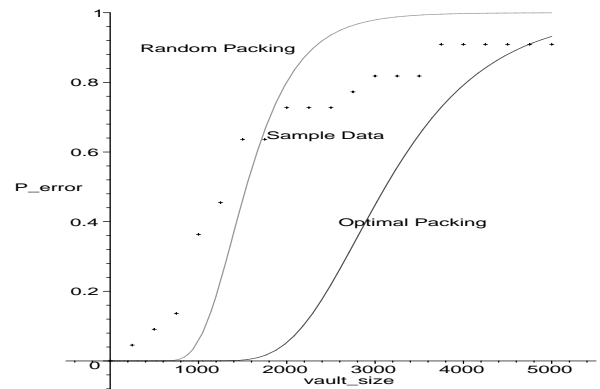


Figure 4: Probability of error as a function of r , using $\sigma^2 = 9$ and $p = 251$, with vault parameters $t = 40$ and $\delta = 12$, plotted with error probabilities using random packing in sample data sets.

Figure 4 shows the error probability as a function of the vault size for the optimal and randomized packing methods. Also included is error probabilities from actual fingerprint data using the random packing method. For each person, five fingerprint scans were available. The first four were used to create a vault, and the fifth was used to try and unlock it. The plotted probabilities represent the fraction of data sets that had enough true points in the unlocking set to successfully unlock the fingerprint vault. The experimental results are quite close to the theoretical, which is impressive given our simplified noise model.

Note that this is the probability of being able to successfully decode, and does not deal with the complexity of actually performing that decoding.

4.6 Unlocking Set

The unlocking set \mathcal{U} starts off initially as some set U of minutiae locations from a single scan of the user's finger. To select the elements of the unlocking set, for each point in U the user finds the closest point in the vault. If the fingerprint capture process did not introduce noise features, we should have $\mathcal{U} \subseteq \mathcal{L}$. However, there is some probability that a minutiae will by chance be closer to a chaff point than a true point. Also, there is a chance that U will contain

spurious minutiae which was not included in \mathcal{L} .

This is where Reed-Solomon codes comes into play. For any n points, we can determine the polynomial, or Reed-Solomon codeword in this case, if at least $\frac{\delta+n}{2}$ of those points are correct. If n is reasonably close to r , then very few attempts will be required to compute the polynomial.

5. UNLOCKING COMPLEXITY

Vault unlocking can be viewed in two contexts. The first is the complexity of a valid user unlocking a vault with a matching fingerprint image. One goal is to minimize this complexity. The second context is the complexity of an attacker without fingerprint information trying to crack the vault. We wish to maximize this complexity while the attacker wishes to minimize it.

There are two obvious techniques for unlocking the vault. The first is the brute-force method, or $\mathbf{bf}(r, t, k)$, where r is the total number of points, t is the number of real points, and k is the degree of the polynomial. For an attacker, r and t are the same as the ones in the vault parameter, however for a valid user, r is the size of their unlocking set and t is the number of non-chaff points in that set.

THEOREM 2. *The complexity of the $\mathbf{bf}(r, t, k)$ problem using a suitable δ to ensure a unique result is $\mathcal{C}_{bf} = \binom{r}{\delta} \binom{t}{\delta}^{-1}$.*

EXAMPLE 1. *Consider a vault over \mathbf{F}_{2512} with $r = 1000$ total points and $t = 40$ real points over a degree $k = 8$ polynomial. Under Corollary 1, $\delta = 12$. Using Theorem 2, the complexity of an attacker breaking the vault is approximately 2^{58} polynomial interpolations.*

EXAMPLE 2. *Consider the same vault, only a valid user intersects their unlocking set with the vault to obtain $r = 30$ points, $t = 22$ of which are real points. With such a small vault, obviously $\delta = k + 1 = 9$. Using Theorem 2, the complexity of a valid user unlocking the vault is approximately 2^7 polynomial interpolations.*

The second example illustrates that a brute-force decoding algorithm is less than ideal a valid user. Another method of unlocking the vault is through the use of a Reed-Solomon decoder. In the $\mathbf{rs}(r, t, n, \delta)$ problem, r , t , and δ have the same meanings as before, and n is the size of the Reed-Solomon codewords involved.

THEOREM 3. *The complexity of the $\mathbf{rs}(r, t, n, \delta)$ problem over \mathbb{F}_{p^2} is*

$$\mathcal{C}_{rs} = \binom{r}{n} \left(\sum_{i=\max(\frac{n+\delta}{2}, n-r+t)}^{\min(n,t)} \binom{r-t}{n-i} \binom{t}{i} \right)^{-1} \quad (4)$$

such that n satisfies $\delta \leq n \leq \min(r, 2t - \delta)$ and $n|(p^2 - 1)$.

COROLLARY 2. *The complexity of $\mathbf{bf}(r, t, \delta) = \mathbf{rs}(r, t, \delta, \delta)$, or taking $n = \delta$ reduces Reed-Solomon unlocking into a brute-force unlocking.*

COROLLARY 3. *For $r \gg t$, $\mathbf{bf}(r, t, \delta) \leq \mathbf{rs}(r, t, n, \delta)$, for all $n \geq \delta$, or unless an attacker can eliminate a significant number of chaff points of a locked vault, he or she can do no better than a brute-force attack.*

Let's examine the previous two examples in the context of Reed-Solomon decoding. Figure 5 illustrates the complexity of a full-scale attack, an attack where partial fingerprint information is known, and a legitimate unlocking of the vault. We can see that depending on the relationship of r and t , the optimal method for unlocking the vault can change. By using a Reed-Solomon decoder, a valid user can now unlock the vault in 1 or 2 tries, while an attacker can still do no better than a brute-force attack.

If an attacker is able to eliminate many of the chaff points, presumably through side knowledge of some of the fingerprint characteristics, finding the optimal attack now becomes more interesting. The minimum complexity is no longer one of the bounding cases.

In this paper, we assume that we can choose chaff points in such a way as to confuse the attacker and force him to consider all points. For the most part we can disregard attacks where an attacker can eliminate certain chaff points based on the probable minutiae configurations. Research [20, 21, 22] indicates that the probability of two people having the same fingerprint is approximately 1×10^{-80} . Assuming fingerprints are equiprobable, this corresponds to $-\log_2(1 \times 10^{-80}) \approx 265$ bits of entropy. Thus, for a complexity-theoretic attack, the minutiae entropy is sufficiently large, thus trivializing attacks which take probable minutiae configurations into account. This also indicates we could never achieve more than 265-bit security, regardless of vault parameters.

6. EMPIRICAL RESULTS

Throughout the paper, the term "reasonable vault parameters" has been used repeatedly. Here, we use actual fingerprint data to determine what "reasonable" is. Each of the vault parameters, p , r , t , and k has limitations placed on it by the behavior of actual fingerprint data.

The locking and unlocking algorithms were implemented in MATLAB, and sample fingerprint data was used to test the error probabilities. Four scans of the same individual were used to create a vault, and a fifth used to try and unlock it. The number of true points in the unlocking sets for these real vaults was used to validate the statistical models.

The field, \mathbb{F}_q , defines the underlying mechanics of our entire system. Throughout, we have been using \mathbb{F}_{p^2} , for prime p . In general, we wish to represent a feature pixel location. For the examples presented so far, $p = 251$ was used. This way, minutiae locations can be stored in 16-bit numbers, and $251^2 - 1 = 63000 = 2^3 \cdot 3^2 \cdot 5^3 \cdot 7$, a very smooth number, yielding many choices for the Reed-Solomon codeword size.

Increasing the fingerprint image resolution and consequently the field size has little effect on the resulting security. As the resolution increases, so does the minutiae variance, σ . These two parameters cancel one another out, making the underlying field selection based more on convenience than security.

The number of real points, t , is the size of the locking set. The algorithm described earlier takes several scans of the same person and locates minutiae that appear in two or more of the scans. This algorithm was implemented in MATLAB and used to create various locking sets. For 5 scans of each person and $S = 1$, we obtain locking sets which ranged from 25 to 60 points, with mean 38 and standard deviation 11.

The degree of the polynomial the vault protects is bounded below by the amount of data we wish to encode in it. Each




Figure 5: Log of complexity for Reed-Solomon decoding as a function of codeword size; (a) complexity of full attack, $rs(1000, 40, n, 12)$; (b) complexity of partial information attack, $rs(120, 40, n, 12)$; (c) complexity of legitimate unlocking, $rs(30, 22, n, 12)$.




Figure 6: Vault performance as a function of vault size, with $k = 14$, $\tau = 20$, and $t = 38$ over \mathbb{F}_{251^2}

coefficient is an element of \mathbb{F}_{251^2} , and can therefore hold 15.9 bits of information. As a result, a 128-bit key can be encoded using 9 coefficients, or in a degree 8 polynomial. Consequently, we shall consider $k \geq 8$.

The total number of points r depends on the number of chaff points added to the vault, and is a function of the desired error probability. Figure 4 gave the probabilities for a particular set of input values. Here, we shall examine this trade-off in more detail.

First, examine how the vault performs as a function of its size. Figure 6 shows both the complexity of a normal user, and the complexity of an attacker for a vault with 38 real points over a degree 14 polynomial. We can see that as the total number of points increases, so do both complexities. In order to keep user complexity to a minimum, we shall select the largest value of r such that the user has zero complexity.

The other key parameter that can be varied to alter our vault performance is k , the degree of our polynomial. Figure 7(a) shows the attack complexities as a function of k . This complexity was computed by first finding the maximum number of points r such that the user has zero complexity, and from there computing δ , the minimum number of points

interpolating our polynomial in order to guarantee success. Using r and δ , the difficulty of a brute-force attack can be computed.

Figure 7(b) is essentially a reality check on our selection of τ , the size of our unlocking set. For a given τ and k , real fingerprint data was again used to compute the probability of successfully unlocking the vault. It can be seen that given $\tau = 20$, approximately 20 to 30 percent error occurs. From a user’s perspective, this means that every couple times they access their smartcard, a second fingerprint scan will be required in order to successfully unlock the vault. This seems reasonable given that we expect that the false positive rate to be infinitesimally small. The corresponding curve in 7(a) indicates that the maximum complexity is 2^{69} for $k = 14$.

Consequently, over \mathbb{F}_{251^2} we have determined the optimal vault to be:

- polynomial: $k = 14$, $\delta = 17$
- chaff points: $r = 313$, $d = 10.7$
- attack complexity: 2^{69}

7. CONCLUSION

In this paper, we have considered the practical implications of using fingerprint information to secure a smartcard. Because fingerprints are often inconsistent, we must resort to a fuzzy scheme for storing the secret key. We show that with real-life parameters, it is impossible to ensure the security envisioned by Juels and Sudan. However, we define a modified scheme called the *fingerprint vault*, provide associated algorithms and a mechanism for finding optimal vault parameters. Parameters are provided which makes retrieving the secret 2^{69} times more difficult for the attacker than a legitimate user.

There are a couple ways by which security may be improved. An obvious way is to use multiple fingerprints to store a longer private key which could be hashed down to the appropriate length. Another way is to improve the detection and extraction algorithms so as to lower σ , allowing us to pack in more chaff points.

8. REFERENCES

- [1] AGRAWAL, D., ARCHAMBEAULT, B., RAO, J., AND ROHTAGI, P. The em-side channel(s). Workshop on Cryptographic Hardware and Embedded Systems, CHES 2002.




Figure 7: Vault performance as a function of k and τ : (a) attack complexities as a function of k for various τ ; (b) decoding failure as a function of k for various τ

- [2] ARORA, S., AND KHOT, S. Fitting algebraic curves to noisy data. ACM Symposium on Theory of Computing, STOC 2002.
- [3] BLAHUT, R. *Algebraic Codes for Data Transmission*. Cambridge University Press, 2003.
- [4] BLAHUT, R. *Modem Theory: An Introduction to Telecommunications*. Cambridge University Press, preprint.
- [5] BLEICHENBACHER, D., AND NGUYEN, P. Q. Noisy polynomial interpolation and noisy chinese remaindering. Advances in Cryptology, EUROCRYPT 2000.
- [6] DAVIDA, G., FRANKEL, Y., AND MATT, B. On enabling secure applications through off-line biometric identification. IEEE Symposium on Privacy and Security, 1998.
- [7] GURUSWAMI, V., AND SUDAN, M. Improved decoding of reed-solomon and algebraic-geometric codes. Symposium on Foundations of Computer Science, FOCS 1998.
- [8] HILDEBRAND, F. B. *Introduction to Numerical Analysis*. McGraw-Hill, 1956.
- [9] JAEGER, H., AND NAGEL, S. Physics of granular states. *Science* 255, 1524 (1992).
- [10] JUELS, A., AND SUDAN, M. A fuzzy vault scheme. ACM Conference on Computer and Communications Security, CCS 2002.
- [11] JUELS, A., AND WATTENBERG, M. A fuzzy commitment scheme. ACM Conference on Computer and Communications Security, CCS 1999.
- [12] KOCHER, P. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. Advances in Cryptology, CRYPTO 1996.
- [13] KOCHER, P., JAFFE, J., AND JUN, B. Differential power analysis. Advances in Cryptology, CRYPTO 1999.
- [14] KUHN, M., AND ANDERSON, R. Tamper resistance: A cautionary note. Workshop on Electronic Commerce, USENIX 1996.
- [15] KUMMERLING, O., AND KUHN, M. Design principles for tamper-resistant smartcard processors. Workshop on Smartcard Technology, USENIX 1999.
- [16] LOOI, M., ASHLEY, P., SEET, L. T., AU, R., AND VANDENWAUVER, M. Enhancing sesamev4 with smart cards. International Conference on Smartcard Research and Applications, CARDIS 1998.
- [17] MASSEY, J. L. Shift register synthesis and bch decoding. *IEEE Transactions on Information Theory* 15, 1 (1969), 122–127.
- [18] MONROSE, F., REITER, M., AND WETZEL, S. Password hardening based on keystroke dynamics. ACM Conference on Computer and Communications Security, CCS 1999.
- [19] NICHOLS, R. K., Ed. *ICSA Guide to Cryptography*. McGraw-Hill, 1999, ch. Biometric Encryption.
- [20] OSTERBERG, J., PARTHASARATHY, T., RAGHAVAN, T., AND SCLOVE, S. Development of a mathematical formula for the calculation of fingerprint probabilities based on individual characteristics. *Journal of the American Statistical Association* 72 (1977), 772–778.
- [21] PANKANTI, S., PRABHAKAR, S., AND JAIN, A. On the individuality of fingerprints. *IEEE Transactions on PAMI* 24 (2002), 1010–1025.
- [22] SCLOVE, S. The occurrence of fingerprint characteristics as a two-dimensional process. *Journal of the American Statistical Association* 74 (1979), 588–595.
- [23] STEINHAUS, H. *Mathematical Snapshots*, 3 ed. Dover, 1992.
- [24] VANDENWAUVER, M., GOVAERTS, R., AND VANDEWALLE, J. Overview of authentication protocols: Kerberos and sesame. IEEE Carnahan Conference on Security Technology 1997, pp. 108–113.
- [25] VERIFINGER. Neurotechnologija ltd. <http://www.neurotechnologija.com>.
- [26] YLONEN, T. Ssh secure login connections over the internet. Security Symposium, USENIX 1996, pp. 37–42.

A1. PROOFS

Proof of Corollary 1: A value of δ satisfying the above requirement is

$$\delta \geq \left\lceil \frac{\log \frac{1}{3} p^{2k}}{\log \frac{kp^2}{r}} \right\rceil \quad (\text{A1})$$

The requirement on the number of spurious polynomials can be rewritten as

$$\log \left(\frac{1}{3} p^{2(k-\delta)} \left(\frac{r}{\delta} \right)^\delta \right) < 0. \quad (\text{A2})$$

After expanding and rearranging, we have

$$\delta > \frac{\log \frac{1}{3} p^{2k}}{\log \frac{\delta p^2}{r}} \quad (\text{A3})$$

Unable to isolate the δ , on the right hand side substitute k for δ . In general, δ is larger than k , and the logarithm of the two is relatively close. For reasonable vault parameters, this only increases δ by approximately 0.3, and either does not affect the result or adds a small probabilistic safety net. ■

Proof of Lemma 2: Given elements of \mathbb{F}_{p^2} have pairwise Euclidean distance no less than d , the total number of elements r with packing density ρ is less than $\frac{4\rho p^2}{d^2\pi}$.

The problem reduces to packing circles within a rectangle. The square of possible locations has area p^2 , and the circles have area $\left(\frac{d}{2}\right)^2 \pi$. As a result, the number of circles is bounded above by

$$r \leq \rho \left(\frac{p^2}{(d/2)^2 \pi} \right) = \frac{4\rho p^2}{\pi d^2} \quad \blacksquare \quad (\text{A4})$$

Proof of Lemma 3: The probability of successfully decoding a single point at distance at least d from all others using the maximum likelihood rule is

$$P_s \leq 1 - \exp \left(\frac{-d^2}{8\sigma^2} \right) \quad (\text{A5})$$

The success probability is bounded below by integrating the Gaussian of distance $d/2$ from its mean, a frequently made simplification in communication theory [4]. Here, we use polar integration on the multivariate Gaussian distribution to compute the probability. Consequently, the computation is

$$P_s \leq \int_0^{2\pi} \int_0^{d/2} \frac{1}{2\pi\sigma^2} e^{-r^2/2\sigma^2} r dr d\theta \quad (\text{A6})$$

which simplifies to the above expression. ■

Proof of Theorem 1: The probability of error for decoding points at least δ out of t points in the fuzzy vault

$\mathcal{V}(\mathbb{F}_{2p}, r, t, k)$ is bounded below by

$$P_e \geq \sum_{i=\delta}^t \binom{t}{i} \exp \left(-\frac{\rho p^2}{2r\pi\sigma^2} \right)^i \left(1 - \exp \left(-\frac{\rho p^2}{2r\pi\sigma^2} \right) \right)^{t-i} \quad (\text{A7})$$

for a given point variance σ^2 .

The result above is essentially a combination of the first two lemmas. The vault is designed to allow some decoding errors, since there are $t \geq \delta$ valid points available, any combination of at least δ successful decodings is necessary. Also note that we are ignoring the probability of an error yielding another real point, rather than a chaff points. The real points used in the vault can be chosen to minimize the probability of this event.

The distance d is defined by the first lemma as

$$d^2 \leq \frac{4\rho p^2}{\pi t}. \quad (\text{A8})$$

Substituting for d in the P_s , we can compute P_e as

$$P_e \geq \sum_{i=\delta}^t \binom{t}{i} (P_s)^i (1 - P_s)^{t-i} \quad (\text{A9})$$

which expands to the given expression. ■

Proof of Theorem 2: The complexity of the $\mathbf{bf}(r, t, k)$ problem using a suitable δ to ensure a unique result is $C_{\mathbf{bf}} = \binom{r}{\delta} \binom{t}{\delta}^{-1}$.

The proof is a fairly straight-forward combinatorics argument. In the brute-force method, we must find δ points that interpolate a degree k polynomial. There are $\binom{r}{\delta}$ sets of any δ points. Of those sets, $\binom{t}{\delta}$ will yield successful results, as all δ points will exist on the degree k polynomial. The quotient of the two is the expected number of trials required to open the vault. ■

Proof of Theorem 3: The complexity of the $\mathbf{rs}(r, t, n, \delta)$ problem over \mathbb{F}_{p^2} is

$$C_{\mathbf{rs}} = \binom{r}{n} \left(\sum_{i=\max(\frac{n+\delta}{2}, n-r+t)}^{\min(n,t)} \binom{r-t}{n-i} \binom{t}{i} \right)^{-1} \quad (\text{A10})$$

such that n satisfies $\delta \leq n \leq \min(r, 2t - \delta)$ and $n|(p^2 - 1)$.

The argument here is similar to the proof for brute-force. We select and try codewords of size n . There are $\binom{r}{n}$ such codeword selections. The number of such codewords that succeed is a more difficult question to answer.

In general, our Reed-Solomon code requires $\frac{n+\delta}{2}$ elements to successfully produce the degree k polynomial interpolated by at least δ points. As a result, $\frac{n+\delta}{2} \geq t$ (or r if r is relatively small). Since $n \geq \delta$, the overall condition on n is derived: $\delta \leq n \leq \min(r, 2t - \delta)$.

How many sets of n points will succeed? Well, there must be at least $\nu = \frac{n+\delta}{2}$ real points and no more than $n - \nu$ chaff points. Given we have i real points where $\nu \leq i \leq n$, there are $\binom{r-t}{n-i}$ ways to choose the chaff points and $\binom{t}{i}$ ways to

select the real points. This results in the summation

$$\sum_{i=\max(\nu, n-r+t)}^{\min(n,t)} \binom{r-t}{n-i} \binom{t}{i}. \quad (\text{A11})$$

The additional constraints on i guarantee that we never select more chaff points or real points than we actually have. ■

Proof of Corollary 2: *The complexity of $\mathbf{bf}(r, t, \delta) = \mathbf{rs}(r, t, \delta, \delta)$, or taking $n = \delta$ reduces Reed-Solomon unlocking into a brute-force unlocking.*

Since $\delta \leq t \leq r$, the summation ranges can be determined: $\min(\delta, r) = \delta$, $\max(\delta, \delta - r + t) = \delta$. Since the summation only operates with $i = \delta$, the combinations become $\binom{r-t}{0} \binom{t}{\delta} = \binom{t}{\delta}$. Thus, the overall expression simplifies to $\binom{r}{\delta} \binom{t}{\delta}^{-1}$. ■

Proof of Corollary 3: *For $r \gg t$, $\mathbf{bf}(r, t, \delta) \leq \mathbf{rs}(r, t, n, \delta)$, for all $n \geq \delta$, or unless an attacker can eliminate a significant number of chaff points of a locked vault, he or she can do no better than a brute-force attack.*

We wish to select an n , such that $k \leq n \leq 2t - k$ that minimizes \mathcal{C}_{rs} . To do this, we shall expand the equation for \mathcal{C}_r , and examine the terms that most significantly contribute to the overall complexity.

First, for $r \gg t$, the terms can be slightly simplified.

$$\mathcal{C}_{\text{rs}} = \binom{r}{n} \left(\sum_{i=(n+\delta)/2}^{\min(n,t)} \binom{r-t}{n-i} \binom{t}{i} \right)^{-1} \quad (\text{A12})$$

Now, expanding the summation:

$$\binom{r}{n} \left(\binom{r-t}{(n-\delta)/2} \binom{t}{(n+\delta)/2} + \dots \right)^{-1} \quad (\text{A13})$$

However, notice that we really are only interested in the first term of the summation. For $r \gg t$, the first combination is significantly larger than the second. Additionally, the second term in the summation will be approximately r times smaller. As a result, if we write out the dominant terms:

$$\binom{r}{n} \binom{r-t}{(n-\delta)/2}^{-1} \quad (\text{A14})$$

To minimize our complexity, we wish to minimize the numerator and maximize the denominator. However, these are conflicting goals. Fortunately, the first term contributes approximately r^t times more to the result than the second term. Consequently, to minimize $\binom{r}{n}$, select the smallest n possible. ■