Authentication Using Pulse-Response Biometrics

Kasper B. Rasmussen
University of California, Irvine
Computer Science Dept.
kbrasmus@ics.uci.edu

Marc Roeschlin
University of California, Irvine
Computer Science Dept.
mroeschl@uci.edu

Ivan Martinovic
University of Oxford
Computer Science Dept.
ivan.martinovic@cs.ox.ac.uk

Gene Tsudik
University of California, Irvine
Computer Science Dept.
gts@ics.uci.edu

***Abstract*—We propose a new biometric based on the human body's response to an electric square pulse signal, that we call pulse-response. We explore how this biometric can be used to enhance security in the context of two example applications: (1) as an additional authentication mechanism in PIN entry systems, and (2) as a continuous authentication mechanism on a secure terminal. The pulse-response biometric is effective because each human body exhibits a unique response to a signal pulse applied at the palm of one hand, and measured at the palm of the other. Using a prototype setup, we show that users can be correctly identified, with high probability, in a matter of seconds. This identification mechanism integrates very well with other well-established methods and offers a reliable additional layer of additional security, either on a continuous basis or at login time. We build a proof-of-concept prototype and perform experiments to validate the feasability of using pulse-response as a biometric. Our results are very encuraging: we achieve accuracies of 100% over a static data set and 88% over a data set with samples taken over several weeks.**

I. Introduction

Many modern access control systems augment the traditional two-factor authentication procedure (something you know and something you have) with a third factor: "something you are", i.e., some form of biometric authentication. This additional layer of security comes in many flavors: from fingerprint readers on laptops used to facilitate easy login with a single finger swipe, to iris scanners used as auxiliary authentication for accessing secure facilities. In the latter case, the authorized user typically presents a smart card, then types in a PIN, and finally performs an iris (or fingerprint) scan.

In this paper, we propose a new biometric based on the human body's response to a square pulse signal. We consider two motivating sample scenarios:

The first is the traditional access control setting described above where the biometric is used as an additional layer of security when a user enters a PIN, e.g., into a bank ATM. The pulse-response biometric facilitates unification of the steps of PIN entry and biometric capture. We use PIN entry as a running example for this scenario throughout the paper. This is because PIN pads are often made of metal, which makes capturing pulse-response biometric straightforward: a user would place one hand on a metal pad adjacent to the key-pad, while using the other hand to enter a PIN. The metal pad would transmit the pulse and a sensor in the PIN pad would capture the biometric.

The second scenario corresponds to *continuous authentication*. One example is verifying that the user, who securely logged in earlier, is still the same person currently present at the keyboard. To address this problem, we need a mechanism that continuously monitors the user's biometrics. However, for obvious usability reasons, this must be done **unobtrusively**. The pulse-response biometric is particularly well-suited for this setting. Assuming that it can be made from – or coated by – some conductive material, the keyboard would generate the pulse signal and measure response, while the user (remaining oblivious) is typing. The main idea is that the user's pulse-response is captured at login time and identity of the person currently at the keyboard can be verified transparently, at desired frequency.

The continuous authentication problem is particularly difficult to solve using traditional biometrics. For example, if fingerprints are used instead of pulse-response, the user would have to interrupt work to periodically swipe a finger on a scanner, which would be very disruptive. There have been some attempts to solve this problem using a webcam and face recognition [12], [21], [23]. However, such systems can be fooled by a photo of the legitimate user and they also require the user to keep the head in a more-or-less constant position, unless a more advanced head tracking system is used.

To assess efficacy and feasibility of the pulse-response biometric, we built a platform that enables us to gather pulse-response data. Its main purpose is to verify that we can identify users from a population of test subjects. We also used it to test the distinguishing ability and stability of this biometric over time. We also explored two systems that apply the pulse-response biometric to the two sample scenarios discussed above: one to unobtrusively capture the biometric as an additional layer of security when entering a PIN, and the other – to implement continuous authentication.

The rest of the paper is organised as follows: Section II provides some background on biometrics and presents our design goals. Section III describes the pulse-response biometric in detail. Sections IV and V present the PIN entry and continuous authentication systems, respectively. Section VI describes the biometric data capture setup and Section VII presents experimental results. Related work is overviewed in Section VIII and the paper concludes with Section IX.

NDSS '14, 23-26 February 2014, San Diego, CA, USA

http://dx.doi.org/doi-info-to-be-provided-later

II. Background

This section provides some biometrics background and summarizes the terminology used throughout the paper. Then, design goals are presented.

A. Biometrics

The meaning of the term *biometric* varies depending on context. The US National Science & Technology Council's (NSTC) Subcommittee on Biometrics describes its two valid meanings: (1) a measurable biological (anatomical and physiological) and behavioral characteristic that can be used for automated recognition of individuals, (2) an automated method of recognizing an individual based on measurable biological (anatomical and physiological) and behavioral characteristics [6]. Throughout the rest of this paper we use *biometric* in the former sense, i.e., as a *characteristic* of a particular individual.

The US National Institute of Standards and Technology (NIST) divides biometric measurements into two categories [25], physiological and behavioural. The former relies on the physiology of a person and includes: fingerprints, hand geometry, facial recognition, speech analysis, and iris/retina scans. Behavioral biometrics is based on user behaviour and includes: keystroke timings, speech pattern analysis, gait recognition, and analysis of stylus pressure, acceleration and shape in hand-writing.

Physiological biometrics can help identify an individual from a large pool of candidates. However, there are some caveats. In general, physiological biometrics are considered moderately difficult to circumvent. For example, although hand geometry is very stable over the course of one's adult life, it does not provide enough distinguishing power to be used as the only means for identification [6]. Also, some facial recognition systems can be fooled by an appropriately-sized photo of a legitimate user. This is certainly a weakness if facial recognition is used to unlock a smartphone. On the other hand, the failure might not be due to the biometric itself but to inadequacy of current technology.

Behavioral biometrics measure user actions over time, meaning that, for each action, there must be a beginning, an end, and a duration. Consequently, behavioural biometrics indirectly measure characteristics of the human body. Behavioral biometrics are learned and, therefore, can be also re-learned. However the consensus in the literature seems to be, that after reaching a certain age, changes in behaviour become more difficult to achieve, even with specific and sustained effort [31]. Behavioural biometrics can therefore be regarded as valid biometric identifiers, even though they are neither as unique, nor as permanent as physiological biometrics. In most cases, behaviour biometrics are used to discern a user from a small(er) pool of candidates. One advantage is that they are less invasive and therefore more user-friendly. For example, a system that analyses keystroke timings or speech patterns can usually do so in the background. Whereas, an iris or fingerprint scan mandates specific user actions.

There is an ongoing debate about whether DNA constitutes a valid biometric. As a measurable part of the human physiology, it can very accurately identify an individual. In that sense, it is certainly a physiological biometric. However according to most definitions, e.g., [6], a biometric must be a characteristic that can be used for *automated* recognition of individuals. Thus, DNA's labeling as a biometric is questionable, at least for the time being. Albeit, one could imagine a future technology whereby DNA samples are continuously taken (and analyzed) from a user typing at a keyboard, e.g., by sampling body oils secreted by fingertips.

B. Biometric Authentication vs. Identification

Authentication refers to identify confirmation or verification. When a user claims a certain identity (e.g., by inserting a card into an ATM or entering a userid into a terminal, and then typing in a PIN or a password) authentication entails deciding whether the claim is correct. The goal of the biometric classifier is to compare the current sample to the known template for that user. The classifier returns the likelihood a match. We refer to this kind of comparison as 1:1.

Authentication differs from identification, where the current sample comes from an unknown user, and the job of the biometric classifier is to match it to a known sample. We call this a 1:n comparison. Identification is further divided into two types: open-set and closed-set.[1] We say that an identification is closed-set, if it is known *a priori* that the user is in the classifier's database, i.e., the classifier must choose the best match from a pool of candidates. Otherwise, we refer to it as open-set identification.

C. Design Goals

When designing a new biometric system it is important to take into account lessons learned from past and current systems. There are good discussions of design goals for biometric systems in the literature, e.g., [14]. Ours are as follows:

Universal: The biometric must be universally applicable, to the extent required by the application. For example, if a fingerprint reader is added as an additional level of access control, what to do about people that are missing all or some fingers? It is important for the biometric to apply to everyone intended to use the system.

Unique: The biometric must be unique within the target population. Measuring someones height would not work as an identification mechanism on a large scale. At the same time, (adult) height alone *can* usually identify individual family members.

Permanent: The biometric must be consistent over the time period where it's used. Very few biometrics will stay constant over a lifetime, e.g., face geometry, voice, gait and writing. However, as long as the biometric is consistent over the lifetime of the system, these biometrics work well.

Unobtrusive: A good biometric should be maximally unobtrusive. If the user can be identified passively, without interference, the biometric is much more likely to be accepted.

Difficult to circumvent: This is essential for a biometric in any security context. Ideally, a user should be unable to change the biometric at all. Moreover, it must certainly be the case that a user can not modify the biometric to match that of another user.

[1] See, for example, http://www.biometrics.gov/documents/biointro.pdf

An additional common design goal that we achieve "for free" is *Collectability*. It means that the biometric can be measured quantitatively. Since pulse-response biometric is based on measuring an electronic signal, no extra features are needed to achieve this goal.

There are a few other goals commonly found in the literature that we do not emphasize here:

Acceptability: The biometric is one that users are likely to feel comfortable with. It is hard to predict what users will or will not be comfortable with. Clearly, acceptability is a sensible design goal but it is one that we are not able to make significant claims about, so we have chosen not to present it as a requirement.

Cost Effectiveness: The relationship between the distinguishing power of the biometric and its deployment and maintainance costs. Since we focus on assessment of a new biometric and building a prototype, it is premature to seek insights about costs of a possible commercial system.

The same argument applies for *Performance*: the biometric should require minimal resources.

III. Pulse-Response Biometric

The pulse-response biometric works by applying a low voltage pulse signal to the palm of one hand and measuring the body's response in the palm of the other hand. The signal travels up through the user's arm, across the torso, and down the other arm. The biometric is captured by measuring the response in the user's hand. This response is then transformed to the frequency domain via Fast Fourier Transform (FFT). This transformation yields the individual frequency components (bins) of the response signal, which form raw data that is then fed to the classifier. Working in the frequency domain eliminates any need for aligning the pulses when they are measured. Details of our measurement setup and experiments can be found in Section VII.

The main reason for this biometric's ability to distinguish between users is due to subtle differences in body conductivity, at different frequencies, among different people. When a signal pulse is applied to one palm and measured in the other, the current has to travel through the body tissue – blood vessels, muscle, fat tissue, cartilage and bones – to reach the other hand. Differences in bone structure, muscle density, fat content and layout of blood vessels, result in slight differences in the attenuation of the signal at different frequencies. These differences show up as differences in the magnitude of the frequency bins after the FFT. This is what allows us to distinguish between individuals.

Pulse-response is a physiological biometric since it measures a person's physiological characteristics, rather than how that person behaves. However, it has an attractive property normally associated with behavioral biometrics: it can be captured in a completely passive way. Although some other biometrics also have this passive capture property, e.g., face recognition, pulse-response is not as easily circumventable. The combination of unobtrusiveness and difficulty to circument makes it a very attractive identification mechanism. Essentially it offers the best properties of both physiological and behavioral biometrics.

At the same time, pulse-response requires special-purpose hardware. The same is true any other physiological biometric. For example, fingerprints need a fingerprint reader, face recognition requires a precision camera and hand geometry – a scanner. Since pulse-response is captured using electrical signals, there are few restrictions on the exact construction of the biometrics capture hardware. We explore this issue in Sections IV and V.

A. Liveness and Replay

A common problem with many biometric systems is liveness detection, i.e., determining whether the biometric sample represents a "live" user or a replay. For example, a fingerprint reader would want to detect whether the purported user's fingerprint was produced by a real finger attached to a human, as opposed to a fingerprint mold made of putty or even a severed finger. Similarly, a face recognition system would need to make sure that it is not being fooled by a user's photo. More details and concrete examples are given in Section VIII).

In traditional biometric systems, liveness is usually addressed via some form of active authentication, e.g., a challenge-response mechanism. In a face recognition system a user might be asked to turn his head or look at a particular point during the authentication process. Although this reduces the chance of a photo passing for the real person, the user is forced to take active part in the process, which can be disruptive and annoying if authentication happens on a continuous basis. Also, a good 3-D model of a human head can still fool such measures.

Fingerprint scanners often include some protection against replay. This might be accomplished by detecting other characteristics normally associated with a live finger, e.g., temperature, or presence of sweat or skin oils. Such counter-measures make it more difficult to use skin-tight gloves or a "cold dead fingers" to fool the biometric system. Still, replay remains a major challenge, especially, for low-end fingerprint readers.

In the context of the pulse-response biometric, unlike fingerprints or face recognition, it is difficult (but not impossible) to separate the biometric from the individual to whom it belongs. If the adversary manages to capture a user's pulse-response on some compromised hardware, replaying it successfully would require specialized hardware that mimics the exact conductivity of the original user. We believe that this is feasible: the adversary can devise a contraption that consists of flat adhesive-covered electrodes attached to each finger-tip (five for each hand going into one terminal) with a single wire connecting the two terminals. The pulse response of the electrode-wire-electrode has to exactly replicate that of the target user. Having attached electrodes to each finger-tip, the adversary can type on the keyboard and the system could thus be effectively fooled. However, the effort required is significantly harder than in cases of facial recognition (where a photo suffices) or fingerprints, which are routinely left (and can be lifted from) numerous innocuous locations.

Furthermore, in contrast to face or fingerprint biometrics, pulse-response can be made to depend on the capture platform. Thus, even if the adversary captures this biometric on one piece of hardware, it would not mach the user's measurements on a different measurement (capturing) system. One way to achieve this is to add a specific (frequency-dependent) resistance to

the measurement platform, i.e., electrodes and/or wiring. If the adversary uses its own capture system to measure the user, there is an additional signature which is actually part of the pulse-response reader.

Finally, the real power of the pulse-response biometric is evident when used for continuous authentication (see Section V). Here, the person physically uses a secure terminal and constantly touches the keyboard as part of routine work. Authentication happens on a continuous basis and it is not feasible to use the terminal while at the same time providing false input signals to the authentication system. Of course, the adversary could use thick gloves, thereby escaping detection, but the authentication system will see input from the keyboard without the expected pulse-response measurement to accompany it, and will lock the session.

B. Ethics and User Safety

As mentioned above, the pulse-response biometric is captured by applying low voltage to one hand of the user and measuring the resulting signal in the other. This involves current flowing through significant portions of the human body. This process naturally raises questions about user safety and ethics. We believe that these are important issues that need to be addressed. The issue of safety might be compounded by users having undocumented or undisclosed medical conditions, including implantable medical devices, e.g., pacemakers, or other devices that may be adversely affected by applying an external signal to the body.

Two primary causes for concern are voltage and current levels that are applied to a user. An average healthy human being can easily withstand farily high voltage levels ($\geq 500/\ V$) provided that the current level is low. Strength of current sent through the body is of greater importance to human safety. Studies have shown that currents as low 1 mA can be perceived as slight tingling and currents as low as 5 mA are uncomfortable [28]. For this reason we chose to add a 10 kΩ output resistor to our signal generator, to act as a current limiting device. The 10 kΩ resistor insures that – even if the output terminals of the signal generator are shorted out – the maximum current strength will not exceed 0.1 mA per volt of input signal. Our initial experiments were done using three different voltage levels: 1V, 5V and 10V. The 5V and 10V levels where used on a small set of volunteers as a parameter search, in order to identify the minimal voltage level for the biometric to work. It quickly became apparent that very good results could be obtained using only a 1V signal.

The amount of current that a particular voltage induces in the human body varies from person to person and depends on external conditions. For example, if a subject's hands are wet, conductivity is significantly higher (i.e., resistance is significantly lower) than with dry hands. The same is true if the subject's hands have cuts or broken skin close to where the signal is applied. If resistance is lowered, current strength increases according to Ohm's law. Normal resistance of the human body is between $1,000$ and $5,000\ \Omega$. However, even in extreme conditions, resistance does not drop bellow $500\ \Omega$. With our current limiting resistor on the signal generator, the worst case current (with 10V test signal) is $10V/10.5k\Omega = 0.95\ mA$, which is bellow the sensitivity limit. The vast majority of subjects were only exposed to a 1V signal, which translates into the worst case current strength of $0.095\ mA$.

All subjects were given detailed information about the nature of the experiment beforehand and all were given the opportunity to opt out. None expressed any discomfort or, in fact, any perception of the current during the experiments.

We note that many commercial systems and products involve applying a similar (or higher) voltage to humans. For example, so-called "touch lamps" (popular since 1970-s) turn on and off whenever the user touches the metal frame. The lamp's touch detection mechanism works by having the user close an electric circuit between the lamp and the ground, i.e., the current takes a path similar to that in our pulse-response capture system. The magnitude of the signal (1-6V) used in touch lamps is similar to our case. Such lamps are commercially available and are known to be safe for people, even those who have a heart condition or implanted medical devices. Another related household product example is a wall-mounted touch-based light switch (although some of those are capacitive touch sensors).

Another point of comparison is a regular 9V battery. The internal resistance of a 9V battery varies depending on the type (e.g., zinc carbon, lithium, alkaline) between 1 Ω and 20 Ω. Consequently, the safest of these (with the highest internal resistance) can deliver $9V/20\Omega = 450mA$ current if the terminals are shorted out. This is a much stronger current than in our setup, even in the worst case. Meanwhile, the terminals of a typical 9V battery are not protected. The reason is that such voltage and current levels are considered to be completely safe for humans.

IV. Combining PIN Entry with Biometric Capture

This section describes how to use pulse-response to enhance security of PIN entry systems without inconveniencing the user.

A. System and Adversary Models of PIN Entry Scheme

We use a running example of a metal PIN key-pad with an adjacent metal pad for the user's other hand. The PIN key-pad has the usual digit (0-9) buttons as well as an "enter" button. It also has an embedded sensor that captures the pulse-signal transmitted by the adjacent metal pad. We can envisage this setup in the setting of a bank ATM allowing authorized users to withdraw cash.

The goal of the adversary is to impersonate an authorized user and withdraw cash. We assume that the adversary can not fool the pulse-response classifier with probability higher than that found in our experiments described in Section VII.

We assume that the ATM is equipped with a modified authentication module which, besides verifying the PIN, captures the pulse-response biometric and determines the likelihood of the measured response corresponding to the user identified by the previously inserted ATM card and the entered PIN. This module works as depicted in Figure 1. We assume that the ATM has access to a database of valid users, either locally or over a network. Alternatively, the user's ATM card can contain data needed to perform pulse-response verification. If stored on the card, this data must be encrypted and authenticated using a key known to the ATM; otherwise, the adversary (who can be assumed to be in possession of the card) could replace it with data matching its own pulse-response.

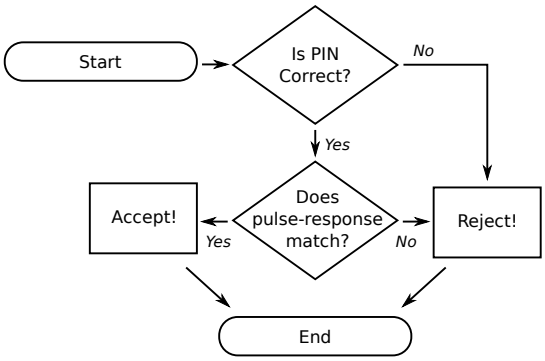

Fig. 1: ATM decision flowchart.

B. PIN Entry Scheme

The ATM has to determine whether data sampled from the user while entering the PIN, is consistent with that stored in the database. This requires the use of a classifier that yields the likelihood of a sample coming from a known distribution. The likelyhood is used to determine whether the newly measured samples are close enough to the samples in the database to produce a match. Using our prototype, we can make such decisions with high confidence. (See Section VII-D.)

Before we look at the security of the pulse-response PIN entry system, we need to make sure that it meets our design goals.

Universal. A person using the present PIN entry system must use both hands, one placed on the metal pad and one to enter the pin. This requires the user two actually have two hands. Whereas, a normal PIN entry system can be operated with one hand; thus, universality of our system is somewhat lower. This is a limitation of the biometric, although a remedy could be to store a flag on the user's ATM card indicating that a disability, thus exempting this person from the pulse-response check. This would allows our approach to gracefully degrade to a generic PIN entry system.

Unique and *Permanent*. In Section VII-D we show that our prototype can determine, with high probability, whether a subject matches a specific pulse-response. Thus, it is extremely unlikely for two people to exhibit exactly the same pulse-response. We also show that an individual's pulse-response remains fairly consistent over time.

Unobtrusive. The proposed scheme is very unobtrusive. From the user's perspective, the only thing that changes from current operation is the added requirement to place the free (not used for PIN entry) hand on a metal pad. There can even be two such pads accommodating both left- and right-handed people. Also, the ATM screen could display system usage instructions, even pictorially to accommodate people who can not read. Similarly, audio instructions could be given for the sake of those who are vision-impaired.

Difficult to circumvent. Given that pulse-response is unique, the only other way to circumvent it is to provide the sensor (built into the PIN pad) with a signal that would correspond to the legitimate user. Although this is very hard to test precisely, assuming that the adversary is unaware of the target user's pulse-response measurements, the task seems very difficult, if not impossible.

C. Security Analysis of PIN Entry Scheme

The additional layer of security provided by the pulse-response biometric is completely independent from security of the PIN entry system alone. For this reason, we model the probability P_{break} that the proposed PIN entry system can be subverted, as follows:

$$P_{break} = P_{guess} \cdot P_{forge}$$

where P_{guess} is the probability of the adversary correctly guessing the PIN and P_{forge} is the average probability that the adversary can fool the classifier. We model this as the false positive rate divided by the number of users. If a PIN consists of n decimal digits and the adversary has t guesses then $P_{guess} = \frac{t}{10^n}$. The false positive rate is the complement of specificity [30] In Section VII-D, we determine specificity to be 88%. Thus $P_{forge} = (1 - 0.88)/5$, which yields the combined probability:

$$P_{break} = \frac{(1 - 0.88)t}{5 \cdot 10^n}$$

For example, if the adversary is allowed 3 guesses with a 4-digit pin, $P_{break} = 7.2{\cdot}10^{-6}$, whereas a 4-digit plain-PIN system has a subversion probability of $3 \cdot 10^{-4}$. Though this improvement might not look very impressive on its own, it is well known that most PIN attacks are performed by "shoulder surfing" and do not involve the adversary guessing the PIN. If we assume that the adversary already knows the PIN, $P_{break} = 2.4\%$ with our system, as opposed to 100% without it.

V. Continuous Authentication

We now present a continuous authentication scheme. Its goal is to verify that the same user who initially (and securely) logged into a secure terminal, continues to be physically present at the keyboard. Here, the pulse response biometric is no longer used as an additional layer of security at login time. Rather, the user's pulse-response biometric is captured at login time and subsequent measurements are used to authenticate the user using the initial reference.

A. System and Adversary Models for Continuous Authentication

We continue using the example for continuous authentication introduced in Section I. This example entails a secure terminal where authorized users can login and access sensitive data. We use this example throughout this section to make it easier to present the details of our system. However, applicability of continuous authentication via pulse response is not limited to this specific scenario.

The system consists of a terminal with a special keyboard that can send out pulse signals and capture the pulse-response biometric. This requires that the keyboard must be either made from, or coated by, a conductive material. Alternatively, the pulse signal transmitter could be located in a mouse that the user operates with one hand and the keyboard could then contain the mechanism that captures the pulse-response. Without loss of generality, we will assume that the keyboard contains both the pulse transmitter and the receiver. Otherwise, the keyboard

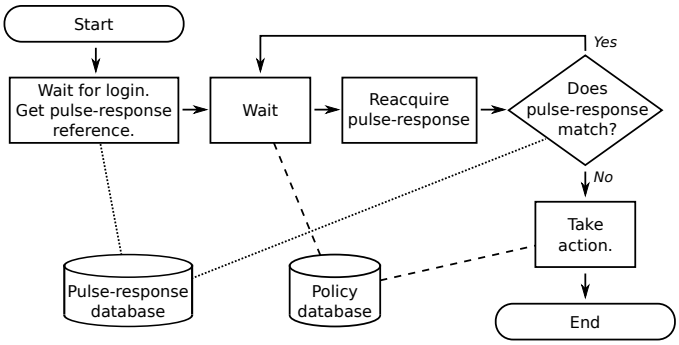

Fig. 2: Flowchart of the *Continuous Authentication Process* decision procedure.

operates normally and is used for both login and routine activity at the terminal.

The adversary is another person who, with or without consent of the authorized (at login time) user, physically sits down at the unattended terminal and attempts to access resources within the already-open session. We assume that the adversary at the keyboard has complete access to the active session, and that this happens some time after the original user logged in. The goal of our system is to detect that the original user is no longer present, and that the keyboard is being operated by someone else. If a different user is detected, the system consults a policy database and takes appropriate actions, e.g., locks the session, logs out the original user. raises alarms, or notifies administrators.

In addition to the peripherals required to capture the pulse-response signal, the continuous authentication system consists of a software process that manages initial login and frequency of reacquisition for the pulse-response biometric. This process is also responsible for displaying warnings to the user and notifying administrators in case of a violation. We refer to it as the *continuous authentication process* (CAP) and assume that neither the legitimate user nor the adversary can disable it.

B. Continuous Authentication Scheme

At login, while the user is entering a password, CAP captures the user's pulse-response biometric and stores it locally. Periodically, e.g., every few seconds, CAP reacquires a pulse-response from the user by sending and receiving a pulse signal through the keyboard. The newly acquired measurement is checked against the value acquired at login. If the likelihood that the new measurement is sampled from the original user is too low, CAP consults its policy database and takes appropriate actions, as discussed above. Figure 2 shows the CAP decision flowchart. The decision policy can be further refined. For example, in a corporate setting, all employees could have their pulse-response biometrics stored in central database. In this scenario, CAP could make a distinction as to whether the new (detected) user is a genuine employee who is authorized to use the terminal.

The envisaged continuous authentication system is also useful for training (e.g., corporate) users to adopt security-conscious behaviour. For example, users can be trained to log out when they leave a terminal, either by seeing a warning every time they forget, or by having a centralized system whereby the employee gets a reprimand if she either forgets to logout, or allows someone else to take over her session. Another positive side-effect is that, in order for anyone to use another persons credentials, that person will have to actually give out their username and password, rather than just logging in and leaving the session. We suspect that most users are much more reluctant to give away their login credentials, a opposed to just starting a session for someone else.

Before considering security of the continuous authentication system, we assess it with respect to the design goals.

Universal. The users of the system must have two hands in order for the pulse-response biometric to be captured. The same arguments, as in the case of PIN entry, apply here.

Unique and *Permanent*. In Section VII-D, we show that our prototype can match a pulse-response to previous samples (taken immediately beforehand) with 100% accuracy. The fact that the pulse-response reference is taken at the beginning of the session and is used only during that session, makes it easier to overcome consistency issues that can occur when the reference and test samples are days or months apart.

Unobtrusive. Users do not need to modify their behaviour at all when using the continuous authentication system. Thus, user burden is minimal.

Difficult to Circumvent. With a true positive rate of 100% it is unlikely that the adversary can manage to continuously fool the classifier. Even if the adversary happens to have a pulse-response biometric similar to the original user, it must evade the classifier on a continuous basis. We explore this further in the security analysis section below.

C. Security Analysis of Continuous Authentication Scheme

The adversary can subvert the continuous authentication system by managing to use the secure terminal after another user has logged in and (possibly) left. In the analysis below, we assume that the initial user and the adversary are collaborating. This eliminates any uncertainty that results from the original user "discovering" that the adversary is using its terminal, which is very hard to model accurately. The result of our analysis is therefore a worst-case scenario and the detection probability is a lower bound on security provided by the continuous authentication system.

One parameter in our security analysis is the number of times biometric acquisition is performed since the time when the adversary initially appeared at the keyboard. The longer the period between each acquisition, the longer it takes for the system to measure the adversary a fixed number of times, and therefore (potentially) longer to detect adversary's presence. Policy plays an important role in the practical security of the system. For example, suppose that the policy is to just display a warning whenever a mismatch in pulse-response is detected. Such a system will offer little, if any, security against a determined adversary. Therefore, for the purpose of security analysis, we consider the attack thwarted as soon as the continuous authentication process detects a problem.

We assume that the adversary cannot evade our classifier with a probability higher than that in Section VII-D.

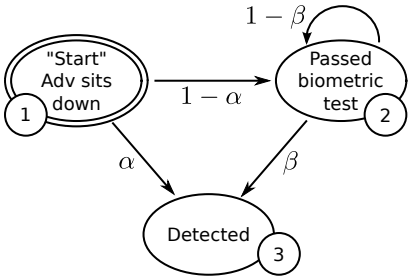

Fig. 3: Markov model of the continuous authentication detection probability. States are numbered 1 to 3 for easy reference in text.

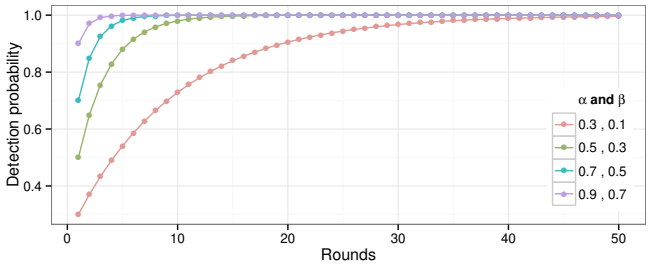

Fig. 4: Detection probability of our continuous authentication scheme as a function of the number of biometric acquisitions performed (rounds), for selected values of α and β.

We model the continuous authentication scenario using two probabilities. The first is the probability that the adversary is detected immediately, i.e., the first time its pulse-response biometric is captured. This corresponds to *sensitivity*, i.e., true positive rate reported in Section VII. We use 99% (rather than the 100% found in our experiments) in order to model the posibility of making a clasification mistake at this point. On average, according to our experiments, the biometric of the adversary differs enough from the original user to be detected easily. We refer to this probability as α.

If the adversary's biometric is very close to that of the original user, it might not be detected every time biometric capture is performed. If the adversary manages to fool the classifier once, it must be because its biometric is very close to that of the original user. Given that the user and the adversary have a similar pulse-response the adversary's subsequent detection probability must be lower, i.e.,

$$P[X_i = adv | X_{i-1} = usr] \leq P[X_i = adv]$$

We call this decreased probability β. We build a Markov model (illustrated in Figure 3) to calculate the probability that an adversary is detected after i rounds. The model uses α and β. When the adversary first accesses the keyboard, it is either detected with probability α or *not* detected, with probability $1 - \alpha$. If the adversary is not detected, its pulse-response biometric must be close the original user's. Thus, β is used for the subsequent rounds. In each later round, the adversary is either detected with probability β or *not* detected, with probability $1 - \beta$. To find the combined probability of detection after i rounds, we construct the state transition matrix P of the Markov model as follows:

$$P = \begin{bmatrix} 0 & 1-\alpha & \alpha \\ 0 & 1-\beta & \beta \\ 0 & 0 & 1 \end{bmatrix}$$

In matrix P each row and each column corresponds to a state. The number in row q and column r, p_{qr}, is the probability of transitioning from state q to state r. To find the probabilities of being in each state we start with a row vector ρ that represents the initial probability of being in state 1, 2 and 3. In this case, $\rho = [1, 0, 0]$, indicating that we always start in state 1. The probability of being in each state after one round (or one transition) can be represented by the inner product ρP. The probabilities for each subsequent round are found by another multiplication with P. Therefore, the probabilities of being in each state after i rounds (state transitions), is found as follows:

$$[1, 0, 0] \cdot P^i = [0, (1-\alpha)(1-\beta)^{i-1}, 1 - (1-\alpha)(1-\beta)^{i-1}]$$

As expected, the probability of being in state 1 (the initial state) is 0, because the first state transition forces a transition from the initial state and there is no way back. (See Figure 3.) The probability of being in state 2, i.e., having escaped detection for i rounds is given by the second element of ρ: $(1-\alpha)(1-\beta)^{i-1}$. The probability of being detected is thus: $1-(1-\alpha)(1-\beta)^{i-1}$. Using the numbers from our experiments (see Section VII-D) $\alpha = .99$ and $\beta = .3$, we the detection probability after 10 rounds is:

$$\begin{aligned} 1 - (1-\alpha)(1-\beta)^{i-1} &= 1 - (1-0.99)(1-0.3)^{10-1} \\ &= 1 - 0.01 \cdot 0.7^9 = 0.99959 \approx 99.96\% \end{aligned}$$

There is a 99.96% chance of detecting the adversary after 10 rounds. This grows to 99.99999997% after 50 rounds. Thus, the frequency of biometric acquisition clearly determines the time to detect the adversary. Figure 4 shows the detection probability as a function of the number of rounds for various values of α and β.

D. Handling False Negatives

False negatives refer to incorrectly detecting the presence of an adversary, i.e., when the original user is still at the terminal. In a scenario where biometric identification is used as an additional layer of security during the authentication procedure, this problem can be managed simply by restarting the login procedure, if the first attempt fails. In a continuous authentication system where a single detection event might cause the system to lock up, false negatives have to be dealt with in a more organized manner.

One way of dealing with false negatives in a continuous authentication system, is to specify a policy that allows a certain number of adversary detection events every n-th round, without taking any action. For example, allowing one adversary detection event every 100 rounds corresponds to a false negative rate of 1%.

Another option is to combine the continuous authentication mechanism with a less user-friendly biometric to deal with ambiguous detection events. For example, after a few adversary

detection events, the user is asked to confirm its identity by swiping a thumb on an adjacent fingerprint scanner. Without pulse-response, the user would have to do that every ten seconds or so, which would render the system quite unusable. However, combined with our continuous authentication system, such confirmation might need to ocur much less frequently.

Finally it is possible to gradually ramp up the severity of actions taken by the continuous authentication process, every time an adversary detection event occurs. For the first time, displaying a warning might be the most appropriate action. If detection re-occurs, more and more severe actions can be taken. It is very unlikely, with a reasonably low false negative rate, to have multiple consecutive adversary detection events if the original user is still at the terminal. Although the false positive rates we achieve are quite low, they cold certainly be improved with a more advanced biometrics capture system. In conjunction with a sensible policy, our continuous authentication system might be appropriate for any organization with high security requirements.

VI. Biometric Acquisition System Design

In this section, we describe decisions and parameters that went into the design of our final classifier. We conducted several experiments during to test different signal types, voltage levels, and frequencies. To support choices made in Section VII, we present some of those results below.

A. Signal Type

We start out with the hypothesis that the biometric signature will vary, depending on the frequency of the signal transmitted through the body. If this is in fact true it makes sense to test the performance of various frequency sweeps. Our initial test signals are three different linear 0.6-second sine-wave sweeps from 1 Hz to 250, 500 and 980 Hz. We also test the performance of square-wave sweeps from 1 Hz to 250, 500 and 980 Hz, respectively. For a few specific values of voltage and frequency we get decent results using Linear Discriminant Analysis (LDA), but at this point our results are not very robust. We continue to experiment with different signal types and it turns out that, contrary to our initial assumption, single pulse signals have significantly higher distinguishing power. We experiment with different pulse widths between 100 ns and 1 ms, and voltage levels of 1, 5 and 10 volts.

The box plots in Figure 5 summarize our results. We present the results from the four classifiers that performed the best in our application: Support vector machines (SVM), Euclidean distance, linear discriminate analysis (LDA) and 3-nearest neighbors (3nn). On the x-axis are the most promising of the signals we tested. The signal name is composed of a signal type, a voltage and a maximum frequency (or width for pulses). The signal types are: single pulses (Pulse), a linear sine sweep (SineLin) and a linear square wave sweep (SquareLin). The voltage is either 1, 5 or 10 volts, and the frequency is 250, 500 and 980 Hz. The frequency information for the pulse signals indicate the width of the pulse (in hundreds of nanoseconds) rather than maximum frequency. The y-axis is the binary detection error rate, i.e., the amount of times the classifier failed to classify a sample correctly, normalized by the number of samples and converted to value in percent. The distribution denoted by the box plots themselves are the results of the classifiers achieved by five times 5-fold cross-validation. We show the box plots rather than just the mean to clearly show the variance in performance for each classifier.

We see that the narrow pulse signal outperforms every other signal type by a remarkable margin. We get consistent error rates close to zero for a pulse signal of 1 volt and a width of 100 nanoseconds. Wider pulse signals also give decent results but the quality of the result seems to decrease with the width of the pulse. For the sine and square wave sweeps the results vary significantly with the choice of classifier. Using LDA, some sine sweeps look interesting but nowhere near as good as the narrow pulse signal.

B. Signal Voltage

There are several factors besides the distinguishing power of the resulting biometric, to consider when choosing voltage levels. It is very important that the users of our system do not experience any discomfort when their biometric information is captured. That requires the voltages to be reasonably low. We test three different voltage levels for all signal types: 1, 5 and 10 volts peek-to-peek (Vpp).

For sine and square signal sweeps the 10 Vpp and 5 Vpp provides better separation between the subjects but also higher noise levels. For example, in Figure 5, using the LDA classifier, we see that the *SineLin-5-500* signal has a lower detection error rate (i.e., better performance) than the *SineLin-1-500* signal, but the latter has less variance. For pulse signals there is no significant correlation with voltage level. Since the pulse signal is clearly the best choice for our final classifier we chose 1 volt pulses to minimize any potential discomfort that users of our biometric system might feel.

C. Signal Frequency

We initially thought that (almost) all frequencies would contribute to the distinguishing power of our classifier but our experiments show that the classifier mainly uses the lower frequencies to distinguish between users. In fact, we see an increase in the true positive rate when we only use the first 100 frequency bins of the FFT. This suggests that most of the high frequency content is noise when operating at such low power levels.

D. Choice of Classifier

Although we apply an FFT to the data before the classification step we can think of our task as time series classification. This is because an FFT is a reversible linear transformation so the euclidean distance metric is preserved. Thinking of the problem as a time series clustering problem, there are many known approaches that work well. One common method is to compare the first n frequency components by using appropriate distance- or similarity metric. We compare several different classification techniques to see which ones provide the best results for our application.

Euclidean Distance (Euclidean) A new measurement is treated as an n dimentional point and classified according to the euclidean distance to the centroid of each class. This

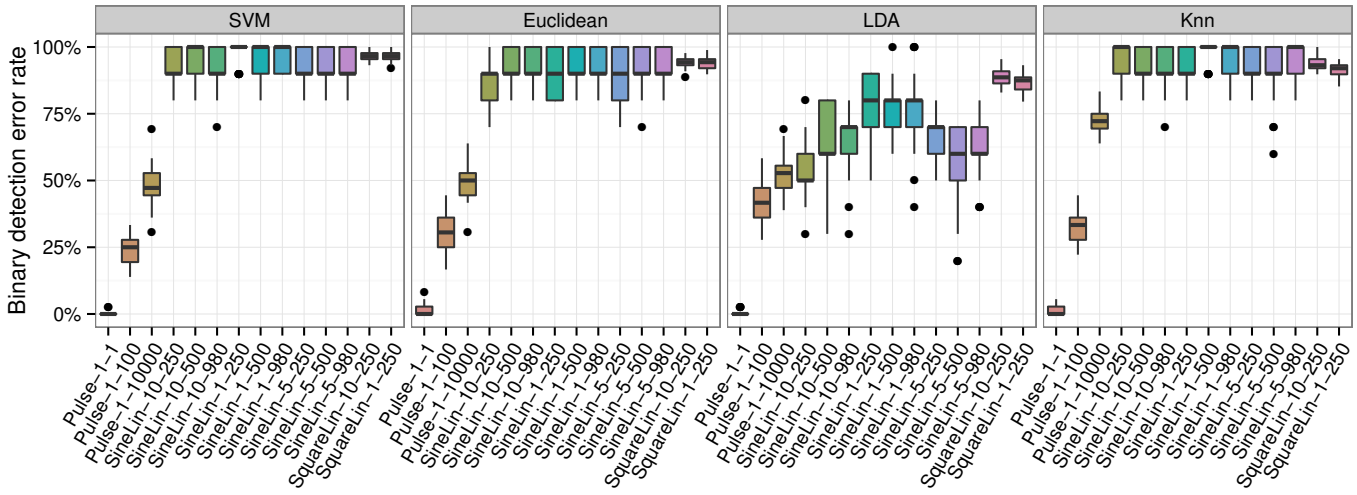

Fig. 5: Box plots of the binary detection error rate for four different classifiers. The distribution shown by each box plot is the result of applying stratified 5-fold cross-validation to the data set five times in a row. We test several different signal types, voltage levels and frequencies for each classifier. We see that narrow pulse signals are consistently performing well.

classifier is conceptually very simple but still offers reasonably good results.

Mahalanobis Distance (MH) Rather than assuming uniform and orthogonal dispersion among the frequency components (as in the Euclidean classifier) the covariance matrix for each class is taken into account in the distance calculation. This allows for a distance metric that is proportional to the shape of the class (in n dimensional feature space). The performance of this classifier did not differ significantly from the Euclidean, suggesting that the shape of each class is not significantly skewed.

Support Vector Machine (SVM) For each pair of groups we train one binary classifier (one-against-one). The final prediction is found by voting. The inverse kernel width for the Radial Basis kernel is determined by the 0.1 and 0.9 quantile of the pairwise Euclidean distance between the samples. This classifier gives consistently good results and is our final choice of classifier.

Linear Discriminant Analysis (LDA) LDA seeks to reduce the dimensionality of the input data while preserving as much of the class distinguishing power as possible. Our LDA classifier performs the linear analysis on all the classes in our database, then compares the position of new samples in the resulting lower dimension feature space. The overall performance of this classifier degrades more gracefully than many of the other methods but ultimately it did not prove as powerful as the SVM method.

K Nearest Neighbor (Knn) We tested the k nearest neighbors classifier for $k = 1$ and $k = 3$, using euclidean distance. It is a simple classifier that often works very well in practice. In our case though the performance of Knn was still not as good as SVMs.

VII. Experiments

In this section we will describe our experimental setup and present the results of our experiments with our final classifier.

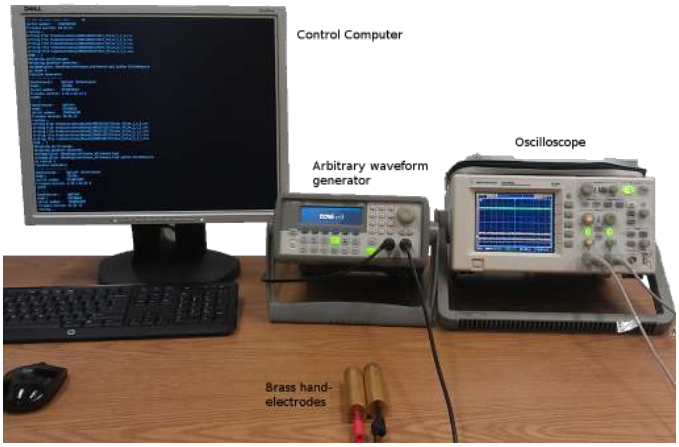

Fig. 6: Our proof-of-concept measurement setup. The test subject is holding the two brass hand electrodes [20] and the pulse signal is generated by an Agilent 33220A (20 MHz) arbitrary waveform generator. The receiver is an Agilent DSO3062A (60 MHz), 1 GSa/s digital storage oscilloscope.

The design decisions and motivations behind our final classifier are described in detail in Section VI.

Any names from test persons appearing in this section's figures have been anonymized through pseudonyms.

A. Proof-Of-Concept Measurement Setup

In order to be able to gather stable and accurate pulse-response data we build a data acquisition platform consisting of an arbitrary waveform generator, an oscilloscope, a pair of brass hand electrodes and a desktop computer to control the apparatus. Figure 6 is a photo of our setup. We use an Agilent arbitrary waveform generator as the source of the pulse signal. The flexibility of the waveform generator is useful during the

initial design phase and allows us to generate the required pulse waveforms in our final classifier. To measure the pulse waveform after the signal passes through a test subject we used an Agilent digital storage oscilloscope which enabled us to store the waveform data for later analysis. The output of the waveform generator is connected to a brass handle that the user holds in the left hand. The other brass handle is connected to the oscilloscope's signal input. When a test subject holds one electrode in each hand the signal travels from the generator through the test subject and into the oscilloscope. To ensure exact triggering, the oscilloscope is connected to the synchronization output of the waveform generator.

We use polished brass hand electrodes to ensure optimal electrical contact between the measurement setup and the user. This reduces contact resistance and increases the stability of the measurements.

The function generator and oscilloscope are controlled by a desktop computer that is connected via USB. We wrote a custom software library to set measurement parameters and retrieve the desired waveform data. This software is available upon request.

B. Biometric Capture Procedure

We had each subject follow a specific procedure during the biometric capturing process. This ensures that only minimal noise is introduced by the process itself. The test subjects are given a brief explanation of the setup and purpose of the experiment and then told to grab a hold of the brass hand electrodes. The red lead in the left hand and the black in the right hand.

The test subjects could choose to either stand or sit in a chair during the experiment as long as they did not touch the sides of their body with their elbows or upper arms. We did this to ensure that the current of the pulse signal had to go through more or less the same path, for all samples and all users. Before each new test subject was measured, the brass handles where wiped down with a disinfectant, both for hygienic reasons and to ensure good electrical contact between the electrode and the user's palms.

The capture process itself lasts about eleven minutes and each subject was given the opportunity to take a break three times during that period. In the initial design phase each test subject was sampled ten times for each of the three signal types, for each voltage level and for various frequencies. Once a decision had been made that the pulse signal gave us the best results, we acquired samples for two different data sets. The first one consists of 22 samples from each test person, taken in one measuring session, i.e. at one point in time. The second one encompasses a total of 25 samples per test person, obtained in five different sessions over time.

Our subject population consists of both men and woman between the ages of 24 and 38. We sampled all our test subjects at different times during the day, over the course of several weeks. We tried to sample subjects in such a way that we would end up with sampling conditions as diverse as possible, for each user. The interval between measurements sessions with the same user varied between a few hours up to several weeks. This was done in order to try to eliminate any effect that sampling at a specific time of day might have on our results, i.e., that our biometric would remain more or less permanent over time, and across different periods of the day.

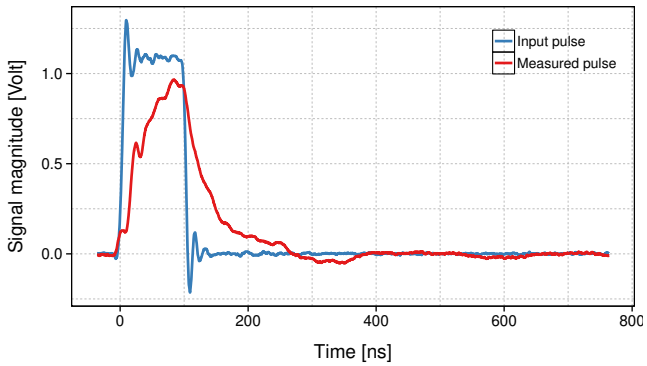

Fig. 7: The input and output waveform. One measurement consists of 4000 samples with a sample rate of 500 MSa/s. It is clear that the measured pulse has been modified by passing through the user.

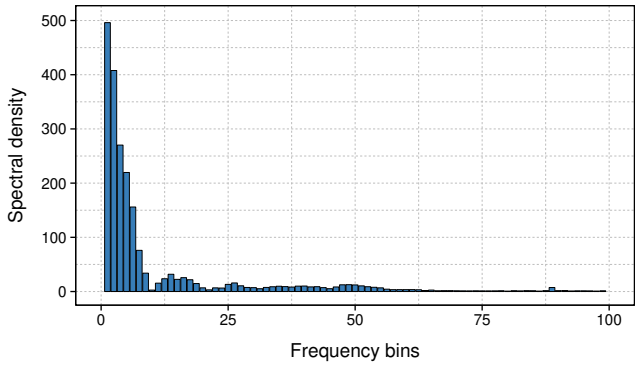

Fig. 8: The raw FFT data of the measured pulse, extracted from our measurement setup. The data consists of the first 100 frequency bins of the measured waveform.

C. Feature Extraction

The data we extract from our measurement setup is in the form of a 4000 sample time-series describing the voltage variation as seen by the oscilloscope. Figure 7 shows the input pulse sent by the waveform generator and the pulse measured by the oscilloscope.

The time series measurements are converted to the frequency domain using FFT and the first 100 frequency bins of the FFT data is used for classification. Operating in the frequency domain has several advantages. First we do not have to worry about the alignment of the measured data pulses when computing metrics like euclidean distance between pulses. Second, it quickly became apparent that only the lower frequency bins carry any distinguishing power. The higher frequency bins were mainly noise. This means that we can use the FFT to do a dimensionality reduction of the original 4000 sample time-series to vector of 100 FFT bins. Figure 8 shows an example of the raw data we end up with after the FFT. This data is then fed into the classifier.

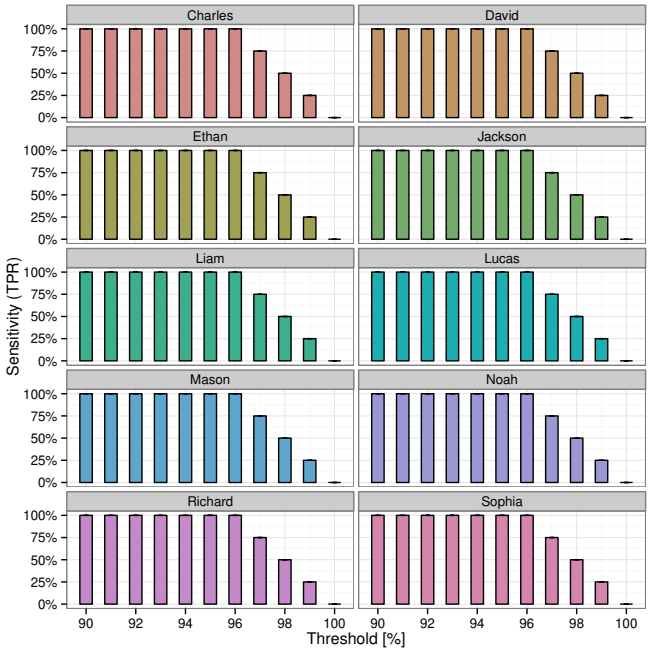

Fig. 9: The results for our authentication classifier based on the single data set. We obtained the true positive rate by performing five times 5-fold cross-validation for each test subject. The x-axis describes the discrimination threshold for assigning the classifier's prediction output to a *positive* or a *negative*.

D. Results

We present two different classifiers, one for authentication and one for identification. The authentication classifier is based on support vector machines (SVM) and solves the problem of verifying a 1:1 match between a sample form an unknown person and the requested person's stored biometric template. The identification classifier, also based on SVM, verifies a 1:n match between a sample from a known person against all the samples in a database. Our identification classifier is a closed-set classifier. Refer to Section II for a more detailed description of open- and closed-set classifiers.

We further divide our findings into results on a single test-set, which shows the inherent distinguishing power of our pulse-response biometric, and results of our classifier when applied to data sampled over time. The samples taken over time show the stability (permanence) of our biometric over a longer time period.

1) Authentication Classifier: Our authentication classifier is a 1:1 classifier based on SVM. The results of running this classifier on our single-session data set can be seen in Figure 9. Each bar is the classifiers performance for different threshold levels, for each of the ten test subjects. The threshold is a measure of how sure you want to be that the identification is correct. If you can accept a small false positive rate a better sensitivity can be achieved. The classifier performance is the result of 5-fold cross validation to ensure statistical robustness. We see that all subjects are being recognized with a very high probability as the true positive rate confirms.

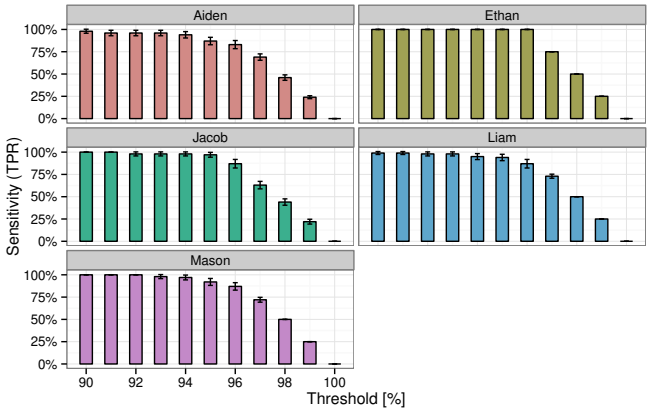

Fig. 10: True positive rate for each test subject when our authentication classifier is fed with the data sampled over time. The error bars show the 95% confidence interval. As in Figure 9, the x-axis depicts different discrimination thresholds.

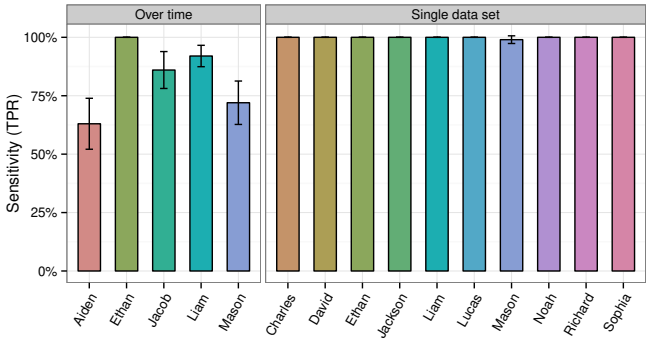

Fig. 11: The results of our identification classifier. The true positive rate for each test subject is obtained by applying five times stratified 5-fold cross-validation. The error bars show the 95% confidence interval.

Figure 10 shows the stability performance of our authentication classifier. In this figure the classifier has been applied to a data set collected over several weeks. Here we see a similar picture. If we can accept, e.g., 10% false positives, we achieve a sensitivity of almost 100%.

2) Identification Classifier: Identification is a multi-class classification problem, our classifier consists of multiple SVMs and follows a one-against-one approach (aggregation by voting). Due to this increased complexity we expect a slight drop in performance in comparison to authentication, which is a binary classification task.

When we run our identification classifier on the two different data sets we get the results shown in Figure 11. Even with the increased complexity we see that our identification classifier still performs very well on both data sets. In the single data set we have ten people and the goal of the classifier is to identify each person as accurately as possible. In other words, distinguish each person from everybody else. We see a slight decrease in performance in the data set containing samples being taken

					in [%]		
	TP	FP	TN	FN	Sensitivity	Specificity	Accuracy
Authentication							
– Single set	2.0	0.0	18.0	0.0	100	100	100
– Over time	4.4	2.4	17.6	0.6	88	88	88
Identification							
– Single set	2.0	0.0	18.0	0.0	100	100	100
– Over time	3.4	1.6	18.4	1.6	68	92	87.2

TABLE I: Summery of our results of the authentication and identification classifiers, averaged over all users. This performance figures have been assessed on the basis of test data not involved in any development or training phase of the classifiers. Values for true/false positives/negatives are at the equal error rate of $EER = 0.00$ on the single data set and $EER = 1.12$ over time. For a more detailed view on the performance of the classifiers see the ROC curves in Figure 12.

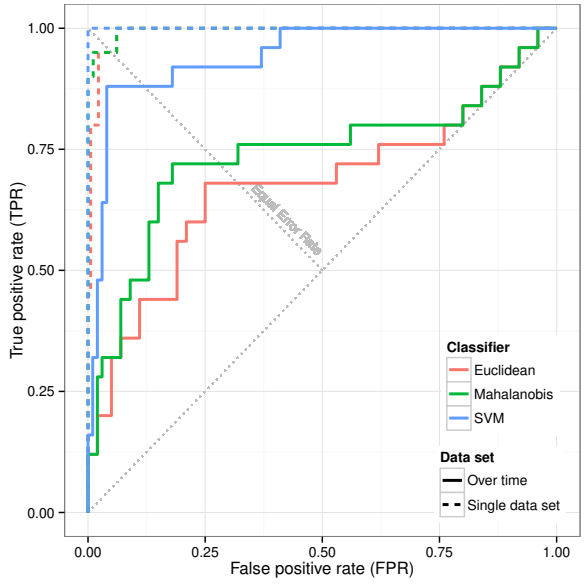

Fig. 12: ROC curves for our authentication classifier based on unseen test data. We show results for three different classification methods. The dashed lines are for the results on the single data set, the solid lines are over time. For the exact number corresponding to the Equal Error Rate see Table I.

several weeks apart. The reason for the performance decrease over time is that samples taken far apart are influenced by very different conditions. There might be physiological changes such as a loss or gain of weight, or there might be differences in the ambient temperature, clothing, or a number of other factors.

Table I summarizes our results for the two classifiers, authentication and identification, on both, the single data set and the data set taken over time.

Both classifiers can be tuned by selecting a specific false positive rate, acceptable for the scenario in which the classifier is being used. For example if the classifiers are being used in a continuous authentication application and false negatives are a problem, the classifiers can be tuned for a lower false negative rate, by accepting a higher false positive rate. Figure 12 shows the relationship between false positives and true positives. The classifiers can operate on any point on this curve, if desired.

VIII. Related Work

Biometrics, as a means of recognizing an individual using physiological or behavioural traits, has been an active research area for many years. A comprehensive survey of conventional physiological biometrics can be found in [13]). While physiological biometrics tend to be relatively stable over time, they are sensitive to deception attacks. These include attacks on: (1) fingerprint identification, e.g., using mock fingers made of glycerine, gelatine or silicon [1], [2], (2) facial recognition, e.g., using photographs or 3D models of an actual user [3], [22], and (3) iris scan, e.g., using patterned contact lenses that replicate a genuine user's iris [8].

In contrast, behavioural biometrics are much harder to circumvent. However, performance of systems that implement behavioral biometrics, in terms of false rejection rates (FRR) and false acceptance rates (FAR), is much lower and can require re-calibration due to varying and often erratic nature of human behaviour. Initial results on behavioral biometrics were focused on typing and mouse movements, see, e.g., [4], [24], [29]. In particular, keystroke dynamics gained lots of popularity through [18], where it was used to augment password authentication similarly to our pin-entry scenario. Keystroke dynamics is another method that could be combined with our PIN entry scenario, but it requires longer sampling duration to work well. A survey on the large body of literature on biometrics using keystroke dynamics is given in [16]. In contrast to keystroke dynamics, some studies on mouse movements argue that it should not be used as biometrics, as it is too unreliable [26], while others report high accuracies [9], [19], [33]. Recently, [33] achieved EER as low as 1.3% using successive mouse actions between clicks. The best accuracy has been reported in [19] with a FAR of 0.36% and a FRR of 0%, although it has been suspected that this result was influenced by recording the data on a different computer for each user [15].

The work in [17] uses multi-modal biometrics composed of voice, face, and signature data for authentication on mobile phones. The goal is to enable legally binding contracts to be signed. According to [17], the face verification shows very high Equal Error Rate (EER), around 28%, the EER of voice and signature are around 5% and 8%, respectively. The fusion of the three biometrics decreases the EER to 2%, yet the price to be paid is the highly intrusive procedure where the user needs to sign, read, and enter a PIN-based password. The work in [7] is related to multi-modal biometrics. It investigates users' touch screen gestures captured by their smart phones. The study shows low error rates, e.g., EERs between 0% and 4% when using SVM and k-NN classifiers. Although not in the area of system security, the work in [11] describes a similar approach based on Swept Frequency Capacitive Sensing, which measures the impedance of a user to the environment across a range of AC frequencies. Finally, a comprehensive survey on multi-modal behavioral biometrics can be found in, e.g., [32].

[27] covers recent papers on biometrics based on the electroencephalography (EEG), the electrocardiogram (ECG),

and the skin conductance, also called electro-dermal response (EDR). Probably the most related to this paper is the work in [5], where bioimpedance is used as a biometric. A wearable sensor is designed to passively recognize wearers based on a body's unique response to the alternating current of different frequencies. Experiments were conducted in a family-sized setting of 2 to 5 subjects, where a person wears a bioimpedance sensor on the wrist. They achieve recognition rate of 90% when their impedance measurements are augmented with hand geometry. Our biometric solves a different problem but it still uses the body's response to a signal. We achieve an achieve recognition rate of 100% when samples are taken in one session and 88% when samples are taken weeks apart. We also do not require any augmentation.

Although not directly related to our work, it is interesting to mention a cryptographic key generation scheme described in [10]. It introduces a key generation resistant against coercion attacks. The idea is to incorporate skin conductance measurements into the cryptographic key generation. They experimentally show that the skin conductance measurement will help to reveal user's emotional states and recognize the attack as a stressful event (significantly different from the state when the keys were generated). This way, the generated keys include a dynamic component that can detect whether a user is forced to grant an access to the system.

IX. Conclusion

We have proposed a new biometric based on the human body's response to an electric square pulse signal. We used our new pulse-response biometric as an additional authentication mechanism in a PIN entry system, enhancing the security of the PIN entry mechanism without adding additional inconvenience for the user.

We also applied our new pulse-response biometric to the problem of continuous authentication. We designed a continuous authentication mechanism on a secure terminal, ensuring that the user that started the session continued to be the person physically at the keyboard.

We showed through experiments on our proof-of-concept prototype system, that each human body exhibits a unique response to a signal pulse applied at the palm of one hand, and measured at the palm of the other. Using our prototype setup we where able to identify users with high probability in a matter of seconds. This identification mechanism integrates very well with other well established methods, e.g., PIN entry, to produce a highly reliable additional layer of security, either on a continuous basis or at login time.

Acknowledgements

We thank the anonymous reviewers and the assigned "shepherd" for their comments and all the help in improving this paper. We also thank Srdjan Ĉapkun for his help during the early phases of this work.

References

[1] C. Barral and A. Tria, "Fake fingers in fingerprint recognition: Glycerin supersedes gelatin," in *Formal to Practical Security*, ser. Lecture Notes in Computer Science, V. Cortier, C. Kirchner, M. Okada, and H. Sakurada, Eds. Springer Berlin Heidelberg, 2009, vol. 5458, pp. 57–69. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-02002-5_4

[2] V. Biometric. (2009, Feb.) How to make the fake fingerprints (by VIRDI). Last accessed 03.08.2013. [Online]. Available: http://www.youtube.com/watch?v=-H71tyMupqk

[3] A. Boehm, D. Chen, M. Frank, D. Huang, C. Kuo, T. Lolic, I. Martinovic, and D. Song, "Safe: Secure authentication with face and eyes," in *In Proceedings of International Conference on Security and Privacy in Mobile Information and Communication Systems*, Jun. 2013.

[4] N. Clarke and S. Furnell, "Advanced user authentication for mobile devices," *Computers & Security*, vol. 26, no. 2, pp. 109 – 119, 2007.

[5] C. Cornelius, J. Sorber, R. Peterson, J. Skinner, R. Halter, and D. Kotz, "Who wears me? bioimpedance as a passive biometric," in *Proceedings of the USENIX Workshop on Health Security and Privacy*, August 2012.

[6] N. S. . T. Council, "Biometrics frequently asked questions," 2006. [Online]. Available: http://biometrics.gov/Documents/FAQ.pdf

[7] M. Frank, R. Biedert, E. Ma, I. Martinovic, and D. Song, "Touchalytics: On the applicability of touchscreen input as a behavioral biometric for continuous authentication," *Information Forensics and Security, IEEE Transactions on*, vol. 8, no. 1, pp. 136 –148, 1 2013.

[8] 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 *White paper for Black Hat USA 2012*, Feb. 2012, last accessed 03.08.2013. [Online]. Available: https://media.blackhat.com/bh-us-12/Briefings/Galbally/BH_US_12_Galbally_Iris_Reconstruction_WP.pdf

[9] H. Gamboa and A. Fred, "A behavioral biometric system based on human-computer interaction," in *Proc. SPIE 5404*, 2004, p. 381.

[10] P. Gupta and D. Gao, "Fighting coercion attacks in key generation using skin conductance," in *Proceedings of the 19th USENIX Conference on Security*, ser. USENIX Security'10, 2010, pp. 30–30. [Online]. Available: http://dl.acm.org/citation.cfm?id=1929820.1929860

[11] C. Harrison, M. Sato, and I. Poupyrev, "Capacitive fingerprinting: exploring user differentiation by sensing electrical properties of the human body," in *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology (UIST'12)*, 2012, pp. 537–544.

[12] S. V. Inc., "Facial recognition provides continuous system security," 2013. [Online]. Available: http://www.sensiblevision.com/en-us/fastaccessanywhere/overview.aspx

[13] A. Jain, A. Ross, and S. Pankanti, "Biometrics: a tool for information security," *Information Forensics and Security, IEEE Transactions on*, vol. 1, no. 2, pp. 125 – 143, June 2006.

[14] A. Jain, A. Ross, and K. Nandakumar, *Introduction to Biometrics*, ser. SpringerLink : Bücher. Springer, 2011. [Online]. Available: http://books.google.com/books?id=ZPt2xrZFtzkC

[15] Z. Jorgensen and T. Yu, "On mouse dynamics as a behavioral biometric for authentication," in *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*, ser. ASIACCS '11. New York, NY, USA: ACM, 2011, pp. 476–482. [Online]. Available: http://doi.acm.org/10.1145/1966913.1966983

[16] R. Joyce and G. Gupta, "Identity authentication based on keystroke latencies," *Commun. ACM*, vol. 33, no. 2, pp. 168–176, Feb. 1990. [Online]. Available: http://doi.acm.org/10.1145/75577.75582

[17] J. Koreman, A. C. Morris, D. Wu, S. Jassim, H. Sellahewa, J. Ehlers, G. Chollet, G. Aversano, H. Bredin, S. Garcia-salicetti, L. Allano, B. L. Van, and B. Dorizzi, "Multi-modal biometric authentication on the SecurePhone PDA," 2006.

[18] F. Monrose, M. K. Reiter, and S. Wetzel, "Password hardening based on keystroke dynamics," in *Proceedings of the 6th ACM conference on Computer and communications security*, ser. CCS '99. New York, NY, USA: ACM, 1999, pp. 73–82. [Online]. Available: http://doi.acm.org/10.1145/319709.319720

[19] Y. Nakkabi, I. Traoré, and A. A. E. Ahmed, "Improving mouse dynamics biometric performance using variance reduction via extractors with separate features," *Trans. Sys. Man Cyber. Part A*, vol. 40, no. 6, pp. 1345–1353, Nov. 2010. [Online]. Available: http://dx.doi.org/10.1109/TSMCA.2010.2052602

[20] L. Nara, "Hand electrodes brass (1 pair)," 2013. [Online]. Available: http://www.lyranara.com/hand-electrodes-brass-1-pair/

[21] J. F. Nevenka Dimitrova, "Continuous face recognition with online learning," US Patent US 20 090 196 464 A1, 08 6, 2009. [Online]. Available: http://www.google.com/patents/US20090196464

[22] M. D. Nguyen and Q. M. Bui, "Your face is not your password: Face authentication bypassing - lenovo - asus - toshiba," in *In briefings of 2009 Black Hat Conference*, Feb. 2009, last accessed 03.08.2013. [Online]. Available: http://www.blackhat.com/presentations/bh-dc-09/Nguyen/BlackHat-DC-09-Nguyen-Face-not-your-password-slides.pdf

[23] K. Niinuma and A. K. Jain, "Continuous user authentication using temporal information," in *Biometric Technology for Human Identification VII*, B. V. K. V. Kumar, S. Prabhakar, and A. A. Ross, Eds., vol. 7667, no. 1. SPIE, 2010.

[24] M. S. Obaidat and B. Sadoun, "Keystroke dynamics based authentication," in *Biometrics*, A. K. Jain, R. Bolle, and S. Pankanti, Eds. Springer US, 2002, pp. 213–229.

[25] I. T. L. N. I. of Standards and Technology, "The biometrics resource center," 2013. [Online]. Available: http://www.nist.gov/itl/csd/biometrics/index.cfm

[26] M. Pusara and C. E. Brodley, "User re-authentication via mouse movements," in *Proceedings of the 2004 ACM workshop on Visualization and data mining for computer security*, ser. VizSEC/DMSEC '04. New York, NY, USA: ACM, 2004, pp. 1–8. [Online]. Available: http://doi.acm.org/10.1145/1029208.1029210

[27] K. Revett and S. T. Magalhes, "Cognitive biometrics: Challenges for the future," in *Global Security, Safety, and Sustainability*, 2010, vol. 92, pp. 79–86.

[28] P. H. Service, "Worker deaths by electrocution a summary of NIOSH surveillance and investigative findings," National Institute for Occupational Safety and Health, Tech. Rep., May 1998.

[29] R. Spillane, "Keyboard apparatus for personal identification," *IBM Technical Disclosure Bulletin*, vol. 17, no. 3346, 1975.

[30] Wikipedia, "Sensitivity and specificity," 2013. [Online]. Available: http://en.wikipedia.org/wiki/Sensitivity_and_specificity

[31] J. Woodward, N. Orlans, and P. Higgins, *Biometrics*, ser. RSA Press Series. McGraw-Hill/Osborne, 2003. [Online]. Available: http://books.google.com/books?id=j-o_btaFK6wC

[32] R. V. Yampolskiy and V. Govindaraju, "Behavioural biometrics; a survey and classification," *Int. J. Biometrics*, vol. 1, no. 1, pp. 81–113, Jun. 2008. [Online]. Available: http://dx.doi.org/10.1504/IJBM.2008.018665

[33] N. Zheng, A. Paloski, and H. Wang, "An efficient user verification system via mouse movements," in *Proceedings of the 18th ACM conference on Computer and communications security*, ser. CCS '11. New York, NY, USA: ACM, 2011, pp. 139–150. [Online]. Available: http://doi.acm.org/10.1145/2046707.2046725