Authentication and Pairing Using Human Body Impedance

Marc Roeschlin

Kellogg College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2018

Acknowledgements

First and foremost, I would like to thank my supervisor, Kasper B. Rasmussen. He has always provided scientific advice which helped me form my academic thinking and hone my problem solving skills. Working with Kasper has been a great joy and his inventive mind has impressed me many times.

I would also like to whole-heartedly thank my co-supervisor, Ivan Martinovic. The discussions with Ivan were always fruitful and inspiring. He has always motivated me to tackle the next challenge.

Thanks to my thesis examiners Andrew Martin and Emiliano De Cristofaro for their constructive feedback. My thanks also go to Cas Cremers and Andrew D. Ker, who were the intermediate assessors I had during the course of my research at the University of Oxford.

I would also thank the reviewers of my published research. Your feedback guided my thoughts and transformed my research into this thesis.

I thank the Department of Computer Science at the University of Oxford and Kellogg College for helping with travel and conference funding.

Most of all, I thank my parents, family and close friends, who are always there for me and have supported me throughout my whole life. Their support and love has made this journey possible. Many thanks go out to Hannah for lending an ear to me and encouraging me at all times.

Many thanks to all my friends in Switzerland, the UK and the US. You made sure I stayed sane and healthy during this time.

Last but not least, a shout-out to the crew in RHB 101! We have spent so many hours together and you have made this ride worthwhile.

Abstract

The electrical properties of the human body are an unexplored source of much potential in Computer Security and Biometrics. The body is an electric conductor with very interesting features: at lower frequencies it can be treated like a resistive cable and at higher frequencies it behaves similar to an electric antenna. These physical phenomena allow the construction of electrical circuits that interface the human body or even build on it as the core element. In particular, one can use the human body as a transmission medium for electromagnetic signals. Using so-called body channel communication techniques, it is possible to send information wirelessly through the body of a person. In addition, the electrical properties of the body can be measured and constitute physiological characteristics unique to the every individual.

Although these properties could prove very useful in the context of Computer Security, their potential for security applications is largely unexplored. We find, however, that extracting the body's distinctive impedance characteristics or using the body as a communication medium entails many interesting application scenarios. The human body can serve as an additional layer to augment the security in existing protocols or enable the design of entirely new methods.

In this thesis, we propose solutions to two challenging problems in the area of System Security: user authentication and secure device pairing. In order to address these problems, we conceptualize and develop solutions that make use of the electrical properties of the human body.

Keeping simplicity and universality in mind, we primarily focus on applications where the human body can be interfaced in ways that require little to no user involvement and are completely noninvasive.

We find that it is most natural, even for novice users, to interface the body through the person's hands. Our proposed solutions therefore connect to the body through (capacitive) electrodes that the user has to touch with their hands. In case of user authentication, we propose a method that measures the human body impedance from one hand to the other, and for device pairing, we present an approach that allows two devices to communicate with each other when a person touches them both, one with each hand.

Contents

List of Figures

List of Tables

List of Algorithms

1 Introduction

The introduction presents the aims and objectives of this thesis, outlines the overarching research questions, and identifies the contributions to the research fields of Computer Security and Biometrics.

Contents

1.1 Motivation

Cryptography and Computer Security have come a long way. Not many people vividly remember reading newspaper articles on stolen Enigma Machines and the eventual announcement that the code had been broken. Most are only aware of how cryptography has affected the Second World War thanks to movies, books and documentaries. While the Enigma code had cryptographic weaknesses and is regarded as completely outdated by today's standards, Allied cryptologists were able to crack the code not only due to their immense efforts at Bletchley Park, but also due to mistakes caused by German operators [1]. The encryption hardware was complex for the time and anecdotal evidence suggests that many German radiomen

were not familiar with all the intricacies of the Enigma [2]. Some might blame the lack of proper training, but far more likely is the fact that the technology was simply not straightforward enough to be used in the heat of battle. It must have been difficult to handle key tables and change parameters of the encryption protocol while under pressure.

The take-away from this detour into the history of cryptography is that even if the underlying mathematical model of encryption is sound, we need to observe how the designed encryption mechanism is interfaced by humans and we must strive to make this process as robust as possible. Even though the context in which we rely on cryptography and security protocols has changed significantly from Enigma's times, the element of human interaction is still required in many cases. Fortunately, a multitude of technologies has revolutionized user input and simplified the way we interact with electronic devices, but the need for explicit user action to execute a cryptographic protocol persists. In most cases, a human is at least required to initiate the protocol and, in doing so, plays a major role in contributing to the protocol's correctness and soundness. Thus, the human factor can not be neglected, which past events, such as the breaking of the Enigma code, clearly demonstrated.

A promising strategy is to devise methods that do not heavily rely on human interaction, but still provide authority and control to the user. This is a challenging problem, especially since the resulting mechanism still has to be secure from a cryptographic standpoint [3].

With the introduction of public-key encryption and public-key infrastructures, the reliance on the user has shifted to trusted third parties. Devices that are interconnected via the Internet, or have other infrastructure support, largely do not require user interaction in order to run security protocols between them. They rely on trusted certificate authorities that verify the identities of devices and thereby provide a root of trust. As a result, the human interaction has become less relevant for secure machine-to-machine communication. However, there are security protocols where user interaction is absolutely necessary. Two of these are user authentication and device pairing.

User authentication is the process of verifying a person's identity claim in an automated way. This procedure naturally requires some form of user cooperation and while failure of the user to execute the required action is not crucial for the security of the protocol, it inconveniences the user if authentication is not successful. Therefore, it is desirable to achieve user authentication that is noninvasive and only demands a minimum of user cooperation.

Secure device pairing, on the other hand, is concerned with the problem of having electronic devices establish secure communication between them if they do not have any shared knowledge or association to a trusted party. In this case, a human needs to mediate between the devices to establish a communication channel and the security of pairing protocol is directly dependent on the user's actions.

In this thesis, we seek to find solutions to those two challenging security problems, where user interaction is unavoidable and constitutes a major component of the security mechanism. While prior research has yielded many approaches to user authentication or device pairing, each comes with their own advantages and drawbacks. We focus on solutions that are particularly simple and easy from a user's perspective and thus can minimize errors that could jeopardize security.

Our goal is to design a new means by which users can interact with devices to perform user authentication and device pairing in a completely effortless fashion. This new method should neither require the users to carry an objects, such as a key or security token, nor should it rely on memorized secrets, such as passwords. Additionally, we want our approach to be fast and computationally inexpensive, but most of all, it needs to be secure. This means that the mechanism's underlying protocol, as well as the necessary man-machine interaction, cannot be subverted. If the protocol terminates successfully, it must provide certain guarantees. In case of user authentication, we want the mechanism only to validate an identity claim, if it is legitimate; and in case of device pairing, the mechanism should establish a secure communication channel between the devices.

The new mechanism we design in this thesis relies on the electrical properties of the human body and tries to meet all described requirements. Our method surpasses many existing techniques in simplicity, capitalizing on the user-specific information extracted from the body without any cognitive effort or demanding specific user action. All that is needed is to touch two electrodes.

1.1.1 Research Questions

With aims to design novel user authentication and device paring mechanisms that rely on the the electrical properties of the human body, we can break down the required research work into the following questions that need to be answered:

- Can we interface the electrical phenomena of the human body in a safe way to enable security mechanisms, such as user authentication and device paring? What are the technical requirements for the physical implementation of such a mechanism?

- What are suitable application scenarios for security mechanisms based on the electrical properties of the human body and how secure are they?

User authentication:

- How can we exploit the electrical properties of the human body to obtain user-specific information that identifies an individual? Is this information obtainable without the need to carry an object or memorization? What is the required user action?
- How do we define unique (biometric) features extracted from the electrical properties that facilitate user recognition?
- How does such a recognition mechanism compare with other methods?

Secure device pairing:

- Can we send data through the body in a secure and safe way? Can two electronic devices that are in physical contact with the body establish such a data link?
- What are the characteristics of the human body as a transmission medium for electrical signals?
- Can the communication be interfered with and if so, what are the circumstances under which this is possible?

At this moment, there is virtually no existing research literature that to answers these questions. In order to find answers, we need to acquire a profound understanding of the electrical properties of the human body, the concept of biometric recognition and secure device pairing. Only this will allow us to reformulate these questions as research goals that we then can pursue to develop novel security mechanisms.

1.2 Contribution

This thesis is one of the first works to present an extensive look at techniques which exploit the electrical properties of the human body for applications in System Security. Body impedance, i.e., how the human body opposes electrical current flow, is a concept that has been thoroughly researched and is well-understood in the medical field. Body impedance analysis and impedance tomography are medial methods that are routinely used for health monitoring and diagnostic purposes. The

success of these methods is rooted in the fact that the body's impedance provides a wealth of information and encodes the composition and state of each individual.

It is useful to attempt tapping into this potential outside of the medical sciences and to investigate possible applications in other fields.

However, medical procedures often do not directly translate to other contexts since they are invasive and need to be performed in a controlled and supervised environment. It is therefore challenging to modify and potentially simplify these techniques in order to make them suitable for general use.

In this thesis, we show how body impedance readings can be acquired without medical-grade gel or needle electrodes, but with the simple touch of two dry conductive pads. While this approach might not be precise enough to diagnose diseases, we demonstrate that our technique can successfully extract information from the human body and utilize the body as a medium for data communication.

In particular we show that, measuring body impedance in such a simple fashion is a way to extract idiosyncratic features from the body to facilitate the recognition of individuals. Using the exact same interface, i.e., two conductive pads, we show that the body can also be used to establish a secure communication channel by sending electrical signals through a person.

These findings address two interesting and timely problems in Biometrics and Computer Security. Firstly, the harnessing of unique features present in body impedance can be leveraged to implement unobtrusive user authentication protocols. Secondly, the human body can serve as an authenticated channel to enable secure device pairing.

The fact that our interface for body impedance measurements is so effortless makes this a particularly attractive approach to tackle those problems and justifies detailed research on this phenomenon. We outline below how human body impedance provides effective solutions to user authentication and secure device pairing.

User authentication This thesis proposes a novel biometric trait based on the electrical impedance of the human body. The electrical impedance of a person's body allows the extraction of physiological characteristics that enable biometric recognition. Due to their idiosyncrasies and persistence, the extracted characteristics represent viable traits for automated biometric verification and identification. In fact, body impedance measurements can facilitate user recognition with remarkable performance.

In addition to high recognition rates, measuring a person's body impedance has many appealing properties. For instance, the biometric capture process is

fast and unobtrusive to in comparison to existing biometric characteristics. From a user's perspective, no specific interaction is required except for touching two electrodes for a short period of time. The nature of body impedance as a biometric modality and the fast acquisition time make this biometric trait applicable to continuous authentication, that is, when a user's identity is verified repeatedly and in quick succession.

In this thesis, we present a way to distill and quantize a person's body impedance and extract the features needed for recognition. We test the feasibility of an impedance-based biometric recognition system in a number of user experiments, with a prototype measurement set-up that follows the design of our biometric modality. Using this proof-of-concept, we determine the contexts in which body impedance can be applied to user authentication and we identify its strengths and weaknesses. We also assess the resilience of the biometric trait against attacks under different adversarial models.

Secure device pairing Device pairing is the problem that arises every time when two devices have to establish a secret in order to secure subsequent communication. Research has suggested many ways to address this problem and proposed a number of approaches to perform device pairing without relying on shared knowledge or a trusted third party. Some of these methods are successfully implemented in practice, but require various levels of user involvement, most often in the form of a manual transfer of information from one device to the other.

Our approach to device pairing does not require the user to perform any particular action other than touching an electrode on both devices, one with each hand. By performing this task the user establishes a physical channel for the devices to communicate. As long as the user is in contact with both devices, his body acts as a transmission medium for intra-body communication.

The devices use this channel as part of a device pairing protocol that allows them to agree on a mutual secret. At the same time, they can extract physical properties of the messages they receive on the body channel during the pairing process. Using this knowledge, the devices can detect with high accuracy if they are being held by the same person. This mechanism gives the user control to authorize a pairing operation by simply touching them simultaneously.

We prove that the device pairing protocol we present in this thesis is secure in our threat model and, to verify the idea in practice, we build a proof-of-concept set-up and conduct experiments with a notable sample of people.

1.3 Peer-reviewed Work

This section provides a list of the scientific publications that have resulted as part of this thesis and explains how they have been incorporated in this work.

1.3.1 Main Publications

Publications that provide the foundation of this dissertation are listed in the following.

- Chapters 2 and 3 are largely based on our three publications on biometric recognition using body impedance:

  1. The paper **Authentication Using Pulse-Response Biometrics** presented at *The Network and Distributed System Security Symposium (NDSS), 2014* [4].
  2. The same-titled magazine publication **Authentication Using Pulse-Response Biometrics** published in *The Communications of the ACM (CACM), 2017* [5].
  3. The journal publication **Pulse-Response: Exploring Human Body Impedance for Biometric Recognition** published in *The ACM Transactions on Privacy and Security (TOPS), 2017* [6].

  The original publication on pulse-response recognition [4] received a Distinguished Paper Award and allowed us to present our work to a broader audience through the ACM magazine [5].

- Chapter 4 extends body impedance based recognition to key generation and is mainly based on the paper **Generating Secret Keys from Biometric Body Impedance Measurements** published in *The Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society (WPES) at the Conference on Computer and Communications Security (CCS), 2016* [7].

- Chapter 5 presents an application of body impedance recognition in the context of Vehicular Ad-hoc Networks and is based in large parts on the paper **Bionyms: Driver-centric Message Authentication using Biometric Measurements** which has been accepted for publication in *The Proceedings of the IEEE Vehicular Networking Conference (VNC), 2018* [8].

- Chapters 2 and 6 build up on the publication **Device Pairing at the Touch of an Electrode** presented at *The Network and Distributed System Security Symposium (NDSS), 2018* [9].

Authorship Apart from the few exceptions that are specifically mentioned below, all results and research papers included in this thesis are Marc Roeschlin's own work that was conducted with guidance from his supervisors, Kasper B. Rasmussen and Ivan Martinovic.
Marc Roeschlin would like to point out the following exceptions and thank his co-authors for their contributions:

- Kasper B. Rasmussen contributed the original idea of using body impedance as a trait to perform user authentication and he suggested to conduct a feasibility study.

- Gene Tsudik provided valuable high-level suggestions for the body impedance recognition system presented in Chapter 3 and the key generation mechanism proposed in Chapter 4.

- Ivo Sluganovic, Ivan Martinovic's D.Phil. student, worked in collaboration with Marc Roeschlin when developing the feature extraction method based on Siamese convolutional networks and the key guessing algorithm (both in Chapter 4). Ivo Sluganovic wrote code to fine-tune the feature extraction (see Appendix A) and helped calculate the run-time of the key guessing strategy.

- Christian Vaas, Ivan Martinovic's D.Phil. student, and Marc Roeschlin both equally contributed to the publication "Bionyms: Driver-centric Message Authentication using Biometric Measurements" (see above) which is the basis for Chapter 5. While Marc focused on the biometric aspect of driver recognition and the authentication protocol, Chris explored the effect of the protocol on the VANET infrastructure and network performance. Parts which are mainly Chris' work are not included in this dissertation.

1.3.2 Research Collaborations

Major research publications that Marc Roeschlin was involved in as a co-author during his D.Phil. and thus tangentially shaped this dissertation are mentioned below.

- A collaboration with Ivo Sluganovic focused on the feasibility of using eye movements for fast user authentication in place of a password or fingerprint. The resulting paper **Using Reflexive Eye Movements for Fast Challenge-Response Authentication** was published in *The Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS), 2016* [10]. An extended journal article based on the original paper has been accepted for publication in *The ACM Transactions on Privacy and Security (TOPS)* [11].

- Together with Simon Eberz (Ivan Martinovic's D.Phil. student) and two other D.Phil. students at the University of Oxford, we managed to devise a method to successfully trick biometric systems based on ECG (i.e., heartbeat) with artificially generating ECG signals. The results were summarized in the paper **Broken Hearted: How To Attack ECG Biometrics** and were presented at *The Network and Distributed System Security Symposium (NDSS), 2017* [12].

1.4 Thesis Outline

The structure of this thesis is as follows:

- **Chapter 2** familiarizes the reader with the two most important concepts that this thesis builds up on. It provides background on the electrical properties of the human body and biometric recognition. We also provide an overview and comparison of existing biometric methods.

- **Chapter 3** verifies if body impedance is a unique trait that can be used for biometric recognition. It also looks at two use cases of body impedance based authentication and analyses the security of such applications.

- **Chapter 4** proposes a mechanism by which user-specific keys from body impedance measurements can be generated. We discuss, implement, and evaluate an approach to extract features from impedance measurements and condense them into cryptographically secure secrets.

- **Chapter 5** presents a particularly interesting application of body impedance in the context of driver authentication and Vehicular Ad-hoc Networks and makes use of the techniques devised in the preceding chapters.

- **Chapter 6** proposes a secure device pairing protocol that exploits the electrical properties of the human body. Instead of measuring body impedance to distinguish between individuals, we use the body as an authenticated channel for the transmission of keying material.

- **Chapter 7** finally concludes this thesis by discussing the key findings and possible directions for future work.

2 Background

This chapter provides the necessary background to formulate the research goals of this thesis in technical terms. It also helps the reader understand the subsequent chapters and gives an overview of existing biometric modalities found in literature.

Contents

2.1 Electrical Properties of the Human Body

Bioimpedance and *bioelectricity* both describe two very related electrical phenomena of the human body [13]. Bioimpedance covers the passive electrical properties of tissue, whereas bioelectricity describes the ability of tissue to generate measurable

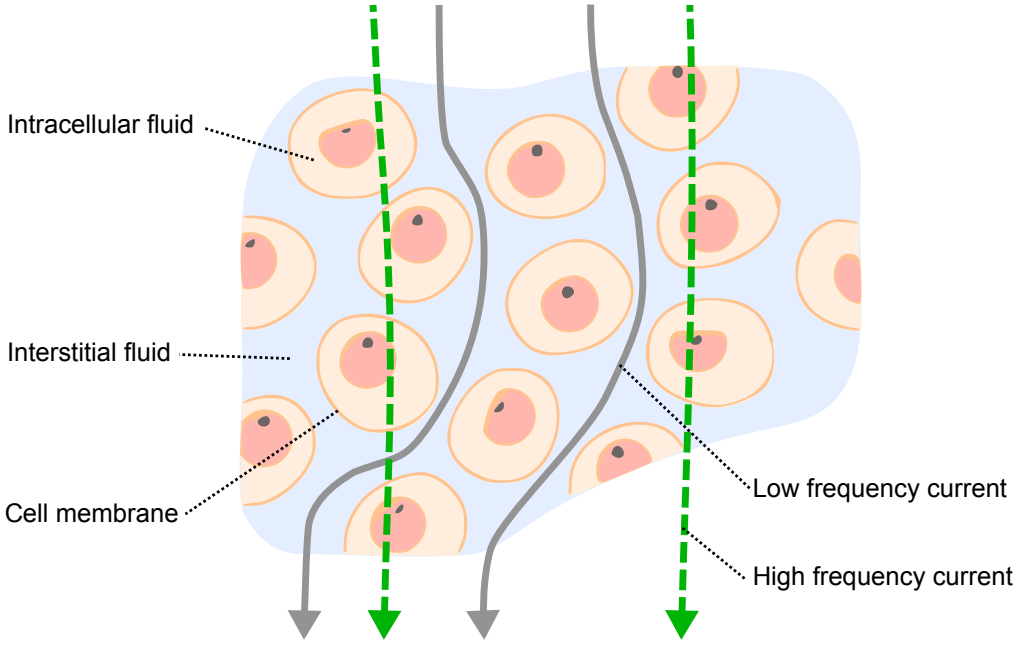

Figure 2.1: Current flow at the cellular level. Currents at lower frequencies will travel through the fluids around the cells, whereas higher frequencies will penetrate cell membranes.

electricity on its own, e.g., the human heart, and how tissue can be affected by externally applied electricity, e.g., stimulating cell growth.

In this work we mainly focus on *bioimpedance*, i.e., the passive electrical properties. We apply an external electrical signal to the body and observe how this current flow is altered by the body. We do not consider any effects caused by what is categorized as *bioelectricity*, i.e., electricity generated by the body itself or non-passive elements.

Biological tissue In general, biological tissue behaves as either an electrical conductor or a dielectric[1]. At lower frequencies of around 100 kHz or less, most tissue exhibits the behavior of an electrically conducting solution with capacitive components that exist due to the cell membranes. However, at frequencies of 50 kHz or more, the dielectric properties can prevail and tissue essentially becomes an insulator. At very high frequencies, i.e., in the Giga-Hertz range, insulation is comparable to that of pure water.

The frequency-dependent behavior of biological tissue originates from the fact that most tissue consists of intra- and extra-cellular fluid, which is separated by the plasma membrane. In Figure 2.1 a cross-section of (generic) tissue at the cellular level is shown. If a current is applied at lower frequencies, the majority

[1] An electrical insulator that can be charged/polarized.

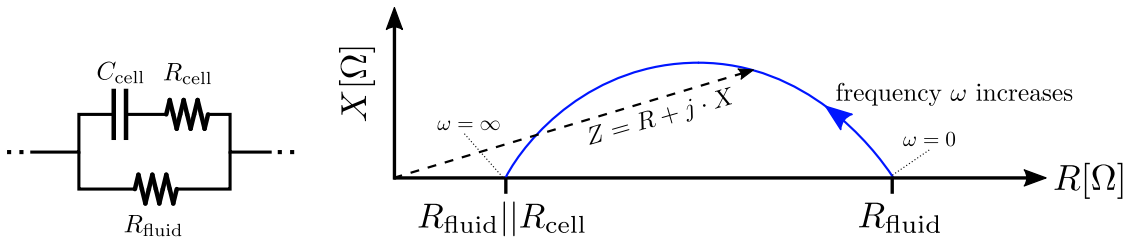

Figure 2.2: Equivalent circuit for biological tissue and corresponding cole-cole plot.

of the electrons travel through the interstitial fluids surrounding the cells. Higher frequencies on the contrary, penetrate the cell membranes and are not restricted to the fluids in the tissue.

This relationship can be modeled and illustrated graphically. A simple theoretical model for the passive electrical properties of biological tissue is a circuit containing three passive elements, as shown in Figure 2.2; capacitor C_{cell} describes the cell membrane, resistor R_{cell} models the intracellular fluid, and resistor R_{fluid} represents the fluids surrounding the cells. The qualitative cole-cole plot in the complex plane depicts the behavior of such an electrical circuit. In Figure 2.2 the impedance is broken down into the real and the imaginary part, i.e., the resistive component and the capacitive/dielectric component. The half circle shows the resulting impedance and is parameterized through the frequency ω of the current. When applying direct current only ($\omega = 0$), capacitor C_{cell} blocks the current flow and therefore the resulting impedance is defined by R_{fluid}, which is purely resistive. As ω increases, C_{cell} starts to conduct electricity and the dielectric component becomes more prevalent. At the theoretical maximum of $\omega = \infty$, the capacitor is short-circuited and the impedance amounts to that of both resistors being in parallel, i.e., $R_{\text{fluid}} \| R_{\text{cell}}$.

While the equivalent circuit captures the general impedance profile, the parameters can vary significantly depending on the type of tissue, since the composition is directly reflected in the electrical impedance: tissues with higher water content, such as blood and muscles, generally act as a resistive conductor, whereas adipose and cornual tissues exhibit a mostly capacitive impedance.

By exploiting these differences in electrical impedance, it is possible to not only determine the water content of a tissue sample, but also to identify the composition and structure of tissue. Procedures like impedance biopsies have become remarkably accurate and have been used in the medical field with success for quite some time now.

Research in bioimpedance The research areas that deal with bioimpedance are for the most part biomedical engineering and biophysics. Electrical Impedance Tomography [14] and Bioelectrical Impedance Analysis [15] have become well-established and noninvasive techniques to measure and monitor tissue composition over a period of time. In humans, the medical imaging of body parts using bioimpedance can deliver an accurate estimation of the percentage of fat, muscle and fluid levels [16].

In the last few years, many more specific applications have been proposed, such as the impedance imaging of individual organs, e.g., heart and lungs [17, 18], or the detection of early tumors [19, 20] and skin cancer [21]. Impedance measurements can also be a reliable indicator for diseases [22], such as meningitis [23] and asthma [24].

Being such an unobtrusive technique to observe and study the physiological parameters of the human body, bioimpedance spectroscopy has even been proposed as a means for the medical examination of newborns [25]. The study in [26] suggests that impedance imaging can detect insufficient blood supply to the brain of newborns and can help monitor the cardiac activity of neonates that need to be supported by mechanical ventilation [27].

Distinguishing power Since bioimpedance has been successfully used to detect minute differences in tissue, we hypothesize that, if acquired correctly, bioimpedance measurements can be utilized to distinguish between individuals. The structure of blood vessels, muscle, fat tissue, cartilage and bones differ widely across individuals, and consequently, it would be logical if impedance measurements could be used to differentiate between people.

We try to investigate and explore this hypothesis in the following chapters, where we focus on the external application of bioimpedance in a non-medical scenario. For our studies, we consider the electrical impedance of the entire human body, which can be conveniently acquired by coupling electricity to the human skin through electrodes or conductive pads. This type of impedance measurement provides insight into the electrical properties from a macroscopic perspective, capturing extant differences in the entire human body.

2.1.1 Body Impedance

When discussing *bioimpedance* of the human body, we refer to it as *body impedance*. Consequently, body impedance analysis is concerned with the measurement and quantification of the electrical impedance of (parts of) the human body. The

electrical impedance Z can be computed according to Ohm's law by taking the ratio between voltage and current:

$$Z = \frac{U}{I} \tag{2.1}$$

Impedance is a complex unit that consists of a real and an imaginary part. The real part is called resistance and the imaginary part is the reactance, resulting in the following mathematical definition, $Z = R + j \cdot X$. Alternatively, electrical impedance can be defined as $Z = |Z| \cdot e^{j \cdot arg(Z)}$ where $arg(Z)$ is the phase difference between voltage and current.

For the purpose of the studies conducted in this thesis, we assume that body impedance can be modeled as linear and time-invariant[2], a common assumption in Electrical Engineering when reasoning about lumped electrical circuits. "Lumped" refers to networks that only consist of resistors, capacitors and inductors, i.e., passive components. Since bioimpedance is concerned with the passive electrical properties of tissue, this assumption is directly applicable and allows us to use readily available frequency analysis techniques. They, in turn, can break down body impedance measurements into separate frequencies ω and compute the frequency-dependent impedance $Z(\omega)$, which can help reveal the resistive (at lower frequencies) and dielectric properties (at higher frequencies) described above. We exclude any behavior that changes over time, or the ability of the human body to generate electricity, such as the heartbeat.

Acquisition of impedance readings An important aspect of body impedance is its acquisition. As it is clear from the definition of impedance in Equation (2.1), in order to compute body impedance, either a known voltage U must be applied to the body and the resulting current I is measured, or I is applied and U is measured. In Chapter 3, we use *pulse-response* as an instance of a particular body impedance measurement, which is acquired by emitting a low current in the shape of a short square pulse. The resulting voltage can then be correlated with the known (emitted) current to estimate the impedance over a wide range of frequencies.

[2] "Time-invariant" means impedance is not a *direct* function of time. Impedance does depend on frequency.

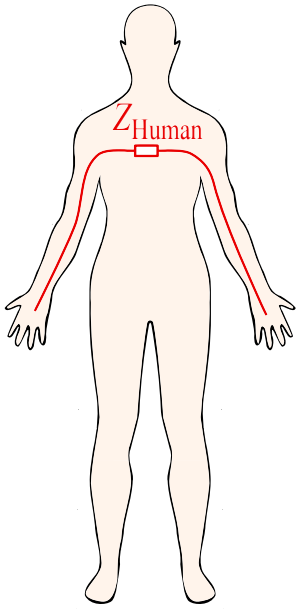

Figure 2.3: Horizontal impedance of the human body Z_{Human} is measured from one hand to the other.

Horizontal body impedance The type of body impedance we focus on in this thesis is captured between the two hands of a human. We call this the *horizontal* body impedance since it spans the two arms and the upper body as depicted in Figure 2.3. The rationale behind using the *horizontal* impedance is its simple acquisition that directly translates to many potential applications in the context of system security. During our everyday activities the hands are for the most part uncovered (unless wearing gloves), a fact that can be exploited by accommodating electrodes or conductive pads on objects that we regularly touch with our hands. For the duration where a physical connection between hands and electrodes exists, a small imperceptible current can be applied to the body and a reading for the impedance can be acquired.

Measurement methods and electrodes Electrodes are a crucial component of any bioimpedance measurement. In order to send current through the body, at least one source and one sink electrode are required. If only two electrodes are connected to the body, both electrodes are used for the current injection and the voltage measurement at the same time. This is often called a two-point, or bipolar configuration. On the contrary, a four-point, or so-called Kelvin method consists of four electrodes: one pair emits and sinks the current, and the other pair measures the voltage. This comes with the advantage that the voltage measurement is less affected by the impedance of the electrodes. In a four-point configuration, only a negligible

amount of current flows through the two dedicated voltage sensing electrodes, which essentially removes the impedance of the electrodes from the measurement. However, the two voltage measuring electrodes need to be placed in-between the two current emitting/sinking electrodes in order to obtain meaningful measurements.

In this thesis, we opt for a two-point measurement method to simplify the measurement process as much as possible and to foreclose electrode alignment issues. In a medical application, electrodes can be carefully placed on the body, but for the applications we envision, there is no supervision and the process needs to be as simple as possible.

For the same reason, we only consider "dry" capacitive touch-electrodes at the interface to the body. Bioimpedance measurements captured at the surface of the human body are often conducted using sticky gel electrodes that attach a thin silver foil to the skin with the help of a contact gel. Such electrodes are common in Electrocardiography (ECG) as the gel quickly builds up a highly conductive salt bridge from the metal to the skin [13], allowing uninterrupted measurements.

Dry touch-electrodes that do not feature gel, i.e., simple metal plates or conductive surfaces, can lose contact and the measurement might be interrupted. Nevertheless, they are much more suitable in cases where electrodes are touched and released frequently. Attaching gel electrodes to the skin is a time-consuming task that is undesirable especially in non-medical applications. Another drawback of gel electrodes is their potential to cause skin irritation when applied for extended periods or when attached to the same place repeatedly [28].

Capacitive touch-electrodes exhibit a higher impedance than gel-coated electrodes, which can, under some circumstances, adversely affect the measurement. The capacitive element of the electrode-skin contact is in series with the human body and therefore forms a high-pass filter. This can lead to the attenuation of lower frequencies that are relevant for the acquisition of body impedance readings.

However, in the applications we propose in this thesis, having non-sticky conductive pads as electrodes clearly outweighs this disadvantage.

2.1.2 Intra-Body Communication

Unlike bioimpedance and body impedance analysis, body communication does not harness minute differences in tissue composition, but rather makes use of the fact that the body is a good electrical conductor. Intra- and on-body communication are techniques by which the human body is used as a signal transmission guide [29].

If the right frequencies and modulation are chosen, electrical signals can be transmitted through the body in a way that the generated electromagnetic field

is largely contained by the skin. In certain scenarios, characteristics like these are superior to those of other methods based on radio frequencies, such as Bluetooth, Wi-Fi and Infrared. In Personal Area Networks (PANs) [30], for example, intra-body communication has gained a lot of popularity. A PAN tries to provide an infrastructure for sensors and devices in close proximity (less than 10 meters) to exchange information in a wireless fashion using as little power as possible. Intra-body communication effectively takes PANs to the scale of the human body and gives rise to so-called Wireless Body Area Networks (BANs) in which implants, medical equipment and wearable devices communicate wirelessly through the body. This motivated the definition of intra-body communication as a physical communication layer in the relatively new IEEE 802.15.6 standard [31], which is the latest international standard covering BANs.

Even though the IEEE 802.15.6 standard mentions medical and non-medical target applications for intra-body communication, the main drivers for the development of electrical near-field communication in and around the human body have been the biomedical sciences and the medical field. Utilizing the body as a transmission medium for electrical signals is key to achieve low-power wireless sensors for (real-time) health monitoring [32, 33].

The main advantages for the use of intra-body communication are the high conductivity of the human body compared to air and the fact that most electromagnetic energy is confined through the body's surface and not radiated into the environment, resulting in very low energy consumption [33]. Since most of the signal is restricted to the body area, external (radio frequency) interference does not affect the communication channel and robust data transmission can be implemented without a large antenna.

Although these features could prove very useful for applications in the context of Computer Security, the use of the human body as a communication channel for security applications is largely unexplored. The possibility to transmit electrical signals through the human body, whilst most energy is confined to the transmission medium, should be of particular interest and is a property normally not found with other wireless communication techniques, such as Wi-Fi or Bluetooth.

Since its first appearance in 1996 [30], intra-body communication has been covered in a large body of research literature. Numerous proposals on different transmission methods, receiver and transmitter types, as well as modulation techniques, have been published (see [34–36]). We briefly cover existing techniques for body channel communication to support understanding of our design choices in Chapter 6 of this thesis. Body-channel communication can be divided into roughly three groups: galvanic coupling, capacitive coupling and waveguide methods.

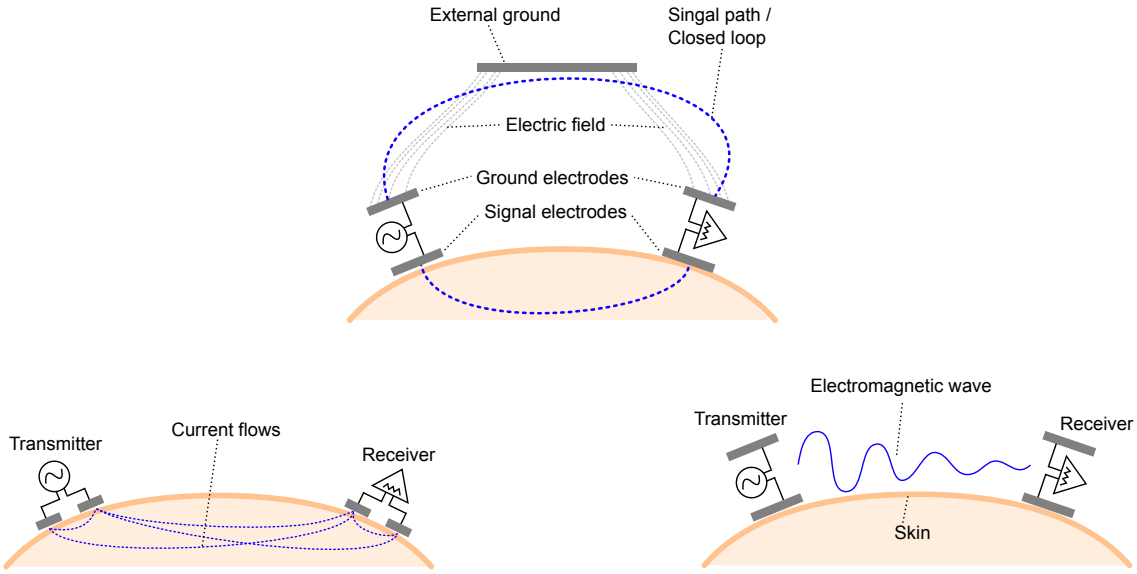

Figure 2.4: Three main methods for intra- and on-body channel communication. From left to right: galvanic coupling, capacitive coupling and surface waves.

Galvanic coupling The concept of galvanic coupling is to induce alternating current into the human body. It was first proposed for intra-body communication in [36, 37] and it works by differentially applying a signal over two electrodes that are attached to the body. Both transmitter and receiver each have two electrodes that are coupled to the human body as shown in Figure 2.4. Most of the induced current flows directly from one sender electrode to the other, but a small portion propagates through the body to the receiver, where it is then detected as the voltage differential between the two receiver electrodes. The carrier of the information is the ionic fluids in the body that form a closed loop for signal transmission [38]. One advantage of galvanic coupling is that the electrical field does not leak outside of the body. This is due to the fact that galvanic coupling relies on electron flow, rather than electromagnetic transmission. A second advantage is that no external ground reference is needed; the return path of the signal transmission is the human body.

Capacitive coupling Capacitive coupling uses an electromagnetic signal for data transmission. The transmitter emits the signal through an electrode that is in touch with the human body. After having traversed the body, the signal is picked up by a receiver that is also coupled to the body (see Figure 2.4). The return path of the signal between transmitter and receiver is established through the environment by electrostatic coupling to external conductive objects, most often earth ground.

This type of communication is enabled by two physical properties:

1. At a frequency of less than 100 MHz, the wavelength λ of an electromagnetic signal is far greater than the size of the human body, i.e., $\lambda > \frac{c}{100\ \text{MHz}} \approx 3$ meters. As a result, the electrical field around the body can be approximated as constant over time, i.e., the phase of the signal remains uniform anywhere close to the human body [39].

2. The human body can be modeled as a conducting wire at low frequencies and capacitive near-field coupling establishes a return path that allows signal transmission [30].

Surface wave techniques Surface techniques are often referred to as "on-body" or "near-body" transmission. They use higher frequencies than capacitive coupling and galvanic coupling. Most often frequencies greater than 100 MHz are used. While some electromagnetic waves propagate through the body in a similar way as with capacitive coupling, usually a significant amount radiates into the air [40]. In addition, as the signal travels through the body, it is attenuated considerably [39]. Unlike with capacitive coupling or galvanic coupling, there is no closed loop for signal transmission; the receiver just measures the intensity of the electromagnetic signal, which is analogous to conventional radio frequency transmission.

2.2 Biometric Recognition

This section provides background knowledge on biometric recognition and describes how it facilitates biometric verification and identification. We summarize the terminology and introduce common requirements and design goals for systems relying on biometric recognition.

The terminology used throughout this thesis is in conformity with the ISO/IEC JTC 1/SC 37 [41] standards on biometrics derived by the Information technology committee and the Biometrics subcommittee of International Organization for Standardization. Wherever possible we use vocabulary and methodologies that are suggested in the respective (sub-)standards. Furthermore, terms "sample" and "measurement" are used interchangeably throughout this thesis.

2.2.1 Definitions

Biometric recognition or just *biometrics* refers to the automated recognition of individuals based on their biological and behavioral characteristics [42]. The general meaning of biometrics usually includes the counting, measuring and statistical analysis of any kind of biological or medical sciences; however, in the context of information security, the meaning is restricted to *biometric verification* and *biometric identification* of individuals, i.e., human beings.

Automated recognition means that a machine based system is used for the biometric recognition either for the full process or in an assisted fashion. The system performing the recognition is usually called the *biometric system* and encompasses both biometric and non-biometric components.

Biometric characteristics are the biological and behavioral characteristics from which distinguishing, repeatable *biometric features* can be extracted and enable the biometric system to perform biometric recognition.

2.2.2 Physiological and Behavioral Characteristics

Although biometric characteristics cannot be completely separated into two distinct categories, the most common classification divides them into physiological and behavioral characteristics, as suggested by the US National Institute of Standards and Technology (NIST), for instance [43].

Physiological characteristics rely on the physiology of a person and include the following: fingerprints, hand geometry, facial recognition, vein patterns, and iris/retina scans. Behavioral traits are based on user behavior and include keystroke timings, speech pattern analysis, gait recognition and analysis of stylus pressure, acceleration and shape in hand-writing.

Why the separation is not completely definitive can be seen using the following example: a fingerprint image results from the physiological traits of the finger ridge patterns, but presenting the finger to a fingerprint reader constitutes a behavioral act [42].

Physiological biometric characteristics can help identify an individual among a large pool of candidates. Under normal conditions, 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 [44]. Also, facial recognition systems that do not employ liveness detection can be fooled by an appropriately-sized photo of a legitimate user. This might pose a weakness

if facial recognition is used to unlock a smartphone. On the other hand, the failure might not be caused by the biometric characteristic itself, but due to the inadequacy of current (sensor) technology.

Behavioral characteristics constitute user actions over time, i.e., for each action, there must be a beginning, an end, and a duration. Consequently, behavioral characteristics indirectly measure properties of the human body. Behavioral characteristics are learned processes and, therefore, can be also re-learned. However, the consensus in the literature seems to be that after reaching a certain age, changes in behavior become more difficult to achieve, even with specific and sustained effort [45]. Behavioral characteristics can therefore be regarded as a valid means of identification, even though for the most part they are neither as unique nor as permanent as their physiological counterparts. An advantage is that they are less invasive and therefore more user-friendly. For example, a system that analyzes keystroke timings or speech patterns can usually do so in the background. In contrast, an iris or fingerprint scan requires specific user actions.

2.2.3 Biometric Verification and Identification

Biometric verification is the process of confirming a biometric claim by a system through a biometric comparison. It is the primary component for user authentication[3] based on biometrics and helps to verify or reject a user's identity claim based on the biometric comparison. The goal of the biometric system is to compare a sample of the provided biometric characteristics to the known reference of that user. The system then yields the decision outcome of whether the biometric probe and the biometric reference have the same biometric source. We refer to this kind of comparison as 1:1.

Sometimes, biometric authentication is used synonymously for biometric verification. If the system leverages biometric verification to decide whether an identity claim is correct or not, we refer to this as biometric user authentication.

Biometric identification differs from biometric verification, in the sense that the biometric system compares a sample from an unknown user against a database to find a stored reference attributable to a single individual and thereby identifying the user. We call this a 1:n comparison.

Biometric identification can be further divided into two types: *open-set* and *closed-set*. We say that an identification is *closed-set* if it is known a priori that the user is in the system's database, i.e., the system must find the best match from a pool of candidates. Otherwise, we refer to it as *open-set* identification.

[3] the act of proving to be of undisputed origin or veracity.

2.2.4 Requirements and Goals for Biometric Recognition

In this thesis, a novel biometric characteristic is explored. In particular, we investigate if the electrical impedance of the human can serve as a reliable method for biometric recognition. When assessing a new biometric modality and envisioning systems that use this characteristic for verification and identification, it is advantageous to consider lessons learned from past and current systems. General design goals for biometric systems can be found in the literature, e.g., [46].

The requirements and goals that we use to test against our envisaged system are described in the following:

Universal The biometric mode must be universally applicable, to the extent required by the application. The recognition method should apply to everyone intended to use the biometric system.

Unique The biometric trait must be unique within the target population.

Permanent The biometric trait must remain consistent over the period of use. Few biometric characteristics stay constant over a lifetime, but they work well if they are consistent over the lifetime of the biometric system.

Unobtrusive Biometric recognition should be as unobtrusive as possible. If users of a system that features biometric recognition can be identified passively, without requiring a lot of interaction, a biometric system is more likely to be accepted.

Difficult to circumvent Users of a biometric system should be unable to change the characteristic that is captured for biometric recognition. At a minimum, it should be difficult for a user to modify the biometric characteristic to match that of another user.

Other important, non-technical goals are:

Acceptability The biometric recognition should be one that users are likely to feel comfortable with. Clearly, acceptability is a sensible requirement. Depending on their functional principle, new biometric methods that have not established themselves for general use can raise concerns about safety. Before biometric methods are ready for widespread use, it needs to be shown that they are harmless to health, as well as their acceptability and perception must be discussed.

Cost effectiveness The relationship between distinguishing power of the biometric system and its deployment and maintenance costs. A new biometric mode requires building a prototype which is then refined to the level of a possible commercial system. In this thesis, we focus on the technical assessment of a new biometric mode and although it is premature to seek insights about deployment and maintenance costs, we can compute estimates on production and unit costs.

2.2.5 Performance Evaluation Measures

The performance of biometric authentication systems is most commonly assessed based on the false accept rate (FAR) and the false rejection rate (FRR). FAR specifies the percentage of unauthorized users incorrectly recognized as legitimate users, whereas FRR measures the percentage of legitimate users who are mistaken for unauthorized users and the system incorrectly detects the presence of an adversary.

An ideal system therefore has 0% FAR and 0% FRR, which means that no unauthorized user is accepted and no legitimate user is rejected. Unfortunately, such a system does not exist in reality and a trade-off between FAR and FRR has to be made.

Often the algorithm deployed to perform the recognition outputs a confidence value for every authentication attempt. False accept rate and false reject rate are determined by the (internal) discrimination threshold of the biometric system. The authentication attempt is only accepted if the confidence value lies above the threshold and it is rejected otherwise. The discrimination threshold can be chosen such that FAR and FRR are equal. In this case, they attain the equal error rate (EER) and the percentage of false rejects and false accepts are the same: FAR = FRR = EER. The equal error rate (EER) can be a measure to quantize the overall accuracy of the biometric system.

2.2.6 Continuous User Authentication

Most user authentication systems in use today are "one-shot" identity verifiers. Identity claims are only validated at specific points in time, for instance, every time a user logs on to the system. Unfortunately, one-time verification entails the deficit of no instantaneous authentication guarantees: in an ideal situation, the identity of the user is unaltered and can be taken for granted by the system over the entire span of user interaction. Additionally, the system should not have to make any assumptions about trustworthiness or reliability of the user.

A user authentication framework that extends beyond verification at log-in phase is therefore highly favorable in many real-world applications, especially those where a user interacts with a computer or device over a period of time.

Continuous user authentication addresses these shortcomings by adding the dimension of time to the authentication process, which can verify a user's identity seamlessly over a period of time. Often biometric verification is used for this purpose to keep the authentication process as unobtrusive as possible. For example, repeatedly entering a password during a terminal session is very likely to be perceived as obtrusive. Biometrics suitable for continuous user authentication are based on characteristics that can be measured effortlessly and remain stable during the required time span. This can be several weeks, months, or just a workday; the duration entirely depends on the biometric recognition method and the intended use.

Having a certain application scenario in mind, biometrics suitable for continuous user authentication can be divided into *integrated* and *detached* methods.

Detached biometrics can be acquired at any point in time as they are not directly coupled to the application or work flow, but they require significant user effort, e.g., swiping a fingerprint reader during a computer terminal session. *Integrated* methods allow biometric verification at successive points in time with very little to no user cooperation, e.g., face detection with a web cam. *Integrated* recognition is thought to be part of the interaction process with the system and therefore does not inconvenience the user. Compared to *detached* methods, *integrated* ones are likely to be less invasive, but they come with the drawback that they can only be acquired at times when the user performs the required action, e.g., key-stroke characteristics, that can only be used as a form of authentication when the user actually types on the keyboard.

Performance evaluation In continuous user authentication, the decision whether a user can be authenticated is based on multiple biometric readings over time, and therefore user recognition must occur in a specific (moving) time window. Consequently, measures such as false acceptance rate and false rejection rate computed on discrete biometric samples (as in the case for "one-shot" user authentication) have to be adapted for continuous authentication. This is especially true when the biometric recognition methods generate a high number of false rejects. In continuous verification, a false accept is a security breach, while a false reject inconveniences the legitimate user, who is de-authenticated by the system and re-authentication is required. Therefore, most approaches to continuous user authentication, based on biometric traits, use a metric for the similarity between

samples. Then, an aggregated score is computed to define false reject rate and false accept rate over a time window.

Another measure to assess performance of a continuous authentication system is the time between the moment an adversary decides to take over the system and the moment when adversarial presence is detected. The adversarial presence can be measured in a unit of time or the number of actions an attacker can perform before detection.

Moreover, due to the mode of biometric acquisition, continuous authentication is always bound to latency. During the time when no new biometric reading is available, the system is left unsecured and no decision about the user's identity can be made. Hence, the security of a continuous authentication system depends on the time between successive biometric measurements. The sampling time should be kept as low as possible, but not inconvenience the user, a trade-off that is especially relevant for biometric traits that have to be actively measured, i.e., detached recognition methods.

2.3 Overview of Existing Biometric Methods

Biometric characteristics, as a means of recognizing an individual using physiological or behavioral traits, have been an active research area for many years.

Despite the promising results and increasing popularity in academia, behavioral biometrics have not yet become prevalent in forensics or governmental agencies [47]. This is mostly due to their higher variability and the lack of longitudinal studies to confirm feasibility over time. As a comparison, fingerprints have been an official form of identification since the twentieth century and research has created a wealth of studies on the applicability of physiological identification over the years [48]. A comprehensive survey of established physiological biometrics can be found in [49].

Physiological biometrics do not come without drawbacks, however. While physiological biometrics tend to be relatively stable over time, they can be sensitive to deception and presentation attacks. These include, for instance, attacks on: (1) fingerprint identification, e.g., using mock fingers made of glycerin, gelatin or silicone [50, 51], (2) facial recognition, e.g., using photographs or 3D models of an actual user [52, 53], and (3) iris scan, e.g., using patterned contact lenses that replicate a genuine user's iris [54].

In contrast, behavioral biometrics are thought to be harder to circumvent. However, the performance of systems that implement behavioral biometrics, in terms of false rejection rates (FRR) and false acceptance rates (FAR), is usually

lower and can require re-calibration due to varying and often erratic nature of human behavior. Initial studies on behavioral biometrics were focused on typing and mouse movements (see [55, 56]).

In particular, keystroke dynamics gained much popularity, when it was proposed to augment password authentication [57]. Keystroke dynamics make use of the typing cadence and timings of an individual while typing on a keyboard. However, as recognition rates of keystroke dynamics greatly improve with longer sampling durations, it would be better suited to continuous authentication. Keystroke dynamics could serve as an alternative or as a complement to body impedance recognition. We compare keystroke dynamics and body impedance in Section 3.7.1.

In contrast to keystroke dynamics, some research studies on mouse movement biometrics argue that it should not be used as biometric for authentication, as it has too high intra-class variability, is highly device-dependent [58], and requires a long sampling duration, while others report high accuracies [59–61]. The authors of [61] achieved equal error rates (EER) as low as 1.3% using successive mouse actions between clicks. Some of the best results have been reported in [59] with a FAR of 0.36% and a FRR of 0%, although it is suspected that this result was influenced by recording the data on a different computer for each user [62].

An evaluation of keystroke dynamics, mouse movements, application usage and system footprint can be found in [63]. A total of 99 users participated in the study and biometric data covering 20 hours per week during a span of 10 weeks was gathered. In addition to the biometric readings, the system acquired CPU and RAM usage, and an interaction profile of the participants. The study comes to the conclusion that keystroke dynamics is most useful for continuous authentication.

2.3.1 Comparison of Behavioral Biometrics

Due to the fact that body impedance recognition—the novel biometric method we devise in this dissertation—is based on a physiological trait, although it has some of the desirable properties of behavioral traits, we provide further discussion of modalities to compare our proposed method with later on.

We cover a broad range of behavioral biometrics found in research literature and commercial products. In Tables 2.1 to 2.3, we list and compare biometrics methods alongside three high-level criteria: usability, offered security, and performance. We subdivide these three factors to reflect the requirements and design goals in Section 2.2.4. For each group of biometrics that we analyze we present examples and cite corresponding literature. We also show the performance metrics reported in those works.

The categories for the comparative evaluation in Tables 2.1 to 2.3 are further explained below:

- **(Technical) Usability** is crucial for the acceptance of biometric systems. Ideally, biometric systems are self-explanatory and do not require the user to take a particular action. We thus investigate how usable a biometric system is from a technical standpoint. We, for instance, examine how unobtrusive the acquisition is, whether the user needs to carry an object for the system to work, if recovery from loss of the biometric trait is possible, and how long enrollment and verification takes. We further examine if the characteristic supports biometric identification and what the potential costs for a deployed biometric system would be. Our assessment is purely based on technical aspects; it does not cover user acceptance and user perception.

- **Security** We investigate the security guarantees that the presented biometrics offer. In particular, we look into whether a biometric trait can be easily forged and if fake biometric material can be fabricated. Moreover, we investigate if a biometric method is computationally hard to circumvent; an important property for behavioral biometrics, as attacks on behavioral biometrics often do not require a physical counterfeit. In addition, we examine if a biometric is resilient to physical observation, i.e., an attacker cannot impersonate a user after visually or audibly observing the authentication phase, and if there are any guarantees against targeted impersonation, i.e., it is not beneficial for an attacker to impersonate a user by exploiting knowledge of personal details.

- **Biometric trait** We make a comparison based on properties inherent to the modality. We show if a biometric relies on motor-skills to recognize individuals and whether it is purely behavioral. Almost always, these two properties are mutually exclusive. Motor-skills denote the ability of humans to contract and extend muscles in order to move (part of) the skeleton. Muscles depend on functions of the brain and the nervous system. Motor-skills therefore constitute intrinsic, but mostly learned behavior. A purely behavioral biometric does not rely on muscle movement, but only on thought processes. Broadly speaking, these can be emotional states, knowledge, skills, and habits. We also investigate if the biometric method allows the continuous acquisition of readings and—if so—whether it can be integrated into the work flow (as explained in Section 2.2.6). Finally, we state the requirements in terms of hardware and/or sensors to capture a biometric reading.

Table 2.1: Comparative evaluation of behavioral biometrics

Usability							Security				Biometric trait					Performance [%]			
Unobtrusive	Nothing to carry	Recovery from loss / leakage	Verification time	Enrollment time	Identification	Cost	Resilient to physical observation	Resilient to targeted impersonation	Computationally expensive	Difficult to forge	Motor skill	Purely behavioral	Continuous measurements	Can be integrated into work flow	Hardware / sensors required	Detection rate	False accept rate	False reject rate	Equal error rate
Keystroke dynamics																			
✓	✓	✗	s	m	○	.	✗	✗	✗	✗	✓	✗	✓	✓					
SVM Based Pass-phrase Enrollment [64]															Keyboard				6.96
Free-Text keystroke dynamics [65]																			2.13
Mouse dynamics																			
✓	✓	✗	s	m	✗	.	○	✗	✗	✗	✓	✗	✓	✓					
Mouse actions between successive clicks [61]															Mouse	1.30			
Variance reduction via extractors with separate features [59]																	0.36	0.00	
(Multi-) Touch gestures and touchscreen biometrics																			
✓	○	✗	s	m	✗	.	✗	✗	✗	✗	✓	✗	✓	✓					
Swiping / Scrolling [66]															Touchscreen				4.00
Finger movements (sensor glove) [67]																	4.66	0.13	
Touch and walking behavior [68]																99.00			
Instant smartphone authentication [69]																57.00			
Eye movement biometrics																			
✓	✓	○	s	m	○	⋱	○	✗	✗	✗	✓	✗	✓	✓					
Task-independent continuous authentication [70]															Eye tracker				3.98
Eye movements during scene understanding [71]																85.70			
Estimation of familiarity levels for intrusion detection [72]																			N/A
Signature verification and handwriting																			
✗	✗	✗	s	m	○	:	✓	✗	✗	✗	✓	✗	✓	✗					
Online signature verification with Support Vector Machines [73]															Digital pen/ Touchscreen				0.41
Online signature verification on mobile devices [74]																			2.89
Offline signature verification using distance statistics [75]																	23.18	20.62	

✓ the method supports the property ○ the method supports the property partially
✗ the method does not support the property h hour(s) m minute(s) s second(s)
⋱ high cost : medium cost . low cost

Table 2.2: Comparative evaluation of behavioral biometrics (continued)

Usability							Security				Biometric trait					Performance [%]			
Unobtrusive	Nothing to carry	Recovery from loss / leakage	Verification time	Enrollment time	Identification	Cost	Resilient to physical observation	Resilient to targeted impersonation	Computationally expensive	Difficult to forge	Motor skill	Purely behavioral	Continuous measurements	Can be integrated into work flow	Hardware / sensors required	Detection rate	False accept rate	False reject rate	Equal error rate
Gait and stride dynamics																			
✓	✓	✗	s	m	✗	:	✗	✗	✗	✗	✓	✗	✓	○					
Comparison of different techniques on 4007 subjects [76]															Accelero-meter/ Gyroscope/ Camera				0.03
Voice and speech																			
✓	✓	✗	s	m	○	:	✗	✗	✗	✗	✓	✗	✓	○					
Text-independent speaker recognition (120 seconds) [77]															Microphone				2.63
Text-dependent speaker verification [78]																			0.00
Stylometry and authorship																			
✓	✓	✗	m	h	✗	.	✓	✗	✗	✗	✗	✓	✗	✗					
Text authorship [79]															CS/SP		0.10	0.00	
Authorship attribution for Twitter in $\leq$ 140 characters [80]																			0.73
Biometric sketch and graphical passwords																			
✗	✓	✓	s	m	✗	.	✗	✗	✗	✗	✓	✗	✗	✗					
Multi-factor biometric sketch authentication [81]															Mouse				7.20
Screen interface recording																			
✓	✓	✗	m	m	✗	.	✗	✗	✗	✗	✗	✓	✓	○					
Optical flows of screen recording when scrolling [82]															CS/SP		20.67	12.38	
User search patterns																			
✓	✓	✗	m	m	✗	.	✗	✗	✗	✗	✗	✓	✓	✓					
Masquerade detection and search-behavior in file system [83]															CS/SP		1.10	0.00	
User behavior patterns as seen from the operating system																			
✓	✓	✗	m	m	✗	.	✗	✗	✗	✗	✗	✓	✓	✓					
Masquerade detection based on user's tasks [84]															CS/SP		16.23	9.43	
Cognitive fingerprints																			
✓	✓	✓	h	h	✗	.	✓	✗	✗	✗	✗	✓	✓	✓					
Web-enabled cognitive fingerprints [85]															CS/SP		25.83	1.13	

✓ the method supports the property ○ the method supports the property partially
✗ the method does not support the property h hour(s) m minute(s) s second(s)
⋱ high cost : medium cost . low cost PC/SP = Computer/Smartphone

Table 2.3: Comparative evaluation of behavioral biometrics (continued)

Usability								Security				Biometric trait					Performance [%]			
Unobtrusive	Nothing to carry	Recovery from loss / leakage	Verification time	Enrollment time	Identification	Cost		Resilient to physical observation	Resilient to targeted impersonation	Computationally expensive	Difficult to forge	Motor skill	Purely behavioral	Continuous measurements	Can be integrated into work flow	Hardware / sensors required	Detection rate	False accept rate	False reject rate	Equal error rate
Behavioral web analytics																				
✓	✓	✗	s	h	✗	.		✗	✗	✗	✗	✗	✓	✓	✗					
Browsing habits and browser history [86]																Computer/ Smartphone		1.10	13.80	
Blinking																				
○	✓	✓	s	m	✗	.		✗	✗	✗	✗	✓	✗	✓	✗					
Biometric identification using song-based blink patterns [87]																Camera	82.02			
Lip movement																				
✓	✓	✗	s	m	✗	.		✗	✗	✗	✗	✓	✗	✓	✗					
Lip motions during utterance of a word [88]																Camera				8.00
Dynamic facial features																				
✓	✓	✗	s	m	✗	.		✗	✗	✗	✗	✓	✗	✓	✗					
Holistic feature extraction covering entire face [88]																Camera				48.00
Movement when answering/placing phone call																				
✓	✓	✗	s	m	✗	.		✓	✗	✗	✗	✓	✗	✗	✗					
Transparent user authentication on smartphone [89]																Smartphone		2.50	8.00	

✓ the method supports the property ○ the method supports the property partially
✗ the method does not support the property h hour(s) m minute(s) s second(s)
⋱ high cost : medium cost . low cost

In Chapter 3, we identify keystroke dynamics and touch(-screen) biometrics to be closest to the method we propose in this work. We present a detailed comparison of those biometric methods in Section 3.7.

2.4 Research Goals and Open Challenges

Having provided the necessary background, we can now formulate the research questions listed in the introduction (see Section 1.1.1) as the main goals that we attempt to achieve in the following chapters. As there is little related research, this dissertation is one of the first works to pursue those goals.

Body impedance as a unique trait We need to investigate if body impedance measurements can be used to identify individuals. To this end, we have to design a (hypothetical) biometric system that can acquire measurements and process them in order to extract user-specific features for automated recognition. Then, we can assess the results with common performance metrics used for the analysis of biometric systems. Further, we need to rigorously define how these features can be obtained from a technical standpoint.

Application scenarios for impedance-based methods It is crucial to envision and design practical applications based on body impedance to show that such a method is not just a lab experiment but can be practical. We need to implement application scenarios and evaluate their security guarantees to justify the use of body impedance as a biometric modality.

Compare body impedance with other modalities To obtain a notion of how useful the methods that we develop are, we need to compare them with other existing biometric modalities. For this purpose, we can use performance metrics as well as the high-level criteria introduced earlier in this chapter (see Section 2.3). A comparative analysis can, for instance, provide insight on what modality is best suited for a given application scenario.

The body as communication channel We know that intra-body and on-body communication can facilitate the transfer of information through the body using various techniques based on electromagnetic waves. To enable device pairing, we have to adopt a technique which can establish a channel whenever two devices are in touch with the body. For this purpose, a modulation technique and a transceiver that can send and receive information on the body channel have to be developed.

Characteristics of the body as transmission medium We need to examine how the body can be characterized as a transmission medium. To use body channel communication in a security context, we have to understand the physical properties of this channel and the circumstances under which data transfer is possible and secure. This determines whether the human body can constitute an out-of-band channel for device pairing.

Secure pairing protocol We need to develop a secure pairing protocol that makes use of the body as a communication channel. This protocol must be secure on the logical and the physical layer and facilitate the pairing of two devices that are in physical contact with the same person. For the protocol to be secure, it needs to withstand man-in-the-middle attacks and the devices should only be allowed to complete a pairing when they can conclude that they are held by the same human.

2.4.1 Open Challenges

As the methods and technologies presented in this work are novel, it is not feasible to address all open questions. We therefore mention below two challenging problems that we believe are important, but are not addressed in this dissertation.

Long-term performance assessment The mechanisms we propose need to be evaluated over a long time period to draw conclusions about sustainability. The technologies we use rely on the electrical properties of the human body which are directly coupled with the composition and shape of the body. These inevitably change throughout one's lifetime. As a consequence, a long-term assessment is needed to verify continuity of our proposed methods. This is especially true for user recognition.

Refinement of the proposed technology The security technologies presented in this thesis need to be refined to the extent where they could become viable as a consumer product and go beyond a feasibility study. This process would, for instance, require engineering challenges and user acceptance studies that are not within scope of this work.

3

Body Impedance as a Biometric Modality

This chapter documents the research that has been conducted to test if body impedance can serve as a biometric modality. It presents two applications for user authentication where body impedance is a particularly well-suited biometric trait.

Contents

3.1 Introduction

Many modern access control systems augment traditional two-factor authentication (something you know and something you have) with a third factor: "something you are", i.e., biometric verification. The biometric characteristics which can be used as an additional layer of security come 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 high security 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 thesis, we propose a new biometric characteristic based on the human body's electrical impedance as described in Section 2.1.1. We develop a method which can be used to measure the body's impedance as the response to a square pulse signal. For this new biometric recognition method, we consider two motivating scenarios:

The first is an access control setting where this new biometric mode is used as an additional layer of security when a user enters a PIN, e.g., into a bank ATM. Biometric recognition based on body impedance (body impedance recognition in short) facilitates unification of PIN entry and biometric capture. We use PIN entry as a running example for this scenario throughout this chapter. This is because PIN pads (e.g., in ATMs) are often made of metal, which makes capturing the biometric trait 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. This conductive pad would transmit the required signal for the measurement and a sensor in the PIN pad would capture the response.

The second scenario corresponds to continuous user authentication (see Section 2.2.6) at a stationary computer terminal, e.g., verifying that the user, who logged in earlier, is the same person currently at the keyboard. For this scenario, we

need a mechanism that periodically samples one or more biometric characteristics. However, for obvious reasons, this should ideally be done unobtrusively. Body impedance recognition is particularly well-suited for this setting. Assuming that it can be made from a conductive material, the keyboard would generate the pulse signal and measure the electrical response, while the user (remaining oblivious) is typing. The main idea is that the user's body impedance is captured at login time and identity of the person currently at the keyboard can be verified transparently, at desired frequency.

Continuous user authentication at a stationary computer can be a challenging problem to solve using static biometric modalities. For example, if swipe fingerprint sensors are used, the user of such an authentication system would have to periodically stop and swipe a finger on the scanner, which could be disruptive. Such recognition process is *detached* from the actual work flow. Less obtrusive approaches try to solve this problem using automated video monitoring or continuous face recognition [90, 91] with a video camera. However, depending on the context, such systems can be perceived as invasive. Unlike many static traits, behavioral characteristics can allow for very noninvasive continuous authentication, most notably keystroke timings and mouse dynamics [92]. By continuously measuring and quantizing the interaction with mouse and keyboard one can verify if an originally logged in user is still present at the computer terminal, or if someone else took over an open session. We present a possible solution to continuous authentication that is equally transparent and unobtrusive as a keystroke dynamics but is based on a physiological trait.

To assess efficacy and feasibility of body impedance recognition, we build a prototype platform for gathering impedance data. Its main purpose is to assess whether we can identify users from a population of test subjects. The same platform can test the distinguishing power and stability of this trait over time. We also explore two hypothetical systems that apply this new recognition technique to the two sample scenarios discussed above: one to unobtrusively capture the biometric characteristic for an additional layer of security when entering a PIN, and the other to implement continuous user authentication at stationary computer terminal.

3.2 Combining PIN Entry with Body Impedance

This section describes the envisaged use of body impedance recognition to unobtrusively enhance the security of PIN entry systems.

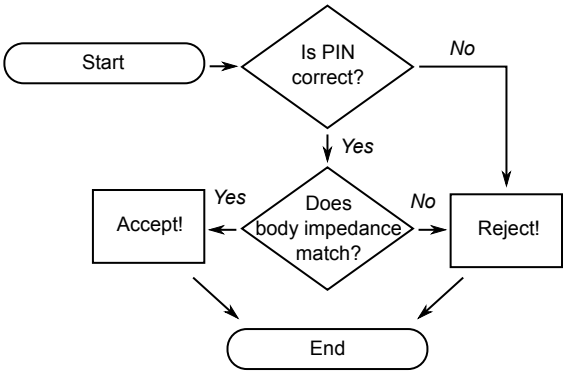

Figure 3.1: ATM decision flowchart for PIN entry scenario.

3.2.1 System and Adversary Model

We use a running example of a metal PIN key-pad with an adjacent metal pad for the user's other hand. The key-pad has the usual digit (0-9) buttons as well as an "enter" button. It also has an embedded sensor that captures the transmitted signal by the adjacent metal pad. This set-up corresponds to a bank ATM or a similar setting.

The adversary's goal is to impersonate an authorized user and withdraw cash. We assume that the adversary cannot fool body impedance recognition with a certain probability that we determine in our experiments described in Section 3.5.

We also assume that the ATM is equipped with a modified authentication module which, besides verifying the PIN, captures body impedance and performs biometric verification, i.e., computes the likelihood of the measured impedance corresponding to the user identified by the previously inserted ATM card and the just-entered PIN. This module works as depicted in Figure 3.1. We assume that the ATM has access to a biometric reference database of valid users, either locally or over a network. Alternatively, the user's ATM card can contain a biometric reference needed to perform impedance 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 body impedance signature.

3.2.2 PIN Entry Scheme

The ATM has to determine whether a biometric sample acquired from the user while entering the PIN is consistent with the reference in the database. This requires a classifier that yields the likelihood of a sample coming from a known distribution. The likelihood is used to determine whether the newly measured samples are close

enough to the reference in the database to produce a match. Using our prototype in Section 3.4.6, we test if we can make such decisions with high confidence.

Before discussing security of the body impedance enhanced PIN entry system, we check whether it meets our requirements stated in Section 2.2.4.

Universal. A person using the modified PIN entry system must use both hands, one placed on the metal pad and one to enter the pin. This requires the user to have two free hands (the user cannot be carrying a wallet, for instance). In contrast, a normal PIN entry system can be operated with one hand. Thus, universality of our system is slightly lower. This is a limitation of the biometric mode, although for people with a disability, a flag could be stored on the user's ATM card and thereby exempting this person from the impedance verification. This would allow our approach to gracefully degrade to a generic PIN entry system.

Unique and *Permanent.* In Section 3.5, we show that our prototype can determine, with high probability, whether a subject matches a specific body impedance signature. Therefore, it is unlikely for two people to exhibit exactly the same body impedance. We also show that an individual's body impedance remains fairly consistent over time.

Unobtrusive. In the envisaged setting, the scheme is very unobtrusive, since from the user's perspective, the the current operation is unchanged apart from the added requirement to place the free (not used for PIN entry) hand on a metal pad. Naturally, some users might have to change their behavior while operating an ATM, as they could be used to holding something in one hand, e.g., their wallet, or shielding their PIN entry. However, shielding could be provided for by such a modified ATM. Also, there could be two conductive pads accommodating both left- and right-handed people. In addition, the ATM screen could display system usage instructions, even pictorially to serve people who cannot read. Similarly, audio instructions could be given for the sake of those who are vision-impaired.

Difficult to circumvent. Given that body impedance 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 hard to test precisely, assuming that the adversary is unaware of the target user's impedance measurements, the task seems difficult. We will discuss biometric attacks in Section 3.6.

3.2.3 Security Analysis of PIN Entry Scheme

The additional layer of security provided by body impedance recognition is completely independent from security of the PIN entry system alone. Therefore, we model the probability P_{break} that the proposed PIN entry system can be subverted, as:

$$P_{break} = P_{guess} \cdot P_{successful-impostor}$$

where P_{guess} is the probability of an adversary correctly guessing the PIN and $P_{successful-impostor}$ is the average probability that the adversary can fool body impedance recognition by presenting his own biometric characteristic. In Section 3.5, we determine the false accept rate to be 9% on average for a zero-effort impostor, i.e., $P_{successful-impostor} = 0.09$.

If a PIN consists of n decimal digits and the adversary has t guesses then $P_{guess} = \frac{t}{10^n}$. Together with $P_{successful-impostor}$ this yields the combined probability:

$$P_{break} = \frac{0.09 \cdot t}{10^n}$$

For example, if the adversary is allowed 3 guesses with a 4-digit PIN, $P_{break} = 2.7 \cdot 10^{-5}$, whereas a 4-digit plain-PIN system has a subversion probability of $3 \cdot 10^{-4}$. Though this improvement might not look very impressive in numbers, it is well known that a lot of PIN attacks are performed by "shoulder surfing" or covertly video-taping the PIN entry sequence (often called ATM skimming). These attacks do not involve the adversary guessing the PIN. If we assume that the adversary already knows the PIN, $P_{break} = 9.0\%$ with our system, as opposed to 100% without it.

3.3 Continuous User Authentication

We present a continuous authentication scheme based on body impedance recognition. 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, body impedance recognition is no longer used as an additional layer of security at login time. Rather, the user's impedance is captured at login time and subsequent measurements are used to authenticate the user by validating against the initial reference.

We continue using the example of a stationary computer terminal throughout this section to make it easier to present the details of the envisioned system. However, applicability of continuous authentication based on body impedance is not limited to this specific scenario. If the measurement apparatus and the electrodes needed to acquire impedance readings are miniaturized, smaller devices such as laptops and smartphones are imaginable.

3.3.1 System and Adversary Models

The system consists of a terminal with a special keyboard that sends out pulse signals and captures body impedance. This requires the keyboard to be either made from, or coated by, a conductive material. Alternatively, the signal transmitter could be located in a mouse that the user operates with one hand and the keyboard captures the impedance. Without loss of generality, we assume the former set-up. The keyboard otherwise operates normally and is used for both login and routine activity at the terminal.

We assume that the adversary, with or without consent of the authorized (at login time) user, physically accesses the unattended terminal and attempts to proceed within an already-open session. In security research literature, this attack scenario is sometimes referred to as "lunchtime" attack (see, e.g., [93]). We assume that the adversary at the keyboard has full access to the active session, and that the attack happens some time after the original user logged in. Our goal is to detect that the original user is no longer present, and that the keyboard is 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 system administrators.

In addition to the peripherals required to capture body impedance, the continuous authentication system consists of a software process that manages initial login and frequency of periodic reacquisition of the biometric characteristic. This process is also responsible for displaying user warnings and reacting to suspicious behavior. We refer to it as the *continuous authentication process* (CAP) and assume that neither the legitimate user nor the adversary can disable it.

3.3.2 Continuous Authentication Scheme

At login time, CAP measures and captures an initial impedance measurement of the authorized user. Periodically, e.g., every few seconds, CAP reacquires the biometric characteristic by sending and receiving a pulse signal through the keyboard. Each newly acquired measurement is checked against the reference captured at login. If the new measurement is sufficiently distinct from the one sampled from the original user at login time, CAP consults its policy database and takes appropriate actions, as discussed above. Figure 3.2 shows a sample CAP decision flowchart. The decision policy can be further refined. For example, in a corporate setting, all employees could have their biometric reference stored in central database to allow for more sophisticated access schema which also include shared resources or devices.

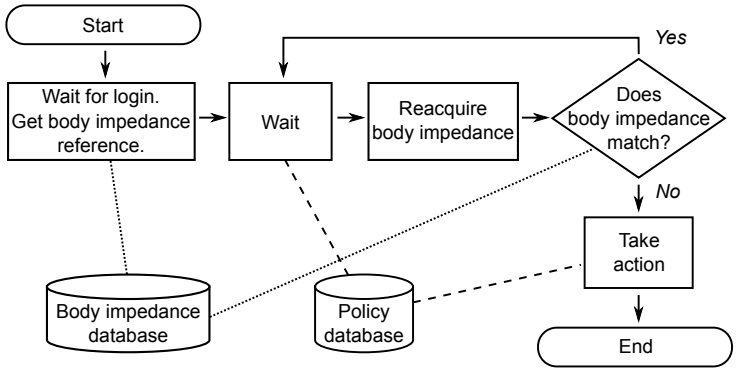

Figure 3.2: Flowchart of the *Continuous Authentication Process* decision procedure.

Before considering the security of the presented continuous authentication system, we verify if our design possesses the desired properties of biometric systems (see Section 2.2.4):

Universal. The users of the system must have two hands in order for their body impedance to be captured. The same arguments, as in the case of PIN entry, apply here as well.

Unique and *Permanent.* In Section 3.5, we show that our prototype can match a body impedance signature to previous samples (taken immediately beforehand) with very high accuracy. Average equal error rate is as low as 2%. The fact that the reference measurement 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. Provided the users of the envisioned system periodically come in touch with the electrodes that emit the measurement signal and measure the impedance, they do not need to modify their behavior and user burden is minimal. In case the electrodes are embedded in a conductive keyboard, this would mean users need to type with both hands. For users who consistently type with only one hand, at least one electrode would have to be incorporated elsewhere, e.g, into the computer mouse the user operates.

Difficult to Circumvent. With a false accept rate of 2% (at equal error rate) it is unlikely that the adversary happens to have an impedance similar to the original user and can manage to continuously fool body impedance recognition. We explore this point further in the security analysis (see Section 3.3.3 below).

3.3.3 Security Analysis of Continuous Authentication

The adversary's goal is to subvert the continuous authentication system by using the secure terminal after the original user has logged in. In the analysis below, we

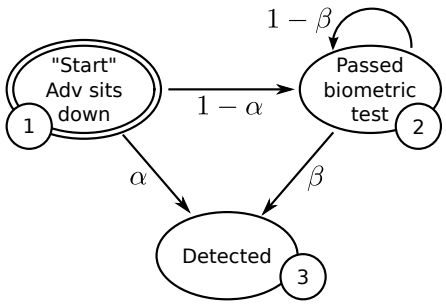

Figure 3.3: Markov model of the continuous authentication detection probability. States are numbered 1 to 3 for easy referencing in text.

assume that the authorized user colludes with the adversary by leaving his session open for the adversary to take over. This eliminates any uncertainty that results from the original user catching the adversary using the terminal, which is hard to model accurately. We consider the worst-case scenario and the detection probability is a lower bound on the security provided by the continuous user authentication system. The exact values for the parameters we use in the security analysis are estimated through experiments (based on our data set) that reflect the worst case our proposed scheme could encounter.

We model the security of the continuous user authentication scenario with two probabilities. The first is the probability that the adversary is detected immediately, i.e., the very first time when his body impedance is measured. This corresponds to the complement of the average false accept rate that we report in Section 3.5 and we call this probability α in the following calculations.

If the adversary's biometric characteristic 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 characteristic is very close to that of the original user. Thus, the probability that the adversary is detected in subsequent measurements is lower:

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

We call this decreased probability β. In Section 3.6.2 we will calculate an experimental lower bound for β based on our gathered data set of body impedance measurements. We measure false acceptance rate in the worst case, i.e., the probability of a successful impersonation attempt for the most promising attacker-victim combination in our data set.

We define a simple Markov model, shown in Figure 3.3, with three states to calculate the probability that the adversary is detected after i rounds. When the

adversary first accesses the keyboard, it is either detected with probability α or *not* detected, with probability $1-\alpha$. In the latter case, its body impedance measurement 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:

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

Each row and each column in P corresponds to a state. The entry in row q and column r, p_{qr}, is the probability of transitioning from state q to state r. To find the probabilities of each state we start with a row vector ρ that represents the initial probability of being in state 1, 2 and 3. Clearly, $\rho = [1, 0, 0]$, 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 $\rho \cdot P$. Probabilities for each subsequent round are determined via another multiplication by 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, since the first state transition forces a transition from the initial state and there is no way back to state 1 (see Figure 3.3). The probability of being in state 2 (i.e., to escape detection for i rounds) is given by the second element of ρ: $(1-\alpha)(1-\beta)^{i-1}$. The probability of detection is thus: $1-(1-\alpha)(1-\beta)^{i-1}$. According to this model, using $\alpha = .98$ and $\beta = .36$ (numbers from our experimental results in Section 3.5 and Section 3.6.2) there is a chance of 98% or more that the adversary is detected after 10 rounds. Not surprisingly, as long as $\beta > 0$, the acquisition frequency largely determines the time to detect the adversary.

Detection could fail in case of a systematic error, i.e., $\beta = 0$, where the distribution of impedance samples the adversary provides is too close to the one of the authorized user. However, we did not observe an attacker-victim combination in our fair-sized data set, which would allow an attacker to constantly impersonate a victim (i.e., $\beta = 0$) and would require extending our model.

3.3.4 Handling False Rejects

False rejects refer to incorrect detection of adversarial presence. If biometric recognition is used as an additional layer of security for user authentication (e.g., the PIN entry scheme in Section 3.2), this can be managed simply by restarting the login procedure, if the first attempt fails. However, in a continuous authentication setting, where a single (and possibly incorrect) detection might cause the system to lock up, false rejects have to be handled more thoughtfully.

One approach is to specify a policy that allows a certain number of detection events every n-th round, without taking any action. Such a mechanism can help mitigate sensor reading errors or short-term environmental changes that could adversely affect body impedance recognition and change the reading, such as overly sweaty hands or arms/hands accidentally touching each other or any non-involved metal object.

Another option is to integrate potentially less user-friendly and more detached biometric recognition to deal with ambiguous events. For example, after a few detection events, the user might be asked to confirm his identity by swiping a thumb on a fingerprint scanner. Such a combined approach would be suitable for a system with very high security requirements. Deploying body impedance recognition could drastically reduce authentication requests from the principal biometric method which might be more obtrusive.

3.4 Biometric Capture of Body Impedance

In this section, we present a method to capture body impedance measurements and realize biometric recognition based on this characteristic. We measure the horizontal impedance of the body, i.e., from one hand to the other, using two capacitive electrodes (see discussion in Section 2.1.1). Such a configuration is applicable to many use cases including the two presented in the previous section: the PIN entry scheme and continuous user authentication.

We first give a brief overview of how our proposed way of capturing body impedance works, we discuss potential ethical concerns, and then highlight some of the design choices that led to this particular method. More specifically, we describe parameters for our prototype set-up that allow us to measure body impedance reliably. We conducted several experiments to test different signal types, voltage levels and frequencies.

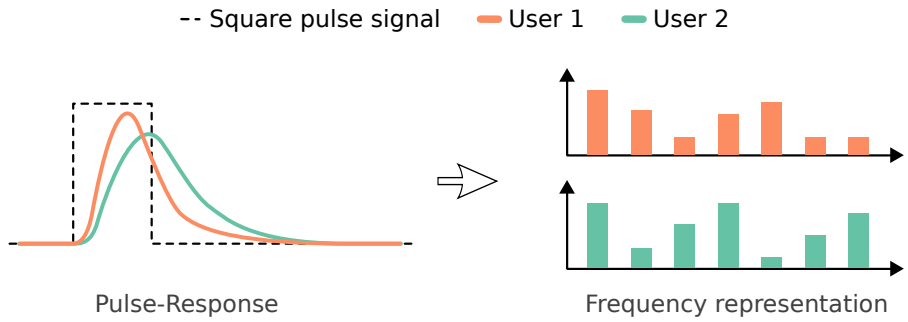

Figure 3.4: Overview of pulse-response recognition. The electrical response is captured and transformed to the frequency domain. Each individual has a distinct pulse-response due to differences in body impedance.

3.4.1 Pulse-Response Recognition

The method we devise is called *pulse-response* recognition or just *pulse-response* in short. The name originates from the fact that it works by applying a low voltage pulse signal to the palm of one hand and measuring the body's electrical 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 emitted signal is modified on its way by the impedance of the body and is transformed to a different signal which can be captured in the user's hand and thus directly reflects the horizontal body impedance. We correlate the injected pulse signal (which is also measured) with the response signal to remove the effect of different current levels. The result is then transformed to the frequency domain via the Fast Fourier Transform (FFT). This transformation yields the individual frequency components (bins) of the response signal, which form the biometric features that are fed to a classifier. Working in the frequency domain has the advantage that there is no need for aligning the pulses when they are measured and the exact electrical impedance can be determined at different frequencies.

The main reason for the ability of pulse-response to distinguish between users is due to two facts: (1) the captured signal contains the person's body impedance and (2) there are subtle differences in body impedance, at different frequencies, among different people. We explained in Section 2.1.1 that body impedance is very likely different enough for every individual such that it can be measured and exploited for biometric recognition. With pulse-response, when a signal is applied to one palm and measured in the other, the current travels through various types of body tissues—blood vessels, muscle, fat tissue, cartilage and bones—to reach the other hand. Differences in bone structure, muscle density, fat content and layout (and size) of blood vessels result in slight differences in the attenuation of the signal at different frequencies.

These differences show up as variations in the magnitude of the frequency bins after the FFT, and this is what facilitates distinguishing among individuals. Figure 3.4 illustrates the concept of how pulse-response recognition measures body impedance and depicts the differences in pulse-response for two different users.

The challenging task is to find those frequencies where body impedance (and thus the response signal) differs the most, but also proves to be stable over time. Only this will facilitate reliable recognition. We decided to determine these key parameters experimentally, which requires experimenting with electricity on live test subjects. In order to do this safely, we need to carefully design the experiments in such a way that they do not pose a threat to human health and we need to address any ethical concerns.

3.4.2 Ethics and User Safety

As mentioned above, capturing a person's body impedance naturally requires electricity. In case of pulse-response, this entails applying a low voltage pulse to one hand of the user and measuring the resulting signal in the other. The voltage will create a current flowing through the human body and this process naturally raises questions about user safety and ethics. Clearly, these are important issues that we must address. The issue of safety might be compounded by users having undocumented or undisclosed medical conditions, including implantable medical devices, e.g., pacemakers, that may be adversely affected by applying an external signal to the body.

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, overall conductivity is significantly higher (i.e., resistance is 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 1000 and 5000 Ω. However, even in very extreme conditions, resistance does not drop below 500 Ω.

It is therefore important to keep the current limited even in a worst case scenario where a subject has very sweaty hands and/or broken skin. With our current limiting resistor of $10k\Omega$ on the signal generator, the worst case current (with 10V test signal) is $10V/10.5k\Omega = 0.95\ mA$, which is below the sensitivity limit. The vast majority of subjects were only exposed to a 1V signal, which translates into the worst case current strength of $0.095\ mA$, less than the current flow induced by touching the terminals of a standard 1.5V battery.

Such a current is on the order of what medical devices and commercially available body-fat scales emit. Body-fat scales determine body impedance at predefined frequencies (usually 50 kHz) by sending an alternating current of up to 0.5 mA through the body. They then estimate body fat percentage based on the measured impedance and additional information such as body height and body weight.

Since body composition monitors are indented for daily use and the methods presented in this work require similar or weaker current strengths, we believe our experiments are also harmless, even if conducted over an extended period of time. We point out, however, that there exist no published studies on the long-term effects of the electrical fields created by such low currents, and therefore our experimental set-up is only safe according to the knowledge available to this day.

All subjects part of our study were given detailed information about the nature of the experiment beforehand and were then asked if they would like to opt out or continue. None expressed any discomfort or, in fact, any perception of the current during the experiments.

Our experimental prototype set-up and its safety and methodology have been reviewed and authorized by the Central University Research Ethics Committee of the University of Oxford, under approval reference MSD-IDREC-C1-2014-156.

3.4.3 Signal Type

Starting out with the hypothesis that body impedance varies across subjects depending on frequency and voltage level of the signal, we conducted a preliminary study to test the distinguishing power of various frequency sweeps, and pulse signals with different widths. Although the sweep signals cover a broad range, short square pulse signals prove to be strongly unique among our test subject population.

The box plots in Figure 3.5 summarize our biometric comparison results with four classifiers that performed well in our application: Support vector machines (SVM), Euclidean distance, linear discriminant analysis (LDA) and 3-nearest neighbors (3-nn). The most promising of the pulses (Pulse), linear sine sweeps (SineLin) and linear square wave sweeps (SquareLin) are listed on the x-axis. The signal name is composed of a signal type, a voltage and a maximum frequency (or width for pulses), i.e., $Type - Voltage - Freq/Width$. The voltage is either 1, 5 or 10 volts. Starting frequency is 1 Hz for all sweeps and the stopping frequency is 250, 500 or 980 Hz, respectively. The width of the pulses (given in hundreds of nanoseconds) is either 100 ns, 10 μs or 1 ms. The y-axis shows the binary detection error rate, i.e., the amount of times the classifier failed to identify a biometric sample correctly,

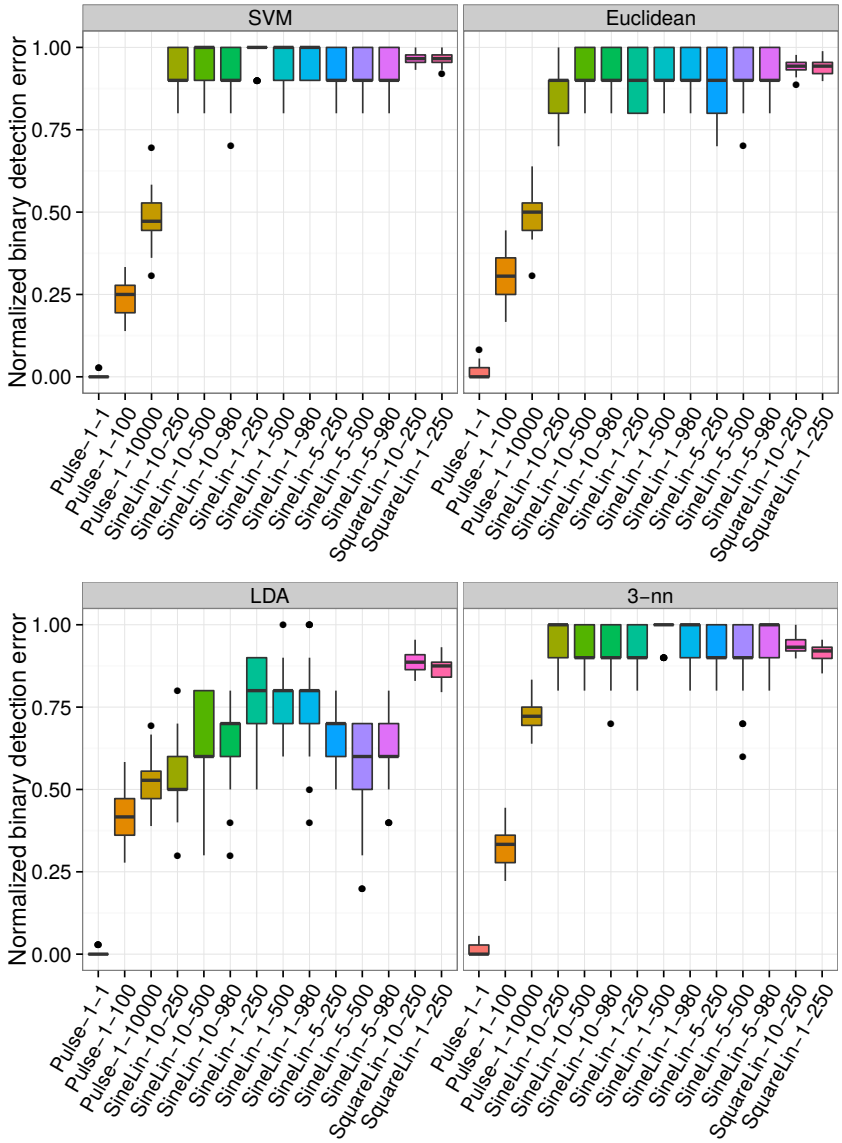

Figure 3.5: Box plots of the binary detection error rate for four different classifiers. The distribution shown by each box plot is the result of applying stratified 5-fold cross-validation to the data set five times in a row. We test several different signal types, voltage levels and frequencies for each classifier. We see that narrow pulse signals are consistently performing well.

normalized by the number of samples. The distribution denoted by the box plots shows the results of the classifiers achieved by five times 5-fold cross-validation.

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. This is due to the fact that the wider the pulse is, the narrower and more constrained the spectrum, i.e., the Fourier Transform, and therefore the response signal contains fewer relevant frequencies.

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.

Besides shape and form of the signal, voltage levels are an important factor to consider. It is important that the acquisition is safe and the users of our system do not experience any discomfort when their biometric information is captured. 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, 10 Vpp and 5 Vpp provides better separation between the subjects due to lower noise levels. For example, in Figure 3.5, using the LDA classifier, we see that the *SineLin-5-500* signal has a lower detection error rate than the *SineLin-1-500* signal, but the former has high variance. For pulse signals, we did not observe significant correlation with voltage level. Since the pulse signal is clearly the best choice for our biometric characteristic, we chose 1 volt pulses to keep the current levels at a minimum.

3.4.4 Signal Frequency

A possible reason for why short square pulses outperform the tested frequency sweeps is because they contain a much wider spectrum and therefore, they elicit a response from the body's impedance at a broad frequency range.

The Fourier Transform of a square pulse is the theoretically infinite spectrum $G(f) = A \cdot W \cdot \text{sinc}(f \cdot W)$, where A is the amplitude, W is the width of pulse, and f the frequency. The first zero-crossing of $G(f)$ is at $f = 1/W$, which means that for a pulse width $W = 100$ ns, frequencies of 10 MHz and beyond are represented in the signal.

In Chapter 5, we show that certain frequency sweeps can be an alternative to pulses and achieve similar performance to pulse-response recognition, albeit at lower acquisition rate due to the time it takes to generate the sweep. A short pulse has the advantage of a wide spectrum and can be generated in a fraction of a second.

3.4.5 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. Approaching the classification as a time series clustering problem, there are many known methods 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 dimensional point and classified according to the euclidean distance to the centroid of each class. This classifier is conceptually very simple but offers reasonable 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 improved significantly from the Euclidean distance metric, suggesting that the shape of each class should be taken into consideration.

Support Vector Machine (SVM) For each pair of groups, we train one binary Support Vector Machine classifier (one-against-one approach). 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 when pulse-response recognition is used for biometric verification.

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. This classifier turns out to be especially useful for biometric identification. It does however not prove as powerful as the SVM classifier for the binary classification task of biometric authentication.

k nearest neighbor (k-nn) We test 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 k-nn is still not as robust as SVMs or LDA, respectively.

3.4.6 Proof-Of-Concept Measurement Set-up

In order to gather stable and accurate pulse-response measurements we build a data acquisition platform consisting of: (1) an arbitrary waveform generator, (2) an oscilloscope, (3) a pair of brass electrode handles, and (4) a desktop computer to control the apparatus. Figure 3.6 shows a photo of our set-up. We use an Agilent arbitrary waveform generator as the source of the pulse signal. Flexibility of the waveform generator is useful for experimenting during the initial design phase and allows us to determine the required pulse waveforms for the final application. To measure the pulse waveform before and after the signal passes through a test subject, we used an Agilent digital storage oscilloscope which allows storage of 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 signal input terminal. When a test subject holds one electrode in each hand the signal travels from the generator through the body and into the oscilloscope. To ensure exact triggering, the oscilloscope is connected to the synchronization output of the waveform generator. The oscilloscope also records the output signal of the waveform generator for reference.

We use polished brass hand electrodes to ensure optimal electrical contact between the measurement set-up 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.

When measuring body impedance we make each subject follows a specific procedure to ensure that only minimal noise is introduced into the measured data. The test subjects are given a brief explanation of the set-up 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 can choose to either stand or sit in a chair while holding the electrodes as long as they do not touch the sides of their body with their elbows or upper arms. We do this to ensure that the current of the pulse signal has to go through more or less the same path, for all samples and all users. Before each new test subject is measured, the brass handles are wiped down with a disinfectant, both for hygienic reasons and to ensure good electrical contact between the electrode and the user's palms.

While our prototype set-up ensures accurate biometric measurements and shows feasibility of body impedance based recognition, it might not translate directly to

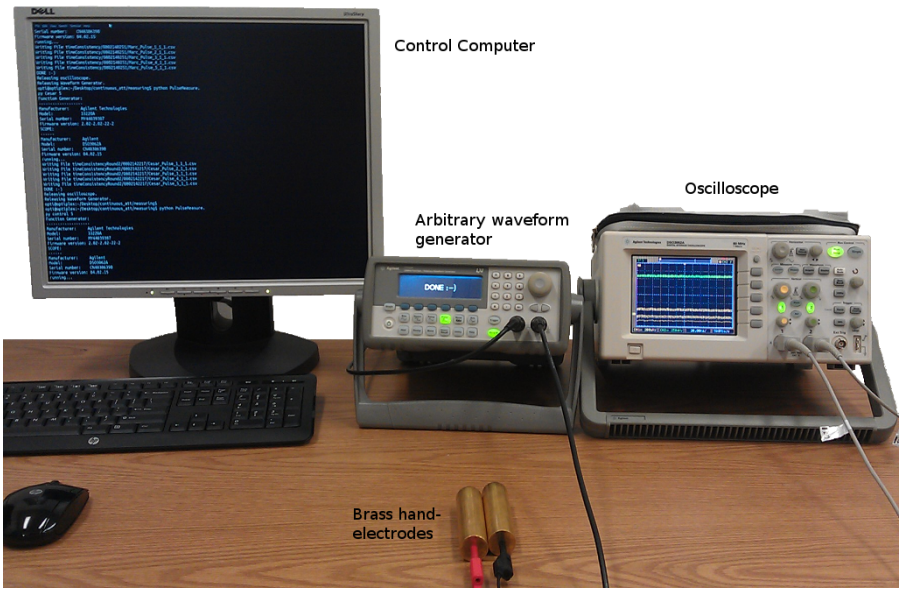

Figure 3.6: Proof-of-concept measurement set-up. The test subject holds two brass electrode handles [94] and the pulse signal is generated by an Agilent 33220A (20 MHz) arbitrary waveform generator. The receiver is an Agilent DSO3062A (60 MHz), 1 GSa/s digital storage oscilloscope.

the described application scenarios of PIN entry and continuous authentication in terms of electrode design and other ergonomic requirements. Obviously, further practical tests would be needed before deploying body impedance recognition, as to find out to what degree soiled electrodes or sweaty hands have an effect on the biometric reading. In Chapter 5, we propose to use body impedance recognition for driver authentication and we determine that movement and change in posture do effect body impedance measurements, but only marginally.

3.4.7 Test Subject Population

In the initial design phase, each test subject was sampled ten times for each of the different signal types, for each voltage level and for various frequencies. Once we selected the pulse signal with the best results, samples were acquired for two data sets.

The first consists of 20 samples for each subject, taken in one measuring session. A total of 30 people were measured for this data set, including 9 women and 21 men. We call it the snapshot data set.

The second data set includes 25 samples per subject from a total of 16 subjects, obtained in five different sessions over time. To assess stability and permanence of impedance measurements, we measured the biometric over a longer period of time.

Table 3.1: Test subject population and sample size of body impedance study

Data set	Test subjects	Females	Males	Samples per subject
Snapshot	30	9	21	20
Over-time†	16	2	14	25

The age band of the subject population ranges from 24 to 38.
†Test subjects were measured in five different sessions over time.

We sampled all test subjects at different times during the day, over the course of several weeks. The median time span between consecutive sessions was 8 days and there was a minimum time interval of at least one day between sessions.

We tried to sample each subject in conditions as diverse as possible in order to capture various other potential factors that might influence body impedance, such as varying body water percentage, body temperature or time of the day.

Table 3.1 summarizes the composition of the test subject population.

3.4.8 Feature Extraction

Data extracted from the measurement set-up is in the form of a $4,000$ sample time-series describing voltage variation as seen by the oscilloscope. Figure 3.7a shows the input pulse sent by the waveform generator and the pulse measured by the oscilloscope. Time series measurements are converted to the frequency domain using the FFT and the magnitude of the first 100 frequency bins are used for classification. Operating in the frequency domain has several advantages. First, there is no need to worry about alignment of the measured data pulses when computing metrics, such as the euclidean distance between pulses. Second, it quickly became apparent that lower frequency bins carry more distinguishing power. Higher frequency bins contained more noise, meaning that the FFT can be used to perform dimensionality reduction of the original $4,000$ sample time-series to a vector of 100 FFT bins. Figure 3.7b shows an example of the raw data we end up with after correlating the two signals and applying the FFT. This data is then fed to the classifier.

3.5 Experimental Results

In this section we present the results of our experiments with pulse-response recognition, a narrow pulse signal, that resulted from our analysis as the final biometric characteristic. The design decisions and motivations behind selecting a short square pulse signal are described in detail in the previous section. We

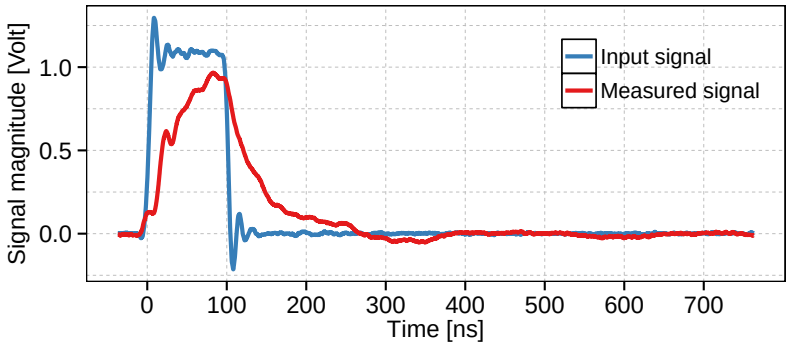

(a) Input and output waveforms

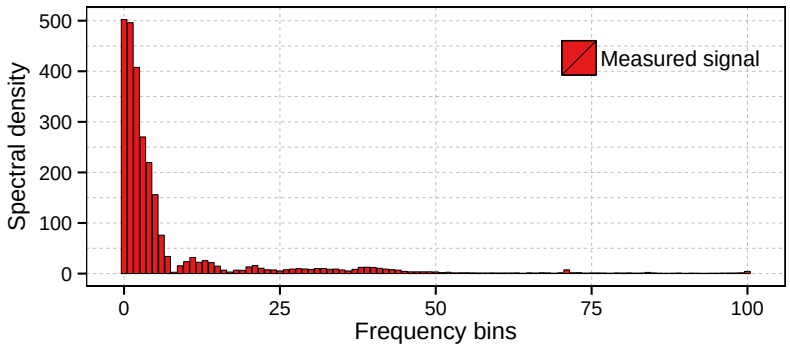

(b) Raw FFT data of the measured pulse

Figure 3.7: One measurement consists of 4,000 samples with the rate of 500 MSa/s. In Figure 3.7a it is apparent that the measured pulse has been modified by passing through the user. The FFT data shown in Figure 3.7b consists of the first 100 frequency bins after computing the FFT of the correlated signal.

report system performance figures of various classifiers when they are applied to pulse-response recognition. To be precise, we divide the results into two different types of classifiers according to the usage scenario of the biometric recognition. We present classifiers for verification and for identification.

We sub-divide the results into the two underlying data sets: (1) those from the snapshot data set, which show the inherent distinguishing power of pulse-response recognition, and (2) those based on the data sampled over time, which assess stability (permanence) of the measurements.

Within our data set and due to our straightforward feature extraction, we did not experience any failure to enroll or failure to capture errors, which means the classifier performance corresponds to the system performance of our prototype set-up. We therefore do not report classifier and system performance separately.

3.5.1 Performance Metrics

We use the metrics introduced in Section 2.2.5 to assess system performance of our prototype set-up: false accept rate (FAR) and false reject rate (FRR).

To illustrate the FAR and FRR graphically we draw the receiver operating characteristic (ROC curve) which shows the relationship between these two performance numbers. The ROC curves shown in the following are vertical averages. We compute a ROC curve for every test person and calculate the average over all false reject rates for given false accept rates (see [95] for an algorithm on vertical aggregation of ROC curves).

We also use the equal error rate (EER), a common performance metric for biometrics. It denotes the rate at which errors for acceptance and rejection are equal and provides a straightforward way to compare different ROC curves. Equal error rates for the best-performing classifiers will be presented in Section 3.5.

To assess performance of pulse-response recognition in identification, we compute the ranking success Rank(N). The ranking success is a metric that measures the ratio of query samples for which the corresponding reference measurement is among the first N references out of all stored biometric reference measurements in the database. The measurements are assumed to be sorted in decreasing order according to their assigned similarity values. Ideally, Rank(1) = 1.0, which means that for all query samples the corresponding reference from the database has been assigned the highest similarity value.

To obtain unbiased and realistic performance measurement numbers, the data sets are partitioned into learning set and test set. We make sure that the test set for the over-time data spans all five measurement sessions. For both data sets, the partitioning into training and test set is repeated multiple times by stratified cross-validation to acquire a robust estimate of the performance of the biometric modality.

3.5.2 Biometric Verification

Biometric verification is a binary classification task. The classifier has to decide whether a presented sample belongs to the group of samples reflecting a specific user or not. A verification classifier for pulse-response recognition is used in the example of Section 3.2 where PIN entry is combined with impedance measurements and the example in Section 3.3 where impedance measurements facilitate continuous authentication.

To simulate an authentication procedure with pulse-response recognition, we separate the samples into two classes: Samples belonging to the legitimate user

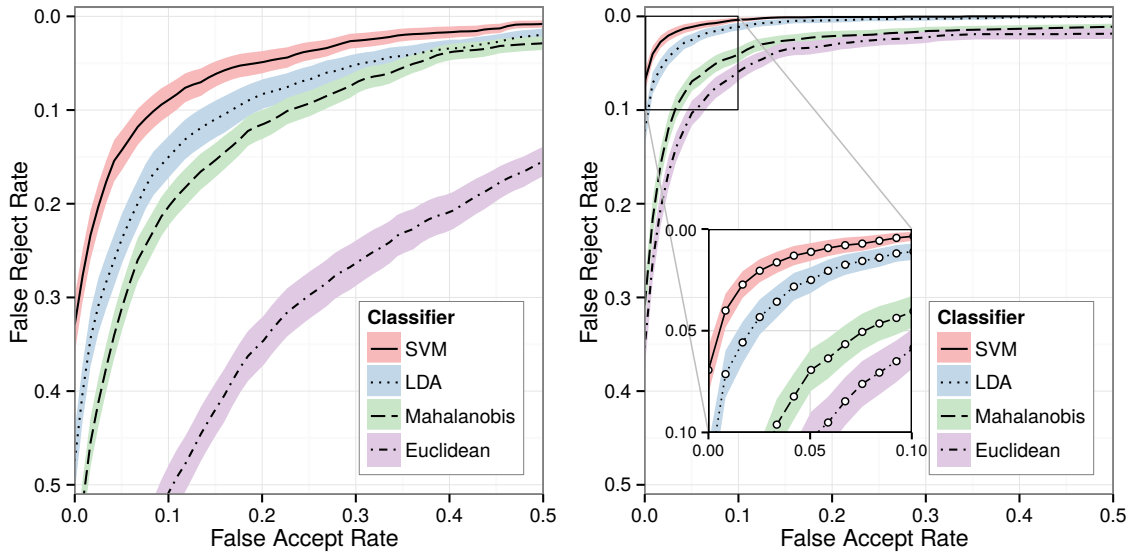

(a) Over-time data set **(b)** Snapshot data set

Figure 3.8: Receiver operating characteristic for authentication. The results presented are averages over all users and obtained by applying 5 times stratified 5-fold cross-validation (ROC curves are vertical averages). Shaded areas show the 95% confidence interval for each classifier.

and samples from all other users. Samples from other users are collected in a large pool and represent potential impostors. Once the classifier is trained, it is presented with unseen samples from both classes. Then, FAR and FRR are computed on the basis of the classifiers' prediction.

In order to solve the binary classification problem, we test four of the classification algorithms described in Section 3.4.5: Support Vector Machines (SVM), Linear Discriminant Analysis (LDA), Mahalanobis distance (Mahalanobis) and Euclidean distance (Euclidean). Figures 3.8a and 3.8b show the performance of each of these methods when applied to the over-time data set and the snapshot data set, respectively. The depicted ROC curves are averages over all test subjects and describe the relationship between the FRR on the y-axis and the FAR on the x-axis. If a higher FRR is acceptable, a lower FAR can be achieved and vice versa. By changing the discrimination threshold the classifiers can operate on any point on the curve if desired.

To ensure statistical robustness, the ROC curves are constructed by performing 5-fold cross-validation and averaging the results vertically. The confidence intervals reveal that there is very little variance in classifier performance even if the data set is partitioned into different training and test sets.

The ROC curves show that all subjects are recognized with high probability, as the FRR and the corresponding 95% confidence intervals confirm. SVM outperforms all other classification techniques, followed by LDA and Mahalanobis. SVM achieves a FRR of less than 10% and a FAR of less than 10% at the same time, i.e., an EER of less than 10%. Given that this assessment is based on the over-time data set, it is a remarkable result.

Applying the classifiers to the snapshot data set yields even better performance as Figure 3.8b reveals. At a FAR of 5%, FRR is close to 0% when using SVM as classifier. This result suggests that body impedance is a very viable biometric characteristic for continuous authentication and shows remarkable distinguishing power. In a continuous authentication system where a certain percentage of false rejects (incorrect rejection of a legitimate user) can be accepted—such as the one described in Section 3.3—pulse-response recognition will, with high probability, detect all adversarial samples.

Moreover, body impedance seems to be especially effective as a biometric trait if the stored biometric reference is fairly recent in relation to the measurement that is being validated. All classifiers show a significant improvement in performance if they only have to deal with samples from a single measurement session, e.g., the snapshot data set. Performance on the over-time data set is likely to be improved with more measurement sessions. The classifiers will gain a clearer picture of the variability of each subject's body impedance if they have access to samples from additional points in time.

3.5.3 Biometric Identification

Biometric identification is a multi-class classification problem. The goal is to identify a person as accurately as possible given unlabeled biometric samples. We consider closed-set identification, which means the system must find the best match from a pool of candidates (see Section 2.2.3).

We test five different classifiers for biometric identification. The conceptual generalization to the multi-class setting is straightforward for all classifiers: The Euclidean and Mahalanobis distance classifiers increase the number of centroids to one centroid per class. LDA generalizes to Multiclass-LDA by introducing one mean per class and measuring between-class variability through the covariance matrix of the class means.

Ranking success rates obtained from the identification classifiers are shown in Figures 3.9a and 3.9b. We depict Rank(1), Rank(3) and Rank(5). The classifiers have been trained on both, the over-time data set (Fig. 3.9a) and the snapshot

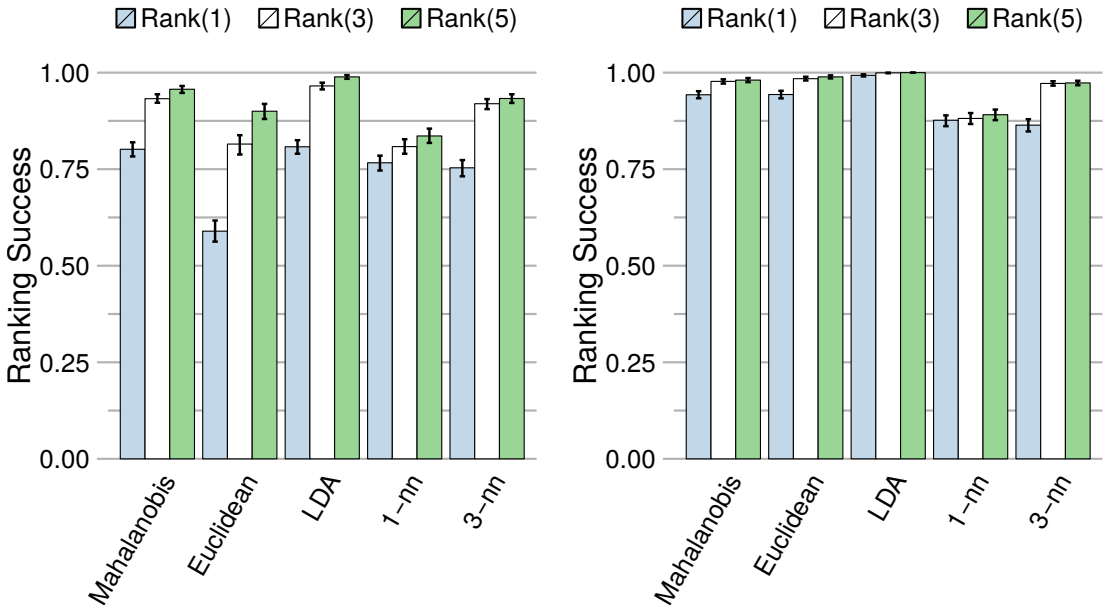

(a) Over-time data set **(b)** Snapshot data set

Figure 3.9: Ranking success rates for identification. The results presented are averages over all users and obtained by applying 5 times stratified 5-fold cross-validation. Error bars show the 95% confidence interval for each classifier.

data set (Fig. 3.9b). The ranking success rates illustrated in the bar plots are averaged over all subjects and obtained by applying 5-fold cross-validation, similar to the authentication scenario.

Even with the increased complexity of multiple classes, all tested identification classifiers perform well on the over-time data and improve on the snapshot data set. For the snapshot data set, a ranking success close to Rank(1) = 1.0% is possible using Multiclass-LDA as classifier. The nearest neighbor classifiers (1-nn and 3-nn) cannot quite match the performance of the Mahalanobis distance method and Multiclass-LDA. Clearly, the conceptually simpler Euclidean distance method cannot cope with the added variability present in the over-time data set (see Figure 3.9a).

All classification methods benefit from measurements that are acquired in a relatively short time frame, e.g., the snapshot data set. They can improve performance significantly if trained and tested on these samples only. Measurements taken far apart are influenced by very different conditions. There might be physiological changes, such as weight loss or gain, or there might be differences in the ambient temperature, humidity, clothing, or a number of other factors. The added uncertainty becomes apparent in the classification performance and, in turn, affects the ranking success rates (compare Fig. 3.9a with Fig. 3.9b).

Table 3.2: System performance of the prototype set-up averaged over all users in [%]

Authentication (SVM classifier)	FAR	FRR	Accuracy	EER
Snapshot set	2	2	96	2
Over time	9	9	87	9

Identification (LDA classifier)	Rank(1)	Rank(2)	Rank(5)
Snapshot set	99	100	100
Over time	81	97	99

Performance metrics are calculated using five times 5-fold stratified cross-validation. Values shown reflect the performance achieved with the best classifier for each scenario.

3.5.4 Summary of Results

Table 3.2 summarizes the results of the best classifiers for verification and identification achieved with our prototype set-up, on both, the snapshot data set and the data set taken over time. For verification, SVM gives the best results whereas for identification Multiclass-LDA proves to be the most suitable classifier.

In a verification scenario, pulse-response recognition achieves a very low EER of 2% on the snapshot data set and an EER of 9% on the over-time data. This makes it clear that body impedance is a viable characteristic for biometric verification.

If pulse-response recognition is used for biometric identification, a ranking success of almost 100% can be achieved for the static snapshot data set. According to our experiments, even if the biometric measurements are captured in sessions that are weeks apart, pulse-response recognition reaches a ranking success rate Rank(1) of 81% and 97% for Rank(3), respectively.

For biometric recognition to work reliably over a long period of time, persistent characteristics are key. The biometric method needs to extract unique and stable features from traits that do not change significantly over the run-time of the system. Otherwise, re-enrollment or the update of the biometric reference is needed. Our results show that EER for identity verification increases from 2% to 9%, when comparing measurements that are captured in quick succession to those that are weeks apart. It remains unclear how recognition rates behave if measurements are even further apart, i.e., months or years. To analyze the characteristic over such long time span, a more extensive study that is not within scope of this work would be required.

3.6 Biometric Attacks Against Body Impedance

In this section, we explore the resilience of body impedance recognition against biometric attacks. We first discuss biometric forgery and presentation attacks and then experimentally determine the probability of success for impersonation attacks.

3.6.1 Liveness and Replay

A common problem with many biometric systems is presentation attack detection. 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. Similarly, a face recognition system would need to make sure that it is not being fooled by a photo or a 3-D artifact

In established biometric systems, presentation attacks are 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.

In the context of body impedance based recognition, unlike fingerprint or face recognition, it is difficult (yet not impossible) to separate the biometric characteristic from the individual to whom it belongs. If the adversary manages to capture a user's body impedance on some compromised hardware, successfully presenting it to a sensor would require specialized hardware that mimics the exact impedance of the original user. We believe that this is feasible: the adversary can devise a contraption that consists of 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 impedance of the electrode-wire-electrode contraption 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 more than in cases of facial recognition or fingerprints, which are routinely left—and can be lifted from—numerous innocuous locations.

Furthermore, in contrast to face or fingerprint recognition, impedance-based biometrics can be made to depend on the capture platform. Thus, even if the adversary captures body impedance on one piece of hardware, it would not directly translate to the user's measurements on a different capturing device. One way to achieve this is to add a specific (frequency-dependent) resistance to the measurement platform. If the adversary uses its own capture system to measure the user, there

is an additional signature which the adversary is unaware of as it is part of the impedance measurement device.

Finally, the real power of the body impedance recognition is evident when used for continuous authentication (see Section 3.3), whereby the person physically uses a secure terminal and constantly touches the keyboard as part of routine work. Biometric verification happens on a continuous basis and thus making it infeasible 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. However, the biometric system will see input from the keyboard without the expected impedance measurement to accompany it, and will lock the session.

3.6.2 Impersonation of Pulse-Response

We experimentally estimate worst-case probabilities for different scenarios where an attacker could impersonate a legitimate user by fooling the biometric system using his own impedance measurements (zero-effort impersonation). In order to estimate the likelihood of such an impersonation attack, we use the pulse-response data set acquired in Section 3.4 and measure similarity of samples.

Attacker Model

We consider four attack scenarios relevant to body impedance recognition. Similar to the previous section, we differentiate between verification and identification. In addition, a potential attacker who tries to impersonate a legitimate user may or may not be known to the system. We refer to an attacker whose biometric reference is known as an *internal* attacker and otherwise we call it an *external* attacker. Thus, an internal attacker has been registered and is enrolled in the system. An external attacker has never used the system and no impedance measurements have been gathered.

Regardless of its type, the attacker's goal is to impersonate a legitimate user of the system. The attacker tries to achieve this by using his own body impedance and trick the classifier.

To give a realistic experimental lower bound on the attack probabilities for zero-effort impersonation, we base our analysis on the over-time data set. The results in Section 3.5 made evident that classifying pulse-response samples with increased variability is more challenging. Consequently, we assume that it is also more difficult to detect an attacker under these conditions.

Table 3.3: Average and worst-case performance for discussed impersonation attacks.

	internal attacker [%]		external attacker [%]	
	average case	worst case	average case	worst case
Authentication (SVM classifier)				
– Sensitivity	98.8	76.0	98.8	88.0
– Specificity	99.9	99.0	95.0	36.0
Identification (LDA classifier)				
– Sensitivity	99.9	76.0	99.6	80.1
– Specificity	99.5	99.0	99.0	92.0

Average and worst-case sensitivity and specificity for four attacks scenarios. We distinguish between authentication and identification and between internal and external attackers. External attackers are not known to the classifiers.

The attack scenarios we present in the following are limited by the scope of our pulse-response data set, but they nevertheless provide an accurate view on the behavior of body impedance recognition.

Internal attackers

We take our best-performing classifiers from Section 3.5 and estimate their performance for all possible attacker-victim pairs in our test subject population. We first train the classifiers on the entire data set and then ask them to classify biometric samples from a predefined attacker and a predefined victim only. We thereby measure sensitivity (that corresponds to the complement of FAR, i.e., $1 - FRR$) and specificity (which denotes the complement of FAR, i.e., $1 - FAR$) for a specific combination of attacker and victim.

This performance assessment is repeated for all possible attacker-victim combinations, which lets us compute average as well as worst-case performance of sensitivity and specificity. The results can be found in Table 3.3 in the column labeled *internal attacker*. Not surprisingly, average sensitivity and specificity attain very high numbers and confirm our previous findings about the classification power of pulse-response recognition. Specificity of almost 100%—on average and in the worst case—guarantees that an internal attacker is very likely to be detected, whether pulse-response recognition is used for verification or identification.

Sensitivity seems to vary more than specificity. For certain attacker-victim combinations, sensitivity only reaches 76%. This means that a particular legitimate user is recognized in 76% of the tested samples. In all remaining cases he was

rejected because the classifier mislabeled him for the attacker. These numbers are congruent with our results from Section 3.5 where we discover that more variability in the pulse-response measurements affects average sensitivity to a greater extent than average specificity.

External attackers for verification

To model an external attack on body impedance verification, we pursue a similar procedure as outlined above for internal attackers. The main difference is that no attacker samples are included in the training phase of the classifier. The classifier should be able to identify adversarial samples without knowing a template/reference describing the attacker's measurements.

We exploit the nature of the binary classification problem for biometric verification and form two classes of samples. Having set aside all measurements from the attacker, we define a class containing samples from the victim and a second class consisting of samples from all other users. Although the actual attacker is not represented in this pool of training samples, the classifier can gain a accurate understanding of what measurements other than those from the victim look like. During classification, an external attacker is likely to fall into the group of "other" users despite the fact that the classifier does not obtain any of the attacker's samples during training.

In Figure 3.10 we show a graphical representation of all possible attacker-victim combinations for an external attack on the verification classifier. Sensitivity and specificity are shown for each attacker-victim pair, as well as average sensitivity and average specificity. The resulting matrix appears to be nearly symmetric. If two subjects have similar pulse-response measurements, it is almost equally likely for both of them to be able to successfully impersonate the other. There are a few deviations, however, which for instance include subjects *Remo* and *Mason* (see Figure 3.10). *Remo* has a higher chance of impersonating *Mason* than the other way round, as specificity is lower when *Remo* simulates the attacker. These differences stem from the fact that the class of samples from two different users can have different shape and dispersion in the feature space. The classifier will not necessarily create symmetrical decision boundaries when it is trained on different subsets of the data.

From the results in Table 3.3, we conclude that, on average, the verification classifier performs almost equally well in both attacker scenarios, internal and external. Sensitivity and specificity are above 95% in all cases. Although average performance is very high, a few attacker-victim combinations reveal detection probabilities significantly below average. For instance, if subject *Ethan* wants to

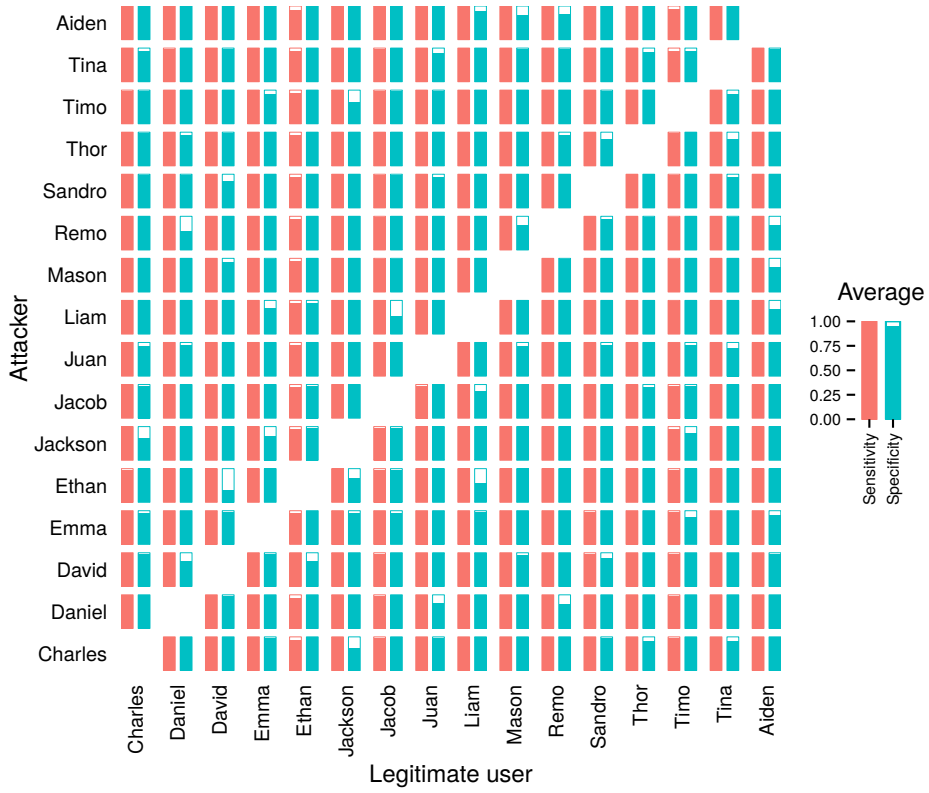

Figure 3.10: An external attacker tries to impersonate a legitimate user. Sensitivity and specificity for every possible attacker-victim combination of the over-time data set based on unseen samples from both, attacker and legitimate user (test persons have been anonymized with pseudonyms).

impersonate *David* then specificity is estimated at 36% which will result in a 64% chance for *Ethan* to go undetected and successfully fool pulse-response recognition (see Figure 3.10). *Ethan* and *David* must have a very similar pulse-response.

The fact that some attacker-victim pairs have similar measurements is what motivated the Markov Model in Section 3.3.3. The model takes into consideration that the measurements of the attacker might be statistically similar to the legitimate user and, as a consequence the attacker successfully passes the biometric test at first and only gets caught eventually.

External attackers for identification

When pulse-response recognition is used for identification, reliable detection of external attackers becomes more intricate. The classifier has to distinguish between multiple classes and detect attacker samples at the same time. It is possible to construct a binary classifier for every single user which decides between legitimate user and attacker. However, this approach requires an aggregation scheme that collects the classifiers' outputs to produce a final decision whether the presented

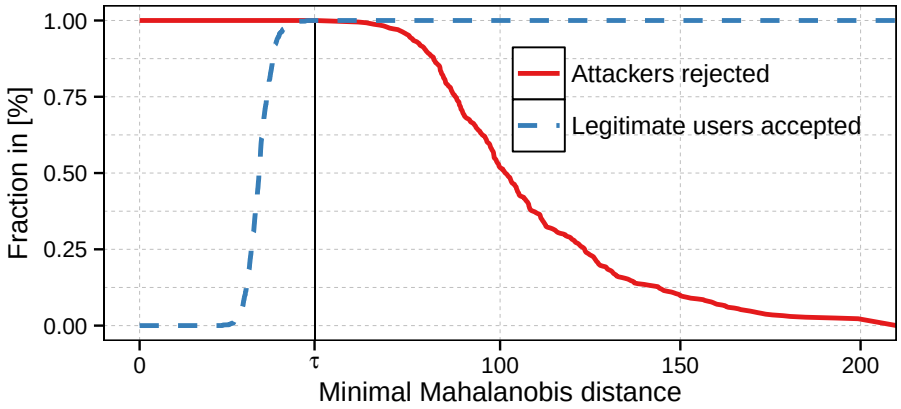

Figure 3.11: Detecting external attackers as statistical outliers with minimal Mahalanobis distance between sample and class means. At discrimination threshold τ the fraction of detected attackers is equal to the fraction of recognized legitimate users.

measurement is indeed an attacker or not. Results in Section 3.5.3 showed that, for pulse-response recognition, pairwise SVM classifiers do not perform as well as conceptually simpler methods, such as Multiclass-LDA. Starting out with this insight, we opt for a less complex model to detect attacker samples. It is based on the assumption that samples from unknown subjects can be detected as statistical outliers. Samples from an external attacker originate from an unknown source as no such samples have been seen by the system before classification. These adversarial measurements might not share any statistical characteristics with the measurements the classifier has encountered during training phase.

An effective approach to outlier detection that we pursue is to determine mean and covariance for each class of samples representing a registered user. This information can then be used to compute the Mahalanobis distance between a new measurement and all stored biometric templates, i.e. the Mahalanobis distance to the class means. Should a new measurement happen to be far from all class means, then the likelihood of it being an attacker sample is high. If the minimal distance exceeds a certain threshold, the sample is declared as an attack and filtered out.

The described method is essentially the same as the Mahalanobis classifier that we tested for pulse-response recognition in Section 3.5.3. This time though, we do not assign class labels to the samples but rather compute the likelihood (i.e., the distance) that a sample belongs to any of the stored templates. The motivation behind choosing Mahalanobis distance for outlier detection is twofold: It showed reliable classification results for pulse-response measurements and its application to outlier detection is straightforward. There is no need to train multiple classifiers and no additional class to accommodate outliers is needed.

The performance graph in Figure 3.11 shows the discrimination threshold for the minimal Mahalanobis distance versus outlier detection. Varying the discrimination threshold not only has an effect on how many legitimate users are recognized but also on how many attackers are detected. Ideally, the system would reject all attacker samples and accept all samples from registered users. The threshold $\tau = 48.5$ used in the experiments is chosen in such a way that the percentage of detected attackers is equal to the percentage of correctly identified legitimate users. τ is found by 2-fold cross-validation and achieves an error rate of almost 0%.

After having filtered out the attacker samples, the system continues to assign class labels to the remaining measurements, i.e., those which have not been found to represent an external attacker. The classification of these samples is analogous to the scenario where only internal attackers are considered. The biometric system employs a Multiclass-LDA classifier to solve the classification task, similar to the identification classifier found in Section 3.5.3.

Since the system now contains two sources of possible misclassification errors (the classifier for the user samples and the preceding outlier filtering stage) the performance assessment must take this fact into account. In particular, we need to consider the rejection of legitimate users during outlier detection. A legitimate user who is incorrectly identified as an external attacker must be treated as a wrong assignment and should impact the sensitivity score.

Table 3.3 lists sensitivity and specificity the identification classifier is able to achieve when samples are pre-filtered by Mahalanobis distance to detect external attackers. Performance experiences almost no decrease compared to the scenario for internal attackers. Average sensitivity and average specificity stay at a very high level of 99%. Worst-case sensitivity is even increased from 76.1% to 80.1%. Worst-case specificity is affected to a marginal extent: it changes from 99% to 92%, which supports our initial assumption that it is more challenging to detect external attackers than internal attackers.

We can still conclude, however, that there is a high chance that impersonation attempts from external attackers are detected. This is mainly due to the effective outlier detection scheme which filters out attacker samples before the measurements are fed to the classifier.

Table 3.4: Comparison with keystroke dynamics and touchscreen biometrics

	Usability							Security				Biometric trait			
	Unobtrusive	Nothing to carry	Recovery from loss / leakage	Verification time	Enrollment time	Identification	Cost	Resilient to physical observation	Resilient to targeted impersonation	Computationally expensive	Difficult to forge	Motor skill	Purely behavioral	Continuous measurements	Can be integrated into work flow
Keystroke dynamics	✓	✓	✗	s	m	○	.	✗	✗	✗	✗	✓	✗	✓	✓
Touch(-screen) biometrics	✓	○	✗	s	m	✗	.	✗	✗	✗	✗	✓	✗	✓	✓
Body impedance recognition	✓	✓	✗	s	s	✓	⋱	✓	✗	✗	○	✗	✗	✓	✓

✓ the method supports the property
○ the method supports the property partially ✗ the method does not support the property h hour(s) m minute(s) s second(s) ⋱ high cost : medium cost . low cost

3.7 Comparison with Other Modalities

Using body impedance for biometric recognition is a novel concept that has largely been unexplored in research literature so far. Only a few proposals exist that exploit the electrical properties of the human body for user authentication.

We provide a comprehensive overview of biometric characteristics in Section 2.3. Here, we cover those biometric methods that are similar to body impedance in terms of modality and biometric capture.

Apart from other proposals that suggest bioimpedance-based traits, we identify keystroke dynamics and touch(-screen) biometrics as two recognition methods that could be used in place of body impedance for both presented examples: the PIN entry scheme in Section 3.2 and the continuous user authentication scheme in Section 3.3. Depending on the input device, the system could measure the respective biometric characteristics: keystroke dynamics in case of an ordinary keyboard and touch(-screen) characteristics in case a touch-screen is available for user input.

In Table 3.4, we show a comparison of the three modalities along the high-level criteria used in Chapter 2. Even though impedance-based characteristics are physiological, they can be measured unobtrusively and do not require the user to carry

anything. Like keystroke dynamics and touchscreen biometrics they do not allow recovery after loss. Body impedance supports identification and has much shorter verification and enrollment times. However, since it needs to be measured with a sensor not commonly found in electronic devices, implementation costs are currently much higher. Unlike keystroke dynamics or touchscreen behavior, impedance-based traits are resilient to physical observation and are more difficult to forge.

We now compare the performance of these three biometrics in detail.

3.7.1 Comparison with Keystroke Dynamics

Keystroke dynamics is one of the most researched behavioral biometrics. Some of the first scientific studies that propose to harness the distinguishing capabilities of keyboard characteristics for identity verification date back to the mid 1970s and can be found in, e.g., [96] and [55]. Since then, many different recognition methods have been proposed. The most straight-forward methods are based on relatively simple statistics, such as mean typing times and their standard deviations [97, 98]. Over the last few years, several pattern-recognition methods have come into vogue and been applied to keystroke dynamics, such as e.g., neural networks [99], fuzzy logic [100], and support-vector machines [64]. A survey on the large body of literature on biometrics using keystroke dynamics is given in [92, 97, 101] and in the comprehensive background section of [102].

A biometric system could be designed such as to measure both biometric traits, keystroke timings and body impedance, at the same time and with minimal user intervention. This would require that the keyboard users are typing on is conductive and can capture body impedance. Both biometric modes do not require the user to change normal work flow when typing on the keyboard, which makes the biometric recognition process very unobtrusive. Clearly, these two modalities could complement each other and result in a more powerful biometric system. However, both methods have the drawback of not being able to acquire biometric measurements during periods when there is no user input. Assuming users of such a combined system do not rest their hands or fingers on the keyboard while inactive, neither keystroke dynamics nor body impedance recognition can bridge the breaks between typing phases. In such cases, other recognition methods, e.g., a video camera for face recognition, could be a better complement and increase security to a greater extent than keystroke dynamics and body impedance in combination with each other. We therefore compare the performance of keystroke dynamics and pulse-response recognition in more detail.

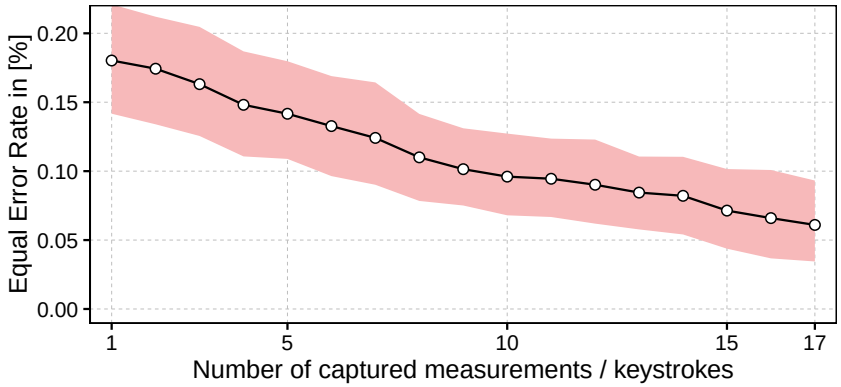

Figure 3.12: Equal error rate (EER) of pulse-response recognition in relation to the number of keystrokes. We assume users are typing on a conductive keyboard and every keystroke results in one pulse-response measurement, e.g., for a five-letter word, five measurements can be captured (measuring errors are omitted). The solid line represents average EER over all users, the shaded area shows the 95% confidence interval.

To this end, we evaluate a scenario specifically designed for this comparison. This allows us to compare pulse-response with performance numbers for keystroke dynamics found in literature. We assume that users type on a conductive keyboard and every keystroke results in only one captured impedance measurement. Since the square pulse used for the capture has a duration of 100 nanoseconds, many more measurements would in theory be possible during a single keystroke.

The enrollment data for this analysis is comprised of five random measurements per user, taken from our over-time data set. The validation data consists of 17 measurements per user, randomly sampled from the snapshot data set. Choosing training and validation data in such a way, we simulate verification of new measurements (captured in quick succession while the user is typing) with the help of a stored biometric reference obtained during enrollment[1]. The final authentication decision is made based on the aggregated classification outcomes of each individual measurement.

Figure 3.12 shows the equal error rate depending on the number of captured measurements, averaged over all users. The 95% confidence interval is depicted as a shaded area. It can be estimated by re-sampling the subsets for enrollment and validation data 25 times for each user. After one single measurement, i.e., after one keystroke, average equal error rate is 18.0% and steadily declines to 6.14% if 17 subsequent measurements can be captured.

[1]This is different from the continuous authentication setting in Section 3.3 because new measurements are not compared to a reference measurement obtained at login time, but validated against a pre-existing, stored biometric reference.

The performance of keystroke authentication systems varies in a similar fashion: If verification consists of a single word, i.e., as it is the case in password augmentation, only a small amount of keystroke data can be captured by the system and recognition rates are consequently lower. The study in [64] which uses a 16 character pass-phrase for both enrollment and verification achieves an equal error rate of 6.96% (vs. 6.14% of pulse-response recognition) whereas free text recognition (users are allowed to type anything for enrollment and verification) can achieve equal error rates as low as 0.95% [103]. Short typing sequences or passwords, however, yield similar results to pulse-recognition. The study in [104] uses passwords between 11 and 17 characters and resulted in 8.1% FRR and 2.8% FAR. [98] operates with a text length of 10 and achieves a FRR of 11.57% and FAR of 1.89%. Finally, the authors of [105] are able to get 6.0% FRR and 0.5% FAR while using a text length of 25.

Research literature appears to be in disagreement as to whether keystroke dynamics are a valid means of user authentication. Most studies report recognition rates high enough such that keystroke dynamics can be considered unique to each individual [98, 102]. However, others show that typing patterns of different individuals can have similar characteristics and the uniqueness of keystroke biometrics must be questioned. The most prominent one is [106] where attackers are shown the typing pattern of their victims and make a conscious attempt to imitate them. The fictional attackers receive training through a textual and graphical feedback interface. After training, false acceptance rate increases from 0.20 to 0.42 if attackers have partial knowledge of the typing statistics of the victim, and from 0.24 to 0.6 if entire typing statistics are known. These results show that keystroke dynamics might not be suitable for high-security environments and existing commercial solutions using keystroke biometrics might not withstand targeted attacks.

3.7.2 Touch(-screen) Biometrics

Nowadays, many modern personal electronic devices, such as smartphones and tablets, are equipped with a capacitive touchscreen as input device and do not rely on keyboard and mouse anymore. As a consequence, research has started to explore the concept of user authentication based on input signals received from a touchscreen. Touchscreen biometrics, i.e., taps, strokes, swipes and gestures executed by one or multiple fingers on a touchscreen, are similar to keystroke dynamics as they can only be measured and evaluated during active user input. They are considered very unobtrusive as they measure users' touchscreen actions which are part of the natural work flow when interacting with a smartphone or similar device. If the biometric capture mechanism needed to measure body impedance can be miniaturized in the

future, impedance-based recognition might be accommodated in smaller devices where it could complement touch(-screen) biometrics similar to keystroke dynamics.

The first work that thoroughly investigates the applicability of touchscreen input as a behavioral biometric can be found in [66]. The authors propose 30 behavioral features that can be extracted from a user's interaction with a touchscreen on a with a smartphone. The paper concludes that touchscreen features might not be applicable to long-term authentication. They could, however, still serve as part of a multi-modal biometric recognition scheme or secure short absences of usage without immediately locking the device. In [69], for instance, touch characteristics are used to unlock a smartphone and to enhance swipe/shape password patterns for instant authentication. The authors achieved a recognition rate of 57% in a two-day user study.

In [67] another framework and a prototype for continuous user authentication on mobile devices is presented. It consists of a sensor glove that delivers fine-grained features, e.g., orientation, direction, rotation, of the finger movements and a smartphone that collects touch gesture data. This augmented approach achieves slightly worse recognition rates, but the authors believe that their system could be used successfully for post-authentication security for a certain amount of time after the user authenticates by some other means, i.e., password or other biometric method.

A similar approach is presented in [107] where a watch-like prototype measures the user's skin impedance profile of the wrist in order to modulate a user-specific signal onto the user's skin that can be picked up by a touchscreen. This allows seamless and transparent authentication on each touch the user makes. The authors recruited 10 participants for a lab evaluation and claim that their classifier produces no false positives when identifying users.

3.7.3 Bioelectricity- and Bioimpedance-based Biometrics

The work in [108] provides a listing of papers on cognitive biometrics based on the electroencephalography (EEG), the electrocardiogram (ECG), and the skin conductance, also called electro-dermal response (EDR). It describes how these biometrics can be harnessed for user authentication. Skin conductance (i.e., EDR) is directly related to body impedance in terms of modality and acquisition method. The main difference is that, unlike body impedance, it captures the emotional state of an individual and not necessarily a physiological trait. The resistance of the skin can vary significantly due to the embedded sweat glands which are controlled by

the nervous system. Body impedance-based biometric recognition methods (such as pulse-response recognition), on the other hand, focus on extracting physiological characteristics independent of emotional state by measuring entire parts of the human body, not only skin conductance.

Probably the most related to our recognition method is the work in [109] where bioimpedance is used as a physiological characteristic. A wearable sensor is designed to passively recognize wearers based on a body's unique response to the alternating current of different frequencies. The authors design a prototype wristband that captures electrical impedance around at the wearer's wrist, as opposed to measuring body impedance from one hand to the other. Experiments in [109] were conducted in a family-sized setting of 2 to 5 subjects, where a person wears the bioimpedance sensor on the wrist. They achieve recognition rate of 90%. In a more recent study [110] the authors improved their prototype and increased the number of test subjects. They report FAR and FRR of 2% for samples taken within a day. Our biometric recognition method solves a different problem—we propose a recognition method that works by temporarily touching two electrodes, not a wearable device—but our technique also uses the electrical response to a signal and we achieve very similar error rates when samples are taken in one session. Both the bracelet in [110] and our method achieve promising error rates over a day, which shows that body impedance is well-suited for continuous authentication. Unfortunately, we cannot compare the results for samples that are taken weeks apart, as we are not aware of any other published study.

3.8 Summary

We show that the horizontal body impedance is a viable biometric characteristic that could serve as an additional authentication mechanism in a PIN entry system, enhancing the security of PIN entry with minimal extra user burden.

The same biometric characteristic is also applicable to continuous user authentication. To this end, we propose a continuous authentication mechanism on a secure terminal, which ensures user continuity, i.e., the user who started the session is the same one who is physically at the terminal keyboard throughout the session.

Through experiments with a proof-of-concept prototype we develop a specific instance of body impedance measurement and demonstrate that each human body exhibits a unique response to a signal pulse applied at the palm of one hand, and measure at the palm of the other. Using the prototype we can identify users in a matter of seconds. This identification mechanism integrates well with other

established methods, e.g., PIN entry, to produce a reliable added security layer, either on a continuous basis or at login time.

We also direct our attention to how likely a legitimate user can be impersonated by an attacker using his own biometric data. We give average probabilities, as well as, experimental lower bounds found through simulations of worst-case scenarios.

Body impedance is a physiological characteristic and therefore it can be seen as distinct from behavioral aspects. However, it has an attractive property normally associated with behavioral characteristics: it can be captured in a completely passive fashion. Although other physiological characteristics used for biometric recognition also have this feature, such as face recognition, body impedance recognition is not easily circumventable.

Body impedance recognition does require special-purpose hardware—which is true for most other physiological traits—, but the combination of unobtrusiveness and resilience to attacks make it an attractive recognition mechanism which offers desirable properties from both worlds, physiological traits and behavioral traits.

4

Generating Secret Keys from Body Impedance Measurements

In this chapter, we focus our attention on the generation of biometric keys from body impedance. After having shown the distinguishing abilities of body impedance in the previous chapter, we devise a method to extract the entropy inherent in impedance measurements to compute user-specific and cryptographically secure secrets.

Contents

4.1 Introduction

Due to its unobtrusiveness, body impedance is well-suited for a wide range of scenarios. Given that the user is in physical contact with two conductive surfaces, multiple biometric measurements can be acquired rapidly and without additional user participation. Thus, embedding the measurement apparatus into a variety of devices can support the use of body impedance in many security contexts, as a means of primary or supplementary user authentication. Examples include: customer authentication on ATMs equipped with conductive PIN-pads (see Section 3.2), continuous user authentication with conductive keyboards (see Section 3.3), and car driver authentication via conductive steering wheels (presented in the next chapter).

In terms of pros and cons, body impedance is similar to other biometric recognition methods. While biometrics offer usability advantages over authentication schemes based on secrets or possession of items, arguably their main drawback is immutability—once a biometric is compromised, it cannot be changed or replaced. Two recent examples are a leak where hackers stole 5.6 million fingerprints of United States federal employees [111], and revelations about insecure storage of fingerprint authentication data by a major mobile phone manufacturer [112]. This is particularly worrying because the research community has shown that original biometric input can be reconstructed from such stored reference data (i.e., the biometric template), if not adequately protected [113, 114].

Many biometric authentication systems therefore resort to protect biometric templates with the help of a Trusted Platform Module to prevent unauthorized access. In an ideal situation, however, biometric templates could be protected in a way similar to best practice in password-based user authentication. When handling passwords, it is the *de-facto* standard to store only hard-to-invert, uniquely salted hashes; anything else is considered unfit for secure authentication.

Biometric template protection [115]—transforming raw biometric input into a representation which does not leak any sensitive information if acquired by an adversary—is an intricate problem and remains one of the main obstacles to widespread deployment of biometric authentication.

One approach to template protection is *biometric key generation*, whereby raw biometric measurements are used to consistently generate the same unique value for each individual user. This allows strong biometric data protection, since in a case of

compromise, an adversary cannot infer information about the user from a stored key or template. Also, biometric keys can alleviate humans from remembering strong secrets, since the keys can be used to either add entropy to an existing cryptographic scheme or serve as a replacement for password- or token-based authentication.

Unfortunately, generating consistent biometric keys from volatile biometric data is a challenging task, both in terms of achieved performance, as well as applying and tailoring existing methods to different biometric modalities. While the previous chapter investigated using body impedance in the context of user authentication and identification, the issue of deriving biometric keys or otherwise protecting this biometric trait is an important step towards practicality.

This thesis represents the first attempt to derive keys and secure biometric templates from body impedance data. As a basis for deriving biometric keys, in Section 4.3, we implement and compare two methods for extracting stable and idiosyncratic features:

1. an approach based on *feature learning* using multilayer convolutional networks in a Siamese architecture, and
2. the feature extraction based on frequency analysis that is used in the previous chapter (see Section 3.4.8).

We then improve the key evaluation approach of Ballard et at. [116] to analyze strength of derived biometric keys, and discuss security guarantees of key generation based on body impedance. Our experimental evaluation in Section 4.6 shows that the electrical properties of the human body can be used to derive biometric keys and sets a baseline for future work in body impedance-based key generation methods.

4.2 Biometric Key Generation

In this thesis, we design and implement a procedure to extract consistent secrets, *biometric keys*, from body impedance measurements and experimentally assess security provided by such keys in the context of protecting users' sensitive biometric data. Stated simply, our goal is to remove the need for storing an individual's biometric measurements or any derivative representation that an adversary could use in case of a breach. To this end, we aim to generate a secret value from the raw biometric measurements and subsequently only use the derived biometric key for user authentication or other protocols.

Before describing our approach, we overview previous work in biometric template protection techniques and state general goals and requirements for biometric key generation schemes.

4.2.1 Background

Biometric template protection techniques can be divided into two main groups [117–120]: biometric cryptosystems and cancellable biometrics. The latter focus on (non-invertible) transforms on biometric features, which allows privacy-preserving template matching of biometric samples. On the other hand, biometric cryptosystems rely on generating (or binding) biometric keys by storing user-specific biometric information as a secure biometric template, in order to retrieve or generate keys.

Randomized Biometric Templates (RBT) [121]—the key generation scheme we use in this work—falls into the subcategory of key-generating schemes, along with other methods such as fuzzy extractors [122] and secure sketches [123, 124]. RBT builds upon quantization schemes which have been successfully applied to different biometric modalities, including iris [125], face [124], voice [126], fingerprint [127] and handwriting [121, 128, 129]. For more background information, we refer to a comprehensive survey of biometric template protection methods [117].

The motivation behind choosing RBT as the key generation method is the close resemblance of impedance measurements to voice recordings in terms of data structure. If body impedance is captured as varying voltage levels over time, such as in pulse-response recognition (see Section 3.4.1), the resulting data consist of a one-dimensional signal with a time component. This is similar to voice recordings (measured as changing pressure levels over time) and RBT has been successfully applied to voice biometrics. We discuss this further in Section 4.7.

4.2.2 Requirements and Goals

A biometric key generation scheme is considered sound if it is: *correct* and *secure* [115].

To achieve correctness, the scheme must output consistent keys, allowing a legitimate user of the biometric system to provide a recent biometric measurement and recreate the biometric key established during enrollment. The rate of failure can be measured by the false reject rate (FRR): the percentage of key generation trials that do not lead to the original key.

To be considered secure, the following three necessary requirements need to hold [115, 116]:

1. **Biometric Uncertainty:** measurement values used to generate the key (i.e., biometric input) must be difficult to guess.
2. **Key Randomness:** the resulting key must appear random.

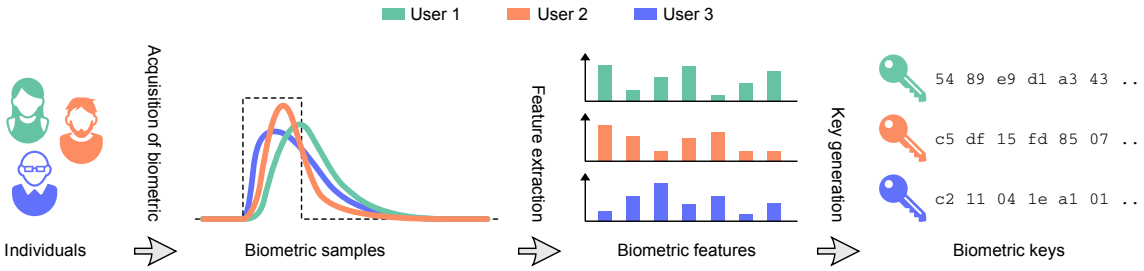

Figure 4.1: Procedure for key derivation from body impedance measurements (compare with Figure 3.4).

3. **Irreversibility:** it must be difficult to deduce information about the biometric input from the generated key, the template and any auxiliary data.

We further elaborate on these requirements in Section 4.4 and experimentally validate our proposed key generation scheme against them in Section 4.6.

4.3 Proposed Approach

The proposed approach, overviewed in Figure 4.1, consists of three parts: *1)* acquisition of biometric measurements, *2)* extraction of a stable feature-set, and *3)* derivation of a biometric key from that set.

The procedure depicted in Figure 4.1 is similar to the overview of impedance-based recognition shown in Figure 3.4 of the previous chapter, but includes important extra steps required for key generation. We now describe each of them in detail.

4.3.1 Acquisition of Impedance Measurements

In order to generate keys, we use pulse-response as an instance of a particular body impedance measurement. Pulse-response recognition measures body impedance from one hand to the other and is acquired by sending a low-voltage electric signal (in the form of a short square pulse) from the palm of one hand to the other. We showed in Section 3.5 that pulse-response recognition is sufficient for biometric verification and identification. These results make pulse-response a good candidate to serve as a basis for our analysis. However, we believe that the proposed method generalizes to other body impedance biometrics.

4.3.2 Feature Extraction

Feature extraction is the process of transforming raw input data into a set of intermediate feature values which serve as input to subsequent machine learning

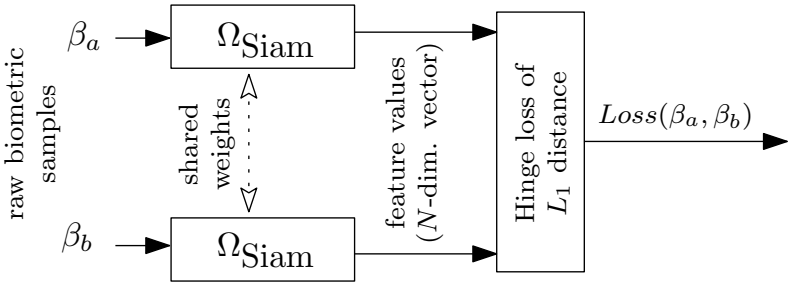

Figure 4.2: Siamese architecture: Two copies of convolutional networks Ω_{Siam} share internal parameters. They are trained on pairs of inputs (β_a, β_b) by optimizing the Hinge loss of the distance of feature value outputs $\Omega_{\text{Siam}}(\beta_a)$ and $\Omega_{\text{Siam}}(\beta_b)$.

tasks, such as classification, identification, or (as in this case) key derivation. For successful key derivation, features extracted from the underlying biometric trait must be: (1) *stable*, i.e., maximally similar for multiple measurements of the same individual, and (2) *idiosyncratic*, i.e., maximally distinct for different users. One traditional approach to feature extraction is to have domain experts manually *engineer* a set of potentially strong features. Feature values are then generated on a subset of data, to allow selection of those features likely to fit their intended purpose.

Feature learning We follow a relatively new approach called *feature learning* that is currently at the forefront of state-of-the-art in machine learning domains, such as speech recognition [130], machine translation [131], and image recognition, where computer performance is for the first time surpassing those of humans [132]. The main advantage of feature learning is in automating manual, labor-intensive and biased tasks of feature engineering using multilayer convolutional networks as feature extraction models. Parameters of these networks are programmatically optimized along clear optimization criteria during a feature learning stage. As a result, the network *learns* a hierarchical nonlinear representation which can map raw input data into a lower-dimensional feature space with desired characteristics [133].

In order to find an optimal mapping, we use a *Siamese* architecture [133, 134], shown in Figure 4.2, where two identical convolutional networks Ω_{Siam} share the same set of parameters. During optimization, the networks accept pairs of raw biometric inputs β_a and β_b (belonging to users a and b) and compute feature values as their outputs $\Omega_{\text{Siam}}(\beta_a)$, $\Omega_{\text{Siam}}(\beta_b)$. Ideally, feature values should be similar (have small distance) if the inputs belong to the same user, and sufficiently different (distance of at least m) if they belong to different users. Such an optimization criterion can

be directly represented with a Hinge loss function applied to a distance measure:

$$Loss(\beta_a, \beta_b) = \begin{cases} |\Omega_{\text{Siam}}(\beta_a) - \Omega_{\text{Siam}}(\beta_b)| & \text{if } a = b \\ \max\left(0, m - |\Omega_{\text{Siam}}(\beta_a) - \Omega_{\text{Siam}}(\beta_b)|\right) & \text{if } a \neq b \end{cases}$$

We optimize the shared parameters of convolutional networks Ω_{Siam} to minimize the loss function of the described Siamese architecture using iterative stochastic gradient descent (SGD) [130]. Pairs of raw biometric measurements are randomly sampled from a pool of data reserved for training (70% of the data set), and the remainder of the data set is used to estimate the Hinge loss on unseen data. We run iterative SGD until each pair of samples has been encountered five times on average.

The number of features extracted by Ω_{Siam} is varied by parameter N that specifies the dimensionality of the resulting nonlinear representation, i.e., a N-dimensional vector. Further details about the structure of used multilayer convolutional network Ω_{Siam} are given in Appendix A.

Baseline feature extraction method As a baseline for comparison, we also test the feature extraction method proposed in the previous chapter where we classify pulse-response measurements. That feature extraction approach, denoted by Ω_{FFT}, consists of transforming the raw signal to the frequency domain by computing the *Fast Fourier Transform (FFT)* to obtain frequency components. We compare performance of the two feature extraction methods in Section 4.6.3 and show that Ω_{Siam} outperforms Ω_{FFT} by a considerable margin.

The intuition behind strong performance of Ω_{Siam} can be found in previous work, which has shown that multilayer convolutional networks generalize remarkably well on many tasks [131, 132], even to the level of supporting *transfer learning*, where the features are trained on one task, and then slightly adapted and used on another [135]. Unlike other dimensionality reduction techniques, such as Linear Discriminant Analysis, Siamese architecture allows us to train the feature extraction model on pairs of training data, thus increasing the number of different inputs seen during training. The model never explicitly takes the number of different classes into account and is therefore well-suited for learning features in open-set applications, such as biometric key generation.

Algorithm 1: Enroll

Global: features $\phi_1, \ldots, \phi_N$, quantization widths $\delta_1, \ldots, \delta_N$

Input: enrollment samples $\beta_1, \ldots, \beta_j$, PIN π

Output: key K, template T_U, token z

$\Psi \leftarrow \mathsf{SelectStrongFeatures}(\beta_1, \ldots, \beta_j)$

$\Delta \leftarrow [\mathsf{Permute}\{\phi_i \in \Psi\},\ \mathsf{Permute}\{\phi_i \notin \Psi\}]$

$T_U \leftarrow [\,],\ \kappa \leftarrow \epsilon$

foreach ϕ_i in the order of Δ **do**

  $\mu_i \leftarrow \mathsf{Median}\{\phi_i(\beta_1), \ldots, \phi_i(\beta_j)\}$

  $\alpha_i \leftarrow \lfloor \mu_i - \delta_i/2 \rfloor \bmod \delta_i$

  $c_i \leftarrow (\ \mathsf{Enc}_1^{\pi}(i), \mathsf{Enc}_2^{\pi}(\alpha_i)\ )$

  $T_U \leftarrow T_U.\mathsf{Append}(c_i)$

  if $\phi_i \in \Psi$ **then**

    $x_i \leftarrow \lfloor \mu_i - \delta_i/2 \rfloor$

    $\kappa \leftarrow \kappa \parallel i \parallel x_i$

$z \leftarrow \mathsf{Hash}^{\mathrm{token}}(\ \pi \parallel \kappa\ )$

$K \leftarrow \mathsf{Hash}^{\mathrm{key}}(\ \pi \parallel \kappa\ )$

return K, T_U, z

Simplified enrollment procedure of Randomized Biometric Templates. Enroll outputs the user's biometric template and biometric key. ϵ denotes the empty string and $\parallel$ represents string concatenation.

4.3.3 Key and Template Generation

We now overview Randomized Biometric Templates [121] (RBT), the key generation scheme used in this chapter. The scheme performs error-correction by quantizing the values of features. It defines two procedures: Enroll derives the randomized biometric template from a set of biometric measurements, and KeyGen uses the biometric template to recreate the key from a single biometric sample. Algorithms 1 and 2 show a simplified version of these two procedures in pseudocode.

Enroll Taking into account that the same feature can have different variability across individuals, Enroll determines which features can be used for consistent key generation and therefore should be included in the set of "strong" features Ψ for a particular user. Then, Enroll computes necessary information to perform error-correction and encodes it in a user-specific *randomized biometric template*.

As input, Enroll requires a set of user's biometric measurements $\beta_1, \ldots, \beta_j$, user's PIN π, access to the biometric features $\phi_1, \ldots, \phi_N$, and the corresponding global

Algorithm 2: KeyGen

Global: features $\phi_1, \ldots, \phi_N$, quantization widths $\delta_1, \ldots, \delta_N$

Input: biometric sample β, PIN π, template T_U, token z

Output: key K or *nil* on failure

$\kappa \leftarrow \epsilon$

foreach c as appearing in T_U **do**

  $i \leftarrow \mathsf{Dec}_1^{\pi}(c[0])$

  $\alpha_i \leftarrow \mathsf{Dec}_2^{\pi}(c[1])$

  $x_i \leftarrow \phi_i(\beta) - (\ (\phi_i(\beta) - \alpha_i)\ \bmod \delta_i)$

  $\kappa \leftarrow \kappa \parallel i \parallel x_i$

  if $\mathsf{Hash}^{\text{token}}(\ \pi \parallel \kappa\ ) = z$ **then**

    $K \leftarrow \mathsf{Hash}^{\text{key}}(\ \pi \parallel \kappa\ )$

    return K

return *nil*

Simplified key generation procedure of Randomized Biometric Templates. KeyGen outputs the user's biometric key if token z can be regenerated. $\parallel$ denotes concatenation.

quantization widths for each feature $\delta_1, \ldots, \delta_N$. As a first step, Ψ is computed using the function SelectStrongFeatures$(\ldots)$, which is illustrated in Figure 4.3; a feature ϕ_i is considered "strong" for a particular user if the range of its values for all enrollment samples $(\beta_1, \ldots, \beta_j)$ is smaller than δ_i, the global width of quantization for feature i. Only "strong" features are used to generate user's biometric key. This ensures that the user can reproduce a key at a later time and results in reduced false reject rates.

After Ψ is determined, Enroll reorders feature indices such that all "strong" features are located before all other features and then both parts are pseudo-randomly permuted:

$$\Delta \leftarrow [\mathsf{Permute}\{\phi_i \in \Psi\},\ \mathsf{Permute}\{\phi_i \notin \Psi\}]$$

Enroll continues to calculate a quantization offset α_i for every feature i. Each α_i is computed such that the median μ_i of the feature values collected during enrollment $\phi_i(\beta_1), \ldots, \phi_i(\beta_j)$ falls in the middle of the corresponding quantization interval.

Finally, the user's biometric template T_U consists of a list of all N pairs of feature indices and corresponding quantization offsets (i, α_i) in the order as they appear in Δ. The pairs are encrypted to $c_i \leftarrow (\mathsf{Enc}_1^{\pi}(i), \mathsf{Enc}_2^{\pi}(\alpha_i))$ by two different pseudorandom permutations, Enc_1^{π} and Enc_2^{π}, determined from the user's PIN π.

Given only T_U with a permutation of pairs (i, α_i), KeyGen does not know the size of the set of "strong" features Ψ. In order to encode the number of "strong"

features, a token z is generated by Enroll to serve as a stopping criterion for KeyGen. z is computed as a hash of the concatenation of feature indices i and quantized feature values x_i of all "strong" features using hash function $\mathsf{Hash}^{\text{token}}$.

KeyGen The key generation algorithm KeyGen requires template T_U, a token z, PIN π, and a new biometric measurement β as input, and outputs a user's biometric key K. After decrypting the template with the PIN (i.e., reversing Enc_1^π and E_2^π), KeyGen computes the quantized feature values for the given β. It subtracts the quantization offsets α_i from the feature values $\phi_i(\beta)$ and uses the global quantization widths δ_i to reconstruct the corresponding quantized value x_i for each feature. The final biometric key K is constructed by applying $\mathsf{Hash}^{\text{key}}$ to the concatenation of i and x_i for all "strong" features according to the order they appear in the decrypted template.

Recall from the description of Enroll that KeyGen does not know the number of features in Ψ, so it iteratively computes the value of the temporary token z' with more and more features, until the values of tokens collide: $z' = z$. It is important to note that even though K and z are generated as hashes of the same values, the scheme uses different hash functions $\mathsf{Hash}^{\text{token}}$ and $\mathsf{Hash}^{\text{key}}$ to ensure that no information about K can be deduced from z.

Using z as the stopping criterion provides additional benefit against an adversary who does not have the user's PIN since the adversary has no way of knowing which features and in which order they should be included when calculating the user's biometric key. Finally, this is a protection even against an adversary who is in possession of user's template and PIN, as he will still not know how many features need to be included when trying to guess the key. We discuss this in Section 4.4.2.

Quantization widths δ_i and level of quantization k The values of δ_i are controlled by the level of quantization k. Since the quantization widths are global and used in the generation of all templates, they should be representative of the overall population statistics. We estimate them on a subset of data from multiple users. For each feature, δ_i is determined by selecting k-th percentile of the spread of the feature's values. Choosing the k-th percentile for δ_i results in quantization intervals that are expected to quantize feature values of k % of the population consistently.

As a consequence, k controls the granularity of quantization: the larger k is, the more features on average are found to be "strong" for a particular user. This affects the chances of consistently replicating the biometric key in two ways. On the one hand, as the error-correction intervals increase, KeyGen is more likely to

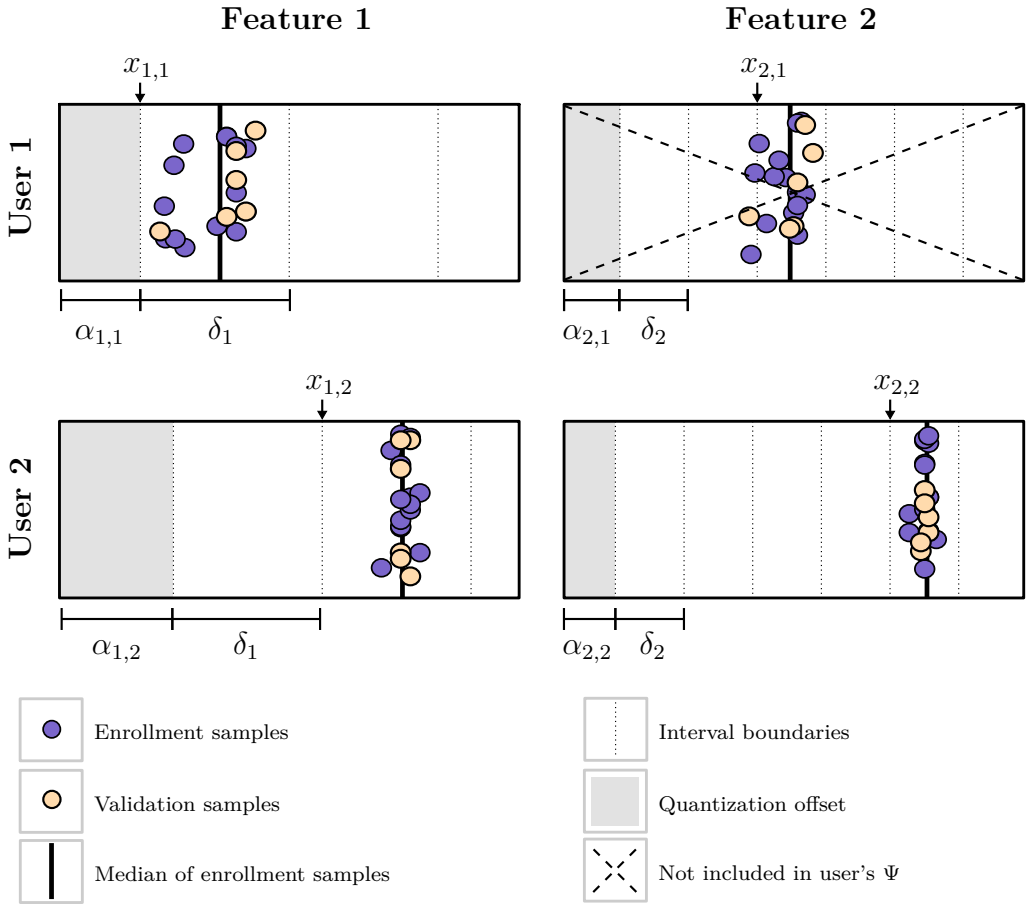

Figure 4.3: Example of quantization and feature assignment for two users and two features. The horizontal position represents raw feature values for enrollment and validation samples. Samples are distributed vertically to improve readability. $\alpha_{i,j}$ denotes the quantization offset of feature i for user j. $x_{i,j}$ represents the quantized value of feature i for user j. Feature 2 is not included in the set of "strong" features (Ψ) for user 1 since the spread of enrollment samples is higher than the quantization width (δ_2).

quantize a feature value properly, but on the other hand, since more features are considered "strong" and used in key generation, the chance of one of the feature values falling into the wrong quantization bucket increases. Finally, larger δ_i results in a lower total number of buckets for a specific feature, which can reduce the number of guessing attempts for an adversary who attempts to directly guess the quantized feature values.

4.4 Security Analysis

We now analyze the proposed scheme against security requirements stated in Section 4.2 that a key generation scheme based on body impedance needs to meet in order to be considered *secure*.

4.4.1 Key Randomness and Irreversibility

As mentioned earlier, we use Randomized Biometric Templates (RBT), which ensures key randomness and irreversibility by construction, under the assumption that biometric uncertainty holds. We now outline our main arguments why key randomness and irreversibility hold and refer to [121] for further details.

Even assuming that the biometric template T_U, the token z and the key K are known, an adversary cannot infer any information about the biometric input a user provided to Enroll during enrollment. This is due to the fact that without knowing the correct PIN π that decrypts the template, an adversary cannot use the quantization offsets α_i in T_U, as he does not know to which features they correspond[1], and the offsets reveal no information about quantized feature values x_i. Hence, in order to learn any information about the biometric input, an adversary can only guess π and the biometric input β simultaneously, by running KeyGen on the guessed values π' and β'. It can verify the guesses by checking whether either the resulting key K' corresponds to K or the computed token z' matches z. Thus, we conclude that if the combined entropy of the biometric and the PIN is sufficiently high, the key generation scheme is irreversible for a computationally bounded adversary.

Similarly, we argue that the key generation scheme yields keys indistinguishable from random. Input to the cryptographic hash function used to compute $K = \mathsf{Hash}^{\text{key}}(\dots)$ consists of π, which is unknown to the adversary, and quantized feature values, which can only be inferred by enumerating all possible combinations of π and β, as described above. Therefore, a computationally bounded adversary cannot distinguish between K and random, even if T_U and z are known.

However, recall that key randomness and irreversibility hinge on biometric uncertainty, i.e., difficulty of predicting biometric input. We focus on it in the following section.

4.4.2 Biometric Uncertainty

The most important requirement for a biometric key generation scheme is that the used biometric trait exhibits high unpredictability. In order to assess the unpredictability provided solely by the biometric input, one must assume that all other parameters of the key generation scheme, i.e., the user's template (including the token) and PIN, are known. This residual unpredictability can be estimated by computing measures such as Shannon entropy or min-entropy. However, estimating

[1] This also assumes that the set of "strong" features Ψ for user U is a random subset of all features.

conditional entropy in a high-dimensional space is usually infeasible due to lack of sufficient data. Moreover, entropy estimation might not capture the variety of different approaches an adversary could pursue to predict the biometric input, such as impersonation attempts or guessing attacks, since entropy measures are summary statistics and do not reflect concrete strategies.

It is therefore important to consider different types of adversaries when quantizing unpredictability of a biometric trait. We distinguish between *imposters* and *algorithmic adversaries*. To show that body impedance can withstand both of these adversaries, we verify whether the biometric trait is resistant against impersonation and automated searches, i.e., guessing biometric input in a structured way.

Resilience to imposters In order to test how well body impedance withstands impersonation attempts, we measure FAR – the probability that the key generation scheme outputs the correct biometric key based on forged biometric material. The lower the FAR, the more secure the key generation scheme. We report FAR for different combinations in the parameter space (N, k) in Section 4.6.4. For certain configurations, the average FAR is as low as 1.9% when measured over the entire test subject population. For comparison, biometric key generation based on voice characteristics achieves FAR between 0.0% and 10.0% (see Section 4.7).

Guessing raw biometric data We assert that, for an adversary in possession of a user's decrypted biometric template, the corresponding token, and the PIN, it is always more beneficial to guess the intermediate set of quantized feature values, than to guess raw biometric measurement data. This is due to the feature extraction process, which strongly reduces the dimensionality of biometric measurements by several magnitudes and arrives at a smaller set of intermediate features which are used in key generation. Taking into account that the final biometric key is a result of a cryptographic hash function, the adversary would also prefer guessing intermediate feature values over trying to guess the key directly. We therefore focus our analysis of guessing attacks on an adversary that tries to predict the inputs to the hash function, i.e., quantized feature values x_i of RBT.

Guessing feature values In the remainder of this section, we estimate the effort of the adversary in guessing a set of quantized feature values x_i, that, combined with the legitimate user's template and PIN (which are at the attacker's disposal), result in generating the correct biometric key. Similar analysis can be made for a scenario where the adversary has the user's biometric template, but not the PIN.

In that case, by construction of RBT, the adversary must simultaneously guess the correct PIN and feature values. This increases guessing complexity by a factor proportional to the number of PIN guessing attempts.

4.5 Key Guessing

To estimate adversarial effort in guessing correct feature values, we adopt and improve the probabilistic search approach described by Ballard et al. [116]. Their approach assumes that an estimate of the strength of biometric keys can be given as $\log_2(\,\mathrm{RunTime}(\mathcal{A})\,)$ where $\mathcal{A}$ is an algorithm that repeatedly guesses the biometric key until it succeeds, with run-time proportional to the number of guesses. The algorithm has access to the user's template (including token), the PIN π used to encrypt it, and a population statistic P_{ϕ_i} for each feature ϕ_i measured over all the feature values for several other users (not including the user under attack).

The algorithm starts by sorting the features according to their information content and then traverses them in that order by recursively guessing values for each feature ϕ_i using the population statistic P_{ϕ_i}. It follows this *depth-first search* pattern until the search space is fully exhausted or one of the guessed feature value combinations yield the correct key.

Considering that the original algorithm was developed to estimate key strength of arbitrary key generation schemes, it does not take into account the fact that in RBT, the "strong" features are stored in user's template before other features. Exploiting this insight, an adversary that is attacking RBT can take a *breadth-first search* approach to deploy a more efficient key guessing strategy, which results in a successful guess in considerably less tries. We describe such improved key guessing strategy in the remainder of this section and then use it in Section 4.6 to compute an estimate of the security guarantees provided by the biometric keys.

4.5.1 Improved Key Guessing Strategy

We improve the approach from [116] and tailor it specifically to RBTs by ensuring that the adversary never attempts to guess the values of those features which are not considered "strong" for a specific user and hence not included in the resulting key. Our approach thus gives a more conservative estimate of security of the evaluated key generation scheme. In Section 4.6.5, we experimentally show that the original approach significantly overestimates the number of guessing attempts. Our guessing strategy reduces – on average – the number of guesses by a factor of 2^{20}, for almost half of the keys in the data set.
We now describe the improved key guessing strategy:

Guessing the number of "strong" features ($|\Psi|$) Even if the adversary has a user's template and the corresponding PIN π, he can only obtain partial knowledge about the user's set of "strong" features Ψ. This is due to how the Enroll algorithm of RBT generates the template T_U. Enroll shuffles the order of the features before they are stored in T_U, keeping apart "strong" features and features used for padding. It then encrypts the index i and quantization offset α_i of every feature:

$$T_U = [\ \{(\mathsf{Enc}_1^\pi(i), \mathsf{Enc}_2^\pi(\alpha_i)) : \phi_i \in \Psi\},$$
$$\{(\mathsf{Enc}_1^\pi(i), \mathsf{Enc}_2^\pi(\alpha_i)) : \phi_i \notin \Psi\}\ ]$$

Recall from Section 4.3 that z serves as a stopping criterion for KeyGen, as it allows evaluating if the key generated from some subset of features is correct, or more features should be included. The exact number $|\Psi|$ of "strong" features included in a user's template is therefore unknown, even if the adversary can decrypt T_U. Still, the earlier a feature ϕ_i appears in the user's template, the higher the likelihood that $\phi_i \in \Psi$. A natural choice for an adversary is therefore to guess features in the order of their appearance in T_U.

Instead of always attempting to guess values for all features (as in [116]), in our algorithm $\mathcal{A}$, the adversary initially assumes that $|\Psi| = 1$ and evaluates all possible values for only the first feature. If no key is successfully generated, $\mathcal{A}$ iteratively increments its current estimate of the size of Ψ, i.e., $|\Psi| = 2, 3, \ldots$, and thereby includes more and more features in its search. This process continues until the adversary assumes the correct number of "strong" features and, while evaluating all possible combinations of feature values for those "strong" features, finally generates the correct key.

Guessing feature values ϕ_i When the adversary is evaluating all possible keys which consist of the first $|\Psi'|$ features, we assume that he guesses the value of the first $|\Psi'|$ quantized features one by one, in a depth-first-search manner. When guessing the value for a specific feature ϕ_i, we assume that the adversary uses the population statistics P_{ϕ_i} acquired from other users' biometric features to guess the values in the order of their decreasing probabilities, thus increasing the probability of an early success. To take this into account, we denote with $\mathrm{G}(\phi_\mathrm{i}, \mathrm{P}_{\phi_\mathrm{i}})$ the number of different values that $\mathcal{A}$ will try when guessing ϕ_i, before it guesses the correct value.

The run time estimate of $\mathcal{A}$ can now be expressed as the summation of W_1 failed attempts when $|\Psi'| < |\Psi|$ and W_2 failed attempts made with the correct estimate of $|\Psi'| = |\Psi|$, plus 1 for the last attempt that yields the correct key:

$$\mathrm{RunTime}(\mathcal{A}) \propto W_1 + W_2 + 1$$

Since generated keys and templates are known in our experimental setting, if $\mathcal{A}$ follows the described search pattern, we can evaluate the number of unsuccessful guesses W_1 and W_2 (and therefore the run time of $\mathcal{A}$) on the used data set:

W_1: number of wrong guesses when $|\Psi'| < |\Psi|$ When assuming a too small number of "strong" features, $\mathcal{A}$ has to enumerate all possible quantized feature values for each feature ϕ_j (denoted by $|\Gamma_j|$), which gives a total of

$$W_1 = \sum_{i=1}^{|\Psi|-1} \left(\prod_{j=1}^{i} |\Gamma_j| \right)$$

wrong guesses, for all possible keys generated from the first $1, 2, \ldots (|\Psi| - 1)$ features.

W_2: number of wrong guesses when $|\Psi'| = |\Psi|$ Once $\mathcal{A}$ assumes the correct number of "strong" features, the key will likely be found before all $\prod_{j=1}^{|\Psi|} |\Gamma_j|$ possibilities are explored as $\mathcal{A}$ can use the population statistics to improve its chances of guessing the correct feature values. However, presuming $\mathcal{A}$ guesses a wrong value for feature ϕ_i, it will unsuccessfully enumerate all $\prod_{j=i+1}^{|\Psi|} |\Gamma_j|$ combinations spanned by the remaining features. Since $\mathcal{A}$ will make $\mathrm{G}(\phi_\mathrm{i}, \mathrm{P}_{\phi_\mathrm{i}})$ unsuccessful guesses before arriving at the correct value for ϕ_i and $\mathcal{A}$ can assume an incorrect value for any "strong" feature (i.e., $i = 1, \ldots, |\Psi|$), we get at a total of

$$W_2 = \sum_{i=1}^{|\Psi|} \left(\mathrm{G}(\phi_\mathrm{i}, \mathrm{P}_{\phi_\mathrm{i}}) \cdot \prod_{j=i+1}^{|\Psi|} |\Gamma_j| \right)$$

unsuccessful key guessing attempts.

As we show and discuss in the experimental evaluation in Section 4.6.5, results indicate that 70% of generated keys require over 2^{30} guessing attempts in the described scenario. Furthermore, half of the keys are even stronger, since $\mathcal{A}$ makes at least 2^{50} guesses before finding the correct key.

4.6 Experimental Results

We experimentally evaluate performance of the proposed scheme and compute error rates for key regeneration. We then examine the strength of biometric keys derived from the experimentally acquired data set by simulating attacks discussed in Section 4.4.

4.6.1 Experimental Data set

We derive biometric keys and compute RBT-s based on the over-time data set of pulse-response measurements (see Section 3.4.7) comprising a test subject population of 16 people. The data set was acquired in five seatings where participants' pulse-response was measured five times in a row, summing up to a total of 25 samples per participant. The median time span between consecutive seatings was 8 days and there was a minimum dwell time of at least one day between seatings.

When acquiring the over-time data set we did not impose any requirements or restrictions on the participants besides ensuring that their hands were not overly sweaty. Since the measurements were gathered over multiple sessions, various other potential factors, such as varying body water percentage or body temperature must have been captured as well. This allows us to evaluate the key generation scheme under conditions as diverse as possible and compute a realistic estimation of its performance.

4.6.2 Experiment Parameters

The proposed key generation scheme has three steps: (1) feature extraction, (2) feature selection (as part of enrollment), and (3) key generation. As described in Section 4.3.2, behavior of feature extraction is characterized by a N – that is the total number of features extracted from biometric measurements. Feature selection, on the other hand, depends on k – the feature quantization level, defined in Section 4.3.3. In the remainder of this section, we show the performance of our scheme by presenting results for a range of parameter pairs: k between 25 and 75 and N between 10 and 40.

4.6.3 False Reject Rate

Correctness of a biometric key generation scheme is primarily determined by measuring FRR – the probability that a user fails to recreate the same biometric key using new measurements. FRR therefore directly reflects user inconvenience, analogous to how FRR is defined for biometric recognition systems (see Section 2.2.5).

The way FRR is computed here coincides with how we obtained FRR for biometric user verification in Section 3.5.2. We partition the data set into training and test set and repeat this process five times by stratified cross-validation to acquire a robust estimate of FRR.

The left plot of Figure 4.4 shows FRR as contour lines in the 2-dimensional plane for various combinations of N and k. Using the nonlinear feature extraction

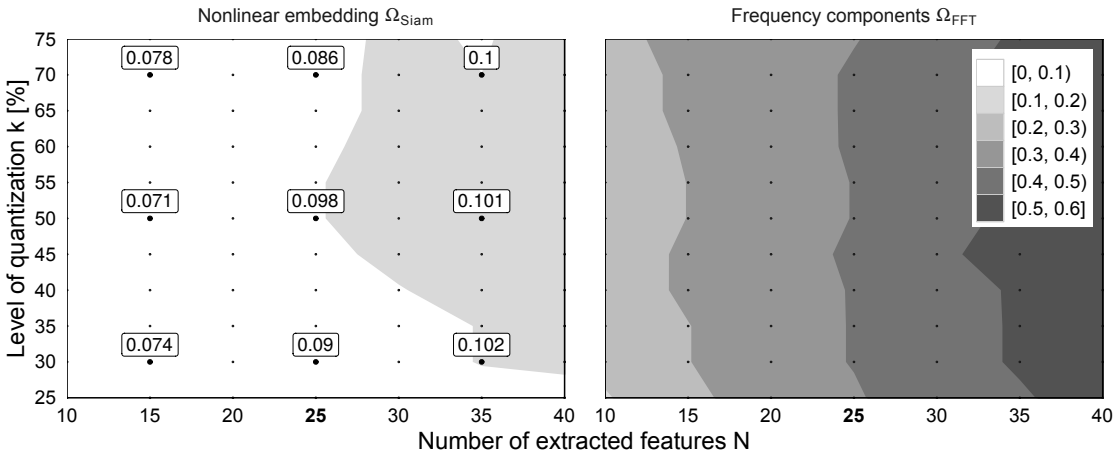

Figure 4.4: False reject rate (FRR) based on the number of extracted features N and the level of quantization k. Contour lines are linear interpolations on the grid of parameters marked with black dots. The white area represents combinations of parameters N and k for which the average FRR for generated keys is lower than 0.1.

method Ω_{Siam}, our key generation scheme achieves FRR between 7.4% and 9.8% for a broad range of parameters. These results are better or comparable to FRR values for quantization-based key generation schemes based on other biometric modalities, such as iris, face and voice [117, 126]. Being the first of its kind, this evaluation also serves as a reference for the future comparison of FRR values of key generation schemes based on body impedance.

Values are computed by running a repeated κ-fold cross validation on the whole data set with $\kappa = 5$ and 5 repetitions to compute average FRR per user. Given a total of ζ user samples, we enroll a user by creating the key and the template on $\zeta \cdot (\kappa - 1)/\kappa$ samples. The remaining ζ/κ samples are used to test the rate of key regeneration based on the previously computed template. We measure the fraction of incorrectly reconstructed keys for all users to arrive at the average FRR.

Comparison with baseline feature extraction As a comparison to our feature extraction method, the right side of Figure 4.4 shows FRR achieved with the baseline feature extraction method which relies on frequency components of the response signal (Ω_{FFT}).

Figure 4.4 shows that independent of quantization level and number of features, Ω_{Siam} achieves significantly lower FRR than the baseline feature extraction method Ω_{FFT}. We thus conclude that the nonlinear embedding performs significantly better in describing idiosyncratic features. For Ω_{FFT}, FRR almost entirely depends on the number of features while the level of quantization only plays a marginal

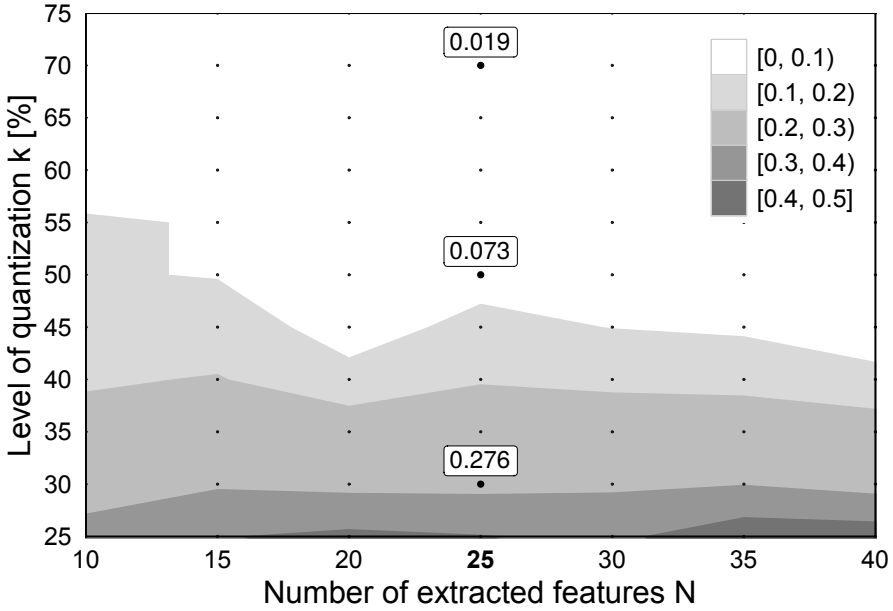

Figure 4.5: False accept rate (FAR). Darker colors indicate higher FAR and therefore higher probability of success for an imposter attack. The white area represents combinations of parameters N and k for which the average FAR is below 0.1.

role, suggesting that most features are less user-specific then those extracted with the Siamese architecture.

Presented false reject rates render performance of the baseline feature extraction method impractical. Therefore, we do not experimentally evaluate the security of that scheme.

4.6.4 Imposter Attacks

We now assess resistance of the proposed key generation scheme to impostor attacks. This is done by measuring the false accept rate (FAR). We compute FAR similarly to Section 3.5.2 where the likelihood for impersonation of pulse-response was estimated. We consider the attackers/imposters to be *external*. Thus, the key generation does not acquire a template or reference of the attack samples before the attack is simulated. Analogous to the FRR calculation above, we use cross-validation to acquire a robust estimate for FAR.

In Figure 4.5, contour lines indicate average FAR, measured on biometric keys generated for different pairs: (N, k). Our key generation method is robust and achieves strong results for a range of N and k (white area), wherein the probability for a successful attack is lower than 0.1.

For each combination (N, k), FAR-s are computed using 5-fold cross validation to generate legitimate user's template from different subsets of data, and then

trying all other users' biometric measurements as inputs, in an attempt to derive the legitimate user's key.

Figure 4.5 also shows that higher quantization levels reduce the likelihood of impersonation attacks. As discussed in Section 4.3.3, this is due to higher number of features assigned to each enrolled user, which reduces the probability of the adversary generating the correct key.

4.6.5 Key Guessing Complexity

We use the same experimental data set to evaluate run time complexity of an attack by an adversary that has the legitimate user's biometric template and PIN. As described in Section 4.4.2, the adversary tries to generate the legitimate user's key by probabilistic search over all possible feature values using our improved algorithm $\mathcal{A}$.

In this evaluation, we focus on the particular configuration of $N = 25$ and $k = 70\%$. On condition that FRR-s remain practical, we set $N = 25$ to maximize the number of extracted features (see Figure 4.4). We set $k = 70\%$ in order to minimize the probability of an impostor attack while maintaining a sufficiently high quantization threshold for stringent feature selection (see Figure 4.5).

Key strength estimates for this particular combination are presented in Figure 4.6. As in Section 4.6.4, results are obtained with 5-fold cross validation. We compute population statistics P_{ϕ_i} using biometric data of all users, except the one being evaluated.

Results of the guessing complexity estimation show that majority of keys withstand an adversary using the approach of algorithm $\mathcal{A}$. Over 70% of generated keys require the adversary to make at least 2^{42} guesses. Also, more than half of generated keys require over 2^{55} attempts.

About 10% of users' keys are predictable due to the feature selection procedure choosing only a small number of "strong" features. In a production system, selection of such a small set of "strong" features would be detected and would result in enrollment failure. In such cases, the user would be required to re-enroll. We included these users and their corresponding keys for the sake of completeness.

Comparison of key guessing algorithms Figure 4.6 compares our improved key guessing algorithm $\mathcal{A}$ with the original guessing algorithm in [116], which $\mathcal{A}$ is built up on. As an additional reference, we show an estimate of Shannon entropy of generated keys from a worst-case perspective by assuming the set of "strong" features to be known for every generated key. To compute the entropy, we assume

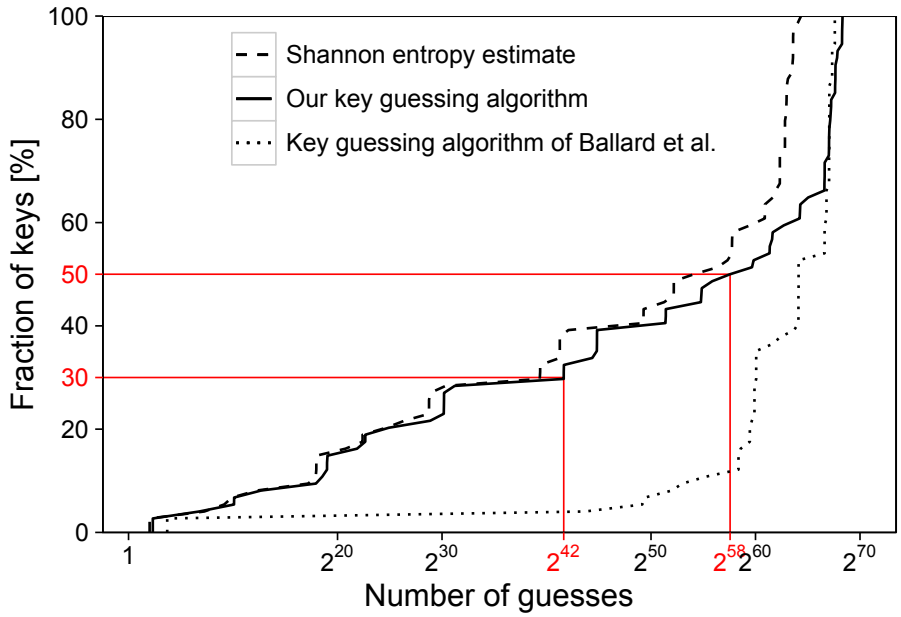

Figure 4.6: Empirical cumulative distribution of the keys found after a given number of guessing attempts estimated using two different algorithms. As our more conservative algorithm shows, for 70% of the generated keys, the adversary has to try guessing a minimum of 2^{42} combinations. For 50% of the keys, the required effort increases to 2^{55} guesses.

feature independence and estimate the probability densities based on all biometric measurements, except those from the user being evaluated.

By design of RBT, a guessing adversary does not learn the number of "strong" features until the key is found, even if the template is decrypted with the correct PIN; see Section 4.3.3. Therefore, this entropy calculation underestimates the strength of the keys. In addition, it is not directly comparable to the run time of a guessing attack, as entropy measures unpredictability in bits and does not describe a concrete strategy to guess keys. However, it serves as an indication of tightness for the estimates derived by our proposed guessing algorithm $\mathcal{A}$ and the original guessing algorithm.

As Figure 4.6 shows $\mathcal{A}$ performs significantly better than the algorithm of Ballard et al. [116], when guessing generated keys from body impedance. This is mainly due to $\mathcal{A}$ being specifically tailored to guessing keys generated by RBT. Although entropy and the number of guesses are not directly comparable [136], the run time of $\mathcal{A}$ reflects the estimated entropy of the biometric keys much more precisely than the algorithm of Ballard et al..

FAR versus key strength A False Accept Rate (FAR) of 1.9% for imposter attacks (computed in Section 4.6.4) can appear high when compared to the entropy

estimate obtained through a guessing algorithm. We estimate that over 70% of generated keys require the adversary to make at least 2^{42} guesses (see Section 4.6.5).

One reason why FAR seems high is due to the fact that the data set for FAR computation can vary wildly depending on what type of FAR is computed. We compute FAR for (zero-effort) imposter attacks, where the adversary has samples at his disposal that have not been encountered by the key generation scheme, but are valid impedance measurements.

In case of key guessing, the adversary only has access to an (aggregated) population statistic from which new features are sampled artificially. Therefore, the entropy estimate depends on how representative the used statistic is of the attacker's target. Additionally, any key guessing algorithm that uses more than the theoretical optimum of one guess only delivers an upper-bound on the number of guesses. It is always possible to design a better algorithm, especially when more information is available to the adversary.

In Section 4.7 below, we see that the estimates we obtained for FAR and key strength lie within the range of the numbers found in related work concerned with other biometric modalities. The difference between FAR and estimated key strength can be even larger—for instance in [137], both FRR and FAR are near 2.0%, but key strength is estimated to be at around 60 bits using guessing entropy.

4.7 Comparison with Other Schemes

Body impedance Despite the growing body of related research, we believe that our work represents the first study of generating keys from body impedance. To the best of our knowledge, the only other work which discusses body impedance in the context of generating biometric keys is the study by Gupta et al. [138], which uses skin conductance as an additional measure to fight coercion attacks during key generation from voice biometrics. The idea is to incorporate skin conductance measurements into the cryptographic key generation from voice. They experimentally show that the skin conductance measurement will help to reveal user's emotional states and recognize an attack as a stressful event (significantly different from the state when the user was enrolled into the system). This way, the generated keys include a dynamic component that can reveal whether a user is forced to grant an access to the system.

Comparison with key generation from voice Since previous research has not explored body impedance for generating biometric keys, we compare our results with other biometrics used in key generation. In terms of modality, we find body impedance most comparable to voice biometrics since both rely on a (one-dimensional) signal with a time component, albeit of different scales. Our experimental evaluation suggests that it is possible to generate keys from body impedance with a FRR of 8.6% and a FAR of 1.9%. Furthermore, entropy estimation from a worst-case point of reference (PIN and template known to the adversary) shows that keys from body impedance yield around 46 bits of entropy, on average.

One of the earliest studies that proposes a technique to reliably generate a key from a user's voice is [137], which achieves both FRR and FAR of near 2.0%. Key strength is estimated to be at around 60 bits using guessing entropy. The work in [139] presents cancellable speech template protection in speaker verification systems. FRR and FAR range from 0.0% to 3.0%, depending on the dimension of the extracted features. While these are remarkably low error rates, no entropy or key strength estimation is reported. Although FAR usually correlates with the amount of information extracted from the biometric data, key strength cannot be directly deduced from error rates, and a separate analysis is needed.

Probably most related to our work is [126] where the authors also make use of Randomized Biometric Templates to generate cryptographic keys from voice, reporting FRR of 7% and FAR of 10% for two utterances and a quantization threshold $k = 50\%$. They ran the original (and less tailored) guessing algorithm proposed in [116] to assess strength of generated keys: At least 2^{30} guesses are needed for 36% of the population and over 2^{40} guesses are needed for 7% of the population.

Finally, a recent result on biometric templates for speech [140] reports equal FRR and FAR rates between 8% and 10% when assuming a similar scenario to our analysis, i.e., adversary knows the PIN or password. They estimate the entropy of the scheme to be 76 bits.

While most of these reported FRR and FAR for voice and speech patterns are comparable to our work, it is important to note the methodologies to estimate entropy and key strengths differ widely and are in most cases not directly comparable. We report a very conservative estimate based on an improved guessing algorithm that also provides a concrete strategy an adversary could follow to guess keys.

4.8 Summary

We explore the use of body impedance for deriving secret keys with the aim of minimizing leakage of sensitive biometric data. Our construction is based on Randomized Biometric Templates [121] and on features extracted using convolutional networks in a Siamese architecture. According to our experiments, the proposed feature extraction method achieves an average false reject rate of 8.6% and a false accept rate of 1.9% for impersonation attempts.

The underlying set of biometric features are found using a recent feature learning technique based on Siamese networks. We adapt this technique to work with body impedance measurements, and achieve significantly better results compared to a feature extraction technique based on frequency analysis.

In order to assess strength of generated keys, we improve the approach proposed in [116] to estimate the required guessing effort. This is far from trivial as an adversary might exploit public and leaked user-specific (auxiliary) information. By taking this into account, our guessing strategy results in a more conservative estimate than commonly used methods, such as summary statistics based on entropy. Using our conservative approach, the presented evaluation shows that strength of generated biometric keys is on a par with keys obtained using similar schemes applied to voice biometrics and speech patterns. We estimate the majority of experimentally generated keys to provide between 30 and 60 bits of entropy.

In summary, this thesis is first to show that biometric characteristics derived from body impedance can be successfully used as a source to generate secret keys and biometric templates that do not leak sensitive information.

More importantly however, the results of this chapter show that body impedance-based recognition can be hardened through key and template generation, which makes body impedance attractive in a wide range of security applications.

5 Driver Authentication Using Body Impedance

This chapter makes use of the techniques we developed in the two preceding chapters. It presents an application of body impedance recognition for driver and message authentication in Vehicular Ad-hoc Networks.

Contents

5.1 Introduction

The technology of self-driving cars and driver-assistance systems has matured to a point where vehicles start to make decisions on behalf of the driver or operate completely autonomously. The introduction of vehicular ad-hoc networks (VANETs) will take this level of autonomy a step further and allow the exchange of information between vehicles while greatly improving efficiency and safety on the road. In fact, the European Union is putting forward a plan which mandates every vehicle to continuously broadcast its coordinates on a VANET. However, when relying on vehicle-to-vehicle communication to make life-critical decisions such as emergency braking, information authenticity and integrity is of paramount importance. Current schemes that satisfy these properties tie the identity of a vehicle's owner to its messages and discourage malicious behavior under penalty of prosecution. Assuming that driver and owner are the same person makes it difficult to identify the source of a message when needed, for example, in case of a rogue network node or in the event of a traffic offense. This is especially relevant for increasingly popular car sharing schemes and transportation as a service where vehicles are owned by mobility providers. In this thesis, we propose a message authentication scheme based on biometric information that provides traceability of each message to the driver, and enables exclusion from the network on a per-individual basis. We analyze the security of our scheme and demonstrate feasibility by implementing our scheme using a driver's body impedance, an unobtrusive biometric modality that can be acquired while holding the steering wheel. Our evaluation is supported by data gathered in a user study with 33 subjects conducted under simulated driving conditions.

5.2 Motivation

Vehicular ad-hoc networks (VANETs) are a cornerstone for the next generation of transportation. Using dedicated short-range communication, these networks allow the wireless exchange of information between vehicles' on-board units (OBU) and stationary roadside units (RSU). The exchanged information facilitates safety applications, such as local danger warning and cooperative collision avoidance, as

well as efficiency functions, e.g., cooperative adaptive cruise control and platoon management [141].

The IEEE 802.11p [142] transmission standard enables vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2X) communication in vehicular networks. Due to its broadcast nature, messages sent on a VANET can be received by anyone within a 300 meter radius. While this provides situational awareness and increases road safety, it also introduces major security and privacy challenges; VANETs have to provide message authenticity and integrity, in addition to protecting passengers' privacy.

If authenticity or integrity are not guaranteed, an attacker can disseminate information that jeopardizes the safety of vehicles and passengers. For example, spoofing vehicle coordinates can create a fictive congestion and trigger the collision avoidance system of individual vehicles. Moreover, messages containing location information such as mandatory basic safety messages [143] facilitate the tracking of vehicles and can infringe on passenger privacy. Consecutive messages from the same vehicle can leak information such as its itinerary or even the identity of the driver [144]. Hence, it is necessary for message signatures to not carry identifying information so that drivers remain anonymous. Meeting these requirements makes it very challenging to conceptualize a secure communication scheme for VANETs. Existing proposals suggest that vehicle-specific signing material be used to secure communication. For privacy protection, these approaches rely on either disguising a vehicle's identity through proxy identifiers or using cryptography to provide sender anonymity. Schemes of the first type have commonly been categorized as pseudonym-based [145], where a trusted authority assigns a set of private/public key pairs to every registered vehicle. The second type, on the other hand, utilizes identity-based or symmetric cryptography, enabling vehicles to generate anonymous signatures [146].

Independent of the underlying scheme, the cryptographic material used to sign messages is always tied to a vehicle and the owner is held accountable for all transmitted messages. This leads to two deficiencies: the vehicle owner could be falsely accused of a traffic offense or the driver could obscure his real identity.

This is particularly problematic for company cars, rental cars, and car-pooling services. Precisely these applications, that provide mobility as a service, are projected to show an annual growth of over 30% [147]. Many cities in Europe aim to make it unnecessary for any resident to own a private car in the next decade [148]. Unfortunately, car sharing schemes face additional administrative effort to keep track of who is using a vehicle and comply with, e.g., data protection regulation. Even more importantly, they require users to provide identifying

personal details for every journey while they could otherwise remain anonymous if they owned a car themselves.

Undoubtedly, a VANET infrastructure that can provide fine-grained message verification while providing accountability would allow to relax privacy intruding procedures as they are currently needed. Therefore, we suggest the integration of biometric measurements into the message authentication mechanism. The automotive industry has already started to adopt biometrics, namely for access control and immobilizers, as well as in-car payments, e.g., at charging stations [149]. Our scheme uses repeatedly measured biometric traits during vehicle operation to bind the identity of the actual driver to the emitted messages. This is achieved by signing the messages with biometric-derived, but anonymous commitments. Remote authentication allows the driver to obtain the signing material from a trusted authority that can check the driver's identity against a revocation list. At the core of the scheme lies an integer commitment scheme—a building block that has not yet been introduced to the context of VANETs. It supports anonymous signatures and authentication based on repeated/continuous biometric measurements collected from the driver.

Before explaining our proposed message authentication mechanism, we give a brief overview on the state-of-the-art in VANET message authentication that is based on vehicle identifiers.

We then discuss the few existing proposals that suggest the integration of biometrics into VANETs in some form or another. The idea of incorporating biometric recognition is not completely new, however, most proposed schemes cannot support the use of VANETs as they are envisioned today.

5.2.1 State-of-the-art in VANETs

In order to protect drivers' location privacy, current VANET message authentication schemes rely on anonymous signatures. These allow a trusted authority to resolve messages to a driver's identity for lawful prosecution and provide anonymity against eavesdropping on the network [145]. Most approaches can be categorized into four groups.

Asymmetric cryptography Using an elliptic curve signature algorithm [150], messages are authenticated through a public key infrastructure [151]. Such authentication schemes [145] aim to protect drivers' privacy using pseudonyms which abstract from their real identity and are frequently changed in order to avoid re-identification. Once a vehicle's local supply is depleted, these ephemeral identifiers

need to be refilled which causes additional overhead [152]. To enable revocation, a certificate authority only provides new pseudonyms if a vehicle's identity is not on a public revocation list.

Identity-based cryptography Approaches incorporating identity-based cryptography offer a reduction in resource requirements [153]. These allow signature verification using only a vehicle's identity which serves as its public key [154]. Based on bilinear pairing, two vehicles with a private key obtained from a mutually trusted authority can verify message signatures without revealing their identities [155]. To evict malicious vehicles, the trusted authority has to assign new keys or provide a revocation list to all vehicles [156].

Group signatures Signatures generated using individual private keys distributed by a group leader can be verified using the group's public key [157] providing group members with an anonymity set. To allow revocation and non-repudiation, the leader stores a mapping between keys and vehicle identities [158].

Symmetric cryptography Hashed message authentication codes change dynamically with the message content and are hence anonymous to an observer. The required symmetric keys are provided by a trusted signature verifier, e.g., a roadside unit [159]. It maintains a mapping between vehicles and short-term keys for non-repudiation and revocation. To decentralize message verification, a delayed key disclosure mechanism can be used [160].

5.2.2 Biometrics in VANETs

Literature that mentions the use of biometrics in conjunction with VANETs can be found in [161–163]. These works pioneer the idea of enhancing driver authentication using biometrics and suggest incorporating biometric measurements in message signatures. Unfortunately though, all proposed schemes suffer from flaws or drawbacks. We briefly outline them in the following and refer the reader to Appendix B for a more detailed discussion.

The authors of [161] propose to enhance user authentication in VANETs using face and fingerprint biometrics. They fuse the two biometric modalities to derive a key and encrypt messages using the Exclusive-Or operation. The computational overhead and authentication time of their approach is evaluated in a scenario with roadside units, authentication servers, and up to 100 simulated vehicles. However, since the protocol uses an XOR cipher it allows an adversary to recover the biometric

material via a chosen plaintext attack. This violates privacy by facilitating tracking but also enables the attacker to sign messages on behalf of a victim. Furthermore, the absence of revocation and traceability combined with a fixed message size renders the protocol unusable.

The scheme proposed in [162] offers mutual authentication using vehicle movement. The authors suggest a Keberos-like method by replacing symmetric encryption keys with biometric keys. Without giving a description of how the biometric information is embedded into the protocol, this system is neither mature nor ready for use.

The biometrics-based anonymous authentication presented in [163] uses biometric encryption and suggests temporary MAC addresses for privacy protection. Every pair of vehicles establishes a separate communication session that makes message broadcasting redundant, but at the same time introduces overhead quadratic to the number of communicating vehicles. In addition to that, establishment and exchange of symmetric keys are not specified and the use of biometrics is not motivated. In a simulation, the authors evaluate privacy on a protocol level without considering implications arising from the proposed biometric encryption. The conclusions that can be drawn for a real-world system are therefore limited.

5.3 Bionym Scheme

We present our approach for biometric-based message authentication and describe how measurements are used for signature generation. Then, we describe our system model and the adversarial environment that we consider for the analysis of our system.

5.3.1 General Idea

Even though current proposals for message authentication (see Section 5.2.1) provide message authenticity and integrity, their reliance on locally-stored credentials only allows the coarse attribution of messages to a vehicle. Our approach differs in the sense that it integrates biometric measurements into message signatures for driver-centric authentication. The ability to sign messages is no longer limited to being granted and revoked per vehicle but per-individual. Similarly, a governmental institution, e.g., the police, can attribute actions on the road and disseminated information on the VANET to the actual driver. In order to realized such a fine-grained message verification, we found the integration of biometric measurements to be crucial. The main advantage of biometrics we identified is that repeatedly measuring

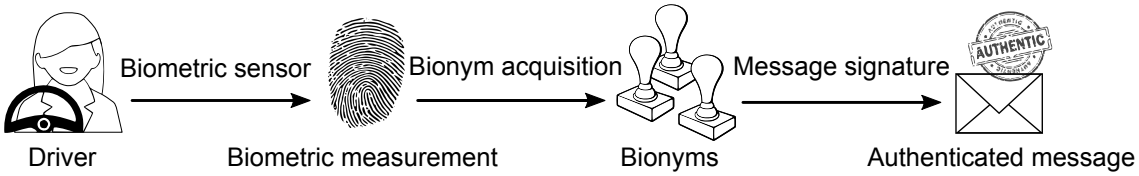

Figure 5.1: Information flow: A driver requests bionyms by providing biometric measurements. The bionyms are anonymized signing material that is used to generate VANET message signatures.

traits while the vehicle is in operation can effectively bind the identity of the current driver to the emitted messages in a VANET and attest the driver's presence.

Based on this insight, we devise a message authentication scheme for VANETs that provides accountability and the revocation of cryptographic signing material on a per-individual basis while still satisfying the security requirements of VANET safety applications.

Our proposed message signature scheme uses anonymous commitments, called *bionyms*, which are unique identifiers designed to authenticate VANET messages. Bionyms are derived from a driver's biometric characteristics following the steps shown in Figure 5.1. The resulting signatures authenticate messages and tie them to the identity of the driver.

In order to acquire bionyms, the driver has his biometric characteristic measured by a sensor built into the vehicle. A request is sent to a trusted authority that contains the biometric measurement in blinded form to anonymously verify the driver's identity. If successful, bionyms are issued to the driver. The bionyms carry the blinded biometric, are unlinkable and have an expiration date after which signatures derived from them are no longer accepted by other vehicles. This forces a vehicle to periodically request new bionyms to refill its repository. During this procedure, the driver's identity can be re-validated and checked against a revocation list.

5.3.2 Strong Driver Authentication

Many car manufacturers offer passive key-less entry and start systems. These systems grant a driver access to the vehicle after an authentication protocol between vehicle and token has been executed successfully. However, a hardware token might be given to another driver to subvert identity verification, similarly to how a car key can be shared.

Thus, in an attempt to personalize access control, car makers have started to adopt biometrics for immobilizers and payment systems. Using a biometric trait

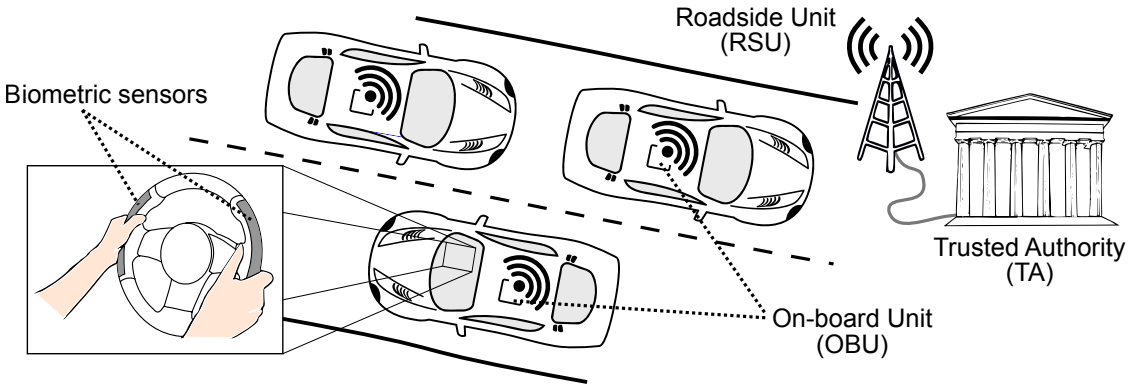

Figure 5.2: System model with the three VANET entities: trusted authority (TA), on-board (OBU), and roadside units (RSU). RSUs relay information between OBUs and the TA. Vehicles can acquire a driver's biometrics via sensors embedded into the steering wheel.

offers non-transferability that cannot be provided by a hardware token alone. In addition, if the biometric modality allows the seamless acquisition of biometric readings, uninterrupted driver verification throughout an entire journey can be guaranteed. Otherwise, if only verified once per trip, additional measures are needed to guarantee the continuity of the driver's identity. For example, a driver's permissions would have to be invalidated immediately if an open door or window is detected. We believe that such additional indicators—especially when relying on simple sensors or switches—are easier to circumvent than continuous biometric recognition.

We propose a combination of a possession-based and biometrics-based mechanism to prevent impersonation while providing privacy to the driver. Our scheme requires a biometric modality and the possession of a simple token (e.g., a smartwatch, bracelet or smartphone). The trust assumptions on the token are minimal, it serves as a storage for a secret which is used to anonymize biometric information. If not combined with biometric information, the secret is meaningless.

We do not consider a solution where the authentication process is executed on an intermediary device for the following reasons: If the token is wireless, special provisions are required to mitigate relay attacks, and if the token is, e.g., a smartcard, the vehicle needs an interface in addition to the biometric sensors. Furthermore, the token would have to be fully trusted and tamper-proof such that it can store biometric information in a secure way.

5.3.3 System Model

Our vehicular network model includes a trusted authority (TA), roadside units (RSU), and on-board units (OBU), see Figure 5.2. A driver provides measurements through a biometric sensor integrated into the vehicle, vehicles exchange information wirelessly.

Trusted authority A driver can enroll in the system at the trusted authority (TA). After verifying a driver's real identity, the TA stores the information necessary to provide a vehicle with driver-specific *bionyms*. This process is supervised and can be combined with the issuance of the driving license. A secure communication channel can be established using the TA's public key, available to every vehicle.

Bionym A bionym is a unique identifier which allows for authentication of messages sent by a vehicle. To compute message signatures, the bionym has to be combined with the driver's biometric information to ensure legitimate use. Although biometric measurements are incorporated in a bionym, the driver's real identity cannot be inferred from this pseudonymous identifier. For privacy protection, bionyms have to be changed frequently to prevent tracking since re-using a bionym makes subsequent messages of the same vehicle linkable. We define a bionym's life-cycle in three stages: (1) acquired, (2) in use, and (3) expired. In case no bionyms are locally available, a vehicle has to acquire them from the TA. Thereafter, these can be used until they reach their expiration time and are no longer accepted by other vehicles.

Driver To participate and benefit from VANET applications, a driver needs to enroll at the trusted authority (TA). This process requires the driver to provide their biometric characteristic and a second factor, e.g., a hardware token, to achieve strong authentication. Once a blinded version of the biometric measurement is stored at the TA, the ability to acquire bionyms is granted. This can be exercised using any OBU as long as the driver provides biometric measurements while operating the vehicle.

Vehicle Vehicles are equipped with an on-board unit (OBU) that enables V2X data exchange via wireless broadcast communication. The OBU includes a tamper-proof device (TPD) that enforces the bionym life-cycle by preventing unauthorized memory access and algorithm interference, e.g., through malware. It is connected to the vehicle's internal bus and can access peripherals like biometric sensors. For authenticity and integrity, the OBU generates signatures by combining a message with a bionym and a recently acquired biometric reading of the driver. A vehicle only accepts messages, if their signatures agree with payload and bionym.

Network The communication of above entities is shown in Figure 5.2. Each vehicle's OBU provides wireless transmission capability for V2X communication with a transmission range of about 300 meters. In order to communicate with the TA, an OBU uses roadside units that are connected to the TA through a wired network. To advertise localized information, such as road and weather conditions, vehicles can communicate with each other via peer-to-peer or geographic routing.

5.3.4 Adversary Model

We consider an adversarial setting that is commonly used when analyzing VANETs: An eavesdropper who aims to reveal the location and identity of drivers by tracking them and an attacker who manipulates traffic by inserting forged messages. In addition to the standard model, we introduce an impersonator who tries to circumvent biometric recognition.

Tracking We assume that a tracking adversary is passive and can observe VANET messages but does not have access to a victim's vehicle. He is also not physically following his target or maintaining an extensive camera network, both of which would make tracking trivial. We assume that the only information available to the adversary is the location periodically broadcast by the victim's vehicle. This information becomes a privacy threat if it can be attributed to a driver's identity. In particular, if an adversary cannot only derive the location of the victim at a single point in time but also learns about entire journeys, chances to identify an individual increase [164]. Generally, a tracking adversary is considered successful in two cases: if he can derive a static identifier including the victim's identity directly from VANET messages or if the observed trip segment is identifying the victim.

Traffic manipulation An adversary can influence traffic by broadcasting fabricated messages that report non-existing accidents or ghost vehicles. These attacks can create a fictive congestion which forces drivers to slow down or even deviate from their original route. An attacker can use this to facilitate car-jacking, robberies, or otherwise cause economic damage.

Impersonation The adversary attempts to impersonate a driver in order to send VANET messages on their behalf. If the adversary is physically present in a vehicle, he can present his own or forged biometric material to the on-board unit (OBU) of the vehicle and claim to be the victim. If the adversary is remote, i.e., does not have access to an OBU, he can try to forge VANET messages using

captured biometric information. Either way, we assume that the adversary has not obtained a valid biometric reading of the victim and cannot break the TPD to gain access to the inner-workings of the OBU.

Independent of the type of adversary, we assume computations are limited to polynomial time and the computational Diffie-Hellman assumption (CDH) holds within a properly constructed cyclic group. Additionally, we base our model on an enrollment phase that cannot be influenced in an adverse way, i.e., the trusted authority detects non-conformity of the procedure.

5.4 Message Authentication Protocol

We describe our message authentication protocol which uses a driver's biometric characteristic to generate digital signatures. Although these signatures are directly linked to the driver's identity, no biometric measurements are transmitted over the network. Moreover, it is not possible to extract the biometric information from message signatures.

The technique we use leverages properties of the integer commitment scheme proposed by Damgard et al. [165]. Each signature is created as a non-interactive proof of knowledge certifying the validity of the biometric measurement provided by the driver. Recipients can verify signatures using an attached commitment. In case a commitment is used for multiple messages, they can be attributed to the same sender and an eavesdropping attacker can therefore reconstruct segments of a driver's journey. To avoid this, our protocol provides OBUs with a set of independent commitments called *bionyms*. Bionyms are indistinguishable from random and hence cannot be linked to the same sender. By changing signing material, a tracking attacker would have to re-identify the victim's vehicle every time an OBU switches to a fresh bionym. The approach by Freudiger et al. [166] proposes to harden such identifier switches using encrypted mix-zones.

Due to the immutable nature of physiological characteristics, unlike passwords, biometric features cannot be revoked or changed once leaked. Therefore, our protocol ensures that biometric information is only transmitted or stored in blinded form.

Our protocol is divided into three phases: *Enrollment* is performed once for every driver, *Bionym Acquisition* is executed periodically when the driver's vehicle requires new bionyms, and *Message Authentication* is invoked every time a message is transmitted on the VANET.

Algorithm 3: BioSign

Global: bases g and h, cryptographic hash function $H(\cdot)$
Input : commitment c, biometric key k, blinding factor γ, payload p
Output : signature σ
$y \leftarrow \mathbb{Z}\,, \quad z \leftarrow \mathbb{Z}$ //draw randomly
$d = g^y \cdot h^z$
$e = H(c, d, p)$
$u = y + e \cdot k, \quad v = z + e \cdot \gamma$
return $\sigma = (d, v, u)$ //pack signature

Generate signature σ for payload p under commitment c with biometric key k and blinding factor γ. Signature σ contains key k only in blinded form.

Algorithm 4: BioVerify

Global: bases g and h, cryptographic hash function $H(\cdot)$
Input : commitment c, signature σ, payload p
Output : *True* or *False*
$(d, v, u) = \sigma$ //unpack signature
$e = H(c, d, p)$
if $g^u \cdot h^v = d \cdot c^e$ **then**
  return *True*
end
return *False*

Verify signature σ of a payload p using commitment c.

5.4.1 Set-up Phase

Before explaining each protocol phase, we briefly outline the onetime set-up procedure to establish protocol parameters. For a detailed description we refer the reader to [165]. The TA constructs an algebraic multiplicative group G such that computing roots of random elements in G is computationally infeasible. The public parameters (g, h) are determined as follows: $h \in G$ is a randomly picked element and $g = h^\omega$ for a random secret ω. The description of G, g, and h are announced to all vehicles in the TA's province and $g \in \langle h \rangle$ is proven without revealing ω (via a Schnorr signature based statistical proof of knowledge). In addition to (g, h), the TA's public key K_{TA}^+ and the two algorithms BioSign and BioVerify are assumed to be public. These algorithms implement our commitment scheme as shown in Algorithm 3 and Algorithm 4.

5.4.2 Enrollment

Before a vehicle can sign messages under a driver's identity, the driver has to be registered with the TA. As shown in Figure 5.3, a driver enrolls in the system by providing his identity D, a biometric measurement q, and a secret s stored on a token.

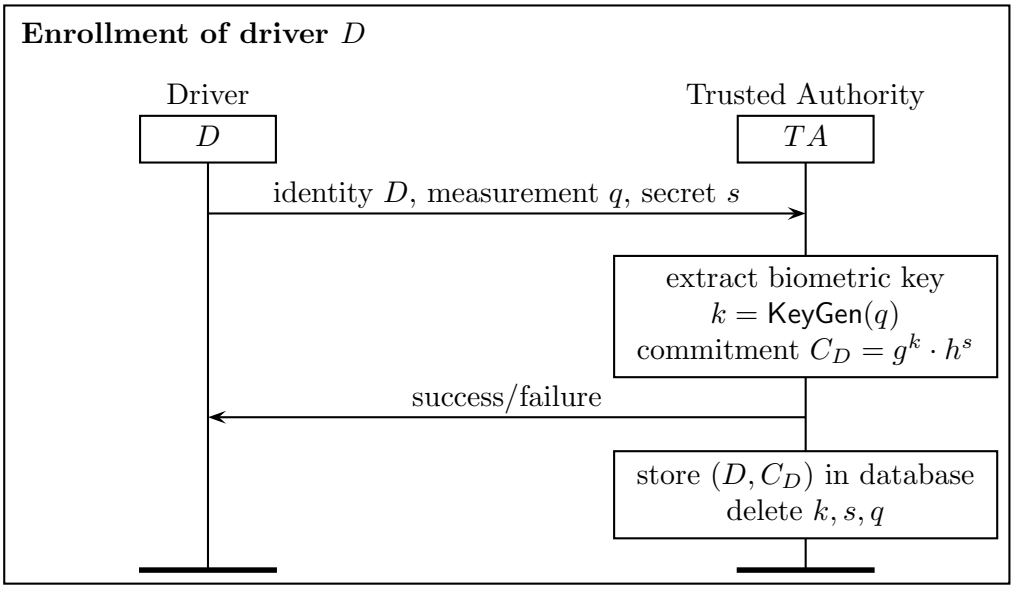

Figure 5.3: Enrollment of driver D with the TA. After identity is confirmed, the enrollment commitment C_D is established.

An officer at the TA verifies D's identity and ensures that D has provided a genuine biometric measurement q by following the intended acquisition procedure. Afterwards, the biometric key $k = \mathsf{KeyGen}(q)$ is derived. $\mathsf{KeyGen}(\cdot)$ is a key derivation function analogous to the one described in the previous chapter. It transforms the biometric characteristic into a cryptographic secret in an irreversible way similar to a hash function. The resulting secret is indistinguishable from random and must not allow the inference of any biometric information, see the requirements for biometric key generation in Section 4.2.2. For instance, if body impedance is chosen as the biometric characteristic, the key k could be generated using the key derivation scheme presented in the previous chapter.

After key k is successfully generated, the TA creates an enrollment commitment $C_D = g^k \cdot h^s$ that encapsulates k and thereafter erases s, k, and q. In order to guarantee that the biometric key cannot be extracted, s is used for blinding. Finally, the driver is notified about the outcome of the enrollment process and the tuple (D, C_D) is stored.

From now on, when interacting with the TA, the OBU computes C_D on-demand to identify as D. As C_D is static, it must not be transmitted in the clear but encrypted with the TA's public key K_{TA}^{+} to prevent tracking.

5.4.3 Bionym Acquisition

In order to send authenticated messages on a VANET, an OBU has to acquire signing material, i.e., commitments in the form of bionyms. Figure 5.4 shows

this process starting with the driver providing his secret s and a new biometric measurement q' to the vehicle. In case the driver prevents the acquisition of q', the vehicle will not start if stationary or not be able to participate in the VANET if already moving. After the OBU has extracted the biometric key $k = \mathsf{KeyGen}(q')$, it uses a fresh asymmetric key pair K_V^-/K_V^+ for bionym retrieval. We note that while the biometric has to be provided for each acquisition, s can be stored on the OBU for the duration of the trip. The tuple of public key K_V^+, a random seed δ_{seed}, the number of bionyms N, and the driver's identity D form the bionym request r. For authenticity, this request is signed under commitment C_D, key k, and secret s using BioSign. We note that for the generation of N bionyms, N randomization values are needed. To keep the communication overhead small, both entities, OBU and TA expand δ_{seed} to a sequence of N numbers using a cryptographically secure pseudo-random number generator.

Upon receiving an acquisition request, the TA loads the commitment C_D for the claimed identity D and authenticates the request using BioVerify. If successful, δ_{seed} is expanded to the randomization values δ_i to compute the bionyms $b_i = C_D \cdot h^{\delta_i}$ and sign them under the public key K_{TA}^+. The TA stores the real identity D with the issued bionyms for non-repudiation. Afterwards, the TA returns the set of bionyms encrypted with the public key provided by the OBU to prevent an attacker from linking the enclosed bionyms. The vehicle then verifies the bionyms using its own sequence of randomization values δ_i. The resulting set $\mathbf{B}$ of bionym and blinding factors γ_i is stored for later use.

5.4.4 Message Authentication

The VANET message authentication is depicted in Figure 5.5. Analogous to an acquisition request, the biometric key k, a bionym b_i, and blinding factor γ_i are required to sign a message m. First, a new biometric measurement q'' is acquired from the driver to extract the biometric key k. To prevent linking attacks, the OBU chooses a fresh bionym b_i and the corresponding blinding factor γ_i. Signature σ can be computed using BioSign and immediately verified through BioVerify locally. If σ is not valid due to an incorrect k, another biometric measurement is acquired and k regenerated. Finally, the message with signature and bionym is broadcast to nearby vehicles.

Similar to TA during the bionym acquisition phase, receiving OBUs use BioVerify to check signatures. However, since the sender committed to the bionym at the TA, the bionym's authenticity can be confirmed by the receiving OBUs using TA's public key K_{TA}^+.

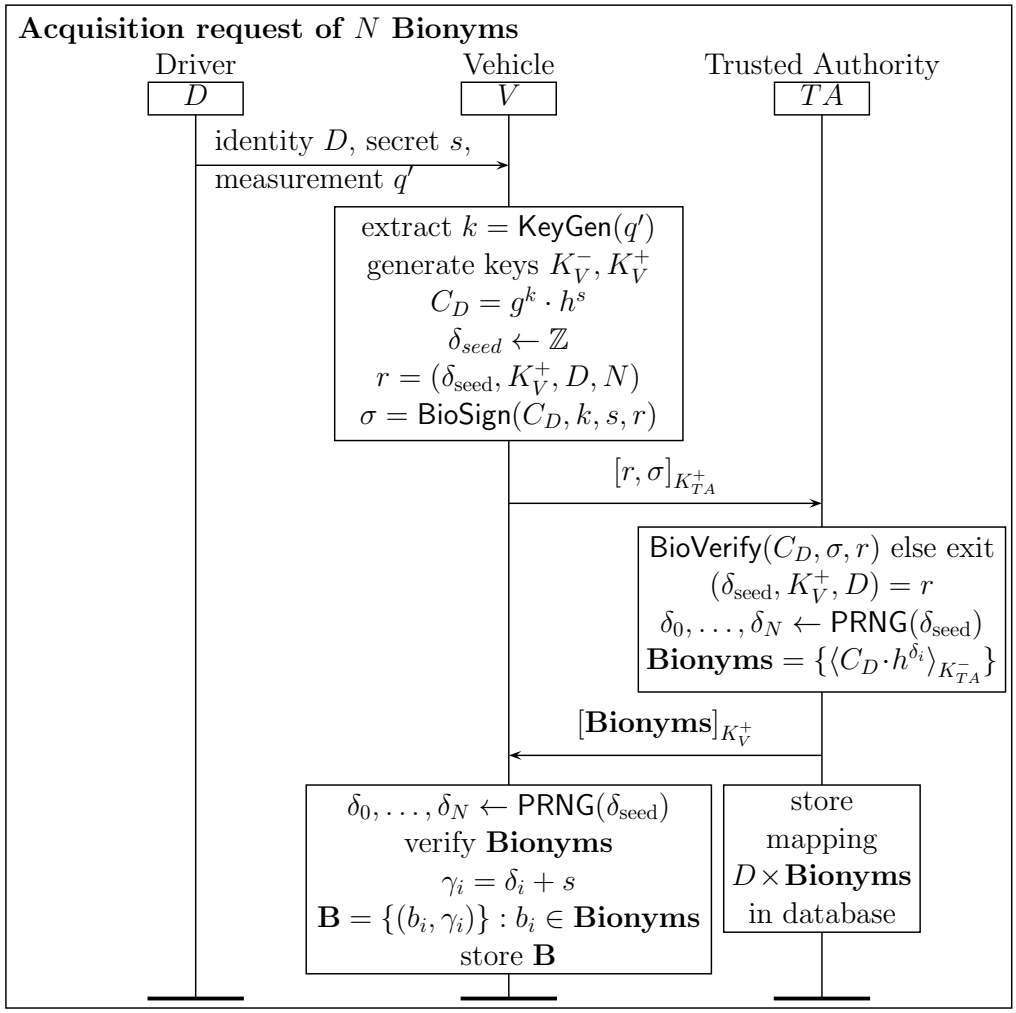

Figure 5.4: Acquisition of N bionyms: $[\cdot]_{K_X^+}$ provides encryption and MAC based message integrity, while $\langle \cdot \rangle_{K_X^-}$ provides signatures using asymmetric cryptography with public/private keys of X.

5.5 Security Analysis

We start this section with a proof of correctness for signature verification during bionym acquisition and message authentication. Afterwards, the system's resilience against protocol level attacks is analyzed. As outlined in the adversary model, these can either be passive eavesdroppers or active manipulators which have full knowledge of the protocol and the parameters g and h. However, we assume that the TA's public key K_{TA}^+ is securely distributed to all OBUs and has not been compromised.

5.5.1 Proof of Correctness

Bionym acquisition and message authentication use the BioSign algorithm to generate signatures based on a commitment, biometric key, and a secret. The selected

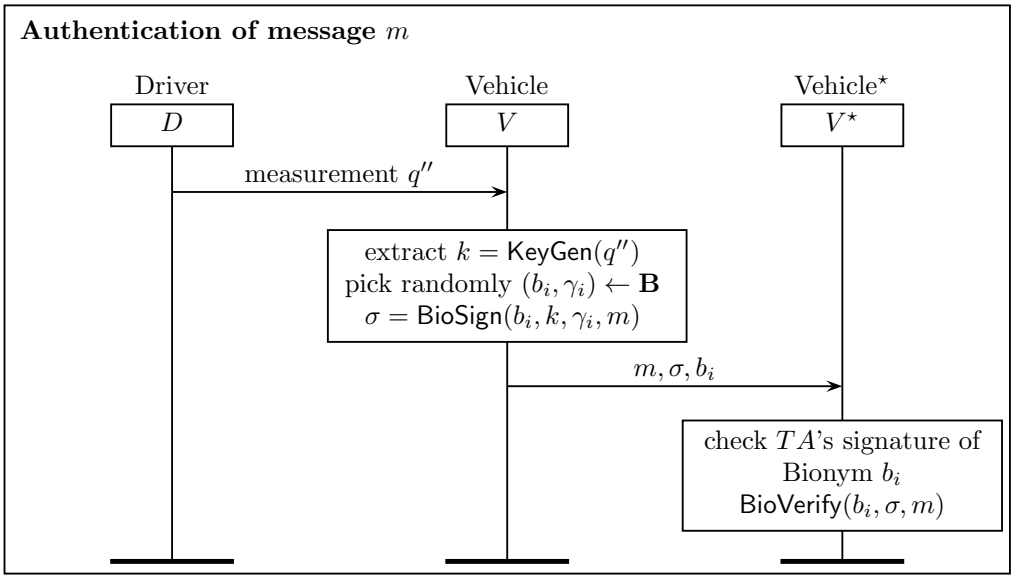

Figure 5.5: Authentication of message m: Vehicle V signs m and transmits the message, signature and bionym over the network. Nearby vehicles $V^\star$ validate the bionym and verify the signature.

commitment depends on the sender's intention: when interacting with the TA, it is necessary to identify the driver for non-repudiation, hence the persistent commitment C_D is used, when signing a VANET message, providing a valid anonymous signature is sufficient and protects the driver's privacy. Therefore the unlinkable bionyms b_i come to use. In either case, Equation 5.1 can be verified using BioVerify and holds if and only if BioSign is provided with the correct inputs. The values for blinding factor γ and commitment c are determined as shown in Equations 5.2 and 5.3 for acquisition and message authentication phase, respectively.

$$g^u \cdot h^v = g^{y+ek} \cdot h^{z+e\gamma} = g^y \cdot h^z \cdot (g^k \cdot h^\gamma)^e = d \cdot (c)^e \tag{5.1}$$

$$\gamma = s, \quad c = C_D = g^k \cdot h^s \ : \text{ when interacting with TA} \tag{5.2}$$

$$\gamma = \delta_i + s, \quad c = b_i = C_D \cdot h^{\delta_i} = g^k \cdot h^{s+\delta_i} \ : \text{ when signing/verifying messages} \tag{5.3}$$

For a detailed explanation of the integer commitment scheme used by the algorithms BioVerify and BioSign the reader is referred to [165].

5.5.2 Passive Eavesdropper

The passive adversary uses observed transmissions to compromise secret s, biometric key k or link consecutive messages of a vehicle for driver tracking and identification.

In general, a driver's secret s is only used in its blinded form $\gamma_i = \delta_i + s$. Therefore, the passive adversary cannot derive any information about its value. The

seed δ_{seed} used in a cryptographic pseudo-random number generator to generate the randomization values δ_i is encrypted with K_{TA}^{+} and hence protected when transmitted. Further, when overhearing a message, the adversary only learns its signature $\sigma = (d, v, u)$ and bionym b_i. The message m and value d are independent of k and γ_i, and therefore, neither can reveal information about the key or the secret. The values of $u = y + e \cdot k$, $v = z + e \cdot \gamma_i$ and the bionym b_i include k and γ_i in blinded form. Hence, the security of k and γ_i relies on the integer commitment scheme used for signature generation and verification. Damgard et al. [165] show that this scheme is unconditionally hiding under the root assumption. Therefore, any algorithm that could with non-negligible probability extract the biometric key k or γ_i from m, d, v, u and b_i would break the computational Diffie-Hellman assumption.

5.5.3 Bionym Linkability

For non-repudiation purposes, the TA can attribute bionyms to drivers using the stored mapping. A passive adversary cannot resolve a bionym directly to an identity. However, he can leverage a bionym's uniqueness in combination with a vehicle's status messages to keep track of its location. As a defense mechanism, an OBU must change bionyms periodically [167] and make use of the fact that they are indistinguishable from random due to the randomization values δ_i. Bionyms by themselves are also unlinkable, since both the acquisition request and the returned set of certified bionyms are encrypted. To guarantee that bionyms cannot be linked due to the context they are used in—e.g., two messages with different bionyms, but close location information might be linked by and adversary—OBUs must adopt a change strategy similar to pseudonym-based schemes mentioned in Section 5.2.1. Bionym changes can be hardened using encrypted mix-zones [166].

5.5.4 Active Manipulation

A Dolev-Yao [168] adversary has full access to the transmission medium, but is not in possession of a victim's secret s or biometric key k. The goal of the adversary is to disseminate messages that are accepted by other vehicles and either appear as if they originated from the victim or cannot be attributed to any real identity. We first focus on the case where the adversary replays or forges messages. Then, we analyze an attacker who aims to obtain or forge bionyms.

Guessing and message forging When modifying existing or creating new messages, the adversary must make the verifier accept the check $g^u \cdot h^v = d \cdot (b_i)^e$. This means that the signature $\sigma = (d, v, u)$ of message m has to satisfy $d = g^y \cdot h^s$ and $e = H(b_i, d, m)$ which is embedded in v and u. If the adversary cannot create a message that passes the verifier's check, message authenticity and integrity is fulfilled. According to the hiding property described above, and assuming that $H(\cdot)$ is resistant against second-preimage attacks, an adversary indeed would have to simultaneously guess either (k, γ_i) or (y, z). While k and γ_i can be directly used for signature generation, compromising y and s allows to "unblind" u and v and obtain values to compute valid signatures. However, the values y, z and γ_i are picked uniformly at random, which means they can only be guessed with negligible probability. The difficulty of guessing k depends on the amount of uncertainty the biometric trait exhibits across a population and how well $\mathsf{KeyGen}(\cdot)$ can extract this entropy.

If we assume that body impedance is used as the biometric modality and run $\mathsf{KeyGen}(\cdot)$ from Chapter 4 to generate keys, results from the experimental evaluation in Section 4.6 apply here, as well. However, the acquisition of body impedance is slightly different from the measurement method in the previous two Chapters 3 and 4. The drivers hold a steering wheel and move their hands. We discuss this further in Section 5.6.3.

Message replay From the adversary's point of view, a completely new set of values d, v, and u is used every time a signature is created. A replaying adversary cannot reuse the captured values v and u, because they depend on $e = H(b_i, d, m)$, which in turn depends on message m. Assuming that $H(\cdot)$ is resistant against second-preimage attacks, the adversary cannot produce a meaningful message $m' \neq m$ that corresponds to a captured σ.

The adversary is therefore left with replaying the same message m, but the ETSI VANET communication standard uses timestamps and sequence numbers (as part of the payload) for replay protection [167], which means such a replay attempt is immediately detected.

Bionym replenishment If the adversary makes an attempt to acquire bionyms on a victim's behalf, he has to send a request to the trusted authority. Signatures for bionym request messages are constructed using secret s, commitment C_D and biometric key k; the adversary cannot generate a valid request without knowing all those parameters. Moreover, if C_D is unknown, i.e., it has not leaked form the TA, the guessing has to be done interactively, which can be hardened through rate limiting.

Bionym forgery When an adversary $\mathcal{A}$ is denied the acquisition of bionyms because his identity is revoked or not enrolled, he can try to forge bionyms. If the attacker can obtain his biometric measurement q, he can use $\mathsf{KeyGen}(\cdot)$ to extract the biometric key k and choose a random secret s to compute $C_{\mathcal{A}}$. From there, he can construct a bionym $b = C_{\mathcal{A}} \cdot h^{\delta}$ using any δ. Message signatures based on b can be validated with $\mathsf{BioVerify}(\cdot)$, but the attacker cannot sign b with the TA's key K_{TA}^{+} and thus, messages signed with this bionym will be rejected regardless.

5.6 Evaluation and Results

In this section, we evaluate an implementation of our message authentication scheme. There are several biometric recognition methods that could qualify for the proposed message authentication scheme. We chose body impedance as the biometric modality since it supports effortless acquisition even when captured in a continuous fashion.

We therefore conduct the first user study that assesses the applicability of body impedance in a simulated driving scenario, and review the suitability of this biometric method for our protocol. Our user study does not involve actual vehicles, but is based on a driving simulator displayed on a screen in front of the participants. The test subjects control a virtual car with a consumer-grade steering wheel used for computer games. Our set-up does not feature a gearshift or dashboard and thus cannot fully reproduce an actual vehicle. Nevertheless, our designed driving scenario allows us to draw conclusions about the feasibility of body impedance-based recognition in the context of driver authentication.

5.6.1 Why Body Impedance?

In Chapter 3, we have established that body impedance can be used for biometric verification and identification [169], and Chapter 4 shows suitability for biometric key generation. Since driving a vehicle involves holding the steering wheel, we envision embedding electrodes into its handles for the unobtrusive acquisition of biometric impedance measurements. Assuming that drivers only infrequently release the steering wheel, a measurement can be captured every time and as long as both hands touch the steering wheel. Acquiring these readings works analogous to the set-up presented in Chapter 3 and is hence nowhere close to having an effect on the driver's body. In our experiments we demonstrate that the short acquisition time of body impedance makes it resilient against occasional interruptions.

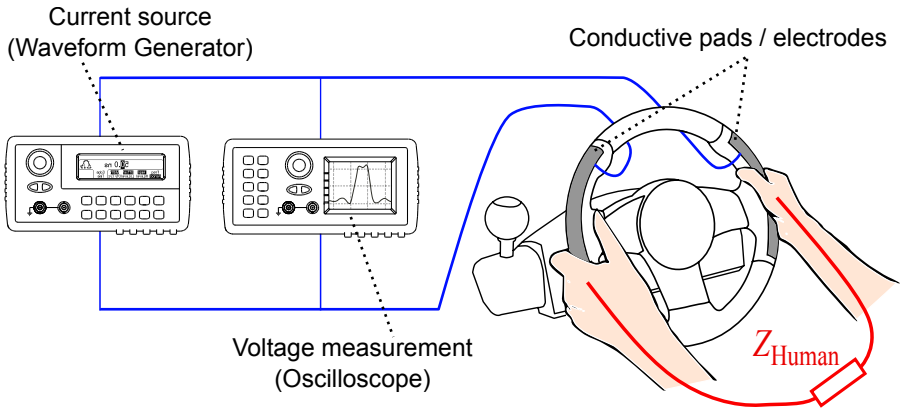

Figure 5.6: Body impedance measurement set-up. When driving, the human subject completes the electrical circuit and the body impedance Z_{Human} can be acquired.

Other biometric methods that could verify a driver's identity and thus constitute viable contenders along with body impedance are face recognition [170], eye movement tracking [171], driver posture [172, 173], and electrocardiography (ECG) [174]. We believe that those modalities could be used as part of our proposed protocol, but we favor body impedance for the following reasons:

A camera directed at a driver could repeatedly verify the face. However, special care has to be taken as the camera has to cope with rapidly changing lighting conditions, e.g., at night time. Furthermore, a camera could record the surroundings and possibly intrude privacy unless head tracking is used and sensitive parts are blurred out.

As for eye tracking, contemporary devices are still expensive, require calibration, and are sensitive to changes in lighting and user positioning.

Preliminary research in posture recognition showed a low level of uniqueness among drivers, making it unsuitable for biometric authentication without further refinement.

To acquire ECG data, three electrodes have to be placed on the subject, one on each forearm and a reference electrode at the leg. This constrains the range of motion while driving and makes (re-)attaching the electrodes cumbersome.

5.6.2 Prototype Set-up and User Study

We improved our proof-of-concept acquisition system described in Chapter 3 and incorporated the electrodes into a steering wheel. A schematic of our set-up is depicted in Figure 5.6. Two electrodes are attached to the sides of an off-the-shelf steering wheel used to control computer games (see Figure 5.8). As soon as the wheel is touched, i.e., short-circuited by a human, one electrode emits a

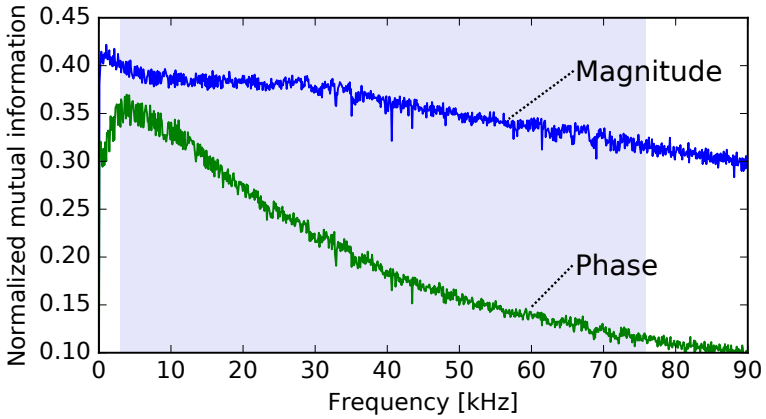

Figure 5.7: Normalized mutual information for magnitude and phase of body impedance. Shaded area highlights the frequency range used for classification.

frequency sweep every 3 seconds. These sweeps range from 100Hz to 100kHz and take 300ms to generate. After the signal has traveled through the human body, it is measured by the second electrode. The emitted signal and the measured signal are then correlated and transformed to the frequency domain in order to compute the complex impedance of the human.

Conceptually, this new set-up is identical to the one presented in Chapter 3 with the exception of electrode arrangement and elicitation signal. The emitted signal is a frequency sweep instead of a pulse. In Section 3.4.3, we concluded that for classification of impedance measurements short square pulses outperform frequency sweeps. The sweeps used here are wider and at much higher frequencies than the ones tested initially. In fact, these wider sweeps include the features extracted by pulse-response recognition (described in Chapter 3.4) and a host of additional features, potentially resulting in even higher distinguishing capabilities.

The additional features come at a cost, however: a sweep from 100 Hz to 100 kHz takes longer to execute than a short square signal—compare 300 ms acquisition time with a 100 ns pulse. Moreover, features extracted from higher frequencies are more prone to noise. Therefore, this modified set-up might not be applicable to continuous user authentication, as described in Section 3.3, where we envision to capture impedance measurements while typing on a conductive keyboard and short measuring times are crucial.

In a driving scenario, we found these adjustments to be advantageous and we conducted a study with 33 participants (26 male, 7 female) using the slightly modified set-up. The participants were aged between 25 and 40, and 94% of them owned a driver's license. The ratio between left-hand side and right-hand side drivers

Figure 5.8: Experimental set-up from the perspective of the test subjects. The steering wheel controls the vehicle in the computer game shown on the screen in front of the participants.

was 17:14. The participants were asked to sit and hold the steering wheel as if they were in a real car. To eliminate potential bias, the goal of these measurements was revealed after the fact[1]. We measured the participants' impedance in two configurations: (*static*) while keeping the wheel still and (*driving*) while controlling a vehicle in a computer game shown on the screen in front of them. For each participant, we acquired 60 measurements over two 4 minute sittings. Having the participants steer a vehicle in first-person view allows us to elicit hand movements close to actual driving and analyze the impact of hand positioning and steering motion on recognition performance.

On average, participants took 15 left and 12 right turns during their session. 95% of the acquisitions during turning succeeded (steering wheel held with both hands) while 20% of the acquisitions during straight segments failed (steering wheel held with only one hand).

Our user study and the modified measurement set-up were approved by the Central University Research Ethics Committee of the University of Oxford, reference R55051.

Features and classification To identify the significant frequencies of the collected impedance measurements, we calculate the normalized mutual information

[1]Some participants found out that the conductive pads were an essential part of the study, as they are attached to the wheel in a very visible way.

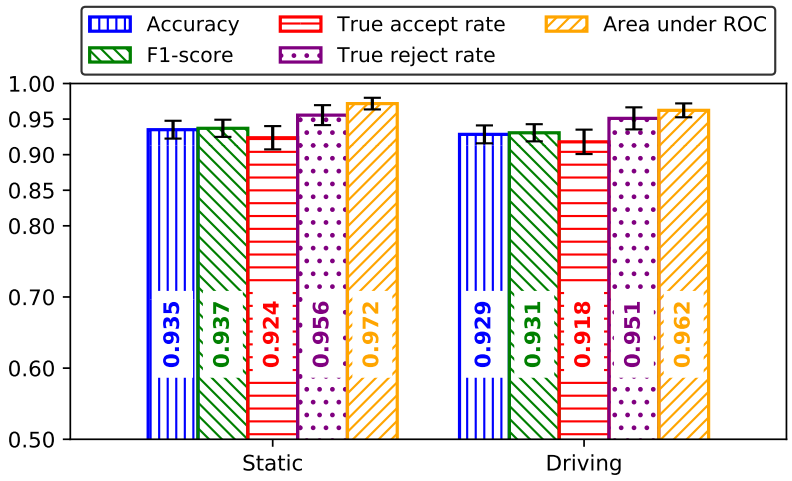

Figure 5.9: Body impedance recognition performance for 5 different metrics. Results are averaged over the test population and obtained through 5-fold cross-validation. Error bars represent 95% confidence intervals.

(see Figure 5.7). Both, magnitude and phase, are most specific at lower frequencies, letting us choose the complex impedance values at the frequencies between 2 kHz and 76 kHz as the features for classification. A feature vector extracted from an impedance measurement is the array of its magnitude and phase values. We feed these arrays into a random-forest classifier that decides whether a measurement matches a reference reading of the claimed identity. In total, we train one binary classifier per participant by dividing the samples into two classes. One class comprises the samples of the participant whose classifier is currently trained and the other class contains all other samples.

5.6.3 Biometric Recognition and Impersonation

We evaluate the performance of biometric recognition based on body impedance using the set-up described above. The obtained results determine how likely an enrolled driver is recognized (i.e., bionym request and message authentication are successful), as well as how likely an attacker can successfully impersonate a victim whose secret s has been compromised, that is when the security hinges on the biometric recognition only.

The results in Figure 5.9 depict recognition performance for the configurations *static* and *driving* using 5 different metrics. In order to minimize statistical overfitting, we conduct 5-fold cross-validation when estimating the performance metrics. We further ensured that a ratio of 1:1 between measurements of an authorized driver and impersonation attempts was used for both, training and testing data. This reduces the effects of class imbalance when training and evaluating the classifier.

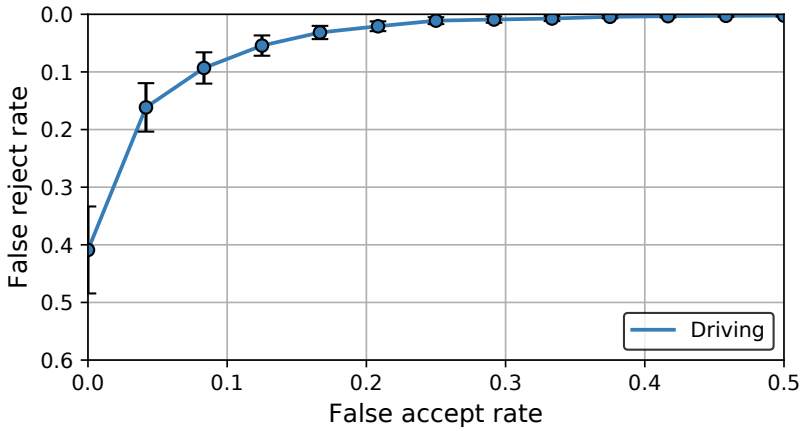

Figure 5.10: Receiver operating characteristic (ROC) for the driving scenario averaged over the test population. Error bars represent 95% confidence intervals.

The results indicate that verification works well; the performance for *driving* is on par with *static*. Even though most of the participants stated that they occasionally let go of the steering wheel when they drive their own cars, in our simulation, steering movement and the way drivers hold the wheel affect authentication only marginally. Unless the driver has a completely one-handed driving style, measurements can be captured in rapid succession and compensate for short interruptions.

Moreover, as can be seen in the ROC plot in Fig. 5.9, confidence for the false reject rate is high, implying a low chance for a systematic error. Hence, in a false rejection event due to noisy data or procedure errors, the reading can be re-acquired and is likely to be recognized. Therefore, a higher false reject rate can be tolerated to achieve a lower false accept rate and reduce probability for impersonation.

Moreover, as can be seen in Figure 5.10, confidence for the false reject rate is high, implying a low chance for a systematic error. Hence, in a false rejection event due to noisy data or procedure errors, the biometric reading can be re-acquired and is likely to be successfully recognized. Since acquisition time in our experiments was less than 3 seconds, measurements can be captured in rapid succession. Therefore, a higher false reject rate can be tolerated to achieve a lower false accept rate and reduce probability for impersonation.

Biometric key generation Our protocol uses biometric keys instead of a binary decision between accepting and rejecting an identity claim. While the experimental results obtained in this section do not allow to draw conclusions about the difficulty of guessing the biometric keys, they adequately quantify the probability of impersonation.

Due to fact that recognition performance in a (simulated) driving scenario proves equal to the results obtained in Chapter 3 and hand movement affects measurements only we do not repeat the key generation analysis. We believe that the strength of keys derived from body impedance while driving is similar to those acquired in a static environment. In Chapter 4, we established that more than 50% of the generated keys exhibit at least 55 bits of entropy and 46 bits on average across the population.

5.7 Summary

This chapter presents a novel scheme for message authentication in Vehicular Ad-hoc Networks. Our approach enables the incorporation of biometric measurements into message signatures, transferring the responsibility from vehicle owners to actual drivers. This provides features not found in existing schemes, such as the exclusion of individuals from the network and the protection of vehicle owners. We offer these guarantees in addition to conditional identity and location privacy of drivers. Our proposed protocol is therefore especially geared to the future use of vehicles in car sharing schemes and mobility as a service where vehicles are owned by third party providers.

In order to assess the feasibility of our scheme, we conduct a user study under simulated driving conditions, with the help of a computer game. The results show that biometric recognition works sufficiently well and readings are available for signature generation in a timely manner.

Although body impedance seems particularly suited for this scenario, the presented protocol is not limited to the chosen biometric modality. Future work might include the assessment of other biometric methods, such as face recognition, and the testing of our message authentication scheme in actual vehicles.

6

Device Pairing Using the Body as a Transmission Medium

In this chapter, we use the human body as an authenticated channel for the transmission of keying material. Instead of extracting user-specific features, we characterize the body as a transmission medium for the purpose of devising a secure device pairing protocol.

Contents

6.1 Introduction

Device pairing is the process of bootstrapping secure communication between two devices that do not share any common secrets. Often the most challenging part of a device pairing protocol is to establish the identity of the other device, i.e., to make sure that one is establishing a key with the intended device and not someone else. For devices on the Internet, this problem is addressed by relying on certificate authorities to certify the identities of hosts, providing a root of trust when establishing the identity of a communicating party. For smaller devices that do not necessarily have (or need) a certified global identity, certificate authorities are not appropriate. Smaller devices instead often use short range radio technology like Bluetooth, and rely on a human to certify the validity to the other device when pairing, e.g., by visually comparing short strings on a screen, or by typing a number displayed by one device into the other. Such schemes require active participation from a human and the security guarantees provided by these protocols rely on the user performing the correct actions at the correct time. If the user makes any mistakes, the security guarantees of these protocols no longer hold.

In addition to human error, device pairing protocols also impose certain hardware requirements on devices. This is not a problem by itself, as all communication requires some form of hardware support, but screens and input devices place restrictions on the size and shape of devices, e.g., a device may have to have a flat surface, and be large enough to support a usable screen.

In this chapter, we propose a device pairing protocol for smaller devices (e.g., phones, headsets, keyboards, etc.) that mitigates these two problems. Our protocol does require human participation but the user never has to make a security relevant decision and the hardware needed for communication can be any conductive surface on the device. This eliminates the possibility of human error and the scheme remains usable regardless of the physical design of the device (as long as the device is large enough to touch with a finger).

Our scheme is based on the core idea that two devices are allowed to be paired if they are both held by the same human, at the same time. The rationale behind this decision is that if a user is physically holding both devices there are very few ways to secure communication between these devices if the user has malicious intentions. For example, a malicious user could run a device pairing protocol involving short string comparison (or any other mechanism), or physically manipulate the devices to achieve his goal. Our scheme enables device pairing by having the user touch a conductive surface on each device. The human body then serves as a transmission medium for capacitive coupling between the devices which can be used for communication. We call this communication channel the "body channel". Devices can distinguish between messages sent on this channel, and messages sent by a remote attacker, and can thus ignore any message that originates from an external source. This means that two devices held by a user effectively have an authenticated channel between them that can be used for key confirmation. Only a small amount of data is sent through the body channel, so device pairing is fast and easy.

In order to test feasibility of our approach, this chapter presents a device pairing protocol that takes advantage of the body channel to quickly and securely establish a shared secret, without the need for certificates or shared knowledge. We then prove the security of the design by reducing the security of the protocol to the security of the underlying primitives under the assumption that the body channel is read-only to the attacker. The read-only assumption models the fact that the receiving device can tell the difference between messages sent by an external transmitter, and a device which is physically connected to the person performing the pairing. The receiving device can thus ignore any message that originates from an external source, which is equivalent to saying that the channel is read-only for the attacker. Using a proof-of-concept prototype, we present thorough experiments to verify this distinguishing ability and document performance as well as user experience of the protocol.

6.2 Device Pairing in Literature

Prior research has yielded a plethora of methods that implement secure device pairing. Most of them work by having the user authenticate information in an interactive way and augmenting the device pairing process with an out-of-band channel to mitigate man-in-the-middle attacks. Various types of auxiliary channels have been proposed, e.g., the visual channel [175], the audio channel [176], or gesture-based channels [177–179]. Some approaches combine different out-of-bound

channels, e.g., the proposal in [180] uses the acoustic and vibration channel to reduce the risk of side-channel attacks. The authors mask the keying material that is transmitted via vibrations by actively injecting noise on the audio channel. Depending on the platform and the sensors available, many combinations of auxiliary channels are possible. In Augmented Reality headsets, for instance, it is feasible to combine the visual channel with a gesture-based channel, as suggested in [181].

We believe that, in terms of usability, gesture-based approaches such as [179] are most similar to the idea presented in this chapter. The authors of [179] present a device pairing solution for smartphones where the user has to perform a hand gesture to point their phone into the direction of the target device. We consider such an approach comparable to touching two electrodes, which is required for our protocol to work. However, most gesture-based solutions need to explicitly capture and understand the user's intention signaled by their gesture. Therefore, the gesture needs to be sensed by the devices, often requiring another auxiliary channel for that purpose (e.g., the audio channel in case of [179]). Our solution on the other hand does not have to record any movements or gestures and has the advantage of only using one auxiliary channel—the human body.

Finally, a comparison and survey of a multitude of secure device pairing methods can be found in [182]. Worth mentioning is also the study in [183] that measures the influence user perception, security needs and other factors can have on device pairing method choices.

6.2.1 Body Area Networks and Medical Sensors

Most research on the security of body channel communication and body area networks focuses on implantable and body-worn medical sensors, e.g., see [184–186]. While our problem statement is substantially different from medical sensors and on-body devices, there are similarities. Most notably, the fact that the human body can serve as a transmission medium. Some on-body or implantable medical devices use body channel communication to transmit and receive sensor readings, either to and from each other or to establish communication with an external device which is used to monitor and program the implantable devices. An extensive survey and overview of literature about the security and privacy of body area networks and implantable medical devices can be found in [187].

The idea of applying body channel communication to device pairing other than for medical sensors or implants is, to the best of our knowledge, a novel concept not documented outside of this work. The paper that is most related

is [188] which proposes a body area network authentication scheme that does not depend on prior trust among the nodes. It is based on variations in received signal strength due to movement. Nodes that transmit on an on-body channel have a distinctive variation behavior of the signal strength. This behavior is different from a transmission on an off-body channel. The authors exploit this fact and perform clustering analysis to differentiate between an attacker and a legitimate node. This approach has similarities to our idea, but it exploits the physical movement of on-body sensors, an artifact our approach cannot rely on, as device pairing should in stationary and moving environments.

The study in [189] is related to our approach as it proposes a method for robust key establishment among (medical) on-body sensors within a body area network using the body channel. Key establishment is directly related to our problem statement, which is secure device pairing. However, their approach is different from ours, as the authors suggest to inject an artificial voltage signal below the action potential level of a human body to construct a covert communication channel that is secure against an outside eavesdropper. The authors determined the effectiveness of their approach with experiments on a dead mouse.

Our solution to device pairing does not construct a covert channel, but aims to generate an authenticated channel that can be accessed by any device that is in physical touch with the person involved in the pairing operation.

6.2.2 Non-medical Applications of Body Channel

There are few proposals in research literature that deal with applications of body communication outside the realm of body area networks.

The proposal most related to device pairing can be found in [190] which presents a system that allows a user to "wear" a unique key and unlock devices by touching them. The described system consists of a watch-like device that acts as a transmitter and houses a signal electrode. The electrode is in permanent contact with the user's skin around the wrist and emits data encoded in an electrical signal every time the user touches a receiver electrode with his finger. Using capacitive coupling, the data is transmitted to the receiver which can be embedded in a door, smartphone or remote control device. Although such an unlocking mechanism bears a lot of potential, the paper draws no conclusions about the security of such a system.

Also not in the field of computer security, but interesting to mention is the work in [191] which describes a near-field-sensing transceiver for intra-body communication between two or multiple devices. The proposed transceiver features an electro-optic sensor to detect the fields generated by the human body when subjected to an

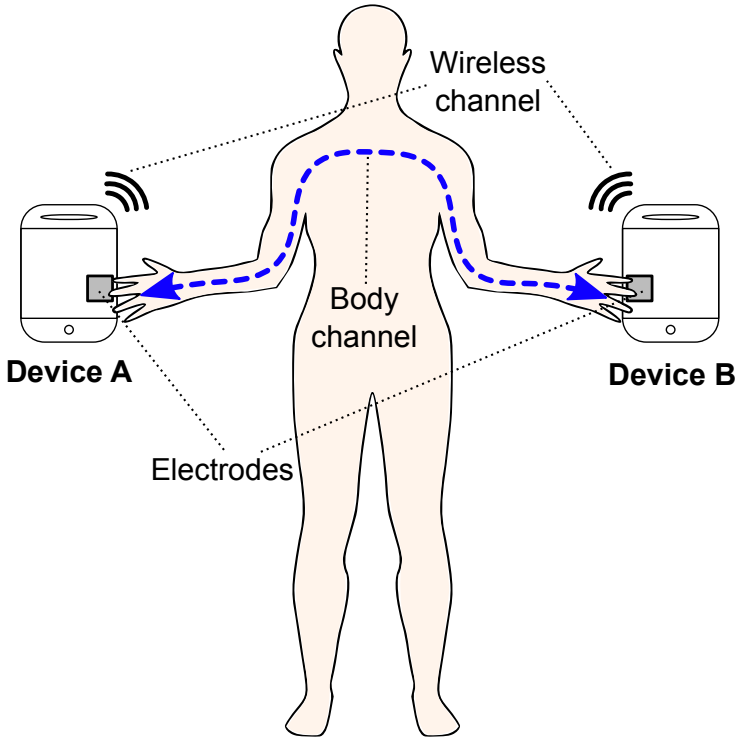

Figure 6.1: A human pairs devices A and B. Both devices can communicate via a wireless channel and use the human body as a transmission medium for a second channel, the human body channel. The body channel is established by touching electrodes on both devices.

electrical signal. The authors' experiments include two transceivers communicating with each other through one and two human bodies.

6.3 Body Channel Pairing

The device pairing method we propose relies on intra-body communication as an out-of-band channel. The basic idea is that two electronic devices should be allowed to perform device pairing if they can successfully communicate with each other through a human body. The fact that two devices can transmit and receive messages using body communication implies that they must be physically close to each other and must be held by the same person. We use this as the criterion for whether two devices are meant to run a pairing protocol with each other and establish a mutual secret. A person can give two devices permission to pair by holding them both at the same time and thereby provides a transmission medium for intra-body communication (see Figure 6.1).

Our proposed device pairing scheme uses capacitive coupling to establish the human body channel, i.e., the "body channel". Our choice to utilize this particular technique is founded on the observations described in the following.

Transmission distance The person pairing two devices should be able to touch them with their hands to perform the pairing. This requires hand-to-hand transmission on the body channel which can over 180 cm in adults. Out of the three techniques for intra-body communication (see Section 2.1.2), capacitive coupling and surface waves are the two techniques that have been reported to cover such a distance reliably. With galvanic coupling only short transmission distances are possible due to the high attenuation of the signal [33, 192]. In addition, the frequency ranges where galvanic coupling operates best are lower than for other techniques, which significantly restricts the data rate for communication [193].

Ease of use and electrode design Capacitive coupling only requires one electrode per device to be in physical touch with the human body, i.e., the person pairing the devices only needs to touch one electrode with each hand (see Figure 2.4). Unlike galvanic coupling, which requires at least two electrodes per device, capacitive coupling can work with a single capacitive touch-electrode per device. This simplifies the implementation of body channel enabled devices and makes the action of pairing two devices straightforward for the user. Additionally, the fewer electrodes there are, the less the effect orientation of transmitter and receiver have on the signal attenuation [194]. We elaborate on the design of the electrodes we use in our experiments in Section 6.7.2.

Electromagnetic interference Surface wave techniques and capacitive coupling can both cover a transmission distance that is sufficient for our application with relatively little signal attenuation. Compared to capacitive coupling, surface wave techniques allow more electromagnetic power to leave the human body during transmission and are more susceptible to external interference. We aim to design body channel communication that is difficult to interfere with from the outside, i.e., with an external radio transmitter. It should require a lot of energy to influence the body channel with a signal source that is not physically connected to the body. Capacitive coupling, which operates at much lower frequencies than surface waves, is therefore better suited for our use case.

6.4 System and Adversary Model

In this Section, we define the system model and the adversarial context that we consider for our device pairing protocol.

6.4.1 System Model

Two devices that do not share any secrets need to bootstrap secure communication. The devices follow the pairing protocol presented in Section 6.5 in order to agree on a mutual secret.

The decision whether two devices should be paired with each other and execute the pairing protocol is made by a human. A person can give the devices permission to run the pairing protocol with each other by physically touching and holding them both at the same time. Only if two devices are held by the same person they are allowed to be paired with each other. If a device is not connected with another device through a person, or if a device is not being held by a person at all, it should not be able to carry out the pairing process.

The devices each have an electrode that when touched by a human enables communication through capacitive coupling. We call this communication channel the *human body channel*. The devices can also communicate with each other on a *wireless channel* (see Figure 6.1). The wireless channel does not have to provide any particular security guarantees for the device pairing to work.

The human body channel is formed when a person is in physical contact with both devices. If a person touches both devices at the same time, one with each hand, the human body acts as a transmission medium for intra-body communication and both devices can send and receive messages on this channel. The human body channel also allows the devices to extract physical properties of received messages to validate if they have indeed been sent over the body channel, i.e., "through" the person who is currently touching both devices.

6.4.2 Attacker Model

We specify three different adversaries: An adversary who *eavesdrops* on the device pairing process, an adversary who tries to perform *remote pairing* with a body channel enabled device, and an adversary who launches a *man-in-the-middle attack* during the pairing of two devices.

- **Remote pairing.** This adversary tries to perform remote pairing with a body channel enabled device. The adversary does not have physical access to the target device and therefore cannot authorize the device to pair by simply touching it. Due to the inability to touch or hold the target device, the adversary cannot establish a body channel for the pairing process, but he can attempt to initiate the the device pairing by sending radio waves from a distance. He might do so while the target device is on its own or while a

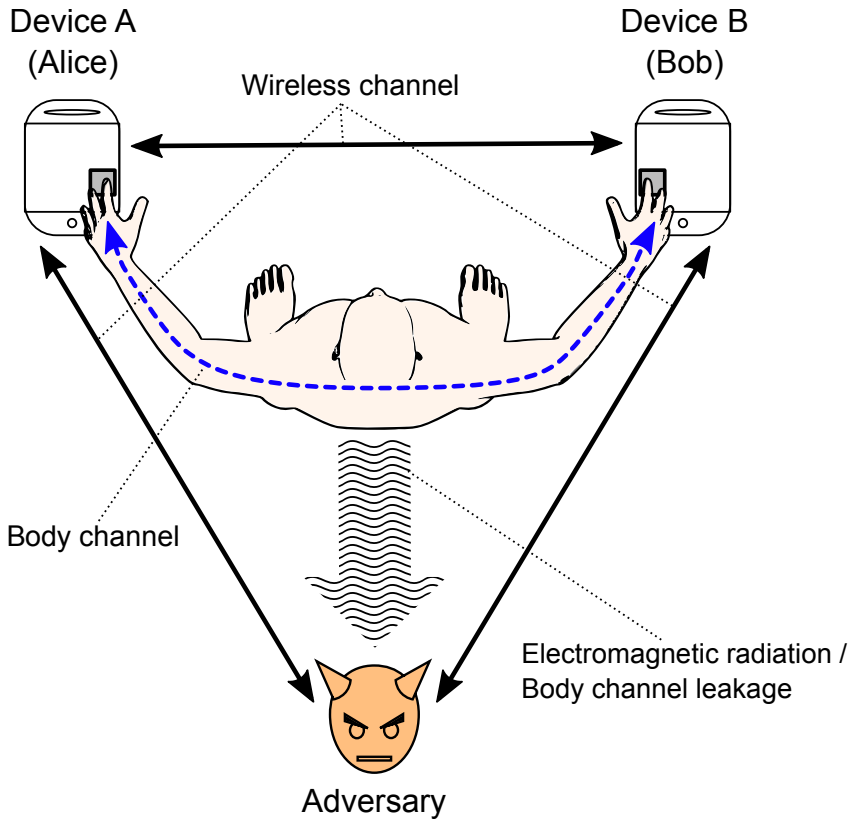

Figure 6.2: An adversary interferes with the wireless channel and records the electromagnetic leakage from the human body channel.

person is in physical contact with the device. It is important to consider such a scenario since a person could be touching the target device accidentally or be part of an ongoing pairing execution.

- **Passive eavesdropping.** This adversary listens on the wireless channel and records the electromagnetic leakage originating from the body channel (see Figure 6.2) in an attempt to learn about the secret that is being agreed on during the pairing of two devices.

- **Man-in-the-middle attack.** This adversary tries to actively participate in the pairing of two devices, with the purpose of making one or both of the devices believe the pairing protocol was successfully completed. We make the assumption that such an adversary can relay, alter and inject messages on the wireless channel as well as record the electromagnetic signals transmitted on the body channel. In addition, the adversary can send electromagnetic signals at the the devices and the person involved in the pairing, but similarly to the remote pairing scenario, we assume that the adversary is not in physical contact with any of the two devices.

For all three adversaries, we assume that they can only establish an actual body channel if they are able to touch the devices or the person involved in the pairing. Further, the devices can successfully extract physical properties of the messages received on the body channel and determine with high accuracy whether a message is an induced radio signal from an outside source or a message send on the body channel.

We thus consider the human body channel as read-only for any signal source other than the devices which are being paired and held by the same person. We show that this is a reasonable assumption in Section 6.8. For the read-only property of the body channel to hold, we state a minimum distance of 50 cm between the adversary and the person involved in the pairing.

Like all other pairing protocols, our proposed pairing mechanism cannot prevent denial of service attacks. Hence, we do not address attacks that have the sole goal of disrupting the communication between the devices.

6.5 Pairing Protocol

Two devices, henceforth referred to as Alice and Bob, jointly agree on a secret using a wireless channel and the human body channel. Alice and Bob follow the device pairing protocol outlined in Figure 6.3. If the protocol terminates, it guarantees that the secret is only known to Alice and Bob, provided they have not revealed it to any other party, of course. The resulting mutual secret can, for instance, be used in subsequent communication between the devices.

The protocol relies on the fact that Alice and Bob can independently verify if the messages they receive on the body channel have traveled through a human body. If they both conclude that the physical properties of the received messages match with the characteristics of the body channel, they must be communicating with each other through the same person. In that case, Alice and Bob must be held simultaneously by the same person and the pairing protocol can terminate successfully.

6.5.1 Protocol Description

The device pairing protocol consists of two steps: key agreement and key confirmation. Alice, who initiates the protocol, chooses a private key a and picks a random nonce N_A. She then sends her identity A, ($g^a \bmod p$) and the freshly picked nonce N_A to Bob on the wireless channel. Bob then picks a private key b and a nonce N_B and sends his identity B together with ($g^b \bmod p$) and the nonce back to Alice. Alice and Bob can now independently construct a mutual secret K and complete the key agreement phase. However, at this point, Alice and

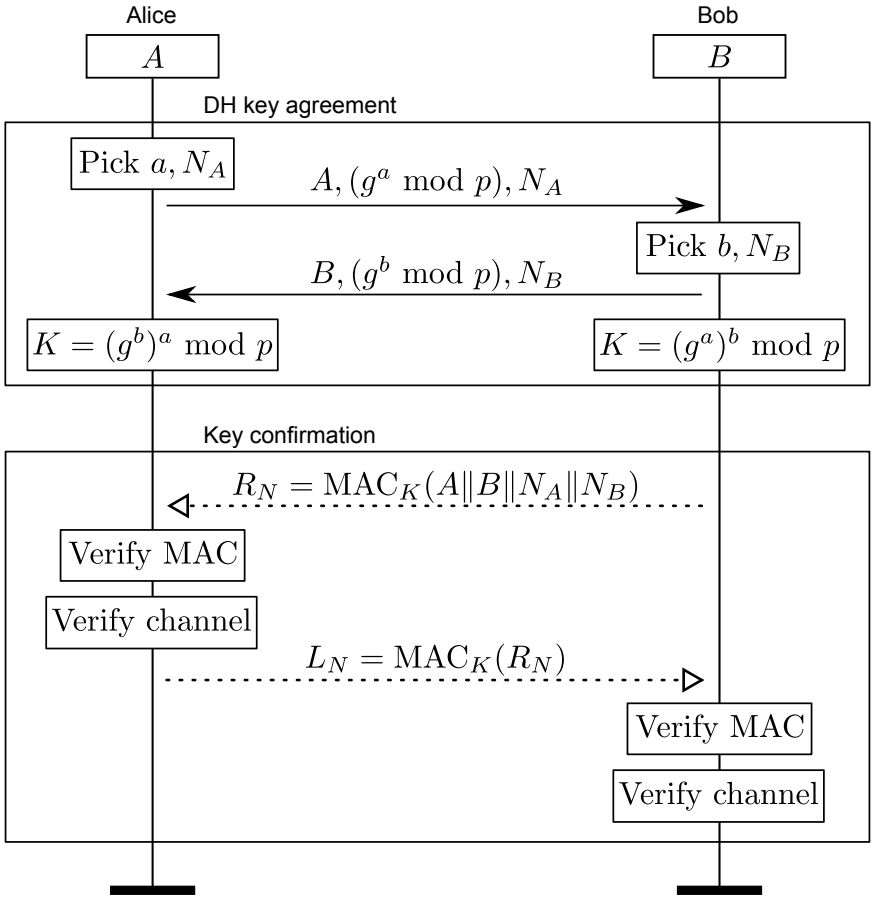

Figure 6.3: The pairing protocol uses the wireless channel (solid arrows) for the key agreement and the body channel (dotted arrows) for the key confirmation.

Bob cannot yet be certain if K is indeed a mutual secret only known by them, since the wireless channel is unauthenticated.

The key confirmation phase follows immediately after the key agreement. Bob computes a message authentication code (MAC) R_N using the newly created key K (or a derivative thereof). The MAC is constructed over the concatenation of the identities and nonces, and is sent to Alice over the body channel. Alice verifies the MAC R_N, and verifies that the message came through the body channel (as described in Section 6.8). If both checks succeed, Alice knows that K is a freshly generated secret shared with Bob. By sending R_N to Alice, Bob demonstrates that he can transmit messages over the body channel and must be connected to Alice through the same human body. He also confirms that he knows K and proves that Alice must have been communicating with him in the preceding key exchange.

Finally Alice computes a MAC of R_N using K, and sends the result to Bob through the body channel. Bob verifies the MAC and the body channel like

Alice did before. This proves to Bob that Alice is in possession of K and can transmit on the body channel.

Termination of the protocol guarantees that the mutual secret K is known to Alice and Bob, and only to them provided none of them revealed it to any other party. Moreover, Alice and Bob can be sure that they were both held by the same person when they ran the pairing protocol. If any of the verification steps fail, the protocol will terminate with an error.

6.6 Security Analysis

The high-level goals of the adversary are to either eavesdrop on the traffic between two legitimate devices, place himself as a man-in-the-middle, or perform remote pairing with a target device.

In this section we show that neither a passive nor active adversary can achieve these goals. We assume that the adversary has full knowledge of the protocol including the public parameters g and p.

6.6.1 Passive Eavesdropping

To show that our device pairing protocol is secure against purely passive eavesdropping, we observe that the only information available to the adversary at the end of the key agreement part of the protocol are the identities of the two devices A and B, the freshly picked nonces N_a and N_B, as well as the public Diffie-Hellman parameters g^a and g^b. The identities are public and do not constitute information leakage. The two nonces are freshly picked independently from the private key, so they cannot reveal any information. If the computational Diffie-Hellman assumption holds for the underlying group, then the adversary cannot get the key K from this information.

Furthermore, we observe that the only additional information the adversary can obtain from the key confirmation part of the protocol are the two different MACs R_N and L_N. The MACs are computed using the key K (or a derived MAC-key), however assuming the MAC scheme is secure against existential forgery, R_N and L_N do not reveal information about the key.

6.6.2 Remote Pairing

In order for a remote adversary (i.e., an adversary that is not in physical contact with the same human as the device) to perform device pairing, the adversary has to execute the protocol with an honest device. Without loss of generality we assume that the adversary takes the role of Alice, i.e., executes the protocol with Bob.

The adversary must proceed according to the protocol otherwise Bob will abort. After the key agreement part of the protocol, the adversary does indeed share a key $K' = (g^b)^{a'}$ with Bob. However, in the key confirmation part, after receiving $R_N = \mathrm{MAC}_{K'}(A\|B\|N_A\|N_B)$ from the body channel, the adversary must send $L_N = \mathrm{MAC}_{K'}(R_N)$ back on the body channel. By the read-only property of the body channel this can only be done with negligible probability (as explained in Section 6.8), thus a remote attacker cannot successfully complete the protocol with Bob (or Alice).

6.6.3 Active Eavesdropping and MITM Attacks

To demonstrate that our device pairing protocol is secure against an active man-in-the-middle (MITM) attack, we observe the following. In order for the adversary to place himself in the middle between Alice and Bob, he must either run the protocol with each of them or interfere in an ongoing pairing session between Alice and Bob. Furthermore, the adversary must replace or modify at least one of the key agreement messages, as this would otherwise be passive eavesdropping.

As we showed above for the remote pairing attack, the adversary cannot successfully complete the protocol alone with either Alice or Bob. The protocol does not terminate in either case, since the body channel is read only for the adversary and thus the key confirmation fails.

Any modification of the public DH contributions g^a or g^b will, except with negligible probability, cause Alice and Bob to disagree on the key. For example, if the adversary replaces g^b with $g^{b'}$, we have

$$K_A = (g^{b'})^a \neq (g^a)^b = K_B,$$

which will result in the verification of R_N to fail in the key confirmation part. Interference with any of the other parameters sent in the protocol, A, B, N_A or N_B, will also cause the verification of R_N to fail, assuming the underlying MAC scheme is second pre-image resistant. By the read-only property of the body channel, the adversary cannot modify or replace R_N. Nor can he replace L_N after Alice has aborted the protocol, as a result Bob will also abort.

The only remaining option for the adversary is to initiate two sessions simultaneously with both Alice and Bob, and then rely on them to complete the key confirmation phase. For this to succeed the adversary must create two sessions where all the nonces, identities and public parameters are the same, since these are inputs to the MAC-function in the key confirmation part of the protocol. If all parameters are identical in the two sessions, and Alice and Bob are both being held by the same human, the protocol would succeed, but the adversary would just have done passive eavesdropping (and learned nothing as shown above).

Table 6.1: Parameters for intra-body communication using capacitive coupling.

Parameter	Value
Frequency bandwidth	0.5 MHz - 3.5 MHz
Transmission distance	Hand-to-hand (180 cm)
Signal electrode	4 cm by 4 cm aluminum plate
Ground electrode	7 cm by 7 cm aluminum plate
Data encoding	Manchester code
Modulation scheme	On-off keying
Sending power	5 mW
Sender voltage	3 Vpp
Current through body	$\sim 10\mu$A

6.6.4 The Human Body Channel

The security of the protocol relies on the assumption that the human body channel is read only for the adversary. This assumption models the fact that the receiving device can tell the difference between messages sent by an external transmitter and a device which is physically connected to the person performing the pairing. The receiving device can thus ignore any message that originates from an external source, which is equivalent to saying that the channel is read only.

In the following sections we will document experiments that verify this particular channel property and we state the assumptions that need to be made in order for the property to hold.

6.7 Implementation

Our design of the intra-body communication channel is inspired by [34]. The authors of [34] are among the first to report reliable intra-body transmission based on capacitive coupling. Their designed receiver front-end achieves a transmission distance that spans the entire body. Our goal is to establish hand-to-hand transmission which typically reaches around 180 cm for adults. We therefore adopted the impedance matching network proposed in [34] and followed the design choices found in Table 6.1.

6.7.1 Measurement Set-up

In order to simulate the pairing protocol between two devices, we designed a proof of concept for a body channel transmitter and receiver. For the purpose of our

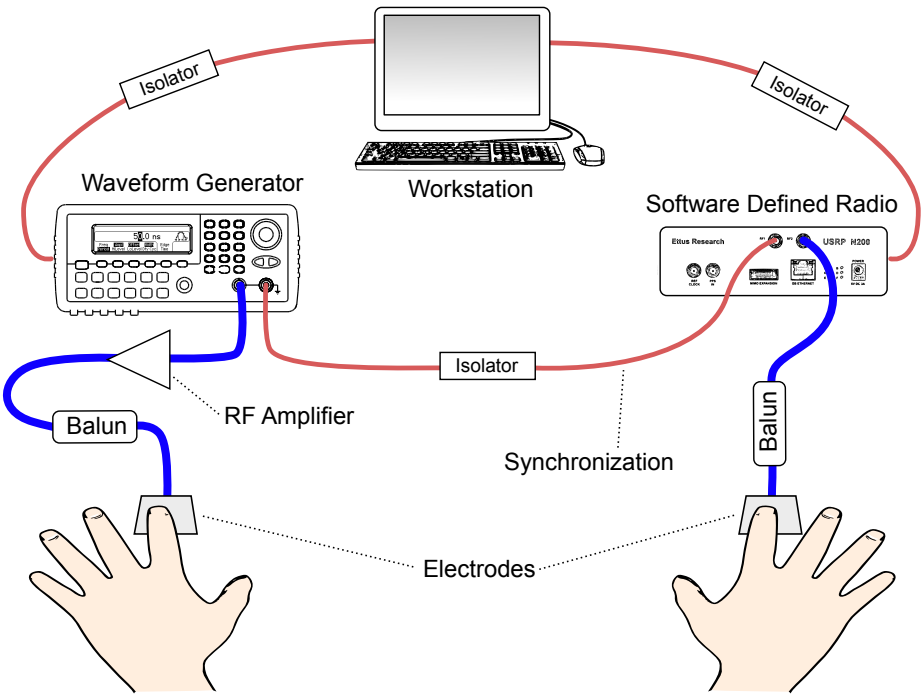

Figure 6.4: Measurement set-up. A waveform generator transforms the message into an electrical signal which is amplified and emitted through the touch-electrode of the transmitter. The touch-electrode of the receiver is connected to a software defined radio which captures the incoming signal.

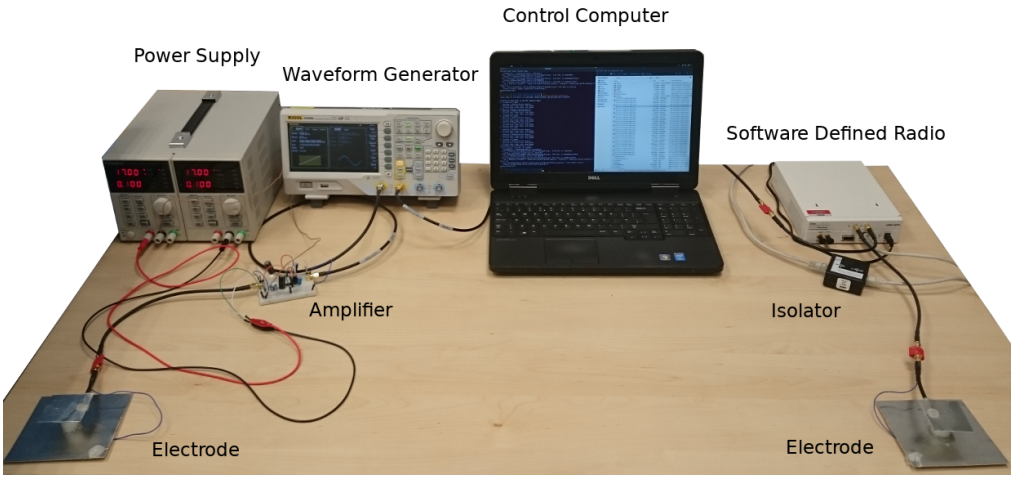

Figure 6.5: Photo of the experimental set-up used to run the pairing protocol.

prototype set-up, we did not implement two transceivers, but a separate transmitter and receiver. A more finished apparatus could combine the circuitry into two body channel transceivers that are capable of sending and receiving messages, i.e., bidirectional transmission.

The front-end of our receiver and transmitter implementation follow the exact same construction, which consists of two electrodes, the ground electrode and the touch-electrode. The person who pairs two devices only touches the touch-electrodes. The ground electrodes are floating. We describe the design of the electrodes in more detail in the following section.

We used lab measurement devices to implement the actual transmitter and receiver (see Figure 6.5). The simplified schematic of the experimental set-up can be seen in Figure 6.4. An arbitrary waveform generator acts as the transmitter and a software defined radio is the receiver. The waveform generator and the software defined radio are both controlled by a workstation computer that is used to specify the messages sent over the body channel and processes the signal received by the software defined radio. The receiver electrodes are directly connected to the software defined radio to record the incoming signal. The transmitter electrodes are connected to the waveform generator through an amplifier to boost the generated signal to the required 5 mW sending power.

For safety reasons and to minimize cross-talk, we made sure that the connections between the measurement devices are optically isolated. We also placed transmitter and receiver in such a way that they are separated by 120 cm and at least 200 cm away from any other electrical conductor. Transmitter and receiver electrodes are also decoupled from earth ground or any other shared potential through a pair of Balun transformers. A Balun transformer converts a single-ended signal (a signal referenced to a known potential) to a balanced signal and thereby eliminates the effect of the shared potential by the grounded measurement instruments. This is absolutely necessary and simulates a realistic scenario for body channel communication, as otherwise the shared ground potential will form a direct return path, yielding an unrealistically strong signal. In a real scenario the transmitter and receiver are not in direct contact with each other and do not have a shared electric potential, such as earth ground. This is especially true if transmitter and receiver are implemented as battery-powered devices (e.g., in mobile devices).

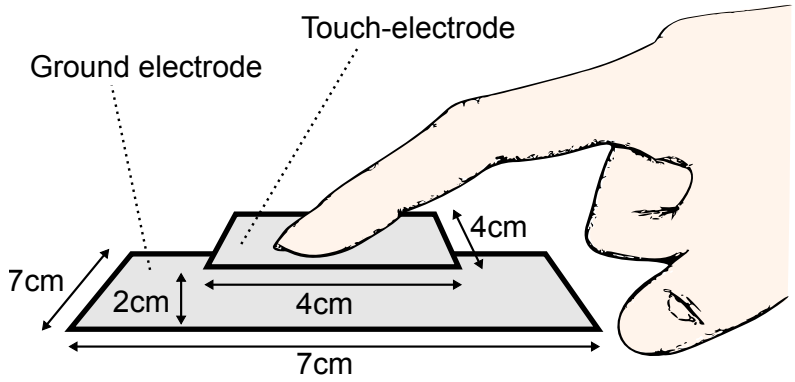

Figure 6.6: Signal and ground electrodes are 2 cm apart and manufactured from two aluminum plates.

6.7.2 Electrode Design

The touch-electrodes, i.e., the electrodes that interface the human body, are 4 cm by 4 cm sized aluminum plates with a thickness of 1 mm (see Figure 6.6). If the touch-electrodes are manufactured from a conductor, the effect of the electrode material on intra-body communication is marginal, see, e.g., [195]. In [33], aluminum and copper electrodes as well as pre-gelled electrodes, such as commercial AgCl electrodes used for electro-cardiogram (ECG) measurements have been tested. Pre-gelled electrodes can have better performance than copper or aluminum plates for capacitive coupling as a body communication method, since the gel enhances conductivity and adherence to the skin. However, analogous to body impedance recognition, gelled electrodes are not an option for our proposed device pairing mechanism for both hygienic and usability reasons. We opted for aluminum plates, as our touch-electrodes should be reusable and a permanent feature of the device.

The ground electrodes of the transmitter and receiver normally do not need to be implemented specifically. In an actual device they would correspond to the ground plane of the circuit board of the transmitter or receiver. For our experiments, we implemented the ground electrodes as square aluminum sheets similar to the touch-electrodes. They measure 7 cm by 7 cm and thus cover an area of 49 cm^2 each. The required surface area of the ground electrodes for reliable body channel communication has been estimated in [196]. The authors of said study developed a distributed RC model to simulate the characteristics of the human body channel when using capacitive coupling in the frequency range of 100 kHz to 150 MHz. According to the authors' empirical formula, 32 cm^2 is sufficient regardless of location of transmitter and receiver on the body if a bit error rate of 10^{-6} can be tolerated. Our ground electrodes cover 49 cm^2 and we achieve similar error rates (see Section 6.7.4).

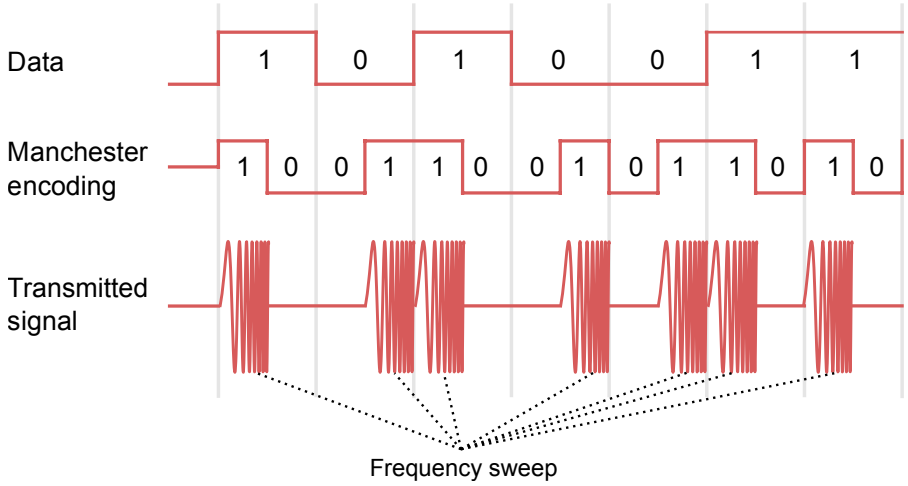

Figure 6.7: Data is Manchester encoded. The transmitted signal follows an on-off-keying modulation. During the "on"-periods a frequency sweep is performed.

6.7.3 Data Encoding and Modulation

We apply Manchester coding to the data before it is sent over the body channel. The Manchester code ensures that it is not possible to change the transmitted data without changing at least one encoded "0"-bit into a "1"-bit, similar to how integrity codes (I-codes) work [197].

The encoded messages are then transmitted using amplitude modulation in the form of on-off-keying. When the bit of the encoding is high ("0"-bit), the power on the channel is "on" and similarly, if the bit of the encoding is low the power is "off". Our scheme differs from a simple on-off-keying in the way that we do not use a single carrier or center frequency for the "on"-period. Instead of transmitting on a single frequency, the sender performs a sweep over a range of frequencies (see Figure 6.7). The frequency sweep is not dependent on the transmitted data. Whenever the power is on the transmitter outputs a signal at a frequency of 0.5 MHz and keeps increasing the instantaneous frequency until it reaches 3.5 MHz and the power is turned off. The purpose of the frequency sweep is to characterize the communication channel. If the sweep is present in the transmitted signal, the receiver can measure the frequency-dependent attenuation over a broad spectrum and verify that the measured characteristics correspond to a human body channel.

6.7.4 Throughput and Error Rate

With a duration of 1 millisecond per "on"-period, one data bit takes 2 milliseconds to transmit. Assuming that there are no bit flips, this results in a theoretical data rate of 500 bits per second. For example, if the message authentication codes R_N

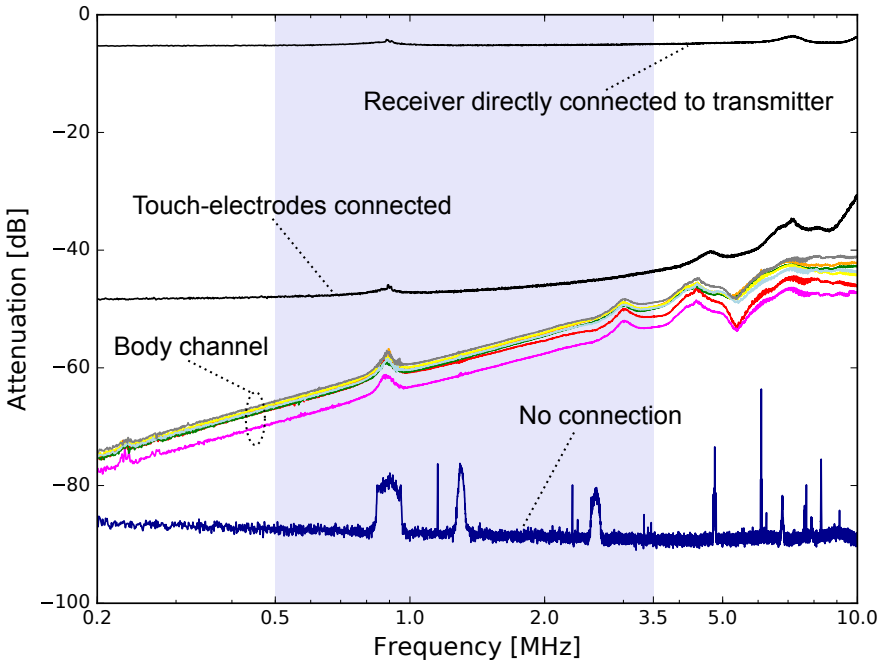

Figure 6.8: Measured attenuation (S_{21} parameters) of the body communication channel. From top to bottom: Both electrodes of transmitter and receiver are directly connected to each other with a wire (first black line), only the touch-electrodes of transmitter and receiver are connected with a wire (second black line), transmitter and receiver communicate through a human body (colored lines for 7 different people), receiver and transmitter are not connected at all (dark blue line at bottom). The shaded area depicts the frequency range we use to distinguish the body channel.

and L_N from the pairing protocol have 56 bit length, just 224 milliseconds are required to transmit both MACs over the body channel.

In all our experiments, the measured bit error rate of the body channel for hand-to-hand transmission was below 10^{-6}. This means that under normal operating conditions, i.e., when the human body is not subjected to external interference, the probability for a flipped bit is very low. The transmission of two 56 bit message authentication codes is errorless with a probability of more than $(1-10^{-6})^{2\cdot 56} = 99.98\%$ if the bit errors are equally likely to happen for every bit. Assuming the MACs have 56 bit length, it is therefore not necessary to compute error correcting codes and introduce redundancy into the messages that are sent over the body channel.

6.7.5 Body Channel Characteristics

Some of the energy transmitted on the body channel is lost due to the effect of capacitive coupling and due to the fact that the human body is not a perfect conductor. As a consequence, the frequency sweeps that are sent by the transmitter

are attenuated. In fact, the attenuation is frequency dependent, which means that not all parts of the frequency sweep are affected to the same extent. Provided the transmitter sends the sweeps at a fixed power level, the receiver can exploit this fact and measure the frequency dependent attenuation. Since there are no active elements in the body channel, the receiver essentially measures the S_{21} scattering parameter of the transmission line through the human body.

By extracting this information from the messages received through the touch-electrode, the receiver can characterize the communication channel. If the receiver knows the attenuation pattern that corresponds to a human body channel, it can verify if the received frequency sweeps have traveled through a human body by matching them with the known pattern.

In Figure 6.8, we show the channel characteristics for 7 different people when they are in physical contact with the touch-electrode of transmitter and receiver. We plot the attenuation over the frequency range from 0.2 MHz to 10 MHz and compare the body channel to the case where the touch-electrodes are either shorted-out (with a cable) or not connected at all. It is apparent that the human body channel exhibits characteristics different from other conductors, such as a cable, for instance. If the touch electrodes are connected with each other through a copper wire, the attenuation is low throughout the entire frequency spectrum. Contrary to that, if the touch-electrodes of the transmitter and receiver are not connected at all, i.e., they are floating, we see that all the frequencies are completely attenuated and are not picked up by the receiver (bottom line in Figure 6.8). In that case, the frequency spectrum highlights only noise and artifacts induced by the measurement set-up.

As explained earlier in Section 2.1.2, capacitive coupling works best in the frequency range of 1 MHz to 100 MHz. However, frequencies higher than 10 MHz are mostly surface waves [40]. We focus on the frequencies between 0.5 MHz and 3.5 MHz to extract the body channel characteristics. Figure 6.8 shows that the higher the frequency, the lower the attenuation, because more power is transmitted through the air. The signal does not travel through or along the human body any more and the channel characteristics become less unique to the human body (i.e., the human body acts as a high-pass filter) and more erratic. We capture the properties of the human body channel where they are most specific and can facilitate the distinction whether the characteristics belong to a human body channel or not.

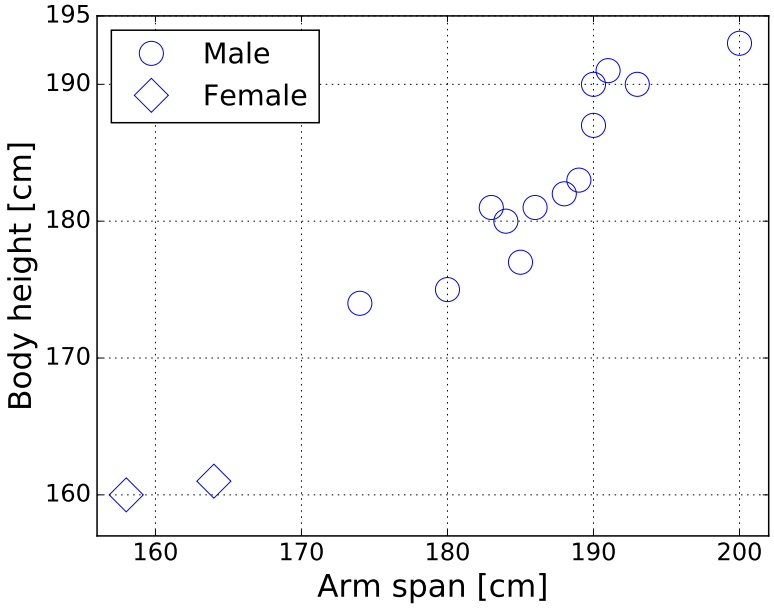

Figure 6.9: Body dimensions of the study participants. Arm span is measured in a T-pose (fingertip to fingertip) and approximately represents the length of the body channel.

6.7.6 Experiment Data set and User Safety

For the experimental analysis of our proposed pairing scheme, we collected data from a total of 15 study participants. The study was approved by the ethics board of the University of Oxford under the reference number R53956/001. The participant group of the study consisted of two women and 13 men who were between 22 to 45 years old. Figure 6.9 shows the body dimensions of the study participants. We collected more than 50 data transmissions per participant and conducted additional experiments to prove that our protocol is secure.

Our implementation of body communication is safe to use and does not pose a risk to human health. The return path for capacitive coupling goes through the air, which results in very high resistance and little current flow [198]. In fact, the current through the body never exceeded 12 micro-ampere (see Table 6.1). This is even weaker than the currents we used for body impedance measurements in Chapter 3 and is nowhere near what commercially available body composition measurement devices emit. Body fat monitors, for instance, pass a current of up to 500 micro-amperes through a person [199].

In addition to the risk of current flow, we have to ensure that the exposure to the electromagnetic field created by the capacitive coupling does not jeopardize human health. We consulted the "Guidelines for limiting exposure to time-varying electric, magnetic, and electromagnetic fields" issued by ICNIRP (see [200, 201])

and concluded that the electric field strength generated inside the human body stays well within the suggested limit of 1.35V/m per 1000 Hz.

Moreover, we verified that the power of our body channel transmitter does not violate FCC regulations [202]. We measured the strength of the radiated electromagnetic field with a rod antenna at a distance of 4 meters for a subset of our participant group. The electromagnetic waves radiated into the air did not exceed the limit of 30 μV/m for the entire frequency range we experimented in, i.e., from 0.2 MHz to 10.0 MHz.

Lastly, we made sure that our lab instruments are isolated from the touch-electrodes, such that even in the very unlikely event of a hardware failure the participants are not exposed to line voltage.

6.8 Experimental Results

In this section, we present experiments that document the properties of the body channel and validate the assumption that the body channel is read-only for an attacker that is not touching the body.

The read only property can be stated in two different ways and we validate both experimentally.

- We verify that a body channel enabled device can detect if a received message has been sent by another device that is physically connected to the same person or an outside signal source. The receiver should be able to classify messages according to their origin; if the message comes from a legitimate body channel or an external transmitter.

- We examine if it is possible to "inject" a message into the body channel in such a way that the physical properties of the message appear at the receiver as if the message was sent on the body channel.

We break the experiments down into these two statements and report the results in the following.

6.8.1 Classification of Body Channel Messages

Our proposed pairing protocol relies on the ability of the body channel receiver to distinguish messages based on their physical properties. This is important, not only from a security standpoint, but also with respect to feasibility. The device pairing protocol does not work if the receiver cannot detect the body channel. To show that a

body channel receiver can identify messages sent on body channel, we performed data transmission through the body channel of 15 test subjects under various conditions.

In order to capture data reflecting the intended use of the device pairing protocol, we asked the participants to touch receiver and transmitter electrodes as if they were paring two devices. With the collected data we establish a baseline of the attenuation pattern of the human body. We then analyze how the channel characteristics change when the body channel is modified, or if there is no human body present. We build and train a classifier that can exploit these differences and decide whether a previously unseen message has been sent on the body channel.

If the classifier is universal enough to distinguish between messages independently of the actual person involved in the pairing, it can be readily deployed in any body channel enabled device. Such a device would not require any user-specific input or enrollment to classify messages and therefore could be taken into operation without in-field adjustments or calibration.

Feature extraction The receiver captures the messages that are transmitted on the body channel in the form of the time varying voltage level at the touch-electrode. The measured electrical signal is transformed to the frequency domain where the channel characteristics become apparent. Much like the feature extraction process for body impedance recognition (see Section 3.4.8), we use the Fast Fourier Transform (FFT) to compute the frequency bins that correspond to the spectrum from 0.5 MHz to 3.5 MHz. The magnitudes of each bin constitute the feature values that are passed to the classifier for training. As an additional step, before we train the classifier, we normalize the extracted feature values to eliminate the effect the power of the transmitter has on classification.

Classifier We use support vector machines (SVMs) to classify the channel characteristics and we treat the classification problem as supervised and binary. The classifier has to decide between two classes; the class of features that belong to the body channel and the class of all unwanted interactions with the body channel receiver, i.e., unwanted interactions are combined to one class for training.

Evaluation We evaluate the classifier on samples that we gathered in multiple scenarios that each fall into one of those two classes. For the intended use of the body channel, we tested two different settings. The participant is either standing or in a seated position when touching the electrodes.

The samples that represent unwanted interactions cover the following scenarios:

1. No connection between the transmitter and receiver electrodes. All electrodes are floating.

2. Transmitter and receiver electrodes are connected to each other through a wire.

3. Transmitter and receiver electrodes are directly facing each other at various distances (5 cm, 10 cm, 30 cm and 50 cm).

4. Transmitter is connected to either a rod antenna of 1 m length or a 25 cm by 80 cm aluminum sheet (a large surface area improves capacitive coupling) directly pointing at the receiver. This scenario represents an external transmitter communicating with the body channel receiver.

5. One of the participant's hand touches the electrode of the receiver, but his other hand is not in physical contact with the electrode of the transmitter. It hovers over the transmitter electrode at various distances (5 cm, 10 cm and 30 cm).

6. The participant only touches the receiver electrode. The transmitter is connected to a rod antenna or an aluminum sheet which is placed at a distance of 30 cm and 60 cm from the participant. This scenario represents an external transmitter communicating with receiver while a person (accidentally) touches the receiver electrode.

The different scenarios listed above are repeated at different transmit power levels. We set the output voltage of the transmitter to 1, 2, 5 and 10 Volts. The data sent in these experiments consists of a random bit-string of 56 bit length.

All experiments are performed twice, once with a frequency sweep containing a sine wave and a second time with a square wave, to determine if the shape of the waveform plays a role in how the channel characteristics are elicited. Sine waves are a straightforward way to measure channel properties, but the study in [203] successfully applied short square pulses to intra-body communication using capacitive coupling and, albeit at different frequencies, we successfully used a short square pulse when designing pulse-response recognition. Since a frequency sweep with a square wave corresponds to a series of pulses of different duration, we also include square waves in our evaluation.

Table 6.2: Classification results for body channel characteristics.

Wave Type	Accuracy	F1 score	AUC (ROC)
Sine wave	0.993 (± 0.018)	0.987 (± 0.033)	0.985 (± 0.037)
Square wave	0.943 (± 0.180)	0.918 (± 0.246)	0.988 (± 0.040)

Shown are the mean values of the metric and the 95% confidence interval.

Results We analyze a total of 1020 instances of the scenarios described above. They encompass data transmissions for every study participant in each of the outlined cases. The balance of the two classes, i.e., the ratio between the number of samples that represent the body channel and those that represent unwanted interactions is 1:1.

Table 6.2 shows the classification performance in terms of three metrics: accuracy, F1-score and the area under the ROC (receiver operating characteristic) curve. The results are obtained by running stratified 10-fold cross-validation. We observe that the SVM based classifier can detect the characteristics of the body channel with high accuracy. If a sine wave is used for the frequency sweep, the probability for a misclassification is less than 2%. All three different metrics are consistently high which suggests that the human body channel is very distinctive even when compared to the various other ways of interacting with the receiver. The results also show that the extracted characteristics are consistent across different people, regardless whether the study participants are sitting or standing. The body pose does not have a significant effect on the body channel. Figure 6.10 shows the receiver operating characteristic curve, representing body channel transmissions as positive samples and unwanted interactions as negative samples. Both curves are very close to each other, with "sitting only" slightly outperforming the other. The classifier can be tuned by setting the discrimination threshold to any point on the curve. Overall, Figure 6.10 shows that the classifier is conservative in assigning a new sample to the class of body channel characteristics and is more likely to reject it as an unwanted interaction.

If a square wave is used for the frequency sweep, the classifier does not perform as well as for a sine wave. An explanation for this behavior is the fact that a square wave generates more spectral components in the high frequency range. These frequencies are mostly transmitted through the air (via surface waves) and therefore do not capture any of the distinctiveness of the human body channel.

In order to understand what scenarios exhibit channel characteristics that come closest to the actual body channel, we list the scenarios according to their likelihood

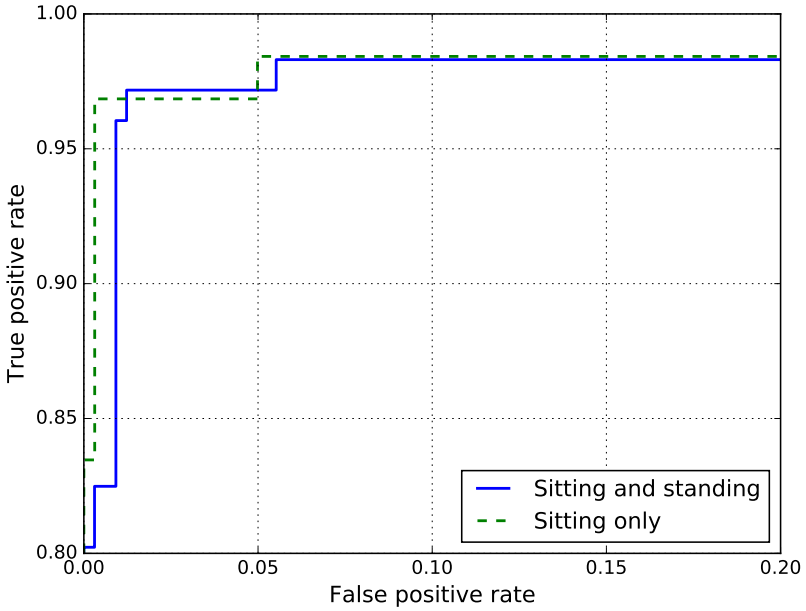

Figure 6.10: Receiver operating characteristic (ROC) for the body channel classifier, zoomed into the upper left area. We compare the effect of two body poses: participants are sitting or standing when touching the electrodes (solid line) or sitting only (dashed line).

Table 6.3: Unwanted interactions with the body channel receiver according to their likelihood for misclassification in [%]

Antenna Type	Participant touches receiver electrode	Receiver electrode is floating
Rod antenna		
at 30 cm	0.40	0.09
at 60 cm	0.01	0.00
Aluminum sheet		
at 30 cm	0.91	0.30
at 60 cm	0.09	0.00
Participant's hand hovers over transmitter electrode		
at 5 cm	70.1	N/A
at 10 cm	55.0	N/A
at 30 cm	1.20	N/A

Shown are mean values obtained by running 10-fold cross-validation.

for misclassification in Table 6.3. We see that, if the participant is in contact with the touch-electrode of the receiver, but only hovers over the transmitter electrode, the channel characteristics are similar to the actual body channel. This is result is not surprising, because the additional distance between the body and the transmitter electrode will increase the capacitance of the channel, but not significantly change other physical properties of the channel. Capacitive coupling still works even if the body is not in direct contact with the emitter of the signal. If an external transmitter is used however, the channel characteristics only match an actual body channel to an extent. Table 6.3 shows that the rod antenna and the aluminum sheet are more successful in establishing a body channel if they are closer to the receiver or the person. At a distance of more than 60 cm the chance of matching the body channel characteristics becomes negligible, assuming the transmitted signal corresponds to what the receiver expects, i.e., a frequency sweep from 0.5 MHz to 3.5 MHz. We investigate the case of an attacker changing the waveform for signal injection in the following section.

6.8.2 External Signal Injection

We have shown that the human body channel can be characterized on the basis of its frequency dependent attenuation pattern. We now approach the question if the body channel is read only from the perspective of the second statement: Can attacker transmit from an external source and by manipulating the signal make it appear as if it was sent on the body channel?

To answer this question we make a number of observations. The first observation is that an attacker has two options, inject his own message on the body channel or modify another message. If he injects an entire message, he has to make sure that all frequency sweeps included in the message match the body channel characteristics. If the attacker's goal is to modify another message, he has to inject at least a part of a message. The messages on the body channel transmitter are Manchester coded and every bit of transmitted data consists of a period where power is on and off. Therefore, even to change a single bit, the attacker has to inject a signal that matches a frequency sweep emitted by a body channel transmitter. Regardless whether the attacker injects an entire message or modifies another message, if the injection of a single sweep fails, then the message is automatically rejected by the classifier, because at least part of the signal has a different signature. We therefore focus on the injection of a single sweep signal in the following.

We also note that changing the overall transmit power does not help an attacker since a constant shift in the attenuation pattern (e.g., achieved by increasing

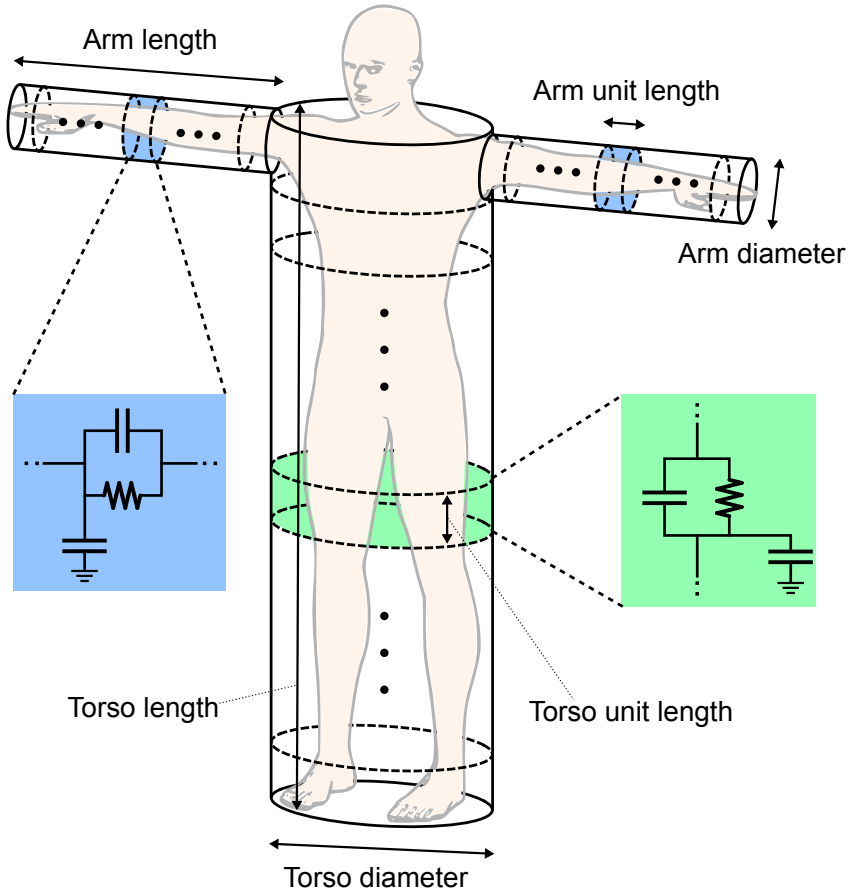

Figure 6.11: Lumped network model for body channel. The human body is simplified to three connected cylinders. The cylinders are further divided into smaller units. Each unit can be modeled with a capacitor and a resistor in parallel, plus coupling capacitance to ground.

the power of the transmitter) is removed during the normalization of the extracted features.

For further analysis, we divide signal injection attempts into the near and far field based on the attacker's distance to the body channel receiver. Near and far field define the behavior of the electromagnetic field around a receiving or transmitting antenna. In the far field "normal" electromagnetic radiation is dominant, whereas in the near field the electromagnetic field is mostly determined by non-radiative and quasi-static effects, such as capacitive coupling. For the purpose of our analysis, we define the boundary between near and far field to be where capacitive coupling becomes ineffective.

Far field An attacker in the far field has to send a signal that matches the body channel signature like an attacker from the near field. However, an attacker in the far field cannot rely on capacitive coupling because the electric field generated by

electrostatic effects falls off with distance cubed [30]. The attacker has to resort to radio frequency transmission, but transmitting on the frequency band of 0.5 to 3.5 MHz at a power level such that the signal is picked up by the receiver electrode (or the human body, provided a person is touching the electrode) is not feasible. The electrodes as well as the person are by far from an optimal antenna for such low frequencies.

The human body does have an antenna effect because of its size [196], but at wavelengths on the order of 100 m it is not viable to induce a field strength at the receiver electrode that would result in a signal greater than thermal noise. Unless the transmitter is highly directional and has an output power in excess of 100 Watts, an attacker cannot inject a meaningful signal, let alone a signal with a signature resembling the characteristics of the body channel. Aiming the antenna at the receiver further increases the complexity of an attack from the far field as well as signal propagation phenomenons such as multi-pathing that cause interference and fading.

Near field Electrostatic coupling, such as capacitive coupling, has the highest chance of success for signal injection. Electrostatic effects diminish with the cube of the distance, but if an attacker is close enough to the receiver (or the person touching the receiver electrode), he can mitigate the attenuation by increasing the output power of his transmitter.

Capacitive coupling works by electrostatically coupling a current into the human body. The air gap between body and transmitter acts as a capacitor and the larger the gap, the lower its capacitance. A low capacitance results in a high-pass filter with a higher cut-off frequency and the lower frequencies are attenuated significantly.

The attacker can overcome this attenuation in two different ways: increase the output power at the transmitter and increase the surface area of the transmitter. This is congruent with the finding in Section 6.8.1, where we show that the channel characteristics are more similar to the body channel when an aluminum sheet with a large surface area is connected to the transmitter instead of a rod antenna. Following this reasoning, an external transmitter needs to have high power, a large surface area and be placed close to the receiver.

Network model for body channel In order to understand if signal injections from the near field are feasible, we build a lumped network model for the body channel which is inspired by [196]. The model approximates the human body as three cylinders, one for the torso and two for each arm (see Figure 6.11). The cylinders are subdivided into units for which an approximation of the electrical circuit can be given. Each unit can be modeled with a capacitor and a resistor in parallel, plus coupling capacitance to ground. The units for torso and arms have the same electrical circuit, but different parameters. The length of a unit is 10 cm for both, arm and torso. The diameter of an arm is 10 cm and the diameter of the torso is 30 cm, respectively. Based on these dimensions, the values for capacitance and resistance per unit can be calculated from the dielectric properties of biological tissues [204].

Using this model the body channel transmitter and receiver can be attached anywhere on the human body, i.e., to any unit block of the model, and the resulting transmission characteristics can be readily computed. If we attach the transmitter to one hand and the receiver to the other hand, we obtain an accurate approximation of the body channel characteristics. Figure 6.12 demonstrates that the computed characteristics (striped area) correspond well with actual body channel measurements.

To simulate an external transmitter that does not directly touch the body, we can attach the transmitter at multiple coherent blocks of the network model to take into account the distance between body and transmitter. The further away the transmitter is the larger the area that is affected by the capacitive coupling. In addition to that, the air coupling capacitance between body and transmitter decreases, as the distance between transmitter and human body increases. Figure 6.12 shows the computed channel characteristics (solid shaded area) when simulating a large aluminum sheet (25 cm by 80 cm) aimed at the person from behind at a distance of 30 cm. Comparing the results with actual measurements, we find that the model approximates the channel characteristics very well.

Both, the computed approximations as well as the actual measurements, make it evident that the channel characteristics for an external source, such as an aluminum sheet, are significantly different from the body channel. For an attacker to successfully inject a signal, he has to change the output power of his transmitter based on the currently transmitted (i.e., instantaneous) frequency. In order to make the injected signal match the body channel signature, the attacker has to constantly vary the power of his transmitter. Taking the example in Figure 6.12, the attacker has to transmit at a low power output at 0.5 MHz and gradually

increase the power until reaching 0.8 MHz. Then he has to back off sharply, only to gradually increase the power again for higher frequencies. We claim that this is not feasible due to two reasons.

- The attacker does not know the exact channel characteristics his transmitter creates and he cannot measure them as this would require physical access.
- The attacker can try to precompute the channel characteristics, but this is likely to be inaccurate, since the attenuation pattern is very volatile.

In order to compute the channel properties, the adversary has to estimate the distance to the body as well as the location and size of the area on the body where capacitive coupling induces a current. Our experiments and the data simulated by the model demonstrate that the channel characteristics are very volatile and become increasingly difficult to approximate the further away the transmitter is placed. At around 30 cm distance, capacitive coupling becomes very weak and unpredictable. We give an example in Figure 6.12 that shows if the aluminum sheet is shifted by only 5 cm, the pattern looks significantly different. The bottom red line depicts an injection attempt where the sheet is placed at a distance of 35 cm instead of 30 cm from the person and the body channel receiver.

Together with the results from Section 6.8.1, these insights let us conclude that our stated read-only assumption for the body channel holds as long as there is a minimum distance of 50 cm between adversary's transmitter and the human body.

6.9 Discussion

Body position Body position and body geometry can have an effect on the measurements as shown in Section 6.8.1. We designed and conducted experiments for two different body positions (seated and standing) to get an estimate of how much the attenuation of the received signal varies. The study in [205] found that for different test subjects the two positions, i.e., seated or standing, exhibited an attenuation of the same magnitude, which is in line with our results. The authors tested several other body poses and even body movement and reported that the attenuation changes by around 5 dB for a transmission distance of 120 cm. Hence, we conclude that body position has an insignificant impact on the use of capacitive coupling for our device pairing protocol.

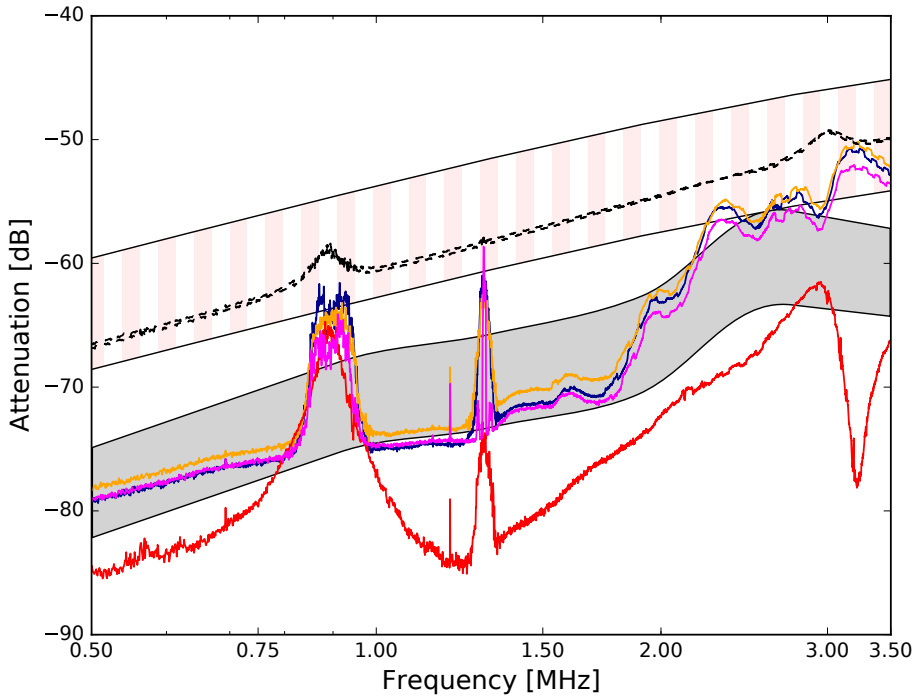

Figure 6.12: Dotted lines represent attenuation patterns of the body channel obtained from two different people. Solid lines depict signal injections with an aluminum sheet. Bottom red line represents an attempt where the sheet is 5 cm further away from the body. Shaded areas show approximations using the lumped network model for a human with a body height between 155 and 195 cm (striped/top area depicts body channel and solid/bottom area represents external transmitter).

Pre-computation of channel characteristics There are environments where pre-computation of the channel characteristics could be feasible. A possible example is an ATM that uses body communication to pair customers' smartphones to an internal module in a similar fashion to the PIN-entry application for user authentication described in Section 3.2. An adversary could install large antennas right next to the ATM which are hidden from the bank's customers. Being so close to the human body and considering that most people assume a similar pose while operating an ATM, an adversary might overcome the capacitive coupling and successfully pre-compute the characteristics of the body channel.

Different pairing protocols The presented body channel communication technique does not have a very high bandwidth and therefore the amount of data that can be transmitted through the body in a reasonable time frame is limited. This is due to the frequency sweeps that take 1 millisecond to complete every time a bit is sent on the medium. However, the measured bit error rate of the body channel is remarkably low (see Section 6.7.4), allowing the use of pairing protocols

other than the suggested one in this work which consists of a (standard) Diffie-Hellman key agreement followed by key confirmation. In principle any protocol that relies on a low-bandwidth, read-only (i.e., out-of-band) channel to bootstrap secure communication is applicable. A survey featuring a comparative evaluation of authentication protocols that are based on an authentic low-bandwidth channel can be found in [206]. Depending on the length of the transmitted messages, the strength of the MAC scheme or hash function can be varied to achieve lower computational cost, for instance. Finally, even group (pairing) protocols [206, 207] that go beyond pairwise applications are feasible, if we consider the scenario of touching more than two devices during pairing.

6.10 Summary

We propose a novel approach to device pairing which builds upon the core idea of using intra-body communication. We present a protocol that allows two devices to securely agree on a mutual secret by sending messages through the body of a person who is in physical contact with both devices. Incorporating the human body as a transmission medium entails an out-of-band channel the devices can utilize to quickly and securely perform key confirmation, without the need for certificates or shared knowledge. Moreover, the human body channel provides the ability for the devices to extract physical properties that are very distinctive of this communication channel. We show that these characteristics are sufficient to determine, with high probability, if a message has traveled from one device to the other via the body channel. Most importantly however, our experiments document that the human body channel cannot be interfered with from the outside as long as there is a distance of at least 50 cm between the external signal source and the person who is pairing the devices.

We believe that the presented device pairing protocol is an attractive solution to this problem and enables even novice users to pair devices with a task that requires very little involvement other than the touch of two electrodes.

Finally, this chapter leaves the interesting question for future work if intra-body communication could be used to enhance security in other protocols as well, such as in user authentication methods. For example, is it possible to combine body impedance recognition and intra-body communication?

7 Conclusion

We conclude this thesis by summarizing the key findings, discussing potential limitations and giving directions for future work.

Contents

Considering the soaring number of electronic devices we use every day, the problems of authenticating users to devices, as well as establishing secure communication between devices, will inevitably arise with increased frequency.

While solutions to these problems have been researched extensively, their implementation often proves difficult in practice or comes with disadvantages. Most often, the drawbacks manifest in the form of increased burden on part of the user who has to make security-related decisions so that the authentication or pairing mechanism works properly.

For example, password-based user authentication can be made arbitrarily complex, assuming the user is willing to memorize a very long secret. The same is true for device pairing, if the user is asked to manually copy a complex character string from one device to the other. Indeed, the increased security comes with more user burden, which is problematic in two ways: first, users who are not security-conscious adopt solutions that require the least amount of cooperation and second, whenever a human factor is involved, mistakes naturally occur. If an

authentication or pairing protocol hinges on the human factor, it cannot provide any security guarantees any more once the user makes a mistake or does not cooperate. Inevitably, user authentication and device pairing (without a trusted third party) will always require some form of user involvement and it is therefore of great importance to develop mechanisms that feel natural and are understandable even to the unskilled user without technical expertise.

This thesis proposed solutions to those challenging problems by presenting novel techniques that only require a very simple interface to work. By providing user authentication and device pairing at the touch of two electrodes, we believe that our proposals are a substantial step in the right direction and can make security easy and accessible to everyone.

7.1 Key Findings

This thesis is one of the first to propose harnessing the electrical properties for two important problems in Computer Security: user authentication and device pairing. We summarize the key results below.

We first established body impedance as a physiological characteristic for user recognition. By exploring different ways of measuring body impedance and evaluating recognition performance based on a user study, we were able to conclude that horizontal body impedance—which is measured from one hand to the other—constitutes a viable method for biometric verification and identification. We showed that horizontal body impedance possesses distinguishing abilities, comparable to methods relying on behavioral traits. Impedance measurements can be acquired almost completely unobtrusively by simply having a person touch two capacitive electrodes or conductive pads.

Leveraging this novel biometric modality for user authentication, we devised and analyzed two hypothetical systems for which body impedance is particularly suited. Besides analyzing the security of the envisioned biometric systems, we used requirements and goals found in literature to assess our proposed systems. We then compared our novel techniques to established and well-researched biometric methods. We identified continuous user authentication as an application where (horizontal) body impedance excels. Taking a reference measurement at log-in time can effectively guarantee that the user currently operating the system is still the same person who was present during log-in phase. All that is needed to implement such a system is a conductive object or item the user repeatedly

comes in touch with, e.g., a conductive keyboard, that allows the acquisition of continuous impedance measurements.

Having studied body impedance from a biometric perspective, we continued to approach the idea of using body impedance as a source to generate user-specific cryptographic keys. Using biometric keys and templates for user authentication can successfully minimize the leakage of sensitive biometric data. However, biometric key derivation is a much more complex problem than extracting features for the purpose of differentiating between individuals. The design and selection of features that meet the requirements for key generation are intricate: the features need to be maximally similar for the same individual and as distinct as possible for different users. We pursued a novel approach based on semi-automated feature learning, resulting in impedance-based secrets that are both unique to an individual and sound from a cryptographic standpoint. For most of the subjects in our test population, the key derivation scheme was able to output keys with 30 to 60 bits of entropy. The generated keys could either serve as a replacement for a password- or token-based component in authentication or add entropy to an existing cryptographic scheme.

After presenting two use cases that are based on the distinguishing abilities of body impedance recognition, we conceptualized and tested an application that builds on biometric key generation from body impedance. We devised a scheme for driver and message authentication in Vehicular Ad-hoc Networks (VANETs). Until now, proposals for VANETs suggested to authenticate messages using signatures that are bound to the owner of a vehicle. However, this is problematic since in the future, passenger cars are likely to be shared among groups of individuals and attributing messages to the actual driver will become increasingly difficult. As a solution, we proposed a biometric signature scheme that supports non-repudiation while at the same time preserving identity and location privacy. We developed a message authentication scheme that relies on biometric measurements captured from the driver every time a message is sent on the VANET. To test feasibility, we simulated driving conditions by having study participants hold a steering wheel to control a vehicle in a computer game. Conductive pads attached to the wheel measured the driver's impedance. The results we obtained showed that impedance-based recognition is a viable candidate for this type of application scenario.

Having successfully extracted static features from the body, we turned our focus to characterizing the human body as a transmission medium. We developed a device pairing protocol that uses intra-body communication as an out-of-band channel. The protocol allows a person to pair two electronic devices by simultaneously touching

an electrode on both devices, one with each hand. The person thereby provides a communication link to the devices that they can use to exchange data as long as they are in physical connection via the human body. We identified characteristics that accurately describe this communication channel and make the transmission medium identifiable to the devices involved in the pairing process. The devices can verify if a message is sent through the body or stems from an outside source, which effectively transforms the human body into an authenticated channel. To verify our findings, we conducted a user study where we performed data transmission through the body channel of the test subjects and gathered transmission samples from external sources. Our results show that the specific characteristics of the body channel enable reliable classification of messages into two groups according to their origin: body channel and outside sources. We were able to validate these results by comparing them to a lumped network model of a human body and simulate the body channel.

7.2 Future Work

We highlight the three most significant ways the work presented in this thesis could be extended and list them according to their relevance.

Build a compact prototype Commercial success of our proposed methods hinges on the design of a compact prototype and on the size to which the prototype can be reduced. The electrodes that interface the human are small enough to be accommodated on a smartphone, but the measurement set-up used in this thesis consists of lab-grade devices and is too expensive to be incorporated in a consumer product. While it is in theory possible to shrink the technology, the smaller the final footprint, the more challenging the engineering tasks will be, such as the design of the electronic circuit. Ideally, a functional prototype that it fits into existing end-user products can be developed.

Once the miniaturized prototype has reached a certain level of maturity, a comprehensive user study involving a large test population needs to be conducted in order to obtain a realistic assessment of usability and acceptance which is currently lacking. Both, user authentication and device pairing could benefit from such a comprehensive assessment. A prototype in a small form-factor would allow us to run user studies closer to real-world scenarios and could prove or disprove feasibility of deployment.

More concretely, a compact set-up for impedance measurements could be accommodated in a keyboard to realize the continuous authentication scheme

proposed in Section 3.3. Furthermore, such a set-up could be installed in actual vehicles in order to test the susceptibility of body impedance to hand movements and steering motion (see Section 5.6) in the real world, without relying on a simulation.

A miniaturized body channel transceiver, on the other hand, could be used to equip small and large devices with the ability to perform device paring over the body channel. We envision our technology to be suitable for small devices or appliances that can be brought close to each other, such that a single person can touch them at the same time. It needs to be investigated what transceiver size and format are ideal in these different scenarios.

Forging of impedance measurements and countermeasures In Chapter 3, we argued that it is difficult but feasible to separate body impedance from the individual to whom it belongs for the purpose of presenting it to a sensor. Since body impedance is a static characteristic, an adversary could launch a targeted attack by capturing a victim's body impedance and re-creating it in hardware, e.g., in the form of an electrical circuit. In a broader, less targeted type of attack, an adversary could collect population statistics on body impedance in order to fabricate a "dummy" body that is likely to impersonate certain individuals in a population. It is important to understand attacks like these in detail and to obtain estimates on how difficult their execution is. To give an example of how significant such attacks can be, we point out the research collaboration on Electrocardiography (ECG) biometrics [12] where we show that with artificially generated ECG signals it is possible to circumvent a commercially available security token based on ECG recognition. We managed to devise a method that can recreate biometric input with the help of a single ECG measurement that has been stealthily acquired from a victim, or with the help of statistics that can be computed across a population[1]. Both attacks can circumvent biometric systems based on ECG with significant probability.

Along with potential attacks on body impedance-based recognition, such as the ones investigated for ECG biometrics, further work should also explore countermeasures. There are two potential remedies for targeted attacks and population attacks. The first approach is to make the impedance measurements sensor- or platform-dependent as suggested in Section 3.6.1. This would guarantee that measurements do not directly translate between different measurement devices.

Another countermeasure is to add randomness to the probe signal. Instead of sending the same elicitation signal, i.e., a short-pulse in case of pulse-response recognition, at every authentication attempt, it could be advantageous to add a

[1]The population on which the statistic is based does not necessarily have to include the victim.

random perturbation to the signal. The change to the signal needs to be distinct enough, such that it can be picked up even after the signal is attenuated by the body impedance, but at the same time, the random contribution to the probe signal should not impact recognition performance. Such a mechanism will guarantee that the adversary can not simply inject a static signal by means of a function generator, for instance. The adversary would have to fabricate a potentially much more expensive "dummy" body in hardware that can react to the random probe signal.

Body channel for wearables In light of the proliferation of wearable technologies and on-body health sensors, it could prove beneficial to extend the body channel to body-worn devices. This would open many new possibilities for the use of the body channel. Electronics that are in constant contact with the body, such as smartwatches, could be paired with larger (stationary) devices by simply touching an electrode embedded in the enclosure. This is even simpler than the pairing process as we presented it in Chapter 6, as only one hand is used. However, having body-worn devices communicate with each other can give rise to new types of attacks. We showed that the body channel is secure if an attacker is at least 50 cm from the body. If an adversary manages to compromise a body-worn device, this statement might not hold anymore and our adversarial assumptions need to be reconsidered. An analysis of how the body channel behaves under such attacks is essential. One needs to determine if it is possible to distinguish between messages sent by devices connected to different parts of the body. The presented network model for the body channel could facilitate such an analysis and give insight into how the channel characteristics change, when transmitter and receiver are attached at different locations on the body.

7.3 Possible Limitations

When proposing novel techniques, it is important to identify and discuss potential obstacles that could hinder further development.

Time-stability of biometric features The user-specific features extracted from body impedance measurements might not be unique or stable enough, resulting in unsatisfactory user recognition rates. As discussed in Section 3.5.4, our results show that equal error rate (EER) for biometric authentication increases from 2% to 9%, when comparing measurements that are acquired in the same seating as opposed to measurements that are weeks apart. It is not clear if recognition errors

increase further when measurements are months or years apart. If the extracted features from body impedance measurements predominantly rely on traits that can be altered substantially, e.g., body weight, impedance-based recognition is not feasible over a longer duration without intermittent re-enrollment. Moreover, effects of external factors, such as sweaty hands, excessive water intake or change in body temperature have not been experimentally studied yet.

Engineering challenges We point out in future work (see Section 7.2) that a compact prototype is crucial for deployment of the technology. This applies to both the measurement device needed for impedance-based recognition and the body channel transceiver. If the technology cannot be shrunk to the needed size, feasibility of deployment cannot be confirmed. The development of a small device could be very involved, since power consumption, radio-frequency design and shielding are just a few factors that become key challenges.

Other wireless communication standards There exist established communication protocols for small devices, such as Near-field Communication (NFC) and Bluetooth. These communication standards are available as highly optimized and dedicated hardware implementations that have been integrated in many commercial products. In their standard version, these technologies do not feature protection against remote pairing attacks or man-in-the-middle attacks [208], unlike the body channel. However, NFC and Bluetooth are so advanced in terms of power-consumption and data throughput that it will be difficult for body channel communication to become prevalent, despite its advantages in security relevant applications.

User acceptability Electricity and radio-frequency transmissions are often perceived as harmful to human health. In many cases this is done rightly so, especially when dealing with high voltages or high-power electromagnetic radiation—in other cases, it is still unclear what effects electricity and radio-frequency waves have, especially over time. There are conflicting opinions on these long-term effects. For instance, the possibility that living next to a mobile network antenna is detrimental to the human body, or whether carrying a mobile phone close to one's body is dangerous. Body impedance measurements and body channel communication could receive an equally negative connotation since electrical signals are sent *through* the body. As a result, this could hurt user acceptance of any product equipped with this technology.

7.4 Conjectures and Closing Remarks

Body impedance recognition and body channel communication are not by any means limited to the scenarios presented in this dissertation. By providing such a simple and natural interface, which barely requires any user interaction, the technologies proposed in this work could serve as a building block in many other systems and security protocols.

Considering the proliferation of wearable devices and the increasing number of electronics we use overall, it is conceivable that the research presented in this dissertation can help enable security across an array of electronic devices, including appliances in smart-homes, augmented reality, and wearable sensors. Most notably, new devices that do not have traditional means to enter data, such as a keyboard and screen, could greatly benefit from impedance-based recognition or body channel communication. In fact, a majority of personal electronic devices, as well as interconnected household appliances, found on the consumer market these days solely communicate in a wireless fashion and hence need special provisions to achieve secure communication. Often, the devices resort to a pairing process where users are asked to manually present tokens or type in codes at each device. We make the conjecture that, provided our technology can be miniaturized, none of these interactions would be necessary. Embedding two conductive pads into the exterior of a device could facilitate user authentication and pairing in a way that is natural and understandable to the average user.

Finally, we believe that the electrical properties of the human body provide much untapped potential and can lead to revolutionizing new technologies in the field of Biometrics and Computer Security.

Appendices

A

Biometric Key Generation from Body Impedance Measurements

In this appendix, we explain in detail the convolutional network presented in Section 4.3. It was trained to extract the biometric features from body impedance measurements and significantly outperformed the baseline extraction method.

Contents

A.1 Structure of the Convolutional Network

We specify and train the multilayer convolutional network for feature extraction (Ω_{Siam}) in Torch [209], an open source machine learning framework.

As shown in Listing A.1, the network consists of six sequential layers, followed by two fully connected linear layers. As input, the network takes an impedance measurement, encoded as a one-dimensional vector with 4000 elements. Each layer of the network consists of three sublayers: one-dimensional convolution, rectified linear unit (ReLU), and maximum value subsampling. The convolution layers act as linear filters with learnable parameters which are optimized during training to detect patterns of increasing complexity. ReLU is a transfer function which adds non-linearity, while max-pooling layers allow temporal generalization of the detected patterns. In each layer, the dimensionality of processed data is reduces, and multiple sequential layers allow the network to learn hierarchical representations

of input patterns. Finally, the number of extracted features (N) is controlled by specifying the sizes of the fully connected linear layers (`nrOutputFeatures`) at the output stage of the network.

Once the network is trained, it serves as a non-linear mapping from a raw biometric measurement to a set of N features that we use as input to the key generation scheme.

```
model = nn.Sequential()

model:add( nn.TemporalConvolution(1, 5, 7, 2) )
model:add( nn.ReLU() )
model:add( nn.TemporalMaxPooling(4, 4) )

model:add( nn.TemporalConvolution(5, 10, 7, 1) )
model:add( nn.ReLU() )
model:add( nn.TemporalMaxPooling(3, 3) )

model:add( nn.TemporalConvolution(10, 20, 5, 1) )
model:add( nn.ReLU() )
model:add( nn.TemporalMaxPooling(3, 3) )

model:add( nn.TemporalConvolution(20, 40, 3, 1) )
model:add( nn.ReLU() )
model:add( nn.TemporalMaxPooling(2, 2) )

model:add( nn.TemporalConvolution(40, 80, 3, 1) )
model:add( nn.ReLU() )
model:add( nn.TemporalMaxPooling(2, 2) )

model:add( nn.TemporalConvolution(80, 120, 3, 1) )
model:add( nn.ReLU() )
model:add( nn.TemporalMaxPooling(2, 2) )

model:add( nn.Linear(480, 200) )
model:add( nn.Tanh() )
model:add( nn.Linear(200, nrOutputFeatures) )
```

Listing A.1: Specification of Ω_{Siam} in Torch

B

Biometric Authentication in VANETs

In this appendix, we present a detailed coverage of the related work that proposes to combine biometric authentication and Vehicular Ad-hoc Networks (VANETs). A more general overview is given in Section 5.2.2.

Contents

B.1 Previous Proposals in Literature

Although the works in [161–163] contain pioneering ideas, we consider these proposals unfit for VANET message authentication by today's requirements. We point out some of the weaknesses below.

Enhancing the Security of User Authentication in VANET Using Biometrics The authors of [161] propose an approach for enhancing the security of user authentication in VANETs based on two layered biometric modalities. They suggest a combination of face and fingerprint biometrics to achieve accurate recognition of the owner of a vehicle and provide secure communication in VANETs. They simulate their approach with a custom-built JAVA application that creates a scenario with RSUs, authentication servers and up to 100 vehicles. The computational overhead and authentication time is analyzed in relation to the number of vehicles. One of the drawbacks of the presented scheme is that the XOR operation is used for

encryption. This implies that all data has to be of the same length as the encryption key which is quite constraining for a real-time critical application like the VANET.

After an initial analysis of their protocol the following issues became apparent. If vehicle i requests to communicate with vehicle j, both have to request a session key from the authentication server (AS). The message that vehicle i receives from AS is $K_{ij} \oplus F_i$ with K_{ij} being the session key and F_i the "fingerprint" vector of the user of vehicle i. Similarly, vehicle j receives $K_{ij} \oplus F_j$. In case this message is intercepted by vehicle i, one can compute $F_i \oplus (K_{ij} \oplus F_i) \oplus (K_{ij} \oplus F_j) = F_j$ and learn user j's fingerprint vector (F_i is known to user of vehicle i). This leaks biometric material which is a privacy issue and can be used to sign messages on behalf of the victim. Additionally, an adversary can launch a chosen-plaintext attack on encrypted messages, if the structure of the message is known, since an encrypted message merely consists of $m_{enc} = m_{plaintext} \oplus K_{ij}$. Also, the scheme offers no means for revocation and traceability.

Authentication Scheme based on Biometric Key for VANET Information System in M2M Application Service The work in [162] proposes a scheme that mutually authenticates vehicles by recording vehicle movement as biometric information. The authors augment a Kerberos-like authentication scheme that makes use of the biometric information. Their construction suggests to replace the symmetric keys in the Kerberos scheme with biometrics-based keys. Unfortunately, the proposal neither includes an implementation, nor a description of how the biometric information is incorporated into the authentication protocol. It also does not cover any details about how the biometric information is obtained in the first place.

A Kerberos-like authentication method for the use in VANETs (not biometric-based) was already proposed in [210]. However, the Kerberos authentication construction was designed without taking the volatile nature of IEEE 802.11p [142] wireless transmissions into account. As a result, such schemes often have complex message routing, which leads to limitations in terms of scalability and communication overhead.

Biometrics-based Data Link Layer Anonymous Authentication The proposal in [163] uses biometric encryption to protect privacy together with temporary MAC addresses. The basic principle is to establish a symmetric key between two vehicles using a trusted third party. Thus, each vehicle pair establishes a separate session key for V2X communication. The scheme does not support message

broadcast and therefore disseminating mandatory safety messages comes with a significant overhead. Even if every vehicle has a session key with neighboring vehicles, broadcasting a message degrades to a multicast that requires n messages for n vehicles in the vicinity. Moreover, the authors claim to protect MAC addresses, a statement which, however, is not explained and investigated in the their analysis. The correctness proof only features a simulation. Anonymity is evaluated based on two vehicles being 100 meters apart from each other, a basis for this assumption is not given unfortunately. Furthermore, the simulation evaluates privacy in general terms without considering a specific system, in particular not the system proposed in the paper.

Additional remarks on the authentication phase of the mechanism: The need for nonces is not motivated. Further, using a symmetric key for the communication makes accountability unfeasible and hence the scheme unusable for VANET applications. Lastly, the use of a biometric reading for message authentication is not fully motivated or justified.

References

[1] R.F. Churchhouse. *Codes and Ciphers: Julius Caesar, the Enigma, and the Internet*. Cambridge University Press, 2001.

[2] J. Copeland et al. *The Turing Guide*.

[3] Ronald Kainda, Ivan Flechais, and AW Roscoe. "Security and usability: Analysis and evaluation". In: *Availability, Reliability, and Security, 2010. ARES'10 International Conference on*. IEEE. 2010, pp. 275–282.

[4] Kasper Bonne Rasmussen et al. "Authentication Using Pulse-Response Biometrics". In: *The Network and Distributed System Security Symposium (NDSS)*. 2014.

[5] Ivan Martinovic et al. "Authentication Using Pulse-response Biometrics". In: *Commununications of the ACM* 60.2 (Jan. 2017), pp. 108–115.

[6] Ivan Martinovic et al. "Pulse-Response: Exploring Human Body Impedance for Biometric Recognition". In: *ACM Trans. Priv. Secur.* 20.2 (May 2017), 6:1–6:31.

[7] Marc Roeschlin et al. "Generating Secret Keys from Biometric Body Impedance Measurements". In: *Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society*. ACM. 2016, pp. 59–69.

[8] Marc Roeschlin et al. "Bionyms: Driver-centric Message Authentication using Biometric Measurements". In: *IEEE Vehicular Networking Conference (VNC)*. To appear. 2018.

[9] Marc Roeschlin, Ivan Martinovic, and Kasper Rasmussen. "Device Pairing at the Touch of an Electrode". In: *The Network and Distributed System Security Symposium (NDSS)*. Feb. 2018.

[10] Ivo Sluganovic et al. "Using Reflexive Eye Movements for Fast Challenge-Response Authentication". In: *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*. ACM. 2016, pp. 1056–1067.

[11] Ivo Sluganovic et al. "Analysis of Reflexive Eye Movements for Fast Replay-Resistant Biometric Authentication". In: *ACM Transactions on Privacy and Security (TOPS)* (2018). To appear.

[12] Simon Eberz et al. "Broken Hearted: How To Attack ECG Biometrics". In: *The Network and Distributed System Security Symposium (NDSS)*. 2017.

[13] Orjan G Martinsen and Sverre Grimnes. *Bioimpedance and bioelectricity basics*. Academic press, 2011.

[14] G. J. Saulnier et al. "Electrical impedance tomography". In: *IEEE Signal Processing Magazine* 18.6 (2001), pp. 31–43.

[15] Rio De Janeiro, C E Neves, and M N Souza. “A method for bio-electrical impedance analysis based on a step-voltage response.” In: *Physiological Measurement* 21.3 (2000), pp. 395–408.

[16] Robert F Kushner et al. “Bioelectrical impedance analysis: a review of principles and applications”. In: *J Am Coll Nutr* 11.2 (1992), pp. 199–209.

[17] Matthias Steffen, Adrian Aleksandrowicz, and Steffen Leonhardt. “Mobile noncontact monitoring of heart and lung activity”. In: *IEEE Transactions on Biomedical Circuits and Systems* 1.4 (2007), pp. 250–257.

[18] S Zlochiver et al. “A portable bio-impedance system for monitoring lung resistivity”. In: *Medical engineering & physics* 29.1 (2007), pp. 93–100.

[19] Christina Skourou et al. “Feasibility studies of electrical impedance spectroscopy for early tumor detection in rats”. In: *Physiological Measurement* 25.1 (2004), p. 335.

[20] Daryl G Beetner et al. “Differentiation among basal cell carcinoma, benign lesions, and normal skin using electric impedance”. In: *IEEE Transactions on biomedical engineering* 50.8 (2003), pp. 1020–1025.

[21] Peter Aberg et al. “Skin cancer identification using multifrequency electrical impedance-a potential screening tool”. In: *IEEE transactions on biomedical engineering* 51.12 (2004), pp. 2097–2102.

[22] Sami F Khalil, Mas S Mohktar, and Fatimah Ibrahim. “The theory and fundamentals of bioimpedance analysis in clinical status monitoring and diagnosis of diseases”. In: *Sensors* 14.6 (2014), pp. 10895–10928.

[23] BK Van Kreel. “Multi-frequency bioimpedance measurements of children in intensive care”. In: *Medical and Biological Engineering and Computing* 39.5 (2001), pp. 551–557.

[24] DG Peroni et al. “Bioimpedance monitoring of airway inflammation in asthmatic allergic children”. In: *Allergologia et immunopathologia* 37.1 (2009), pp. 3–6.

[25] Barbara E Lingwood et al. “Measurement of extracellular fluid volume in the neonate using multiple frequency bio-impedance analysis”. In: *Physiological measurement* 21.2 (2000), p. 251.

[26] Fernando Seoane et al. “Spectroscopy study of the dynamics of the transencephalic electrical impedance in the perinatal brain during hypoxia”. In: *Physiological Measurement* 26.5 (2005), p. 849.

[27] James Tibballs. “A comparative study of cardiac output in neonates supported by mechanical ventilation: measurement with thoracic electrical bioimpedance and pulsed Doppler ultrasound”. In: *The Journal of pediatrics* 114.4 (1989), pp. 632–635.

[28] Charles D Wheelwright. *Physiological sensors for use in project Mercury.* Vol. 1082. National Aeronautics and Space Administration, 1962.

[29] K. Hachisuka et al. “Development and performance analysis of an intra-body communication device”. In: *TRANSDUCERS ’03. 12th International Conference on Solid-State Sensors, Actuators and Microsystems. Digest of Technical Papers (Cat. No.03TH8664).* Vol. 2. June 2003, 1722–1725 vol.2.

[30] T. G. Zimmerman. "Personal Area Networks: Near-field intrabody communication". In: *IBM Systems Journal* 35.3.4 (1996), pp. 609–617.

[31] K. S. Kwak, S. Ullah, and N. Ullah. "An overview of IEEE 802.15.6 standard". In: *2010 3rd International Symposium on Applied Sciences in Biomedical and Communication Technologies (ISABEL 2010)*. Nov. 2010, pp. 1–6.

[32] Chang Hee Hyoung et al. "A novel system for intrabody communication: touch-and-play". In: *Circuits and Systems, 2006. ISCAS 2006. Proceedings. 2006 IEEE International Symposium on*. IEEE. 2006, 4–pp.

[33] Maria Amparo Callejon et al. "A comprehensive study into intrabody communication measurements". In: *IEEE Transactions on Instrumentation and Measurement* 62.9 (2013), pp. 2446–2455.

[34] Hao Wang et al. "A 5.4-mW 180-cm transmission distance 2.5-Mb/s advanced techniques-based novel intrabody communication receiver analog front end". In: *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 23.12 (2015), pp. 2829–2841.

[35] N. Cho et al. "A 60 kb/s - 10 Mb/s Adaptive Frequency Hopping Transceiver for Interference-Resilient Body Channel Communication". In: *IEEE Journal of Solid-State Circuits* 44.3 (Mar. 2009), pp. 708–717.

[36] Marc Simon Wegmüller. "Intra-body communication for biomedical sensor networks". PhD thesis. ETH ZURICH, 2007.

[37] Marc Simon Wegmueller et al. "Galvanical coupling for data transmission through the human body". In: *Instrumentation and Measurement Technology Conference, 2006. IMTC 2006. Proceedings of the IEEE*. IEEE. 2006, pp. 1686–1689.

[38] Yong Song et al. "The simulation method of the galvanic coupling intrabody communication with different signal transmission paths". In: *IEEE Transactions on Instrumentation and Measurement* 60.4 (2011), pp. 1257–1266.

[39] J. Bae et al. "The Signal Transmission Mechanism on the Surface of Human Body for Body Channel Communication". In: *IEEE Transactions on Microwave Theory and Techniques* 60.3 (Mar. 2012), pp. 582–593.

[40] J. Wang, Y. Nishikawa, and T. Shibata. "Analysis of On-Body Transmission Mechanism and Characteristic Based on an Electromagnetic Field Approach". In: *IEEE Transactions on Microwave Theory and Techniques* 57.10 (Oct. 2009), pp. 2464–2470.

[41] International Organization for Standardization. *JTC 1/SC 37: Biometrics*. 2018. URL: `https://www.iso.org/isoiec_jtc1sc37.html`.

[42] International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC). *ISO/IEC 2382-37:2017*. BSI Standards Publication. Feb. 2017.

[43] Information Technology Laboratory – National Institute of Standards and Technology. *The Biometrics Resource Center*. 2014. URL: `https://www.nist.gov/itl/computer-security-division/biometrics-resource-center-website`.

[44] National Science & Technology Council. *Biometrics Frequently Asked Questions*. 2006. URL: `http://www.nws-sa.com/biometrics/biooverview.pdf`.

[45] John Woodward, Nicholas Orlans, and Peter Higgins. *Biometrics*. RSA Press Series. McGraw-Hill/Osborne, 2003.

[46] Anil K. Jain, Arun Ross, and Karthik Nandakumar. *Introduction to Biometrics*. Springer, 2011.

[47] Liang Wang et al. *Behavioral Biometrics For Human Identification: Intelligent Applications*. Hershey, PA: Information Science Reference - Imprint of: IGI Publishing, 2009.

[48] R. Ramotowski. *Lee and Gaensslen's Advances in Fingerprint Technology, Third Edition*. CRC Press, 2012.

[49] Anil K. Jain, Arun Ross, and Sharath Pankanti. "Biometrics: a tool for information security". In: *IEEE Transactions on Information Forensics and Security* 1.2 (June 2006), pp. 125–143.

[50] Claude Barral and Assia Tria. "Fake Fingers in Fingerprint Recognition: Glycerin Supersedes Gelatin". In: *Formal to Practical Security*. Ed. by Véronique Cortier et al. Vol. 5458. Lecture Notes in Computer Science. Berlin: Springer, 2009, pp. 57–69.

[51] VIRDI Biometric. *How to make the fake fingerprints (by VIRDI)*. Video. last accessed 08.08.2018. 2009. URL: `http://www.youtube.com/watch?v=-H71tyMupqk`.

[52] Minh Duc Nguyen and Quang Minh Bui. "Your face is NOT your password: Face authentication bypassing - Lenovo - Asus - Toshiba". In: *Briefings of the Black Hat Conference*. 2009.

[53] Arman Boehm et al. "SAFE: Secure authentication with Face and Eyes". In: *International Conference on Privacy and Security in Mobile Systems, (PRISMS)*. June 2013, pp. 1–8.

[54] Javier Galbally et al. "From the Iriscode to the Iris: A New Vulnerability of Iris Recognition Systems". In: *Briefings of the Black Hat Conference*. 2012.

[55] R. Spillane. "Keyboard Apparatus for Personal Identification". In: *IBM Technical Disclosure Bulletin* 17.3346 (1975).

[56] Nathan L. Clarke and Steven Furnell. "Advanced user authentication for mobile devices". In: *Computers & Security* 26.2 (2007), pp. 109–119.

[57] Fabian Monrose, Michael K. Reiter, and Susanne Wetzel. "Password hardening based on keystroke dynamics". In: *Proceedings of the 6th ACM conference on Computer and communications security*. CCS '99. Kent Ridge Digital Labs, Singapore, 1999, pp. 73–82.

[58] Maja Pusara and Carla E. Brodley. "User re-authentication via mouse movements". In: *Proceedings of the 2004 ACM workshop on Visualization and data mining for computer security*. VizSEC/DMSEC '04. Washington DC, USA, 2004, pp. 1–8.

[59] Youssef Nakkabi, Issa Traoré, and Ahmed Awad E. Ahmed. "Improving mouse dynamics biometric performance using variance reduction via extractors with separate features". In: *IEEE Transactions on Systems, Man, and Cybernetics, Part A* 40.6 (2010), pp. 1345–1353.

[60] Hugo Gamboa and Ana Fred. "A behavioral biometric system based on human-computer interaction". In: *Biometric Technology for Human Identification*. Vol. 5404. 2004, pp. 381–392.

[61] Nan Zheng, Aaron Paloski, and Haining Wang. "An efficient user verification system via mouse movements". In: *Proceedings of the 18th ACM conference on Computer and communications security*. CCS '11. Chicago, Illinois, USA, 2011, pp. 139–150.

[62] Zach Jorgensen and Ting Yu. "On mouse dynamics as a behavioral biometric for authentication". In: *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*. ASIACCS '11. Hong Kong, China, 2011, pp. 476–482.

[63] I. Deutschmann, P. Nordstrom, and L. Nilsson. "Continuous Authentication Using Behavioral Biometrics". In: *IT Professional* 15.4 (July 2013), pp. 12–15.

[64] R. Giot, M. El-Abed, and C. Rosenberger. "Keystroke dynamics with low constraints SVM based passphrase enrollment". In: *Biometrics: Theory, Applications, and Systems, 2009. BTAS '09. IEEE 3rd International Conference on*. Sept. 2009, pp. 1–6.

[65] A.A. Ahmed and I. Traore. "Biometric Recognition Based on Free-Text Keystroke Dynamics". In: *Cybernetics, IEEE Transactions on* 44.4 (Apr. 2014), pp. 458–472.

[66] Mario Frank et al. "Touchalytics: On the Applicability of Touchscreen Input as a Behavioral Biometric for Continuous Authentication". In: *IEEE Transactions on Information Forensics and Security* 8.1 (2013), pp. 136–148.

[67] Tao Feng et al. "Continuous mobile authentication using touchscreen gestures". In: *Homeland Security (HST), 2012 IEEE Conference on Technologies for*. Nov. 2012, pp. 451–456.

[68] Cheng Bo, Lan Zhang, and Xiang-Yang Li. "SilentSense: Silent User Identification via Dynamics of Touch and Movement Behavioral Biometrics". In: *CoRR* abs/1309.0073 (2013). URL: `http://arxiv.org/abs/1309.0073`.

[69] Alexander De Luca et al. "Touch Me Once and I Know It's You!: Implicit Authentication Based on Touch Screen Patterns". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: ACM, 2012, pp. 987–996.

[70] Simon Eberz et al. "Preventing Lunchtime Attacks: Fighting Insider Threats With Eye Movement Biometrics". In: *The Network and Distributed System Security Symposium (NDSS)*. Feb. 2015.

[71] Usman Saeed. "Eye movements during scene understanding for biometric identification". In: *Pattern Recognition Letters* (2015).

[72] Ralf Biedert et al. "Stimuli for gaze based intrusion detection". In: *6th International Symposium on Digital Forensics and Information Security* (2012).

[73] C. Gruber et al. "Online Signature Verification With Support Vector Machines Based on LCSS Kernel Functions". In: *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on* 40.4 (Aug. 2010), pp. 1088–1100.

[74] Napa Sae-Bae and N. Memon. “Online Signature Verification on Mobile Devices”. In: *Information Forensics and Security, IEEE Transactions on* 9.6 (June 2014), pp. 933–947.

[75] Meenakshi K Kalera, Sargur Srihari, and Aihua Xu. “Offline signature verification and identification using distance statistics”. In: *International Journal of Pattern Recognition and Artificial Intelligence* 18.07 (2004), pp. 1339–1360.

[76] H. Iwama et al. “The OU-ISIR Gait Database Comprising the Large Population Dataset and Performance Evaluation of Gait Recognition”. In: *Information Forensics and Security, IEEE Transactions on* 7.5 (Oct. 2012), pp. 1511–1521.

[77] Nicolas Scheffer et al. “Recent developments in voice biometrics: Robustness and high accuracy”. In: *Technologies for Homeland Security (HST), 2013 IEEE International Conference on*. IEEE. 2013, pp. 447–452.

[78] Shi-Huang Chen and Yu-Ren Luo. “Speaker verification using MFCC and support vector machine”. In: *Proceedings of the International MultiConference of Engineers and Computer Scientists*. Vol. 1. 2009, pp. 18–20.

[79] Hans van Halteren. “Linguistic Profiling for Author Recognition and Verification”. In: *Proceedings of the 42Nd Annual Meeting on Association for Computational Linguistics*. ACL ’04. Barcelona, Spain: Association for Computational Linguistics, 2004.

[80] R. Layton, P. Watters, and R. Dazeley. “Authorship Attribution for Twitter in 140 Characters or Less”. In: *Cybercrime and Trustworthy Computing Workshop (CTC), 2010 Second*. July 2010, pp. 1–8.

[81] Arslan Brömme and Stephan Al-zubi. “Multifactor Biometric Sketch Authentication”. In: *In Proceedings of the first conference on biometrics and electronic signatures of the GI working group BioSig*. 2003, pp. 81–90.

[82] Mohammed E Fathy et al. “Screen-based active user authentication”. In: *Pattern Recognition Letters* 42 (2014), pp. 122–127.

[83] Malek Ben Salem and Salvatore J Stolfo. “Modeling user search behavior for masquerade detection”. In: *Recent Advances in Intrusion Detection*. Springer. 2011, pp. 181–200.

[84] J.Benito Camiña, Jorge Rodríguez, and Raúl Monroy. “Towards a Masquerade Detection System Based on User’s Tasks”. English. In: *Research in Attacks, Intrusions and Defenses*. Ed. by Angelos Stavrou, Herbert Bos, and Georgios Portokalidis. Vol. 8688. Lecture Notes in Computer Science. Springer International Publishing, 2014, pp. 447–465.

[85] Myriam Abramson. “Cognitive Fingerprints”. In: *AAAI Spring Symposium Series*. 2015.

[86] L. Olejnik and C. Castelluccia. “Towards Web-Based Biometric Systems Using Personal Browsing Interests”. In: *Availability, Reliability and Security (ARES), 2013 Eighth International Conference on*. Sept. 2013, pp. 274–280.

[87] T. Westeyn and T. Starner. “Recognizing song-based blink patterns: applications for restricted and universal access”. In: *Automatic Face and Gesture Recognition, 2004. Proceedings. Sixth IEEE International Conference on*. May 2004, pp. 717–722.

[88] L. Benedikt et al. “Assessing the Uniqueness and Permanence of Facial Actions for Use in Biometric Applications”. In: *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on* 40.3 (May 2010), pp. 449–460.

[89] Mauro Conti, Irina Zachia-Zlatea, and Bruno Crispo. “Mind How You Answer Me!: Transparently Authenticating the User of a Smartphone when Answering or Placing a Call”. In: *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*. ASIACCS ’11. Hong Kong, China: ACM, 2011, pp. 249–259.

[90] Koichiro Niinuma and Anil K Jain. “Continuous user authentication using temporal information”. In: *Biometric Technology for Human Identification VII*. Vol. 7667. International Society for Optics and Photonics. 2010, p. 76670L.

[91] Sensible Vision Inc. *Facial Recognition Provides Continuous System Security*. 2013. URL: `http://www.sensiblevision.com/en-us/fastaccessanywhere/overview.aspx`.

[92] Salil P Banerjee and Damon L Woodard. “Biometric authentication and identification using keystroke dynamics: A survey”. In: *Journal of Pattern Recognition Research* 7.1 (2012), pp. 116–139.

[93] Simon Eberz et al. “Preventing Lunchtime Attacks: Fighting Insider Threats With Eye Movement Biometrics”. In: *22nd Annual Network and Distributed System Security Symposium, NDSS 2015, San Diego, California, USA, February 8-11, 2014*. 2015.

[94] Lyra Nara. *Hand Electrodes Brass (1 Pair)*. 2013. URL: `http://www.lyranara.com/hand-electrodes-brass-1-pair/`.

[95] Tom Fawcett. “An introduction to ROC analysis”. In: *Pattern Recognition Letters* 27.8 (2006), pp. 861–874.

[96] David Umphress and Glen Williams. “Identity verification through keyboard characteristics”. In: *International Journal of Man-Machine Studies* 23.3 (1985), pp. 263–273.

[97] Rick Joyce and Gopal Gupta. “Identity authentication based on keystroke latencies”. In: *Communications of the ACM* 33.2 (1990), pp. 168–176.

[98] L.C.F. Araujo et al. “User authentication through typing biometrics features”. In: *Signal Processing, IEEE Transactions on* 53.2 (Feb. 2005), pp. 851–855.

[99] Sungzoon Cho et al. “Web Based Keystroke Dynamics Identity Verification using Neural Network”. In: *Journal of Organizational Computing and Electronic Commerce* 10 (2000), pp. 295–307.

[100] Dat Tran et al. “Fuzzy and Markov Models for Keystroke Biometrics Authentication”. In: *Proceedings of the 7th WSEAS International Conference on Simulation, Modelling and Optimization*. SMO’07. Beijing, China: World Scientific, Engineering Academy, and Society (WSEAS), 2007, pp. 89–94.

[101] Pin Shen Teh, Andrew Beng Jin Teoh, and Shigang Yue. “A survey of keystroke dynamics biometrics”. In: *The Scientific World Journal* 2013 (2013).

[102] Kevin S. Killourhy. “A Scientific Understanding of Keystroke Dynamics”. PhD thesis. Pittsburgh, PA: Carnegie Mellon University, 2012.

[103] Daniele Gunetti and Claudia Picardi. "Keystroke analysis of free text". In: *ACM Transactions on Information and System Security (TISSEC)* 8.3 (2005), pp. 312–347.

[104] Saleh Bleha, Charles Slivinsky, and Bassam Hussien. "Computer-access security systems using keystroke dynamics". In: *IEEE Transactions on pattern analysis and machine intelligence* 12.12 (1990), pp. 1217–1222.

[105] Sylvain Hocquet, J-Y Ramel, and Hubert Cardot. "Fusion of methods for keystroke dynamic authentication". In: *Fourth IEEE Workshop on Automatic Identification Advanced Technologies (AutoID'05)*. IEEE. 2005, pp. 224–229.

[106] Chee Meng Tey, Payas Gupta, and Debin Gao. "I can be You: Questioning the use of Keystroke Dynamics as Biometrics". In: *20th Annual Network and Distributed System Security Symposium, NDSS 2013, San Diego, California, USA, February 24-27, 2013*. 2013. URL: `http://internetsociety.org/doc/i-can-be-you-questioning-use-keystroke-dynamics-biometrics`.

[107] Christian Holz and Marius Knaust. "Biometric touch sensing: Seamlessly augmenting each touch with continuous authentication". In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. ACM. 2015, pp. 303–312.

[108] Kenneth Revett and Sérgio Tenreiro de Magalhães. "Cognitive Biometrics: Challenges for the Future". In: *Global Security, Safety, and Sustainability*. Ed. by Sérgio Tenreiro de Magalhães, Hamid Jahankhani, and Ali G. Hessami. Vol. 92. Communications in Computer and Information Science. Springer, 2010, pp. 79–86.

[109] Cory Cornelius et al. "Who Wears Me? Bioimpedance as a Passive Biometric". In: *Proceedings of the 3rd USENIX Workshop on Health Security and Privacy*. Ed. by Carl A. Gunter and Zachary N. J. Peterson. Berkeley, CA, USA: USENIX Association, 2012.

[110] Cory Cornelius et al. "A Wearable System That Knows Who Wears It". In: *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services*. Bretton Woods, New Hampshire, USA: ACM, 2014, pp. 55–67.

[111] Andy Greenberg. *OPM Now Admits 5.6m Feds' Fingerprints Were Stolen By Hackers*. 2016. URL: `http://www.wired.com/2015/09/opm-now-admits-5-6m-feds-fingerprints-stolen-hackers/` (visited on 01/20/2016).

[112] Samuel Gibbs. *HTC stored user fingerprints as image file in unencrypted folder*. 2016. URL: `http://www.theguardian.com/technology/2015/aug/10/htc-fingerprints-world-readable-unencrypted-folder` (visited on 01/21/2016).

[113] Arun Ross, Jidnya Shah, and Anil K Jain. "From template to image: Reconstructing fingerprints from minutiae points". In: *IEEE transactions on pattern analysis and machine intelligence* 29.4 (2007), pp. 544–560.

[114] Javier Galbally et al. "Iris image reconstruction from binary templates: An efficient probabilistic approach based on genetic algorithms". In: *Computer Vision and Image Understanding* 117.10 (2013), pp. 1512–1525.

[115] International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC). *ISO/IEC 24745*. BSI Standards Publication. June 2011.

[116] Lucas Ballard, Seny Kamara, and Michael K Reiter. "The Practical Subtleties of Biometric Key Generation". In: *USENIX Security Symposium* (2008), pp. 61–74.

[117] Christian Rathgeb and Andreas Uhl. "A survey on biometric cryptosystems and cancelable biometrics". In: *EURASIP Journal on Information Security* 2011.1 (2011), p. 3.

[118] Karthik Nandakumar, Anil K. Jain, and Abhishek Nagar. "Biometric template security". In: *Eurasip Journal on Advances in Signal Processing* 2008 (2008).

[119] Nalini K. Ratha et al. "Generating cancelable fingerprint templates". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29.4 (2007), pp. 561–572.

[120] K. Nandakumar and A.K. Jain. "Biometric Template Protection: Bridging the performance gap between theory and practice". In: *Signal Processing Magazine, IEEE* 32.5 (Sept. 2015), pp. 88–100.

[121] Lucas Ballard et al. "Towards practical biometric key generation with randomized biometric templates". In: *Proceedings of the 15th ACM conference on Computer and communications security* (2008), p. 235.

[122] Yevgeniy Dodis et al. "Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data". In: *SIAM Journal on Computing* (2006), pp. 97–139.

[123] Qiming Li, Yagiz Sutcu, and Nasir Memon. "Secure sketch for biometric templates". In: *Advances in Cryptology–ASIACRYPT 2006*. Springer, 2006, pp. 99–113.

[124] Yagiz Sutcu, Qiming Li, and Nasir Memon. "Protecting biometric templates with sketch: Theory and practice". In: *Information Forensics and Security, IEEE Transactions on* 2.3 (2007), pp. 503–512.

[125] Christian Rathgeb and Andreas Uhl. "Privacy Preserving Key Generation for Iris Biometrics." In: *Communications and Multimedia Security*. Vol. 6109. Springer. 2010, pp. 191–200.

[126] Brent Carrara and Carlisle Adams. "You are the key: Generating cryptographic keys from voice biometrics". In: *8th International Conference on Privacy, Security and Trust* (2010), pp. 213–222.

[127] Qiming Li, Muchuan Guo, and Ee-Chien Chang. "Fuzzy extractors for asymmetric biometric representations". In: *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on*. IEEE. 2008, pp. 1–6.

[128] Hao Feng and Chan Choong Wah. "Private key generation from on-line handwritten signatures". In: *Information Management & Computer Security* 10.4 (2002), pp. 159–164.

[129] Clam Vielhauer, Ralf Steinmetz, and Astrid Mayerhöfer. "Biometric hash based on statistical features of online signatures". In: *Pattern Recognition, 2002. Proceedings. 16th International Conference on*. Vol. 1. IEEE. 2002, pp. 123–126.

[130] Geoffrey E. Hinton et al. "Improving neural networks by preventing co-adaptation of feature detectors". In: *arXiv preprint* (2012).

[131] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. “Sequence to sequence learning with neural networks”. In: *Advances in neural information processing systems*. 2014, pp. 3104–3112.

[132] Kaiming He et al. “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification”. In: *Proceedings of the IEEE international conference on computer vision*. 2015, pp. 1026–1034.

[133] S. Chopra, R. Hadsell, and Y. LeCun. “Learning a similarity metric discriminatively, with application to face verification”. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. 2005, pp. 539–546.

[134] Jane Bromley et al. “Signature Verification using a "Siamese" Time Delay Neural Network”. In: *Advances in Neural Information Processing Systems 6*. Ed. by J.D. Cowan, G. Tesauro, and J. Alspector. Morgan-Kaufmann, 1994, pp. 737–744.

[135] Jason Yosinski et al. “How transferable are features in deep neural networks?” In: *Advances in Neural Information Processing Systems*. 2014, pp. 3320–3328.

[136] J.L. Massey. “Guessing and entropy”. In: *Information Theory, 1994. Proceedings., IEEE International Symposium on*. 1994, p. 204.

[137] Fabian Monrose et al. “Cryptographic key generation from voice”. In: *Security and Privacy, 2001. S&P 2001. Proceedings. 2001 IEEE Symposium on*. IEEE. 2001, pp. 202–213.

[138] Payas Gupta and Debin Gao. “Fighting Coercion Attacks in Key Generation Using Skin Conductance”. In: *Proceedings of the 19th USENIX Conference on Security*. USENIX Security’10. Washington, DC: USENIX Association, 2010, pp. 30–30.

[139] Andrew Beng Jin Teoh and Lee-Ying Chong. “Secure speech template protection in speaker verification system”. In: *Speech communication* 52.2 (2010), pp. 150–163.

[140] K. Inthavisas and D. Lopresti. “Secure speech biometric templates for user authentication”. In: *IET Biometrics* 1.1 (2012), p. 46.

[141] Mani Amoozadeh et al. “Platoon management with cooperative adaptive cruise control enabled by VANET”. In: *Vehicular Communications* 2.2 (2015), pp. 110–123.

[142] Daniel Jiang and Luca Delgrossi. “IEEE 802.11p: Towards an international standard for wireless access in vehicular environments”. In: *IEEE Vehicular Technology Conference* 1 (2008), pp. 2036–2040.

[143] Andreas Festag. “Cooperative intelligent transport systems standards in Europe”. In: *IEEE Communications Magazine* 52.12 (2014), pp. 166–172.

[144] Baik Hoh et al. “Achieving guaranteed anonymity in gps traces via uncertainty-aware path cloaking”. In: *IEEE Transactions on Mobile Computing* 9.8 (2010), pp. 1089–1107.

[145] J. Petit et al. “Pseudonym Schemes in Vehicular Networks: A Survey”. In: *IEEE Communications Surveys Tutorials* 17 (2015), pp. 228–255.

[146] S. Zhao et al. “A Survey of Applications of Identity-Based Cryptography in Mobile Ad-Hoc Networks”. In: *IEEE Communications Surveys Tutorials* 14 (2012), pp. 380–400.

[147] Arun Hegde. *Carsharing Market to witness a massive 34%+ growth over 2016-2024*. 2017. URL: http://markets.businessinsider.com/news/stocks/carsharing-market-to-witness-a-massive-34-growth-over-2016-2024-1002207831 (visited on 06/02/2018).

[148] Warwick Goodall et al. "The rise of mobility as a service". In: *Deloitte Rev* 20 (2017), pp. 112–129.

[149] Andrew Krok. *Daimler's new EV charging system takes the hassle out of payment.* 2018. URL: https://www.cnet.com/roadshow/news/daimler-plug-and-charge-hubject-appless-ev-charging/ (visited on 04/19/2018).

[150] Magda El Zarki et al. "Security Issues in a Future Vehicular Network". In: *Symposium A Quarterly Journal In Modern Foreign Literatures* 35 (2002), pp. 270–274.

[151] Maxim Raya, Adel Aziz, and Jean-Pierre Hubaux. "Efficient Secure Aggregation in VANETs". In: *Proceedings of the 3rd International Workshop on Vehicular Ad Hoc Networks*. VANET '06. Los Angeles, CA, USA: ACM, 2006, pp. 67–75.

[152] Zhendong Ma, Frank Kargl, and Michael Weber. "Pseudonym-on-demand: A new pseudonym refill strategy for vehicular communications". In: *IEEE Vehicular Technology Conference* 9 (2008), pp. 1–5.

[153] Dan Boneh and Matt Franklin. "Identity-based encryption from the Weil pairing". In: *Annual international cryptology conference*. Berlin Heidelberg: Springer, 2001, pp. 213–229.

[154] Shushan Zhao et al. "A survey of applications of identity-based cryptography in mobile ad-hoc networks". In: *IEEE Communications Surveys and Tutorials* 14.2 (2012), pp. 380–399.

[155] B. H. Kim et al. "Anonymous and Traceable Communication Using Tamper-Proof Device for Vehicular Ad Hoc Networks". In: *2007 International Conference on Convergence Information Technology (ICCIT)*. USA: IEEE, 2007, pp. 681–686.

[156] Kenneth G Paterson and Geraint Price. "A comparison between traditional public key infrastructures and identity-based cryptography". In: *Information Security Technical Report* 8.3 (2003), pp. 57–72.

[157] Giorgio Calandriello et al. "Efficient and Robust Pseudonymous Authentication in VANET". In: *Proceedings of the Fourth ACM International Workshop on Vehicular Ad Hoc Networks*. VANET '07. Montreal, Quebec, Canada: ACM, 2007, pp. 19–28.

[158] J. Guo, J. P. Baugh, and S. Wang. "A Group Signature Based Secure and Privacy-Preserving Vehicular Communication Framework". In: *2007 Mobile Networking for Vehicular Environments*. USA: IEEE, 2007, pp. 103–108.

[159] C. Zhang et al. "An Efficient Message Authentication Scheme for Vehicular Communications". In: *IEEE Transactions on Vehicular Technology* 57 (2008), pp. 3357–3368.

[160] Yih-Chun Hu and Kenneth P Laberteaux. "Strong VANET security on a budget". In: *Proceedings of Workshop on Embedded Security in Cars (ESCAR)*. Vol. 6. USA: IEEE, 2006, pp. 1–9.

[161] P. Remyakrishnan and C. Tripti. "A Novel Approach for Enhancing the Security of User Authentication in VANET Using Biometrics". In: *Emerging ICT for Bridging the Future - Proceedings of the 49th Annual Convention of the Computer Society of India CSI Volume 2*. Berlin, Germany: Springer, 2015, pp. 299–306.

[162] Keun-Ho Lee and Soo Kyun Kim. "Authentication Scheme based on Biometric Key for VANET Information System in M2M Application Service". In: *Appl. Math* 9 (2015), pp. 645–651.

[163] L. Yao et al. "Biometrics-based Data Link Layer Anonymous Authentication in VANETs". In: *2013 Seventh International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing*. USA: IEEE, 2013, pp. 182–187.

[164] Philippe Golle and Kurt Partridge. "On the Anonymity of Home/Work Location Pairs". In: *Pervasive Computing: 7th International Conference, Pervasive 2009, Nara, Japan, May 11-14, 2009. Proceedings*. Berlin, Germany: Springer, 2009, pp. 390–397.

[165] Ivan Damgård and Eiichiro Fujisaki. "A statistically-hiding integer commitment scheme based on groups with hidden order". In: *International Conference on the Theory and Application of Cryptology and Information Security*. Berlin, Germany: Springer, 2002, pp. 125–142.

[166] Julien Freudiger et al. "Mix-Zones for Location Privacy in Vehicular Networks". In: *ACM Workshop on Wireless Networking for Intelligent Transportation Systems (WiN-ITS)* 51 (2007), pp. 1–7.

[167] European Telecommunications Standards Institute. *Intelligent Transport Systems (ITS); Security; Stage 3 Mapping for IEEE 1609.2*. Vol. 1. Sophia-Antipolis, France: ETSI, 2012, pp. 1–26.

[168] D. Dolev and A. Yao. "On the security of public key protocols". In: *IEEE Transactions on Information Theory* 29 (1983), pp. 198–208.

[169] Ivan Martinovic et al. "Pulse-Response: Exploring Human Body Impedance for Biometric Recognition". In: *ACM Transactions on Privacy and Security (TOPS)* 20.2 (2017), p. 6.

[170] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, et al. "Deep Face Recognition". In: *BMVC*. Vol. 1. Berlin, Germany: Springer, 2015, p. 6.

[171] S Eberz et al. "Preventing Lunchtime Attacks: Fighting Insider Threats With Eye Movement Biometrics." In: *NDSS* 20 (2015), pp. 1–13.

[172] Andreas Riener and Alois Ferscha. "Supporting implicit human-to-vehicle interaction: Driver identification from sitting postures". In: *The first annual international symposium on vehicular computing systems (isvcs 2008)*. Vol. 10. USA: isvcs, 2008, p. 10.

[173] Tyler Kaczmarek, Ercan Ozturk, and Gene Tsudik. "Assentication: User Deauthentication and Lunchtime Attack Mitigation with Seated Posture Biometric". In: *arXiv preprint arXiv:1708.03978* 3 (2017), p. 13.

[174] Fahim Sufi, Ibrahim Khalil, and Jiankun Hu. "ECG-based authentication". In: *Handbook of information and communication security*. Springer, 2010, pp. 309–331.

[175] Jonathan M McCune, Adrian Perrig, and Michael K Reiter. "Seeing-is-believing: Using camera phones for human-verifiable authentication". In: *Security and privacy, 2005 IEEE symposium on*. IEEE. 2005, pp. 110–124.

[176] Claudio Soriente, Gene Tsudik, and Ersin Uzun. "HAPADEP: human-assisted pure audio device pairing". In: *Information Security* (2008), pp. 385–400.

[177] Claude Castelluccia and Pars Mutaf. "Shake them up!: a movement-based pairing protocol for cpu-constrained devices". In: *Proceedings of the 3rd international conference on Mobile systems, applications, and services*. ACM. 2005, pp. 51–64.

[178] Rene Mayrhofer and Hans Gellersen. "Shake well before use: Intuitive and secure pairing of mobile devices". In: *IEEE Transactions on Mobile Computing* 8.6 (2009), pp. 792–806.

[179] Chunyi Peng et al. "Point&Connect: intention-based device pairing for mobile phone users". In: *Proceedings of the 7th international conference on Mobile systems, applications, and services*. ACM. 2009, pp. 137–150.

[180] S. Abhishek Anand and Nitesh Saxena. "Vibreaker". In: *Proceedings of the 9th ACM Conference on Security & Privacy in Wireless and Mobile Networks - WiSec '16*. New York, New York, USA: ACM Press, 2016, pp. 103–108.

[181] Ivo Sluganovic et al. "HoloPair: Securing Shared Augmented Reality Using Microsoft HoloLens". In: *Proceedings of the 33rd Annual Computer Security Applications Conference (ACSAC 2017)*. 2017, p. 13.

[182] Arun Kumar et al. "A comparative study of secure device pairing methods". In: *Pervasive and Mobile Computing* 5.6 (2009), pp. 734–749.

[183] Iulia Ion et al. "Influence of user perception, security needs, and social factors on device pairing method choices". In: *Proceedings of the Sixth Symposium on Usable Privacy and Security*. ACM. 2010, p. 6.

[184] C. Hu et al. "Body Area Network Security: A Fuzzy Attribute-Based Signcryption Scheme". In: *IEEE Journal on Selected Areas in Communications* 31.9 (Sept. 2013), pp. 37–46.

[185] M. Li, W. Lou, and K. Ren. "Data security and privacy in wireless body area networks". In: *IEEE Wireless Communications* 17.1 (Feb. 2010), pp. 51–58.

[186] S. Warren et al. "Interoperability and Security in Wireless Body Area Network Infrastructures". In: *2005 IEEE Engineering in Medicine and Biology 27th Annual Conference*. Jan. 2005, pp. 3837–3840.

[187] Michael Rushanan et al. "SoK: Security and privacy in implantable medical devices and body area networks". In: *Security and Privacy (SP), 2014 IEEE Symposium on*. IEEE. 2014, pp. 524–539.

[188] Lu Shi et al. "BANA: body area network authentication exploiting channel characteristics". In: *IEEE Journal on selected Areas in Communications* 31.9 (2013), pp. 1803–1816.

[189] Sang-Yoon Chang et al. "Body Area Network Security: Robust Key Establishment Using Human Body Channel." In: *HealthSec*. 2012, pp. 5–5.

[190] Nobuyuki Matsushita et al. "Wearable key: Device for personalizing nearby environment". In: *Wearable Computers, The Fourth International Symposium on*. IEEE. 2000, pp. 119–126.

[191] Mitsuru Shinagawa et al. “A near-field-sensing transceiver for intrabody communication based on the electrooptic effect”. In: *IEEE Transactions on instrumentation and measurement* 53.6 (2004), pp. 1533–1538.

[192] MirHojjat Seyedi et al. “A survey on intrabody communications for body area network applications”. In: *IEEE Transactions on Biomedical Engineering* 60.8 (2013), pp. 2067–2079.

[193] Xi Mei Chen et al. “Study of channel characteristics for galvanic-type intra-body communication based on a transfer function from a quasi-static field model”. In: *Sensors* 12.12 (2012), pp. 16433–16450.

[194] Željka Lucev, Igor Krois, and Mario Cifrek. “A capacitive intrabody communication channel from 100 kHz to 100 MHz”. In: *IEEE Transactions on Instrumentation and Measurement* 61.12 (2012), pp. 3280–3289.

[195] Keisuke Hachisuka et al. “Intra-body data transmission for the personal area network”. In: *Microsystem Technologies* 11.8-10 (2005), pp. 1020–1027.

[196] Namjun Cho et al. “The human body characteristics as a signal transmission medium for intrabody communication”. In: *IEEE transactions on microwave theory and techniques* 55.5 (2007), pp. 1080–1086.

[197] Srdjan Čapkun et al. “Integrity codes: Message integrity protection and authentication over insecure channels”. In: *IEEE Transactions on Dependable and Secure Computing* 5.4 (2008), pp. 208–223.

[198] Hao Wang, Jianfeng Wang, and Chiu Sing Choy. “A 2.5-Mbps, 170-cm transmission distance IntraBody communication receiver front end design and its synchronization technique research”. In: *Circuits and Systems (MWSCAS), 2014 IEEE 57th International Midwest Symposium on*. IEEE. 2014, pp. 643–646.

[199] OMRON Healthcare. *Weight Management - Frequently asked questions*. 2017. URL: `https://www.omron-healthcare.com/en-gb/products/weightmanagement` (visited on 06/29/2017).

[200] International Commission on Non-Ionizing Radiation Protection and others. “Guidelines for limiting exposure to time-varying electric and magnetic fields (1 Hz to 100 kHz)”. In: *Health physics* 99.6 (2010), pp. 818–836.

[201] A Ahlbom et al. “Guidelines for limiting exposure to time-varying electric, magnetic, and electromagnetic fields (up to 300 GHz)”. In: *Health physics* 74.4 (1998), pp. 494–521.

[202] Federal Communications Commission. *Electronic Code of Federal Regulations: Title 47: Chapter I*. 2017. URL: `http://www.ecfr.gov/cgi-bin/text-idx?mc=true&node=pt47.1.15` (visited on 06/30/2017).

[203] Željka Lučev Vasić, Igor Krois, and Mario Cifrek. “On a pulse response of a capacitive intrabody communication channel”. In: *EUROCON, 2013 IEEE*. IEEE. 2013, pp. 1785–1789.

[204] S Gabriel, RW Lau, and Camelia Gabriel. “The dielectric properties of biological tissues: II. Measurements in the frequency range 10 Hz to 20 GHz”. In: *Physics in medicine and biology* 41.11 (1996), p. 2251.

[205] Željka Lučev, Igor Krois, and Mario Cifrek. "Effect of body positions and movements in a capacitive intrabody communication channel from 100 kHz to 100 MHz". In: *Instrumentation and Measurement Technology Conference (I2MTC), 2012 IEEE International*. IEEE. 2012, pp. 2791–2795.

[206] L. H. Nguyen and A. W. Roscoe. "Authentication Protocols Based on Low-bandwidth Unspoofable Channels: A Comparative Survey". In: *J. Comput. Secur.* 19.1 (2011), pp. 139–201.

[207] Ronald Kainda, Ivan Flechais, and AW Roscoe. "Usability and security of out-of-band channels in secure device pairing protocols". In: *Proceedings of the 5th Symposium on Usable Privacy and Security*. ACM. 2009, p. 11.

[208] Ernst Haselsteiner and Klemens Breitfuß. "Security in near field communication (NFC)". In: *Workshop on RFID security*. 2006, pp. 12–14.

[209] Ronan Collobert, Koray Kavukcuoglu, and Clément Farabet. *Torch7: A Matlab-like Environment for Machine Learning*. URL: `http://www.torch.ch`.

[210] Hasnaa Moustafa, Gilles Bourdon, and Yvon Gourhant. "Providing authentication and access control in vehicular network environment". In: *Security and Privacy in Dynamic Environments* 201 (2006), pp. 62–73.