Universidad
Carlos III de Madrid

# TESIS DOCTORAL

## Cybersecurity in Implantable Medical Devices

Autor:
Carmen Cámara Núñez

Directores:
Dr. Pedro Peris López
Dr. Juan M. Estévez Tapiador

DEPARTAMENTO DE INFORMÁTICA

Leganés, 19 de Diciembre de 2017

# TESIS DOCTORAL

# Cybersecurity in Implantable Medical Devices

Autor: *Carmen Cámara Núñez*

Directores: *Dr. Pedro Peris López*<br>*Dr. Juan M. Estévez Tapiador*

Firma del Tribunal Calificador:

Presidente:

Vocal:

Secretario:

Calificación:

Leganés, de de

# Abstract

Implantable Medical Devices (IMDs) are electronic devices implanted within the body to treat a medical condition, monitor the state or improve the functioning of some body part, or just to provide the patient with a capability that he did not possess before [86]. Current examples of IMDs include pacemakers and defibrillators to monitor and treat cardiac conditions; neurostimulators for deep brain stimulation in cases such as epilepsy or Parkinson; drug delivery systems in the form of infusion pumps; and a variety of biosensors to acquire and process different biosignals.

Some of the newest IMDs have started to incorporate numerous communication and networking functions—usually known as "telemetry"—, as well as increasingly more sophisticated computing capabilities. This has provided implants with more intelligence and patients with more autonomy, as medical personnel can access data and reconfigure the implant remotely (i.e., without the patient being physically present in medical facilities). Apart from a significant cost reduction, telemetry and computing capabilities also allow healthcare providers to constantly monitor the patient's condition and to develop new diagnostic techniques based on an Intra Body Network (IBN) of medical devices [25, 26, 201].

Evolving from a mere electromechanical IMD to one with more advanced computing and communication capabilities has many benefits but also entails numerous security and privacy risks for the patient. The majority of such risks are relatively well known in classical computing scenarios, though in many respects their repercussions are far more critical in the case of implants. Attacks against an IMD can put at risk the safety of the patient who carries it, with fatal consequences in certain cases. Causing an intentional malfunction of an implant can lead to death and, as recognized by the U.S. Food and Drug Administration (FDA), such deliberate attacks could be far more difficult to detect than accidental ones [61]. Furthermore, these devices store and transmit very sensitive medical information that requires protection, as dictated by European (e.g., Directive 95/46/ECC) and U.S. (e.g., CFR 164.312) Directives [94, 204].

The wireless communication capabilities present in many modern IMDs are a major source of security risks, particularly while the patient is in open (i.e., non-medical) environments. To begin with, the implant becomes no longer "invisible", as its presence could be remotely detected [48]. Furthermore, it facilitates the access to transmitted data by eavesdroppers who simply listen to the (insecure) channel [83]. This could result in a major

privacy breach, as IMDs store sensitive information such as vital signals, diagnosed conditions, therapies, and a variety of personal data (e.g., birth date, name, and other medically relevant identifiers). A vulnerable communication channel also makes it easier to attack the implant in ways similar to those used against more common computing devices [118, 129, 156], i.e., by forging, altering, or replying previously captured messages [82]. This could potentially allow an adversary to monitor and modify the implant without necessarily being close to the victim [164]. In this regard, the concerns of former U.S. vice-president Dick Cheney constitute an excellent example: he had his Implantable Cardioverter Defibrillator (ICD) replaced by another without WiFi capability [219].

While there are still no known real-world incidents, several attacks on IMDs have been successfully demonstrated in the lab [83, 133, 143]. These attacks have shown how an adversary can disable or reprogram therapies on an ICD with wireless connectivity, and even inducing a shock state to the patient [65]. Other attacks deplete the battery and render the device inoperative [91], which often implies that the patient must undergo a surgical procedure to have the IMD replaced. Moreover, in the case of cardiac implants, they have a switch that can be turned off merely by applying a magnetic field [149]. The existence of this mechanism is motivated by the need to shield ICDs to electromagnetic fields, for instance when the patient undergoes cardiac surgery using electrocautery devices [47]. However, this could be easily exploited by an attacker, since activating such a primitive mechanism does not require any kind of authentication.

In order to prevent attacks, it is imperative that the new generation of IMDs will be equipped with strong mechanisms guaranteeing basic security properties such as confidentiality, integrity, and availability. For example, mutual authentication between the IMD and medical personnel is essential, as both parties must be confident that the other end is who claims to be. In the case of the IMD, only commands coming from authenticated parties should be considered, while medical personnel should not trust any message claiming to come from the IMD unless sufficient guarantees are given.

Preserving the confidentiality of the information stored in and transmitted by the IMD is another mandatory aspect. The device must implement appropriate security policies that restrict what entities can reconfigure the IMD or get access to the information stored in it, ensuring that only authorized operations are executed. Similarly, security mechanisms have to be implemented to protect the content of messages exchanged through an insecure wireless channel.

Integrity protection is equally important to ensure that information has not been modified in transit. For example, if the information sent by the implant to the Programmer is altered, the doctor might make a wrong decision. Conversely, if a command sent to the implant is forged, modified, or simply contains errors, its execution could result in a compromise of the patient's physical integrity.

Technical security mechanisms should be incorporated in the design phase and complemented with appropriate legal and administrative measures. Current legislation is rather permissive in this regard, allowing the use of implants like ICDs that do not incorporate any security mechanisms. Regulatory authorities like the FDA in the U.S or the EMA (European Medicines Agency) in Europe should promote metrics and frameworks for assessing the security of IMDs. These assessments should be mandatory by law, requiring an adequate security level for an implant before approving its use. Moreover, both the security measures supported on each IMD and the security assessment results should be made public.

Prudent engineering practices well known in the safety and security domains should be followed in the design of IMDs. If hardware errors are detected, it often entails a replacement of the implant, with the associated risks linked to a surgery. One of the main sources of failure when treating or monitoring a patient is precisely malfunctions of the device itself. These failures are known as "recalls" or "advisories", and it is estimated that they affect around 2.6% of patients carrying an implant. Furthermore, the software running on the device should strictly support the functionalities required to perform the medical and operational tasks for what it was designed, and no more [66, 134, 213].

In Chapter 1, we present a survey of security and privacy issues in IMDs, discuss the most relevant mechanisms proposed to address these challenges, and analyze their suitability, advantages, and main drawbacks. In Chapter 2, we show how the use of highly compressed electrocardiogram (ECG) signals (only 24 coefficients of Hadamard Transform) is enough to unequivocally identify individuals with a high performance (classification accuracy of 97% and with identification system errors in the order of $10^{-2}$). In Chapter 3 we introduce a new Continuous Authentication scheme that, contrarily to previous works in this area, considers ECG signals as continuous data streams. The proposed ECG-based CA system is intended for real-time applications and is able to offer an accuracy up to 96%, with an almost perfect system performance (kappa statistic $>$ 80%). In Chapter

4, we propose a distance bounding protocol to manage access control of IMDs: ACIMD. ACIMD combines two features namely identity verification (authentication) and proximity verification (distance checking). The authentication mechanism we developed conforms to the ISO/IEC 9798-2 standard and is performed using the whole ECG signal of a device holder, which is hardly replicable by a distant attacker. We evaluate the performance of ACIMD using ECG signals of 199 individuals over 24 hours, considering three adversary strategies. Results show that an accuracy of 87.07% in authentication can be achieved. Finally, in Chapter 5 we extract some conclusions and summarize the published works (i.e., scientific journals with high impact factor and prestigious international conferences).

**Keywords:** Security, Privacy, Implantable Medical Devices (IMDs)

# Resumen

Los Dispositivos Médicos Implantables (DMIs) son dispositivos electrónicos implantados dentro del cuerpo para tratar una enfermedad, controlar el estado o mejorar el funcionamiento de alguna parte del cuerpo, o simplemente para proporcionar al paciente una capacidad que no poseía antes [86]. Ejemplos actuales de DMI incluyen marcapasos y desfibriladores para monitorear y tratar afecciones cardíacas; neuroestimuladores para la estimulación cerebral profunda en casos como la epilepsia o el Parkinson; sistemas de administración de fármacos en forma de bombas de infusión; y una variedad de biosensores para adquirir y procesar diferentes bioseñales.

Los DMIs más modernos han comenzado a incorporar numerosas funciones de comunicación y redes (generalmente conocidas como telemetría) así como capacidades de computación cada vez más sofisticadas. Esto ha propiciado implantes con mayor inteligencia y pacientes con más autonomía, ya que el personal médico puede acceder a los datos y reconfigurar el implante de forma remota (es decir, sin que el paciente esté físicamente presente en las instalaciones médicas). Aparte de una importante reducción de costos, las capacidades de telemetría y cómputo también permiten a los profesionales de la atención médica monitorear constantemente la condición del paciente y desarrollar nuevas técnicas de diagnóstico basadas en una Intra Body Network (IBN) de dispositivos médicos [25, 26, 201].

Evolucionar desde un DMI electromecánico a uno con capacidades de cómputo y de comunicación más avanzadas tiene muchos beneficios pero también conlleva numerosos riesgos de seguridad y privacidad para el paciente. La mayoría de estos riesgos son relativamente bien conocidos en los escenarios clásicos de comunicaciones entre dispositivos, aunque en muchos aspectos sus repercusiones son mucho más críticas en el caso de los implantes. Los ataques contra un DMI pueden poner en riesgo la seguridad del paciente que lo porta, con consecuencias fatales en ciertos casos. Causar un mal funcionamiento intencionado en un implante puede causar la muerte y, tal como lo reconoce la Food and Drug Administration (FDA) de EE.UU, tales ataques deliberados podrían ser mucho más difíciles de detectar que los ataques accidentales [61]. Además, estos dispositivos almacenan y transmiten información médica muy delicada que requiere se protegida, según lo dictado por las directivas europeas (por ejemplo, la

Directiva 95/46/ECC) y estadunidenses (por ejemplo, la Directiva CFR 164.312) [94, 204].

Si bien todavía no se conocen incidentes reales, se han demostrado con éxito varios ataques contra DMIs en el laboratorio [83, 133, 143]. Estos ataques han demostrado cómo un adversario puede desactivar o reprogramar terapias en un marcapasos con conectividad inalámbrica e incluso inducir un estado de shock al paciente [65]. Otros ataques agotan la batería y dejan al dispositivo inoperativo [91], lo que a menudo implica que el paciente deba someterse a un procedimiento quirúrgico para reemplazar la batería del DMI. Además, en el caso de los implantes cardíacos, tienen un interruptor cuya posición de desconexión se consigue simplemente aplicando un campo magnético intenso [149]. La existencia de este mecanismo está motivada por la necesidad de proteger a los DMIs frete a posibles campos electromagnéticos, por ejemplo, cuando el paciente se somete a una cirugía cardíaca usando dispositivos de electrocauterización [47]. Sin embargo, esto podría ser explotado fácilmente por un atacante, ya que la activación de dicho mecanismo primitivo no requiere ningún tipo de autenticación.

Garantizar la confidencialidad de la información almacenada y transmitida por el DMI es otro aspecto obligatorio. El dispositivo debe implementar políticas de seguridad apropiadas que restrinjan qué entidades pueden reconfigurar el DMI o acceder a la información almacenada en él, asegurando que sólo se ejecuten las operaciones autorizadas. De la misma manera, mecanismos de seguridad deben ser implementados para proteger el contenido de los mensajes intercambiados a través de un canal inalámbrico no seguro.

La protección de la integridad es igualmente importante para garantizar que la información no se haya modificado durante el tránsito. Por ejemplo, si la información enviada por el implante al programador se altera, el médico podría tomar una decisión equivocada. Por el contrario, si un comando enviado al implante se falsifica, modifica o simplemente contiene errores, su ejecución podría comprometer la integridad física del paciente.

Los mecanismos de seguridad deberían incorporarse en la fase de diseño y complementarse con medidas legales y administrativas apropiadas. La legislación actual es bastante permisiva a este respecto, lo que permite el uso de implantes como marcapasos que no incorporen ningún mecanismo de seguridad. Las autoridades reguladoras como la FDA en los Estados Unidos o la EMA (Agencia Europea de Medicamentos) en Europa deberían promover métricas y marcos para evaluar la seguridad de los DMIs.

Estas evaluaciones deberían ser obligatorias por ley, requiriendo un nivel de seguridad adecuado para un implante antes de aprobar su uso. Además, tanto las medidas de seguridad implementadas en cada DMI como los resultados de la evaluación de su seguridad deberían hacerse públicos.

Buenas prácticas de ingeniería en los dominios de la protección y la seguridad deberían seguirse en el diseño de los DMIs. Si se detectan errores de hardware, a menudo esto implica un reemplazo del implante, con los riesgos asociados y vinculados a una cirugía. Una de las principales fuentes de fallo al tratar o monitorear a un paciente es precisamente el mal funcionamiento del dispositivo. Estos fallos se conocen como "retiradas", y se estima que afectan a aproximadamente el 2,6 % de los pacientes que llevan un implante. Además, el software que se ejecuta en el dispositivo debe soportar estrictamente las funcionalidades requeridas para realizar las tareas médicas y operativas para las que fue diseñado, y no más [66, 134, 213].

En el Capítulo 1, presentamos un estado de la cuestión sobre cuestiones de seguridad y privacidad en DMIs, discutimos los mecanismos más relevantes propuestos para abordar estos desafíos y analizamos su idoneidad, ventajas y principales inconvenientes. En el Capítulo 2, mostramos cómo el uso de señales electrocardiográficas (ECGs) altamente comprimidas (sólo 24 coeficientes de la Transformada Hadamard) es suficiente para identificar inequívocamente individuos con un alto rendimiento (precisión de clasificación del 97% y errores del sistema de identificación del orden de $10^{-2}$). En el Capítulo 3 presentamos un nuevo esquema de Autenticación Continua (AC) que, contrariamente a los trabajos previos en esta área, considera las señales ECG como flujos de datos continuos. El sistema propuesto de AC basado en señales cardíacas está diseñado para aplicaciones en tiempo real y puede ofrecer una precisión de hasta el 96%, con un rendimiento del sistema casi perfecto (estadístico kappa $>$ 80 %). En el Capítulo 4, proponemos un protocolo de verificación de la distancia para gestionar el control de acceso al DMI: ACIMD. ACIMD combina dos características, verificación de identidad (autenticación) y verificación de la proximidad (comprobación de la distancia). El mecanismo de autenticación es compatible con el estándar ISO/IEC 9798-2 y se realiza utilizando la señal ECG con todas sus ondas, lo cual es difícilmente replicable por un atacante que se encuentre distante. Hemos evaluado el rendimiento de ACIMD usando señales ECG de 199 individuos durante 24 horas, y hemos considerando tres estrategias posibles para el adversario. Los resultados muestran que se puede lograr una precisión del 87.07% en la au-

tenticación. Finalmente, en el Capítulo 5 extraemos algunas conclusiones y resumimos los trabajos publicados (es decir, revistas científicas con alto factor de impacto y conferencias internacionales prestigiosas).

**Palabras Clave:** Seguridad, Privacidad, Dispositivos Médicos Implantables (DMIs)

# Acknowledgements

This thesis marks the conclusion of my doctoral studies. An interesting and challenging three and a half years in which I have had the opportunity to learn and specialize myself in a very interesting research area.

I would like to thank my supervisors, Dr. Pedro Peris and Dr. Juan E. Tapiador for their support, helpful advices, friendly inputs and for a large extent of autonomy in this thesis.

I would also like to thank Arturo Ribagorda, the head of our lab, for allowing me to carry out the thesis in his research group and for trusting in me on different occasions.

Finally, my deepest gratitude to my family. To my parents, who have seen and enhance my potential from a very young age. They convinced the academic director of my school that I was able to take computer classes in MS-DOS when I was 6, with students ten years older than me. And I was able to do it, and do it well, impressing the teacher who came from IBM. It was at that moment when I started my link with computer science. Thanks to my father, who used to spent most of the day reading and studying with an insatiable curiosity. We have more books than anything else in the home. It's very probable that I have inherited, or at least has been inspired by his deep passion for knowledge. He has always been an inspiration. And thanks to my mother, who has always pushed us -my sister and me- to do what we really like, trying to get us away from social pressure, present in some moments. I am working on the things that I really love thanks to them. Thanks also to my aunt, my second mother, who has been an irreplaceable support, in many senses, in very difficult personal times.

# Contents

# List of Figures

# List of Tables

# 1

# Security and Privacy Issues in Implantable Medical Devices

## 1.1 Implantable Medical Devices

An IMD is often defined as an electronic device that is permanently or semi-permanently implanted on a patient with the purpose of treating a medical condition, improving the functioning of some body part, or providing the user with a capability that he did not possess before [86]. These devices are often implanted around 2-3 cm under the patient's skin and connected to the organ that needs treatment or monitoring. Cardiac implants (see Fig. 1.1) are possibly the most widely known example of IMDs, but many others are increasingly being used to deal with different medical conditions more efficiently than by traditional methods. The most common types include:

**Cardiac Implanted Devices** These include devices such as Implantable Cardioverter Defibrillators (ICD) and Pacemakers. They are designed to treat cardiac conditions by monitoring the heart's electrical activity and applying electrical impulses of suitable intensity and location in order to make the heart pump at the desired speed [230]. New models are equipped with pressure sensors capable of actively monitoring changes that could lead to a heart failure. This allows to alert the patient or the medical personnel if a pressure increment in the ventricle is detected, as this represents a hazard condition for the patient.

Cardiac implants may also be equipped with accelerometers to measure the patient's physical activity level. This can be set as an input parameter to the IMD controller, allowing to adjust the cardiac stimulation frequency to the one that best suits each moment [209].

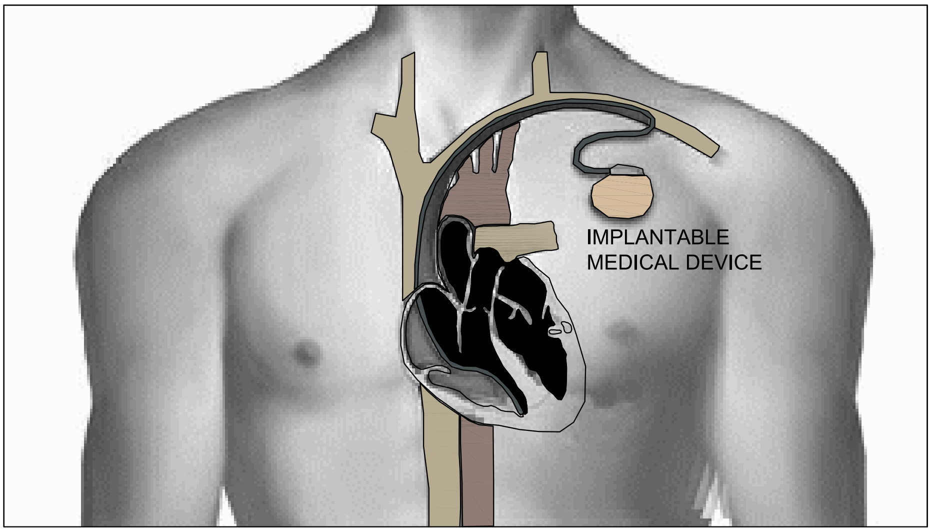

**Figure 1.1:** Implantable Medical Device: Pacemaker

**Neurostimulators** These devices transmit low-amplitude electrical signals through one or more electrodes placed in different locations of the brain. These electrodes are implanted in very specific areas depending on the patient's condition. The process is known as Deep Brain Stimulation (DBS) and allows to treat a variety of pathologies such as Parkinson, dystonia, epilepsy, or even depression that, in some cases, are resistant to medication after several years of treatment [137].

**Drug Delivery Systems (DDS)** A DDS consists of a pump and a catheter that are surgically implanted under the skin. Their function is to supply medication in a controlled, localized, and prolonged way. Since the medication goes directly to the target area, an infusion pump provides a considerable degree of control, which allows to use a lower dose than that required with oral medication. For instance, this type of implants have been successfully used to mitigate pain in cases of cancer where traditional medication does not have good results [127].

**Biosensors** The implant consists of a sensor or a set of sensors placed inside the human body to monitor any part of it. They are capable of measuring certain physiological parameters and use such measures to make decisions. In this sort of implants there exists a special device that acts as a control node, communicating with the sensors and with other external entities (e.g., a programmer). The set of sensors and the control node are often regarded as a wireless biosensor network [37, 54, 222].

### 1.1.1 The New Generation of IMDs with Telemetry

Healthcare systems incorporating numerous communication and networking functions have proliferated over the last years. This has made possible to develop medical sensor networks that, for instance, can monitor patients in their own homes [189, 168, 233, 205]. Doctors, caregivers, or even the patient himself can thus conduct a continuous and more flexible control of his state, as well as access medical data remotely, communicate during an emergency, and even command various household appliances. This also promotes the autonomy of patients who, in many situations, are elderly people or individuals with reduced mobility.

Similar communication and networking capabilities are increasingly being embedded into IMDs. Equipped with a radio transmitter, the IMD can communicate with an external device—generally known as “Programmer” or “Reader”—and send it physiological data such as electrocardiogram (ECG) signals in the case of pacemakers and ICDs, that the doctor can use to track the patient’s pathology. Apart from querying sensed data, the Programmer can also command the IMD to adjust or disable therapies, perform software updates, etc.

Augmenting IMDs with wireless communication and networking capabilities has significant advantages, including:

- It allows to constantly monitor the patient’s physiological parameters and other symptomatology captured by the device, which reduces the time needed to regularly tracking medical conditions and, furthermore, causes less disruptions in the patient’s daily activities.
- Enhanced supervision and management of the IMD operation, which allows to address any problem that might arise and apply adequate correction measures in a shorter time.
- The two previous items also imply a reduction in the overall costs involved in tracking the patient’s condition and managing the operation of the IMD.
- In the case of future IntraBody Networks (IBN) [25, 26, 201], computation and analysis tasks could be shared among different networked devices, which will contribute to the development of new diagnostic techniques.

In their current generation, not all types of IMDs support access to all their available functions through the wireless communication channel. The vast majority of IMDs can be reprogrammed remotely, which allows the doctor to modify therapies as required. The reverse link (i.e., from the IMD to the Programmer) is not present in all of them. For instance, while pacemakers and ICDs can communicate in both directions, current neurostimulators can only receive reprogramming commands but they do not provide any information (e.g., sensed data) back to the Programmer. This fact has caused that most research works on advanced computational and networking issues related to IMDs, including the security and privacy problems addressed in this dissertation, are commonly focused on cardiac implants. Nevertheless, new IMD designs are computationally more complex and are increasingly basing part of their functionality in the ability to communicate externally to perform diagnostic and therapy tasks.

The main standards regulating telemetry for medical devices are:

- Many non-implantable medical devices are compliant with the Wireless Medical Telemetry Services (WMTS) specification, which sets three operating frequency bands: 608-614 MHz, 1395-1400 MHz, and 1427-1432 MHz [35, 85]. This is a U.S. standard defined by the Federal Communications Commission (FCC) in 2000 that is not internationally agreed, hence that its use is often restricted to the U.S. only.
- IMDs operate under the Medical Implant Communication System (MICS) specification, which operates in the 402-405 MHz band. MICS is a low-power (25 microwatt), unlicensed mobile radio service that facilitates data communications between the IMD and an external programmer. The communication range is about 2 m and the bandwidth is very low when compared with wireless communication technologies such as bluetooth or WiFi. The radio signals can go through and be transmitted within the human body due to its conductive characteristics. The purpose of these communications can be accessing the measures taken by the implant or reconfiguring it to, for example, adjust the treatment. MICS compliant IMDs have proliferated in the last years, including pacemakers, ICDs, neurostimulators, hearing aids, and DDSs [32, 56, 193].
- Similarly to the WMTS specification, the Medical Device Radiocommunications Service (MedRadio) defines communication services for both implanted and wearable medical devices. The specification, which was approved by the FCC in 2009 [164], extends the MICS spectrum 1 MHz in both sides, covering a frequency band from 401 to 406 MHz. The use of these frequencies in IMDs is well justified [40]: 1) At those frequencies, radio signals can easily propagate within the human body; and 2) The 401-406 MHz band is compatible with international regulations and does not interfere with

other radio operations in the same band.

Incorporating a wireless communication capability into an implant involves some special requirements that affect their design. One of the most important is that the radio frequency module must consume very little power (e.g., 10 mW and up to 100 mW for a glucose and ECG monitor, respectively [87]) in order to save the implant battery life. Additional design factors include the required communication range (typically from 1 up to 5 meters [224]), the data transfer rate (e.g. 0.1 bps and up to 10 Kbps for a glucose and ECG monitor, respectively [87]), the environmental conditions in which the IMD will operate, and its size and cost [96, 97, 236].

Recently, the FDA has published guidelines for the industry on the design, testing, and use of wireless medical devices [57]. As stated, the security of wireless signals and data is an important issue in order to protect access to patient's data and hospital networks, and to prevent unauthorized communications with medical devices like IMDs or Programmers. Wireless medical devices must use cryptographic techniques (ive., encryption, authentication, secure key storage) to protect communications and accesses. The necessary security level is determined by the sort of threats, and their probability, to which the device is exposed, as well as the operating environment and the consequences/damages on the patient in case of a security incident. For the design of secure solutions, the FDA suggests that wireless medical devices include security measures to protect communications and accesses but also include software protections. Nowadays, the FDA is currently working on the design of recommendations for the management of cybersecurity in medical devices [58]. Apart from the FDA, other organizations are contributing to the elaboration of standards (e.g., X.1120 and X.1139), including tele-biometrics, mobile secure transmissions, secure transmission of personal health information, etc. [107].

## 1.2 Security Assumptions

In this section we first present the system model and then describe the usage scenarios. After that, the threat model is explained and finally the different types of adversaries are introduced.

### 1.2.1 System Model and Usage Scenarios

Fig. 1.2 presents the main entities involved in the system and shows the possible communication interactions (linked to the usage scenarios) be-

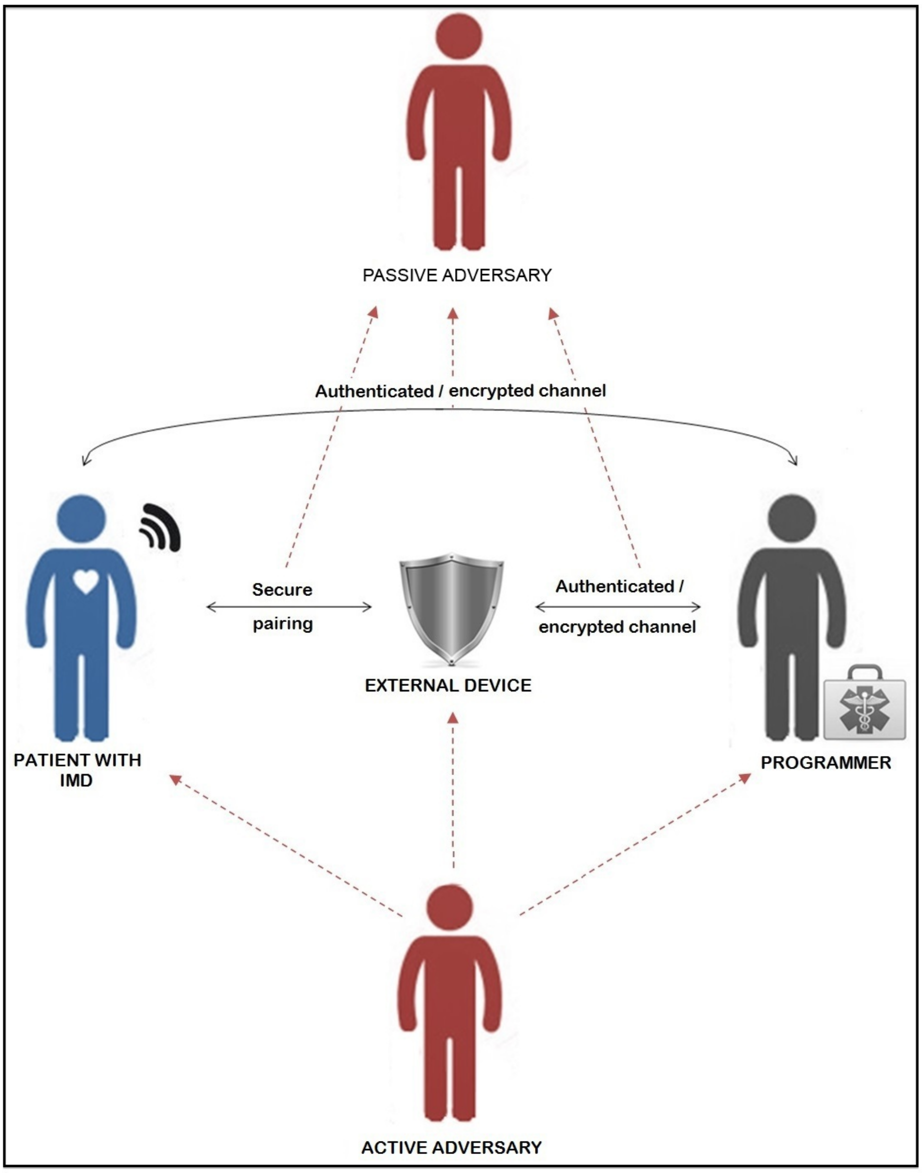

**Figure 1.2:** Typical usage scenario for IMDs

tween these devices. The IMD will communicate with a Programmer, which will be any entity/device authorized to interact with the implant (e.g., medical personnel). In normal operation (i.e., while the implant has not detected an emergency situation [193]), the Programmer has to initiate the communication with the IMD as stated by the FCC regulations. Since the radio channel is a shared communication medium, the programmer will listen to the channel until it detects that is not busy to establish the communication. The goal of this communication is either requesting data (e.g., ECG signals or insulin levels) or sending commands (e.g., treatment modifications). In the case of secure solutions, the IMD and the Programmer are authenticated and sensitive data is passed encrypted on the channel.

Apart from the direct communication between the IMD and the Programmer, some authors have introduced the idea of using an external device (e.g, cloaker [51], shield [73], IMDGuard [234], etc.), which acts as a proxy. In this case, rather than establishing a direct connection with the Programmer, the IMD can delegate this task to an external device that authenticates the Programmer—initially there is a secure pairing between the IMD and the external device. Once the Programmer is authenticated (normal mode operation), this can communicate with the IMD using an encrypted channel via the external device. In emergency mode, the IMD has to answer even if the authentication fails—and, in some cases, the medical personnel must be able to disable the device easily.

As the patient will generally move about different locations and may visit several doctors and hospitals, IMDs will not always communicate with the same, previously known device. Furthermore, the entities authorized to communicate with the implant can vary [82]. Potential attackers must be also considered, as not all signals received by the IMD will actually come from an authorized Programmer and, in many cases, their purpose could be malicious. Under these conditions, guaranteeing the security and privacy of the IMD and its data is essential to protect the safety of the patient.

As described above, an IMD must operate under two different modes: normal and emergency. One major objective is to find a sensible trade-off between these two possible situations:

A. **Security in normal operation mode**. The patient controls what entities can interact with his IMD. In this case, it is necessary to implement both a strong access control mechanism and cryptographic protocols in the communication link to thwart malicious and unauthorized accesses. The IMD must ignore indiscriminate data requests or device reprogramming commands. Ideally, the implant should be undetectable to unauthorized parties. Security mechanisms might be similar to those used in constrained devices like RFID tags or smart sensors (e.g., lightweight hash functions [13, 16, 80] or tiny block ciphers [104, 122, 123]).

B. **Security in emergency mode**. As important as offering strong access control, secure communications, and even undetectability, is the ability of being accessible under an emergency condition. Consider a patient who enters an emergency room in a hospital different to the one he often visits. To further complicate matters, assume that the patient is visiting a foreign county. Even under these circumstances,

> the healthcare staff must be able to communicate with the implant, determine its type (e.g., model and brand), extract physiological data or information about the treatment, and even update its configuration if required. Even in a secure scheme, under an emergency situation such as an urgent surgery of a patient who holds an ICD, in which it is mandatory to deactivate the implant, the IMD should always respond before deactivation.

To understand the importance of emergency conditions, it would be useful to know the frequency of occurrence of these events. Unfortunately, to the best of our knowledge there are no public reports about emergencies involving patients that hold an IMD. Nevertheless, some statistics about pacemakers, which are one the most popular IMDs, may help to shed some light on this matter. For instance, lead complications are one of the principal causes of re-intervention in patients with heart diseases. In a recent retrospective study, Walker et al. [225] reported 1.4 events per 100 patient-years of follow-up for lead-related complications, including vein thrombosis, acute perforation, and dislodgement. This figure doubles if the population under study are children [159]. As for the pocket-related complications (e.g., infection, erosion, or migration of the pacemaker), which is the other main cause of complications for pacemakers [79], the values are slightly higher: 1.9 events per 100 patient-years. Furthermore, the probability of re-intervention increases with every consecutive replacement [18]. We acknowledge that re-intervention due to lead or pocket complications is, strictly speaking, not an emergency condition, since in most cases the surgery is planned. Nevertheless, both situations have in common the need to properly address any security measures deployed in the IMD. It is the job of manufacturers, engineers, and physicians to evaluate the frequency and impact of those events in order to develop a rigorous risk model.

A straightforward solution for emergency conditions that provides the necessary safety to the patient is to force the IMD to disregard authentication and authorization mechanisms and process all incoming commands. Any requester thus becomes an authorized user, possibly with full privileges. This would not be possible if security protocols and strong access control mechanisms are not deactivated, which in turn leaves the implant fully exposed to attackers. Unfortunately, telling apart normal from emergency scenarios is far from trivial for the IMD, and nowadays the best way to provide an adequate security for IMDs is still an open problem. Security tensions between these two conflicting goals are thus created, hence the importance of finding solutions that balance the security requirements to provide security in normal mode while guaranteeing safety during emergencies [185]. Several works (see, e.g., [45, 26, 96]) have proposed schemes in which the IMD can only be accessed by authorized entities and remain invisible for the remaining ones. The prevailing philosophy in most works is that in case of doubt about the patient's safety, security mechanisms should be relaxed and access must be granted. We will discuss in detail the most relevant proposals in this field later in Section 1.4.

### 1.2.2 Threat Modeling

Security threats against the IMD can be categorized using the STRIDE methodology. The acronym stands for six general categories of attacks: Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, and Elevation of privilege. Table 1.1 relates each category with the security service attacked in each case and provides some examples. Generally it is assumed the following set of relations (listing the "security service" versus the linked threat): authentication – spoofing, integrity – tampering, non-repudiation – repudiation, confidentiality – information disclosure, availability – denial of service, authorization – elevation of privileges. Apart from these one-to-one connections, it should be noted that some threats may address various services simultaneously, or that a single attack can be decomposed into individual threats.

The six security services addressed above have their usual meanings, although focused on the IMD domain:

**Authentication** The identity of parties must be correctly established before performing any other operation. Within the domain of implantable medical devices, any device in the system (IMDs, Programmer or External device) can be impersonated. For instance, if the identity of the Programmer is supplanted it might be the starting point for an elevation of privileges attack.

**Integrity** Data, either stored in the device or being communicated through the wireless link, can only be modified by authorized parties. If there is not integrity checking mechanism on the IMD, data could be altered during the transmission over the insecure radio channel. Furthermore the IMD could accept malicious inputs, which could be employed to run a code injection attack [182]. On the other hand the lack of integrity checking would facilitate that the manipulation of the data stored on the IMD memory might be not detected –or be detected in a distant future.

Table 1.1: STRIDE categories and examples in the IMD domain.

| Security Service | Threats | |
|---|---|---|
| Authentication | Impersonate the Programmer<br>Impersonate the IMD<br>Impersonate the external device | Spoofing |
| Integrity | Patient data tampering<br>Malicious inputs<br>Modify communications | Tampering |
| Non-repudiation | Delete access logs<br>Repeated access attempts | Repudiation |
| Confidentiality | Disclose medical information<br>Determine the type of IMD<br>Disclose the existence of the IMD<br>Track the IMD | Information Disclosure |
| Availability | Drain the battery of the IMD<br>Interfere with the IMD communication capabilities<br>Flood the IMD with data | Denial of Service |
| Authorization | Reprogram the IMD<br>Update the therapy of the patient<br>Switch-off the IMD | Elevation of privileges |

**Non-repudiation** Operations performed by/on the IMD are kept securely in an access log. The attacker could focus on delete these inputs in order to cover her traces. On the other hand no all IMDs are equipped with a log system. If this were the case the adversary could repeatability try to gain access to the IMD without leaving any trail. Even if a log system is present the events would be logged but no alarm would be triggered to alert the IMD holder in case of a malicious event.

**Confidentiality** Data, either stored in the device or being communicated through the wireless link, can only be read by authorized parties. In particular, IMDs and the Programmer communicate through the radio channel (401-406 MHz) and these communications are exposed to eavesdroppers. If communications are not encrypted, an adversary could disclose private information such as the IMD model or even medical information of the patient. This would compromise the privacy (data) of the implant holder. Even if communications are encrypted, an attacker could detect the presence of the implant or, even worse, track the movements of its holder. In this case the privacy location would be put at risk.

**Availability** The services offered by the IMD should be available to authorized parties at all times. Availability is crucial for IMDs since these devices are devoted to treat medical conditions of their holders. Unfortunately, an IMD could be rendered inaccessible through the blockage of the radio channel (active jamming). Alternatively the device might be overloaded by flooding the IMD with network traffic over the radio channel. This could be used to block the access to the device or to drain its battery. If the battery runs out of power, the device would become permanently inaccessible and the patient's health could be at risk.

**Authorization** An operation must be executed only if the requester has sufficient privileges to order it. For instance, therapy parameters (e.g., voltage, current, thresholds, operation mode, etc.), cannot be updated by the patient and only doctors should be able to modify these. In this regard, re-programming the IMD must be done under the joint supervision of the doctor and a technician (typically from the manufacturing company of the IMD). On the other hand, the IMD must be kept running at all times and only be switched off under special circumstances that may threat the patient's life (e.g., cardiac surgery with electrocautery devices). In the case of pacemakers, a magnetic field has to be applied near the device (over the patient chest) and this procedure must be authorized by the cardiologist.

#### 1.2.2.1 Types of Attackers

At high level, attackers can be grouped into two main categories: *active* and *passive* (see Fig. 1.3):

**Passive Eavesdropper** A passive attack can only listen to the channel and, therefore, getting access to the messages exchanged between the IMD and the Programmer. Assuming an insecure radio channel, a passive attacker is a direct threat to confidentiality and may threaten authentication. By just reading messages a passive attacker may determine whether a person carries an implant or not; find out

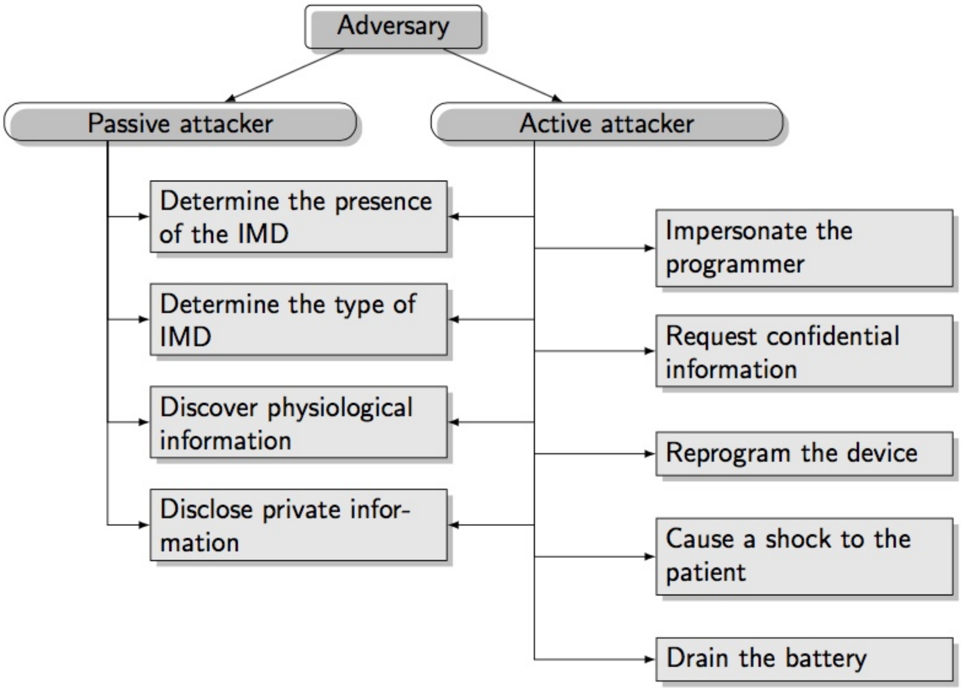
Adversary
Passive attacker
Active attacker
Determine the presence of the IMD
Determine the type of IMD
Discover physiological information
Disclose private information
Impersonate the programmer
Request confidential information
Reprogram the device
Cause a shock to the patient
Drain the battery

**Figure 1.3:** Passive vs active adversaries.

what type of implant and other data such as its model, serial number, etc.; capture telemetry data and disclose private information about the patient, such as the ID of his health records, name, age, conditions, etc. In all cases, the overall result is a serious compromise of the patient's privacy.

**Active Adversary** In this case, the adversary is not only capable of capturing messages exchanged over the radio channel, but also to send commands to the IMD, modify messages in transit before they reach the IMD or the Programmer, or just block them so that they never arrive. Attacks may involve a sequence of interceptions, interruptions, modifications, and generation of messages. The goals pursued by an active attacker are diverse. For example, he could indiscriminately request information from the IMD with the purpose of draining its battery. He could also attempt to modify the configuration of the device, disable therapies, or even induce a shock state to the patient [83].

It must be noted that it is not essential for the attacker (active or passive) to be physically close to the patient to conduct the attack [62]. Depending on the specific communication technology used for the radio link, the IMD could be reachable from a few meters (typically 1-2 meters for MICS and WMTS [166, 224] or even up to 10 meters in case of using advanced communication techniques [88]). Furthermore, communication devices can be acquired very easily nowadays; e.g., certain smartphones can perform this task.

In summary, the technical means needed to carry out most attacks against IMDs are cheap and easy to acquire and use. As a consequence, passive attackers can easily eavesdrop sensitive information about an implant holder without much difficulty. Even if the attacker is not someone who attempts to threaten the patient's safety, the data stored on it might be very valuable for many individuals and organizations.

It is worth mentioning that the attacker could be the patient himself in a deliberate or involuntary attempt to sabotage his own implant. An example of this was reported in [173], where it is described how a patient who sent unauthorized commands to his insulin pump in order to gain unsupervised use of it ended up with a medical condition as his manipulation resulted in the ingestion of a very high dose of medication.

Finally, apart from general system and channel vulnerabilities, attackers can manipulate a number of IMD-specific features to achieve their goals [6, 86]:

- *Manipulation of the distance.* Proximity refers to the distance between the attacker and the IMD. Many current proposals have some form of distance-based access control, allowing access to the IMD only if the Programmer is in short range. The rationale here is to force the attacker to be physically very close to the patient to conduct the attack. In practice, however, the attacker may use a compromised device in the proximity of the patient to launch the attack, including those used in medical facilities.
- *Manipulation of the IMD functions.* IMDs are programmed to perform various activities such as sensing biomedical parameters in the body area where they are implanted, treat a medical condition (actuating), processing gathered data, and communicate with other devices, either external or those in the IBN [240]. These functions can be misused by an attacker, for example by inducing an incorrect sensing to trigger a particular response in the implant.
- *Manipulation of the patient's status.* As we will discuss later, the patient's status plays a key role in the design of many countermeasures. For instance, an implanted biosensor can trigger an alarm if certain parameters fall out of the safety range. In some cases, such an alarm puts the IMD in emergency mode and automatically disables access control mechanisms.

## 1.3 Limitations and Trade-offs

In this Section, we first introduce a number of technological limitations of current IMDs that restrict the sort of security mechanisms that can be implemented on them. We next describe various trade-offs that arise when designing security measures for IMDs and that originate as a consequence of the IMD's computational limitations, the criticality of some of its functions, and the need to support an emergency mode of operation.

### 1.3.1 Limitations

IMDs have restricted capabilities in three separate dimensions: energy, storage, and computing power. All three of them have security implications, either because they can be misused or because they limit the security mechanisms that can be afforded. We next discuss them in more detail.

**Energy** IMDs are powered by an integrated battery that supplies energy to all functions incorporated in the device (i.e., monitoring, treatment, communication, etc.). Once the IMD is implanted, the battery can last from 8 years in the case of neurostimulators [151] up to 10 years in the case of pacemakers [140]. Battery usage has a direct impact over the implant lifetime. Once exhausted, it has to be replaced, which requires a surgical procedure with its associated risks. Some designs support batteries that can be charged wirelessly using magnetic fields, but organs close to the implant could be damaged. Some recent advances in this area can be found in [112, 206, 235].

**Storage** Storage is quite limited in current IMDs. The memory available in the device is used to store historical data from different events and episodes that arise related with the patient's pathology. For instance, pacemakers and ICDs store ECG signals that occurred when the device decided to apply stimulation. The RAM memory of these device varies from 2 KB to 36 KB for the former, and from 128 KB to 1024 KB for the latter. In the case of ICDs, around 75% of this memory is devoted to store ECG signals [105]. Devices with low sensing rate like a Biostator Glucose Controller demand 8 KB for data storage [212]. One consequence of incorporating a reduced memory on-chip is that security mechanisms have to consume as little memory as possible in order to save it for the potential storage requirements required by the medical functions of the device. One may wonder about the possibility of increasing the amount of RAM memory in

IMDs, since this sort of memory is not expensive nowadays. There seem to be two main reasons for keeping limitations on the memory size. On the one hand, an increment on the amount of memory constitutes an increase in the size of the implant. This is a critical feature since IMDs are often located in or over the body of the patient and this parameter (device area) should be kept at a minimum. On the other hand, even if the device size is not an issue, increasing the amount of memory could impact the battery lifetime. Access operations (i.e., reading, writing and erasing) are considered demanding in terms of power consumption [167], so performing them over a large amount of data would decrease the battery life and even exhaust it in a short time.

**Computing and Communication** Both computing and communication capabilities are extremely limited in IMDs due to power restrictions. Communication is the most energetically expensive task for the IMD. Hence, if communications are minimized, the battery life can be extended [192, 197]. As for computation, these are generally supported by a tiny microcontroller. For instance, the micro of a neurostimulator consumes an area of around 5mm$^2$, which is around forty times smaller than the area used for a general purpose microcontroller [101]. In general, the whole chip of the implant occupies an area of around several hundreds squared millimetres.

### 1.3.2 Tensions and Trade-offs

As described in Section 1.2, and IMD can work in two operation modes: normal and emergency. Mechanisms designed to preserve security and privacy properties in both modes must consider various tensions:

**Security vs Safety** Nowadays in a real scenario it is common to assume that all the actors, both the legitimate (new generation of IMDs, external devices and programmers) and the illegitimate ones (active and passive adversaries) will have network connectivity. This should lead to the inclusion of solid security solutions to prevent security incidents. In particular, in normal mode the IMD is vulnerable to a variety of attacks. Attackers could be physically situated at a long distance from the IMD and use its wireless communication capabilities—perhaps relying on a nearby proxy device—to receive data requests and perform update operations. Any proposed solution must guarantee basic security and privacy properties in this case.

Nevertheless, during an emergency the medical personnel must be able to access the implant rapidly and without restrictions. Thus, while the use of strong security measures could provide a high level of protection, it can also put at risk the patient's safety during an emergency situation. The trade-off between safety and security is one of the most critical aspects in the design of security mechanisms for IMDs.

**Battery Lifetime vs IMD Capabilities** As discussed above, IMDs have severe restrictions in terms of energy consumption since extending the battery lifetime is an essential requirement. In turn, this also restricts the amount of computations and communications involved in security functions. This motivates the design of new security and privacy mechanisms that are not very demanding in terms of computation, communications, and storage. An interesting fact in this regard was pointed out in [164]: power consumption increases drastically if the data transfer rate increases. Thus, although it may seem counterintuitive, it is preferable to rely on long transmissions at very low bit rate than on short data exchanges at high speed.

Several solutions have addressed the problem of saving or recharging the battery of IMDs to postpone as much as possible its replacement. For example, in [232] Warwick et al. present an innovative solution to provide higher intelligence to neurostimulators. The idea is to provide the implant with the capability to predict tremor conditions in Parkinsons' activity, so that only in that precise moment an stimulation on the sub-thalamus is triggered. Once the tremor has diminished, the implant stops the stimulation. Intelligent solutions like these could prolong the lifetime of the battery.

Other approaches have suggested techniques to recharge the battery wirelessly. In [223], Arx and Najafi propose to provide the implant with receiver (planar spiral) coils and accompanying circuitry that are capable of receiving transmitted power from a few centimetres away. Another example can be found in [209], where an inductor with a parallel chip capacitor is proposed. In this system, the inductor radiates energy by coupling a signal at the resonant frequency (300 Mhz in this proposal). These systems would allow the IMDs to work without any battery, which would be highly desirable since the battery replacement procedure would be avoided [241].

Using a different approach, Wang and Song proposed to transform mechanical energy obtained from the movement of the patient's muscles into electrical energy [227]. Using this technology, the IMD could be automatically and continuously recharged by the patient's physical activity.

Unfortunately, neither these solutions nor others recently proposed (e.g., [112, 206, 235]) can be nowadays found implemented in commercial IMDs. Therefore, any security measure for implants must take into account existing energy restrictions and potential impacts on the battery life. Furthermore, as there are attacks that pursue to waste the battery of the IMD, security functions should not make this easier (e.g., by allowing the attacker to drain the battery by misusing security mechanisms).

**Answering Time** If the interaction with the implant takes too much time because of the overhead imposed by security controls, the patient's safety could be put at risk. Such controls should be analyzed to guarantee that their worst-case latency is within a reasonable range.

In summary, tensions between safety (i.e., guaranteeing access in critical conditions) and security (allowing access only to authorized entities), coupled with the restrictions present in current IMD platforms, introduce unique challenges in the development of adequate security mechanisms for IMDs. Adapting solutions proposed for other similar environments (e.g., wireless sensor networks) is not straightforward, since questions such as how security mechanisms should behave in emergency mode—and, most importantly, guaranteeing that the existence of this mode is not abused by an attacker—are still open problems.

## 1.4 Protection Measures

In this section, we discuss different security mechanisms that have been proposed to thwart security threats in IMDs. Many of these proposals explicitly address the trade-offs and tensions previously discussed, while others simply focus on counteracting specific attacks. The majority are preventive and attempt to stop attacks from happening in the first place, although detection and correction mechanisms have been also suggested.

Ideally, the inclusion of security measures should not require any modification of the IMD, as this would imply its replacement and, therefore, a surgical procedure. The alternative would be implementing security functions in external devices or independent modules of the IMD chip. Under this approach, the software running on the implant would be exclusively used to treat the patient's medical condition.

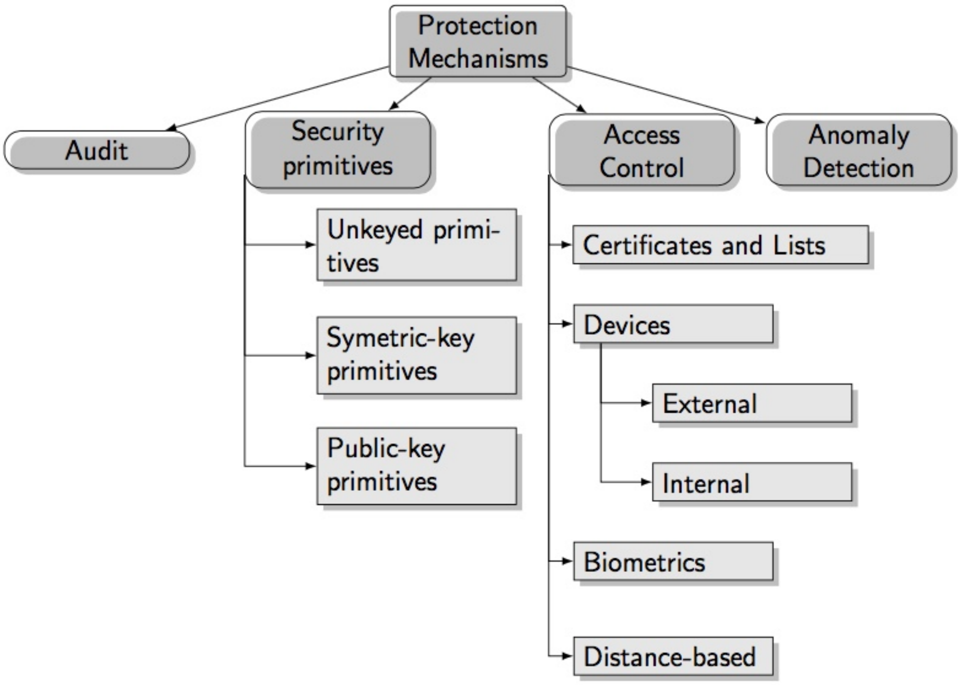

**Figure 1.4:** Protection mechanisms proposed for IMDs.

As discussed above, a major problem with most security measures is that they could put at risk the patient's safety in emergency situations if they cannot be easily disabled. The use of some form of "backdoor" to bypass security could be a straightforward solution, though it is too easily manipulable by an attacker.

Fig. 1.4 provides a classification of the security mechanisms that will be discussed throughout this section.

### 1.4.1 No Security

Many IMDs, particularly the older generations without wireless communication capabilities, have no security mechanisms at all [148, 150]. This is unacceptable for the newest generations of IMDs in which the presence of communication capabilities may jeopardize the patient's safety.

### 1.4.2 Auditing

One of the simplest security mechanisms consists of constantly registering all accesses—authorized or not—together with the patient's status. This is a measure amied at facilitating the detection of non-permitted actions and constitutes a valuable source of evidences to take subsequent actions.

Therefore, auditing helps to combat threats against non-repudiation. Unlikely, it does not prevent the occurrence of attacks, but may act as a deterrent element if appropriately implemented, i.e., if it is not possible for an attacker to compromise the audit log and if it facilitates attribution of the attack. As a consequence of this, this sort of solutions should be be complemented with appropriate mechanisms to detect and block such attacks, as well as measures to prevent them from happening in the first place (e.g., cryptographic or access control solutions).

The main problem that auditing proposals face is the limited amount of memory available in IMDs. For instance, the whole memory of an ICD is less than 1 MB and around 75% of this memory is used for medical functions. In that a case, only a few hundreds kilobytes could be used for logging events, which is extermely restricted. An additional memory could be added to the chip, but this would increase the size of the IMD, which is not recommendable.

To avoid increasing the memory of IMDs, the logging task can rely on an external device without memory and computation limitations. One example in the context of RFID systems is “RFID Guardian” [181], which collects and analyzes evidences of all events that occur in a predetermined range. A similar approach, called MedMon, has been recently proposed for IMDs and e-Health applications [238]. The authors propose the use of an external device that works as a security monitor snooping and analyzing all communications to and from the IMD. The events are locally stored in the external device and an alarm could be raised to alert the patient. A more drastic solution can include blocking the communication channel if a dangerous communication is detected.

### 1.4.3 Cryptographic Measures

Cryptography-based security solutions strongly depend on cryptographic primitives, which can be categorized in three main groups [153], as shown in Figure 1.4. Unkeyed primitives, such as hash functions or one-way permutations, are cryptographic tools that do not use any key. Within the keyed cryptographic tools we can distinguish between symmetric-key and public-key primitives. In symmetric-key primitives a secret key is shared between the trusted entities. The type of primitives in this category is varied including symmetric key ciphers (block and stream ciphers), message authentication codes (MACs), pseudorandom sequences and identification primitives. On the other hand, public-key ciphers and signatures are two examples of asymmetric-key primitives. In this type of algorithms two keys are used, one of them is public and the other one must be kept secret.

In the context of IMDs, cryptographic measures are effective mechanisms to protect the wireless communication channel and the records stored in the device against tampering and information disclosure. Additionally, cryptographic protocols also provide a means to control and manage accesses to the IMD, thus providing protection against spoofing and, in some cases, elevation of privilege attacks. Both symmetric [83, 96] and public-key [55, 210] schemes have been proposed for these applications, although the latter are considerably more expensive in terms of communication, computation, and power consumption. Protocols based on public-key cryptographic schemes often exchange a high number of messages, which makes them quite energy demanding since sending and receiving messages consume power. Furthermore, public-key ciphers result in complex circuits that consume excessive resources (hardware and memory) and are inefficient in terms of power consumption [67, 132]. Due to the resource limitations discussed above for the current generation of implants, solutions based on symmetric-key approaches are the preferred option. Standardized protocols like the one proposed in ISO/IEC 9798 rely on the use of symmetric primitives (i.e., symmetric encryption or keyed hash function) and the encrypted tokens include random numbers (a PRNG is often used for its generation) to guarantee freshness between sessions [103].

Symmetric cryptographic schemes suffer from the key distribution problem. In general, the IMD and other authorized devices such as the programmer need to share a key (or a set of keys) that is used to generate authentication tokens for gaining access to the IMD, and to encrypt communications. The suitability of a particular key distribution scheme depends on the type of IMD, the expected interactions with other parties, and other assumptions about the operational environment. For example, if the programmer and the IMD will have a lasting relationship, a pre-set key can be used. This solution could be valid when the programmer is always a device belonging to the patient or the physician. In these cases, a first approach consists of pre-loading a factor key on the authorized devices. This factor might be renegotiated between the legitimate parties during the first communication session to update the key. We emphasize here that is crucial to protect these keys and guarantee that only authorized entities (i.e., the patient and healthcare staff) have access to the them [199]. Such keys will be used to build various cryptographic tokens (e.g., an authenticated token or an encrypted message) used in the transactions between legitimate entities in the system.

Other solutions suggest that the cryptographic keys used by the IMD can be stored in an external wearable device such as a smart bracelet. Externalizing the key storage incurs a significant risk, as the loss of such a device (e.g., if the patient losses the bracelet or it gets damaged) would render the IMD inaccessible and/or will facilitate access to unauthorized users [64]. Some authors propose to print the key into the patient's skin using ultraviolet pigmentation (i.e., invisible tattoos) that can be read by medical personnel in case of emergency [194]. Note, however, that the keys might be read by an attacker who has physical access to the patient—its presence may be detected due to its proximity.

In the case of sporadic communications with authorized devices that nonetheless do not know the access key, a key agreement protocol must be supported (e.g., RSA-based [210] or using physiological signals [221]).

Providing a confidential channel between the IMD and the programmer is another major goal when using cryptographic solutions. Some approaches suggest to exploit the limited coverage of the physical layer during the initialization phase [136]. Most proposals are based on symmetric ciphers [83], and some of them incorporate a key updating mechanism (e.g, a hash-chain based updating scheme [89]). Recently, Kaadan and Refai have proposed in [116] a novel cryptographic system with claimed military-grade security level that combines a one time pad cipher with a novel key distribution and authentication scheme. Other approaches, like the one discussed in [96], focus on hardware efficiency and propose the use of lightweight ciphers that offer tiny footprints with low power consumption. Recently, a new IMD architecture, evaluated on an artificial pancreas implant, has been proposed. In this case, the implant includes two separates cores (illness treatment and security tasks), and the overhead for the security module in terms of hardware and energy consumption is minimal [212].

The use of standard cryptographic solutions to provide security services in IMDs has been criticized, both for usability reasons and for the lack of rigour in the analysis of many proposals [185]. The main drawbacks that would entail the exclusive use of these solutions are [74]:

**Inalterability** Incorporating cryptographic mechanisms in the device implies that current implants must be re-designed and replaced. This will force patients to undergo a surgery procedure only to get a more secure device, as the treatment functions do not present any problem.

**Patient's safety** The use of cryptographic measures embedded in the device introduces some challenges for emergency situations in which the communication with the IMD is necessary even for unauthorized parties (i.e., programmers who does not know the access key). This problem is not present in solutions based on external devices such as those discussed later in Section 1.4.4.2.

**Maintenance** As security measures are implemented in the device, there is an increment in the amount of software embedded in the implant, which also implies a higher likelihood of errors. Many authors support the idea of restricting as much as possible the software running on the device, keeping just those functionalities needed to treat the medical condition for what it was designed.

## 1.4.4 Access Control

Access control mechanisms prevent unauthorized and inappropriate uses of the IMD functions. Prior to proceed with a particular action (e.g., access, reading, reprogramming, etc.), the privileges of the requester are evaluated with the aim of assessing whether it is authorized to execute that particular action or not. In particular, permitted and forbidden operations are governed through access control policies that establish who can do what, possibly depending on the context in which the access request takes place. Note that access control is fully compatible with other security measures such as cryptographic protocols to protect the communication channel. Furthermore, access control generally requires previous authentication, as decisions on whether an operation is permitted or not are made on the basis on the identity of the requester, who must be previously established.

We next describe a number of access control models suggested for IMDs and discuss their main advantages and limitations.

### 1.4.4.1 Certificates and Lists based solutions

In [64], the authors present two classical authentication mechanisms adapted for IMDs. One is based on Access Control Lists (ACLs)—an implementation of discretionary access control models based on the access matrix—, while the second relies on a Public Key Infrastructure (PKI). The ACL defines which operations an authenticated reader is authorized to execute. Such permissions are permanent once the ACL is programmed. Thus, although it can be reprogrammed in the future, it is intended for providing permanent access to certain readers. Contrarily, in PKI-based solutions the relationship between the IMD and the reader is transitory. In particular, the reader will have to repeat the procedure for obtaining its certificate to authenticate with the IMD in each new session.

In order to optimize the energy consumption in those cases where the reader communicates frequently with the IMD, the PKI and ACL approaches can be combined. For instance, the first time the reader is authenticated with the IMD, it will use its certificate. After that, this particular reader is registered in the ACL, since using this approach is more efficient in terms of energy consumption than PKI-based solutions.

One critical point is that the PKI and the certificate directories should be publicly—and permanently—accessible through the Internet. Consider, for example, a patient suffering an emergency condition while visiting a foreign country or just a different hospital. The medical personnel should be able to obtain the required credentials. Connectivity or authentication problems with the PKI may prevent them from gaining the required credentials to modify or disable the IMD, which in some cases may threaten the patient's safety. Therefore, the needed PKI is very demanding in the sense that is global, a large number of participants are involved, and a huge set of iterations are possible –similar requirements are demanded to the PKI that is used in the borders with the new e-passports [175].

#### 1.4.4.2 Delegation in External Devices

Some authors have suggested to make use of an external device to control accesses to the IMD. Such devices would not be implanted in the patient's body, and part or all of the security functions would be delegated to them. This presents several benefits. On the one hand, the IMD would save battery life since security-related computations are performed externally. On the other hand, a single device can integrate a number of security capabilities, such as auditing, key management, authentication, and access control. Furthermore, as most of these capabilities operate at the physical layer, other sort of solutions can be used at higher layers.

Generally, the role of the external device is to act as a mediator between the programmer and the IMD. When the programmer needs to access the IMD, it first gains access to the external device and then communicates with the IMD. In [51], the authors present a solution based on external devices named "Cloaker". The IMD periodically checks the presence of the Cloaker. While it is detected, the IMD remains silent. Therefore, the Cloaker will provide security to the patient while he holds it. Otherwise, the communications with the IMD are fully open to all readers. Using this approach, in an emergency condition it would suffice to remove the Cloaker from the patient to get full access to the device.

The authors of [51] provide a number of ideas about the role that such an external device could play. Two different possibilities are identified:

- The Cloaker would mediate in all exchanges until the IMD and the programmer successfully finish a key exchange. After that, both parties directly communicate with each other over a secure channel built using the shared key. The external device does not participate in such communications.
- A different possibility is having the Cloaker involved in all communications between the IMD and the programmer. In this case, all packets pass through it, which would allow to record them (for example, for a subsequent forensic analysis) and even implement filtering and attack detection functions. Note, however, that in this setting the Cloaker becomes a single point of failure, so any malfunction or degradation in performance will affect the availability of the IMD.

Solutions based on external devices such as the one presented in [51] attempt to balance the security tensions described in Section 1.3. Security mechanisms are offered only in normal operation and the safety of the patient is guaranteed in emergency conditions. Nevertheless, there are still some open questions that have not been definitely addressed, including:

- The constant detection of the Cloaker by the IMD is not trivial. The authors proposed two ways to do this. In the first case, the IMD sends a "hello" message to the Cloaker whenever an incoming message is detected. Another, more restrictive way consists of the IMD periodically sending "hello" messages to the Cloaker to check its presence. The result is stored in just one bit that indicates whether the Cloaker is present or not.
- Both schemes discussed above are inefficient in terms of energy consumption as a consequence of the messages exchanged to check the presence of the external device. The first solution avoids a continuous flow of requests to the Cloaker, but renders the system more vulnerable since the adversary knows the exact time when the Cloaker would be interrogated. Thus, the attacker could send a fake request to the Cloaker and then impersonate it. Contrarily, the second approach is much more secure but requires the IMD to continuously check the presence of the Cloaker.
- The authors do not address the problem of how deal with an attack

that causes interferences in the communication channel between the IMD and the Cloaker.

- Finally, it is worth mentioning that [51] is not a definite solutions and the authors do not recommend its immediate adoption.

Another solution based on external devices is "RFID Guardian", proposed in [181]. RFID Guardian registers all devices in its range, manages keys, authenticates programmers that request access to the IMD, and blocks all unauthorized entities. Using this approach, all the devices in the neighbourhood of the Guardian (i.e., about 1 or 2 meters according to [181]) are detected and corrective measures could be enforced if needed. Although the solution was originally proposed in the context of RFID systems, the approach can be easily adapted to IMDs. The authors propose to integrate the Guardian into a device that the user (patient) always holds, such as a smartphone or a smart wearable device (e.g., a watch, a bracelet, etc.).

Other approaches are based on the use of hardware tokens. There is a wide variety of these devices, including disconnected and connected tokens, smart cards, bluetooth tokens, etc. In this case, the device stores a password shared with the IMD. The medical personnel would use this key to access the implant. The main drawback is the same as in other solutions based on external devices: if the token is lost or the patient forgets to carry it under an emergency condition, the IMD would be inaccessible [12].

Gollakota et al. proposed the use of an external device, named "shield" [74, 73], that acts as an intermediary so that all communications between the IMD and the programmer pass through it. The shield protects the communication channel by jamming messages sent to and from the IMD in such a way that no other entity can decode them. Similarly, it protects the IMD from unauthorized devices by jamming all messages coming from them. Fraudulent messages are rendered unusable after the jamming and the IMD would discard them simply by its inability to interpret them, thus preventing the execution of malicious actions.

In summary, the main advantage of solutions based on external devices is that they offer a high protection level against unauthorized commands. The IMD will not respond to malicious re-programming commands or attacks to drain the battery. Their main drawbacks include:

- If the patient forgets the external device, the IMD would respond to all incoming (authorized or not) requests, which is only necessary in emergency mode.
- These solutions do not generally consider scenarios in which the ex-

ternal device is replaced by a malicious one. In this case, the security and privacy of the IMD would be highly compromised.

- The external device is fully visible to external entities, which can reveal sensitive information about the patient's medical condition. Moreover, some authors have pointed out that it is a constant reminder to the patient about his medical condition.
- These proposals often assume that the external device is a trusted entity. Nevertheless, this entity can be compromised or act maliciously. For instance, packets can be altered (e.g., flipping certain bits), dropped out, or blocked, which would render inoperative the communication with the IMD.

#### 1.4.4.3 Trusting other Implantable Devices

In [86] it is proposed the idea of using a subcutaneous button that opens access to the IMD only after being pressed. This approach would protect the implant from all incoming communications until the patient deliberately presses the button, which can be done only in controlled environments. Note, however, that this would fail to protect the IMD if the adversary has physical contact with the victim and can press the button, or leave an attacking device in the proximity of the patient waiting for the IMD to be accessible.

The same authors also present the notion of an "IMD Hub", this being an implantable device that works as a network switch for all the devices in the IBN and also plays the role of and authentication center. This approach suffers from an excessive trust on a unique central device, so the use of more connected hubs could be a more interesting approach both for security and performance reasons.

#### 1.4.4.4 Proximity-based Access Control

These solutions base the access decisions on the distance between the programmer and the IMD, allowing only communications with devices located at a short distance from the IMD [181]. In certain applications, it has been suggested that this can be achieved by having the IMD equipped with a magnetic switch. This is just a magnetic sensor that detects the magnetic field generated by programmers in its proximity. Only after this switch has been activated the IMD will become available. After this, the IMD would send to the programmer the key to be used for subsequent communications during this session. Unfortunately, apart from security issues, it is unclear whether these solutions are safe enough, since having a magnetic field close to the implant might alter its functioning [131].

Other solutions are based on classical distance bounding protocols, these being schemes that compute an upper bound for the distance between two devices. In [177], Rasmussen et al. propose a device paring protocol in which the IMD and the programmer obtain a shared key. Messages are sent through an ultrasonic channel and the response times—i.e., the time between sending a request and receiving an answer—are used to estimate the distance between the devices. This process can be repeated several times to upper bound the estimation. If the computed distance is less than a fixed threshold, the communication with the IMD continues; otherwise, it is interrupted. It is also worth mentioning that some authors consider that response times in the protocol could serve as a deterrent against replay attacks [64], as the device could detect old request being replayed and reject them.

Normal and emergency operation modes are considered in [177]. While in normal model, the paring protocol is executed and a session key is established. This process is carried out assuming that the IMD and the programmer initially share a secret key. When in emergency mode, the shared key established above, which is probably stored in an authorization token but the patient could have forgotten it, may be unknown. To address this, the authors propose a mechanism to deal with this contingency. In detail, they propose a scheme to generate a temporary secret key so that the paring protocol can be executed. This is an alternative to the use of the magnetic switch previously described. In this case, the programmer has to be within the emergency range, which is shorter than in normal operation mode.

Distance-based solutions assume that a reader that is close to the IMD is not an adversary. It can be a legitimate programmer with the required credentials and within permitted range for normal mode. Alternatively, it could be a legitimate reader but without the authorization tokens in an emergency condition and located very close (i.e., in emergency range) to the IMD. This leads to two major disadvantages that have not being considered by this sort of protocols:

- The IMD can be compromised if the adversary is within the defined range. It would be desirable to guarantee the security of the patient independently of the distance an attacker can be. There are many daily situations in which the attacker can get very close to the IMD, such as in a public transport vehicle, at the office, etc. In other cases, the attacker can plant a programmer device close enough to the patient’s body and use it as a proxy for conducting his activities. Moreover, the attacker can be the patient himself trying to deliberately manipulate the IMD.

- There are techniques that allow an adversary to simulate being within the permitted range when in reality he is at a longer distance. This is a key limitation for any protocol based exclusively on the computed distance.

#### 1.4.4.5 Biometric Measures

Biometrics refers to a number of identification techniques based on the patient’s physical characteristics, such as his fingerprint, iris pattern, voice, hand, etc. [22, 176]. Interested readers can find in [144] more details about the use of biometrics in the healthcare context.

In [90], Hei and Du propose a solution that restricts access to authorized entities and deals with emergency situations where the patient can be unconscious or not holding his credentials (e.g., an external authorization token). The scheme uses biometric features from the patient in two separate steps or levels. Level 1 employs basic biometric information, such as fingerprints, iris color, the patient’s weight, etc. Once level 1 is passed, level-2 authentication must be passed too in order to finally get access to the device. For that, biometric information extracted from the patient’s iris must be provided. That information is pre-stored as a key in the memory of the IMD. Iris-based authentication has a high accuracy and is very efficient. Furthermore, to obtain a good snapshot of the iris a near infrared camera is needed, and the user has to be at a distance of between fifty and seventy centimetres, which is a very short range. As consequence of this, the patient would easily detect an attacker due to his proximity in many situations.

Similarly, in [234] Xu et al. propose the use of ECG (electrocardiograms) signals to generate the patient’s secret key. By using this the scheme avoids the need for pre-stored keys and the associated key distribution problem. Access control is guaranteed by a cryptographic protocol that employs this ECG-based key. On the other hand, communications between the IMD and the reader are coordinated by an external device named IMDGuard that is very similar to the RFID Guardian proposed in [181]. The presence of the IMDGuard means that the IMD works under normal mode and ECG-based access control is used. When the IMDGuard is absent, communications are not protected and access control is not enforced, which would allow anyone (e.g., medical personnel in an emergency situation, but also an adversary) to interact with the implant. A recent and detailed study about the use of ECG signals for key generation can be found in [186].

In certain cases, biometric techniques can be easier to apply than solutions based on shared keys, since the key distribution problem is avoided and it is harder for the attacker to disclose the keys. In principle, the adversary could not impersonate the programmer or the IMDGuard unless he has physical contact with the patient. Despite this, biometric-based approaches have two main drawbacks:

- Firstly, the physical presence of the patient is needed. This is certainly not a disadvantage in an emergency situation, where the patient is physically located in the emergency room or equivalent. Unfortunately, this is not the case when medical personnel will attempt to access the IMD remotely.
- Secondly, biometric features are never perfect. Two measures taken at different times, or even acquired simultaneously but using two reading devices, could generate different results. A straightforward use of such measures might generate wrong keys, when in reality the user is authorized. Error correction techniques are used to avoid this [10, 38]. This problem is known as truth rejection rate and implies that not all biometric data can be used for key generation (or authentication). Thus, the measure has to be gathered from body parts so that the differences between measures are within a acceptable range and can be corrected [37].

### 1.4.5 Anomaly Detection

The availability of the IMD functions is crucial since the treatment—and even the patient's life—can be compromised otherwise. If an attack is detected, the patient can be informed (e.g., by a notification mechanism) or the device can be rendered inaccessible via switching off the communications (or jamming the channel) while the medical functions are kept running. The difficulty to prevent this sort of attacks mainly arises from the use of the wireless communication channel. Communication between the IMD and the reader starts with the IMD authenticating the reader. If the reader does not pass the authentication step, the communication is interrupted. This consumes resources in the IMD and, therefore, can be exploited by an adversary who, for example, repeatedly attempts to communicate with the IMD. The result would be a classical Denial-of-Service

(DoS) attack in which the battery level could be drastically reduced and memory/storage could be also affected—in each authentication, some registers are used to store security values such as session tokens and logs. In general, this sort of attacks are known as Resource Depletion (RD) attacks and focus on wasting the resources of the IMD [98]. They are very easy to implement and their consequences can be very harmful as the battery life of the IMD could be shortened from several yeas to a few weeks just by sending dummy requests.

Standard cryptographic solutions do not prevent these attacks, and existing studies about RD attacks in sensors networks [179] are not directly applicable to IMDs since implants have more severe resource restrictions. Moreover, there is an extra difficultly for adding new resources—the implant is within the body, which is not the case of sensor networks. This motivates the need for designing solutions that take into account the fact that these devices will be used within a human body.

In the context of IMDs, the combined use of pattern/behaviour analysis and notification systems is the most widely used solution to counter RD attacks. Notification systems inform the patient through an alarm signal (e.g., a sound or vibration) when particular events happen, such as when the IMD establishes communication with a external device [82] or when certain biomedical parameters fall out of the normal range [53]. Such alarms are only informative. Thus, notification does not prevent attacks from happening, although alerting the patient may be valuable to make him aware of unexpected ongoing communication activity. One major drawback of these approaches is that they do not work properly in acoustically noisy environments. Besides, alarms have an associated energy consumption that should not be underestimated. As in the case of auditing, notification mechanisms alone are insufficient and should be complemented with other solutions.

By leveraging the fact that the wireless communications between an IMD and a reader follow a set of observable patterns (e.g., frequency, localization, patient conditions, etc.), Hei et al. propose in [91] a mechanism against RD attacks with an average detection rate over 90%. The scheme uses a Support Vector Machine (SVM), which is assumed to run in the patient's phone. In detail, the authors consider five kinds of input data to carry out detection: reader action type (i.e., the action(s) the reader can execute on the IMD, where the set of actions depends on the type of implant); the time interval of the same reader action; the location (e.g., home or hospital); the time; and the day (e.g., weekly or weekend). Once trained, the classifier will determine whether a pattern is valid or not. For instance, if a particular type of request is always sent from the doctor office, an attempt of the same request from a different location would raise an alarm. The overall system works as follows. Each time the reader attempts to contact with the IMD, the latter sends a message to the mobile phone of the patient with the access pattern. The phone executes the classification algorithm and returns an output that is sent back to the IMD. Depending on that output, three actions are possible: (1) the input vector is considered legitimate. In this case the mobile sends a "1" (true) to the IMD and the communication with the reader continues; (2) the input vector does not correspond with any of the allowed patterns, in which case the phone sends a "0" (false). The request may come from an attacker and the IMD turns into sleep mode to avoid RD attacks; (3) If it is unclear whether the input vector is legitimate or an attack, an alarm is triggered (e.g., an audible alarm) to inform the patient, who must decide if the communication is permitted.

The proposal in [91] has three main drawbacks. Firstly, the scheme assumes that the IMD is always running in normal mode and does not consider emergency conditions in which legitimate access patterns could be certainly anomalous. If that is the case, access to the IMD would be rejected, which could result in severe consequences for the patient's safety. Secondly, the proposal inherits some disadvantages from schemes that base its security on an external device—the mobile phone, in this case—, as discussed in Section 1.4.4.2. Finally, but not less important, the patient has the responsibility of making a decision in case the SVM cannot classify the input data.

Instead of using patterns, Henry et al. have recently proposed in [93] a system to detect malicious/anomalous use of an insulin pump. In particular, the administration of insulin dosages is detected by tracking the acoustic bowel sounds. The events are logged and then used for checking the proper system operation. The proposal is a passive solution and does not offer protection in real-time. Moreover, as in [91], the system is based on the use of an external device needed to measure abdominal sounds.

A new defense method for IMDs based on wireless monitoring and anomaly detection is proposed in [238]. The authors propose the use of a medical security monitor, named MedMon, which eavesdrops communications to and from the IMD. Captured traffic is then passed for analysis to a multi-layer anomaly detection system. If a malicious transaction is detected, the user can be informed (passive response) or alternatively the system can render the IMD inaccessible via active jamming (active response). Jamming refers to the transmission of radio signals with the purpose of impeding communications in the channel by reducing the signal-to-noise ratio. In this case, jamming is used to protect the IMD from being accessible to the adversary. The main drawback of this proposal is that the whole security resides on an external device, but it has the advantage of being applicable on existing devices without any modification. In line with MedMon, Darji and Trivedi have recently proposed a system for detecting active attacks [46]. They propose the use of an external proxy device equipped with several antennas that builds a signature of authorized readers/programmers based on their position. Positions are estimated through triangulation techniques. The proposal seems effective for static scenarios but not for dynamic ones.

### 1.4.6 Overriding Access Control

Although strictly speaking overriding access control mechanisms is not a protection measure, we have included these solutions here for completeness. Furthermore, in an emergency situation keeping the patient alive is more crucial than maintaining the IMD security protections fully functioning.

Access control models are often too inflexible. The access policy is generally established at the design phase, setting what actions are allowed, by which entities and under what circumstances. However, during the system life it is possible that unexpected and unanticipated situations may arise in which access to the implant is vital. For instance, in the context of IMDs and under an emergency condition the usage scenarios are unpredictable. Since guaranteeing the patient's safety is a priority, it is mandatory that access requirements can be removed if it becomes necessary. This type of situations give rise to a family of solutions collectively known as "Breaking-the-Glass" (BTG) that allows to switch the access control requirements off in critical or unknown situations for the system. This would facilitate the access to the implant under a emergency condition, although it also opens the door to a number of security vulnerabilities.

A typical proposal of a BTG policy can be found in [59]. Even though this work is not focused on IMDs, it can be adapted easily. In this case, the access controls requirements can be suppressed even if the entity previously did not have privileges to do that. The BTG is complemented with a non-repudiation mechanism that facilitates a subsequent analysis of the accesses carried out. The authors define a series of steps that must be executed in a precise order to override access control. First, when a user requests access, the system checks if he has the required privileges. If the answer is negative, the system may give him access under the BTG modality provided that the user accepts that all the actions will be recorded. If so, he gets access to the system and assume all responsibilities.

In [184], Rissanen et al. propose a model that distinguishes between allowed actions, forbidden actions, and all those that can be executed (possible actions). The intersection between the sets of possible and forbidden actions represents the actions that can be allowed when overriding the access control policy.

The classical Clark-Wilson access control model for data integrity [39] also provides a reference framework for BTG policies. In this case, the basic steps needed in a BTG system are reduced to [21]:

1. Pre-staging break-glass accounts: emergency accounts are created in advance, so users and passwords are generated for these special cases.
2. Distributing accounts: pres-stages accounts are efficiently managed to guarantee that the required access data is available in appropriate and reasonable manner in case of emergency.
3. Monitoring the usage of the accounts: the system must be audited while being accessed during an emergency condition.
4. Cleaning up: once access in the emergency mode concludes, new access accounts are generated and the old ones are revoked, thus avoiding temporary-authorized users have a permanent access to the system.

Obviously, a set of measures to ensure the proper functioning of the system are required as consequence of bypassing the access control in a BTG system [81, 172, 198]:

1. Users must accept their new privileges, warning them of the possible consequences of their acts.
2. The system must record the actions performed by each user. A posterior analysis will determine if the access was legitimate or not. For that, access requests can be stored together with the system status, which could help to conclude whether the access was justified by the circumstances or not. As in the case of the auditing mechanisms described in Section 1.4.2, this can be considered as a deterrence measure.
3. In an emergency condition, all access operations must be monitored

in real time to grant those privileges needed for them to be executed. Furthermore, the system has to be in emergency mode only while the emergency lasts, returning to the previous access policy as soon as possible.

4. The privileges granted in such situations must be kept to the minimum required to perform the task, but not more.
5. Some proposals like [172] go one step beyond and consider whether the requested action is reversible or not. Thus, actions executed by a user with enhanced privileges due to an emergency condition should be reversible, in such a way that unjustified actions can be undone.

As a practical implementation, the work presented in [21] by Brucker and Petritsch describes the integration of a standard access control model with a BTG policy and discusses the improvements in the architecture. Similarly, a context-based access control mechanism is proposed in [84], which depends on four factors: time, location, identity, and history of events.

In summary, BTG policies extend access control policies to critical situations, dynamically grating privileges to users who require them to execute an essential action. This type of policies are very important since it is unrealistic to assume that all possible situations will be considered at design time. In fact, in the case of IMDs the situations in which an emergency can appear are unpredictable and, in the majority of the cases, a successful management of the emergency situation depends on the access being granted in time. As for the proposals discussed above, some of its properties are difficult to guarantee a priori, such as for example ensuring that the system can recover from the BTG policy by allowing only reversible actions (in most emergency situations the required actions are clearly irreversible). Therefore, although these measures aim at solving the tension between security and patient's safety, its usage can be risky. In general, it would be necessary to carefully define what an emergency situation is and providing the IMD with the means necessary to recognize it. However, as this would have to be done at the design phase, it somehow contradicts the basic BTG motivation, namely that critical situations are unpredictable and must be detected at execution time.

### 1.4.7 Summary

Table 1.2 provides an overview of the proposals discussed throughout this section, detailing for each one of them its main purposes and how it affects three main goals: security, patient's safety, and battery life. Furthermore, in relation with STRIDE methodology, we show the security services addressed.

| Measure | | Type | Safety | Security | Battery Life | Security Services Addressed |
|---|---|---|---|---|---|---|
| **No security** | | ∗ | − | − | ∗ | None |
| **Auditing** | | Detection | − | − | + | Non-Repudiation |
| **Cryptographic Measures** | | Protection | − | + | ± | Authentication, Confidentiality, Integrity |
| **AC** | **Certificates & Lists** | Protection | − | + | ± | Authorization (+Authentication) |
| | **External Devices** | Protection | + | ± | + | |
| | **Internal Devices** | Protection | + | + | ∗ | |
| | **Proximity** | Protection | + | ± | ∗ | |
| | **Biometrics** | Protection | + | + | ∗ | |
| **Anomaly Detection** | | Protection | − | + | + | Authorization, Availability |
| **BTG Policies** | | Protection | + | − | ∗ | Authorization during emergencies |

*Legend: + Positive; − Negative; ± Both positive and negative effects; ∗ No influence.*

**Table 1.2:** Security Solutions for IMDs

# 2

# Human Identification Using Compressed ECG Signals

## 2.1 Introduction

According to [34], the medical sector is the area that has suffered the major number of hacking incidents over the last two years—43 % of the data breaches in US. Medical companies and hospitals have begun to introduce biometric solutions to mitigate attacks and reduce costs. Furthermore, the proper identification of patients when they walk through the door is a major issue nowadays for all the hospitals around the world. Errors in medical records, or even incorrect treatments, are very costly for the medical centres and harmful for the patient. To avoid this, novel solutions propose to maintain a link between the patient's data biometrics and her medical record. Thus, the biometric signature (monomodal or multimodal) is used as an index to recover the patient's medical record in the standard way: the system compares the master template with the one read locally and, if they match, the associated medical record is retrieved. This process is entirely done locally but may be also done remotely, i.e., the user would provide her biometrics data remotely. In this sense, biometrics could accelerate the transition towards home health care [203].

Home health care allows the treatment of a disease at home. On the other hand it is usually as effective as care at the hospital but less expensive and more convenient for the patient. A wide variety of health care services can be offered (e.g., check your vital signs like temperature or blood pressure remotely or have a video conference with a medical staff). Demographic changes (ageing population), social changes (small family units or mobility across countries), and developments in science and technology are some indicators that help to forecast, in a near future, a spread-use of home health care services [217]. The correct and secure identification of each individual is a key-point for the proper operation of these systems. We propose the use of a biometric solution for that purpose. For completeness, we next provide a brief introduction to biometrics.

Biometrics refers to the automatic identification of users based on features derived from their physiological and/or behavioural characteristics. The use of such features for identification (authentication) or verification purposes has been thoroughly explored in the last 30 years. In verification, an identity is provided by the user, which is used to retrieve a master template. The master template is then compared with the verification template (1-to-1 comparison) and a matching score is returned by the classifier. Contrarily, in identification systems like the one proposed in this chapter, the identity of a user rests entirely solely on her biometric information—the classifier performs one-to-many comparisons.

There is a substantial body of knowledge on recognizing subjects by their fingerprint, face, voice, gait, keystroke dynamics, hand, iris, or retina [109]. Depending on the application and operational context, each one of these features can be used separately [63, 102] or combined in a multi-biometrics setting [220]. The accuracy, measured both as the probability of identifying a correct subject and rejecting a false individual, is possibly the single most important feature of a biometric system. However, in practice there are other properties that can severely limit the use of a particular identification technique [178]. The biometric characteristic must be also universal, stable, and unique. Its acquisition has to be easy and without objections by the users. Finally biometrics systems should detect the use of an artefact or substitute. All the mentioned characteristics have been assessed against our proposed system in Section 2.4.

Over the last few years, some works have explored the biometric use of signals that, for different reasons, have traditionally received little attention by the security community. Biosignals—i.e., electrical or chemical signals measuring some activity or parameter of the human body—constitute an important class of such signals, including electrocardiograms (ECGs), electroencephalograms (EEGs), and electromyograms (EMGs). These signals have been thoroughly studied for medical applications, on the hypothesis that they convey information about different pathologies and, consequently, can be used as a valuable diagnostic tool. For example, automatic classification of ECG signals assists cardiologists to diagnose arrhythmias (i.e., tachycardia, bradycardia or atrial fibrillation) [121].

In the last years, several works have demonstrated that many vital signals also contain features unique to the individual and can be used for security purposes. This branch of the biometrics is increasingly referred to as *Intrinsic* or *Hidden Biometrics* [157]. For instance, the electrical activity produced by skeletal muscles can be used for biometrics. EMG is the technique used for measurement and the obtained record is called electromyogram. In [215], Suresh et al. proposed the use of electromyograms to generate a signature for human identification. For that, impulsional electrical stimulation is produced over the muscle and its response constitutes the signature. This proposal has been successfully tested over a population of ten individual.

ECG and EEG signals are by far the most commonly studied signals for Hidden Biometrics. EEG records the electrical activity in the brain through a set of electrodes mounted on the scalp. Existing proposals can be grouped according to the classification algorithm used. Linear discriminant classifiers with auto-regressive feature extraction are demonstrated in [163]. In [33], an LVQ neuronal network with FFT feature extraction is described, while the work in [214] reports results using a neuronal network with energy feature extraction. On the other hand, EEG records the electrical activity of the heart. The algorithms can be classified according to the features extracted from the ECG signal. Fiducial-based methods extract information from the characteristics points of an ECG wave (e.g., amplitude [69], temporal duration [106]). Non-fiducial methods do not use the characteristics points to extract features. Instead, other features like autocorrelation [2], Fourier [188] or Wavelet coefficients [36] are used. Other solutions (hybrid) combined both methods like in [226] or in [207] . The reader is urged to consult for an exhaustive survey of ECG-based biometrics proposals [160].

The rest of the chapter is organized as follows. In Section 2.2 the general architecture of the biometric identification system is presented. After that, we review and explain each of its forming components. The results are presented in Section 2.3. Then, in Section 2.4, we evaluate the main properties of the proposed system.

## 2.2 Methods

In Figure 2.1 we show the general architecture of a biometric identification system. The first step consists of the data acquisition—one or several signals take part depending on whether the system is mono-modal or multi-modal, respectively. Usually a set of sensors are placed over the subject (e.g., chest or head) to read the biosignals. Once acquired, the raw data must be prepared for its analysis. Techniques such as normalization,

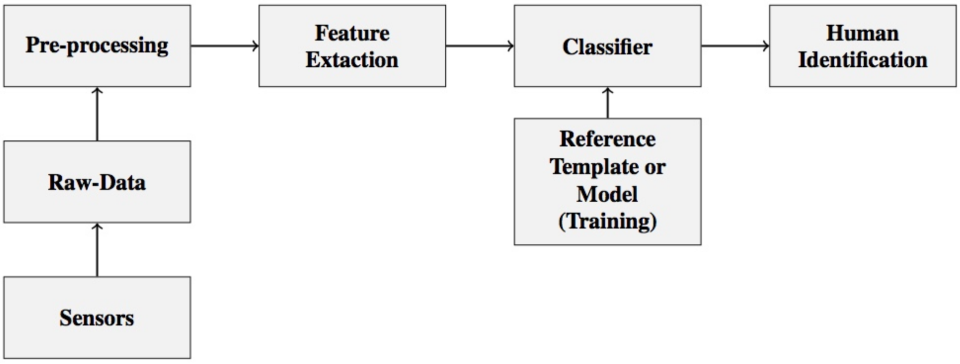

**Figure 2.1:** General structure of a biometric identification system

re-sampling or smoothing are commonly used procedures during the pre-processing step. After that, the more relevant information of the signal is represented by a set of numerical o nominal parameters. This step is usually known as feature extraction and is crucial for the success of the whole process. The generated dataset consists of a number of instances, each one formed by a set of features and a label corresponding to an individual. The aforementioned dataset is split into two subsets for training and testing, respectively. The training set is employed to build the model and the unseen samples (testing set) are used to evaluate the model. That is, for each instance the model outputs a label that is compared to the ground truth. Depending on its success, the classification accuracy will be higher or lower. In the following, each one of these building blocks are explained in more detail taking into consideration the particular procedure used in this proposal.

### 2.2.1 Raw Data and Pre-processing

The electrocardiogram (abbreviated as ECG and sometimes EKG) consists of a measurement over the skin surface to record the electrical activity of the heart. The conduction of ions through the myocardium (heart muscle) change with each heart beat. The ECG represents the sum of the action potentials of millions of cardiomyocyte (heart cells).

For our experimentation, we have chosen a well-known dataset. In particular, the MIT-BIH Normal Sinus Rhythm Database is used [72]. It includes long-term recordings of 18 subjects treated at Boston's Beth Israel Hospital. The decision of using this dataset was motivated for the fact that no significant arrhythmias were detected in the subjects. Therefore, the subjects do not present any bias that could help in the identification task.

The heart rate of a person at resting varies from 60 to 100 beats per minute. In order to pre-process the signal, at the fist step the DC components are eliminated. After that, each ECG signal is filtered using a passband filter. The passband range is often governed by the intended application: for instance, [0.05Hz – 150Hz ] for diagnostic and [0.67Hz – 40Hz] for patient monitoring. In our cause, we use pass-band filter with passband rage between 0.67 and 45 Hz. The lower cut-off frequency is set 0.67 to eliminate the noise introduced by the respiration of the subject. The upper cut-off frequency is set to 45 Hz to keep as much information as possible and to eliminate the power line noise.

Once the signal is filtered, we split the signal in chunks without overlapping. The chunk size is set to 2 seconds, which means that each chunk consists on 2-3 heart beats. We have chosen this size inspired by the fact that algorithms based on fiducial features often use two beats as chunk length. We attempt to obtain similar information by using a compressed version of the signal. The non-fiducial features used in our experimentation are explained in the next subsection.

### 2.2.2 Feature Extraction

Features derived from biosignals are effective in the design of human identification systems [155, 125]. ECG signals are one of the most used for this purpose [119]. Generally, existing algorithms can be classified into two main groups [160]. On the hand hand, the algorithms based on fiducial features use characteristic points (e.g. PQRST peaks) from a ECG trace to extract a set of features (e.g., time intervals between peaks or angles). In Figure 2.2 we show the main characteristic points together with the most common features of an ECG wave. Contrarily, algorithms based on non-fiducial features do not employ characteristic points for generating the feature set.

In this chapter, we propose the use of a non-fiducial based algorithm. In particular, the Hadamard Transform (HT) is used to extract the features of an ECG wave. Figure 2.3 shows the ECG signal in the time domain and in the Hadamard domain, respectively. In the same way as the Fourier Transform (FT) consists of a projection onto a set of orthogonal sinusoidal waveforms, the Hadamard Transform (HT) lies in a projection onto a set of square waves called Walsh functions. In fact, the Hadamard transform is often called Walsh-Hadamard transform, since the base of the transformation consists of Walsh functions.

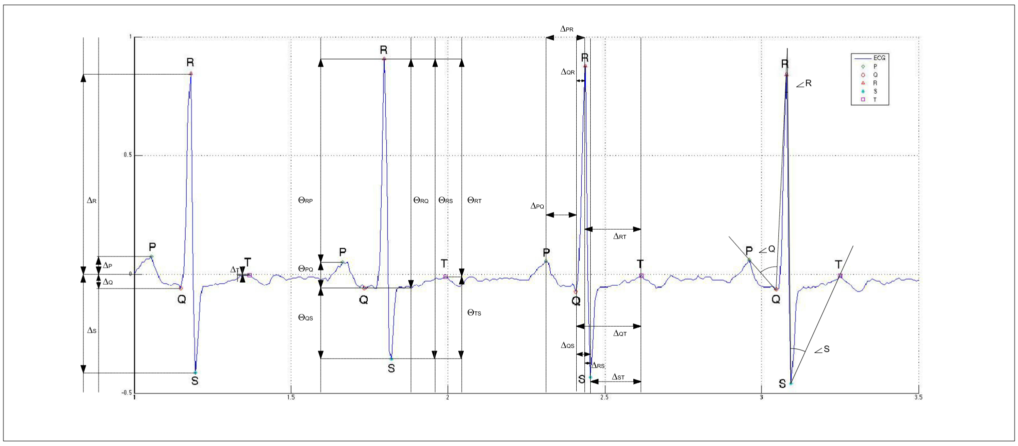

**Figure 2.2:** ECG wave: characteristic points and features

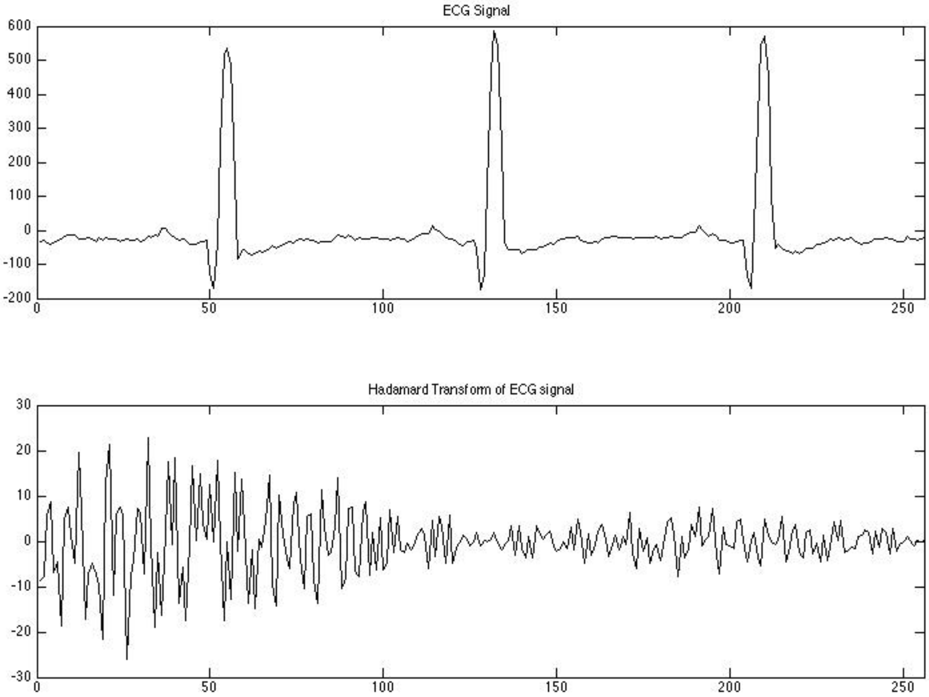

**Figure 2.3:** ECG wave in the time domain and its Hadamard spectrum

The Discrete Walsh-Hadamard Transform (DWT) of a data sequence $x(n)$ and $n = \{1 \cdots N\}$ is given by:

$$X_w(k) = \sum_{n=0}^{N-1} x(n) \prod_{i=0}^{M-1} (-1)^{n_i K_{M-1-i}}, \quad k = 0, 1, \cdots, N-1 \qquad (2.1)$$

where $N$ is the number of samples of the data and restricted to be a power-of-2, and $M = \log_2 N$. Therefore, in a simply way, the transform $(X_w)$ consists on the product of the sequence $(x)$ of length $1 \times N$ by the Walsh matrix $(H)$ with length $N \times N$:

$$X_w = Hx \qquad (2.2)$$

The inverse of the transform can be easily calculated with the next analogous expression that only differs in the constant divisor:

$$x(n) = \frac{1}{N} \sum_{k=0}^{N-1} X_w(K) \prod_{i=0}^{M-1} (-1)^{n_i K_{M-1-i}}, \quad n = 0, 1, \cdots, N-1 \quad (2.3)$$

One advantage of using this transform is that it is computationally more efficient than others, such as the Fourier Transform or the Wavelet Transform. This is important in constrained devices with limited computational capabilities. On the other hand, the usage of this transform facilitates that a compressed version of the signal could be stored, while this compressed signal preserves all the informational of the ECG signal and allows the reconstruction of the signal in the time domain.

To show the effectiveness of the HT with ECG signals, we have studied the effect of compressing the signal. To illustrate this, an ECG wave of 256 samples has been used. The HT is computed over this signal and 256 coefficients are obtained. After that, we have taken fractions of these coefficients (i.e., $\{X_w(0), \cdots X_w(P)\}$ and $P = \{256, 128, 32, 16, 8\}$) and calculated the inverse of the transform to reconstruct the signal. The results of the reconstructed signals are shown in Figure 2.4, which illustrate how the signal can be highly compressed while preserving the signal's main characteristics.

In Figure 2.5 we sketch the feature extraction procedure. As shown, the features used in our proposed ECG-system are mainly based on the coefficients of the HT. In particular, the 24 lower sequencing coefficients has been used—see Section 2.3 for details. Furthermore two additional features have been computed over the whole set of HT coefficients. Shannon

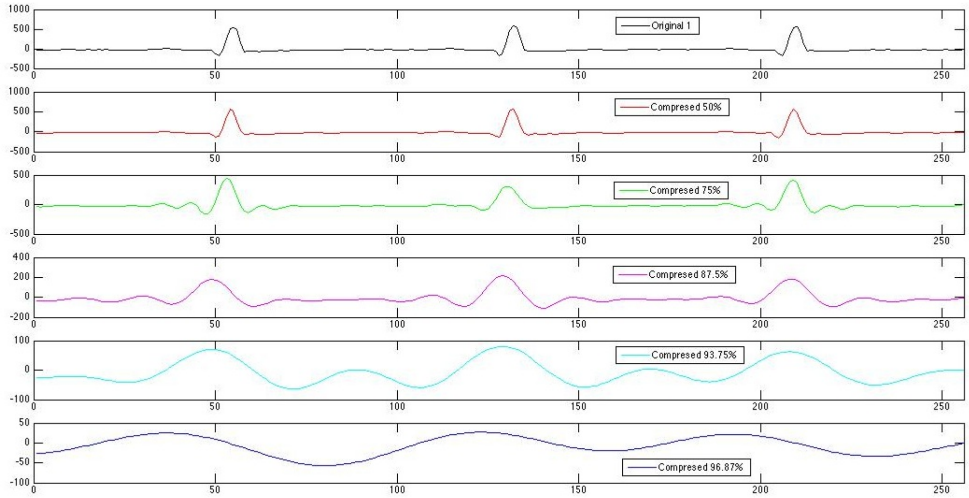

**Figure 2.4:** Reconstructed ECG signal after compression via the HT

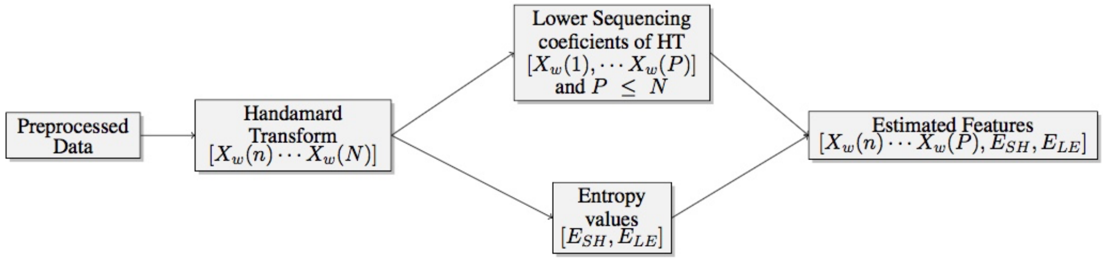
Preprocessed Data
Handamard Transform $[X_w(n) \cdots X_w(N)]$
Lower Sequencing coeficients of HT $[X_w(1), \cdots X_w(P)]$ and $P \leq N$
Entropy values $[E_{SH}, E_{LE}]$
Estimated Features $[X_w(n) \cdots X_w(P), E_{SH}, E_{LE}]$

**Figure 2.5:** Feature extraction procedure

entropy ($E_{SH}$) and Log-Energy entropy ($E_{LE}$) are the two features chosen in our experimentation—other features like statistical metrics were tested but finally discarded. Let $x$ a signal and $X(n)$ the coefficients of $x$ in a orthogonal base, both entropies can be calculated as follows:

$$E_{SH} = -\sum_n X(n)^2 \log(X(n)^2) \tag{2.4}$$

$$E_{LE} = -\sum_n \log(X(n)^2) \tag{2.5}$$

### 2.2.3 Classifier

Inductive machine learning uses the concept of learning by example. A system infers a set of rules from a set of input instances (training set). Once the model is generated, the built model can be used to classify unseen instances (testing set). There is a wide range of classification algorithms and the choice of one or another is determined by the nature of the problem, the dataset characteristics and the application where it will be used. Taking into consideration is function or form, classifiers can be categorized in numerous types, including decision tree learning algorithms, kernel methods, lazy learning algorithms, etc.

In this proposal we use a K-NN algorithm, which fits within the category of non-parametric lazy learning algorithms. Non-parametric refers to the fact that they avoid making assumptions about the data distribution. Lazy means that the training instance are not used to do a generalization, so the training is minimal. The K-NN algorithm makes several assumptions: 1) the instances are in a metric space (i.e., scalars or multidimensional vectors) and distance metrics can be computed between two instances; 2) each instance in the training set is composed of a vector (set features) and a label; and 3) the parameter $K$ determines how many neighbours are considered for classification.

The testing and training phases for the K-NN algorithm are as follows. In the training phase, features vectors with its corresponding class are stored. In the classification phase, let $y_j$ an unseen instance and $\{x_0, \cdots x_k\}$ the $K$ nearest training instances. The label of $y_j$ is determined by majority voting among the labels of its $K$ neighbours.

K-NN has been chosen since it is simple but effective. We have tested several values of the $K$ parameter and finally it has been set to 1. In fact, using higher values for $K$ (i.e., $K = \{3, 5, 9\}$ we do not observe any improvement in the performance while the cost in terms of computational load is significant. Regarding the distance metrics, Euclidean distance ($d_E$) and Manhattan distance ($d_M$) have been evaluated. Let two vectors $x = [x(1) \cdots x(N))]$ and $y = [y(1) \cdots y(N))]$ , both metrics are defined as follows:

$$d_E = \sqrt{\sum_{i=1}^{N} \left[x(i) - y(i)\right]^2} \tag{2.6}$$

$$d_M = \sum_{i=1}^{N} |x(i) - y(i)| \tag{2.7}$$

## 2.3 Results

The algorithm proposed in this work fits within the algorithms based on non-fiducial features. The main difference in comparison with its predecessors is that the algorithm works with a compressed version of the original signal via the Hadamard transform. Furthermore, only a small fraction of all the coefficients are necessary for human identification. Since the number of used coefficients—24 coefficients for each ECG chunk—is effective for identification but insufficient to recover the original signal and to preserve its characteristic points, the proposed system is privacy preserving for the user.

The procedure followed for the analysis of the ECG signal is the one explained in Section 2.2 and sketched in Figure 2.1. For our experimentation, we use the well-known MIT-BIH Normal Sinus Rhythm database. In particular, ECG signals for two electrodes are available and were preprocessed as explained in Section 2.2.1 . Thereafter the same procedure is followed for each electrode signal. The signal is chopped in chunks of 256 samples and for each chunk the DWT is computed. Finally the feature extraction has been evaluated using two approaches:

- Hadamard Coefficients: Only a small fraction of the coefficients are necessary for the identification task. This number has been obtained through experimentation and using the classification accuracy as the metric for comparison between the possible values. After conducting some experiments, we have set this value to 24, which represents less than 10% of the coefficients. Therefore using these 24 lower sequencing coefficients (48 in total considering the two leads) the system can identify an individual with high accuracy. On the other hand, and considering the worst case in which an attacker would capture these coefficients, she could not reconstruct the original signal as shown in the Figure 2.6—only partial information might be retrieved.

- Entropy: Although the system offers a high performance using Hadamard coefficients, we have studied whether additional features are useful for the system. In particular, we have calculated the Shannon and the Log-Energy entropy over the whole set of Hadamard coefficients (i.e., 256 values). It is also worth mentioning that we also tested the inclusion of commonly used statistical metrics (e.g., mean, standard deviation, maximum, minimum, and first derivative, etc.). Nevertheless its benefit over the performance of the system is negligible and for this reason these features were not finally considered in our experimentation.

Once the features are extracted, we have trained and tested a 1-NN classifier. We have used 10-fold cross-validation in order the classifier can accurately predict unknown data. Each instance consists of a set of features

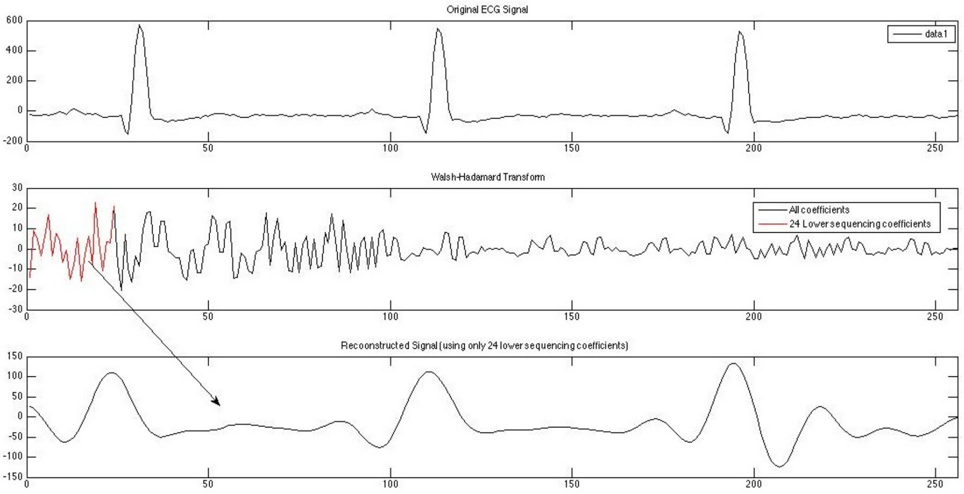

**Figure 2.6:** ECG signal: original, transformed (via HT), reconstructed

| Configuration | | **FNR** | **FPR** | **TPR** | **TNR** |
|---|---|---|---|---|---|
| OP-1 | $d_E$ | 0.0580 | 0.0582 | 0.9418 | 0.9420 |
| | $d_M$ | 0.0570 | 0.0566 | 0.9434 | 0.9430 |
| OP-2 | $d_E$ | 0.0390 | 0.0386 | 0.9614 | 0.9610 |
| | $d_M$ | 0.0340 | 0.0341 | 0.9659 | 0.9660 |

**Table 2.1:** Overall Performance: False Negative (FN), False Positive (FP), True Positive (TP) and True Negative (TN) rates

and a label corresponding to the subject (from 1 to 18). Regarding the features employed we have tested two possible configurations: OP-1 only considers 24 lower sequencing coefficients of the HT—in total 48 features taking into consideration the two leads available; and OP-2 considers the same features as OP-1 plus the Shannon and the Log-Energy entropy (4 additional features considering the two leads). For each configuration, the 1-NN classifier has been evaluated using two distances metrics: Euclidean and Manhattan.

The confusion matrix obtained for each configuration can be summarized through the true positives ($TP$) and false negative ($FN$) rates and its corresponding complementary values, false positive ($FP$) and true negative $TN$ rates, respectively. The obtained values are summarized in Table 2.1. Using these values, the performance of the proposed ECG-based human identification system can be assessed through a number of standard metrics:

- *Classification Accuracy.* Measures the proportion of correct outputs, both positive and negative:

$$CA = \frac{TP + TN}{TP + FP + FN + TN} \tag{2.8}$$

| **Configuration** | | **CA** | **ST** | **SP** | **PPV** | **NPV** |
|---|---|---|---|---|---|---|
| OP-1 | $d_E$ | 0.9419 | 0.9420 | 0.9418 | 0.9418 | 0.9420 |
| | $d_M$ | 0.9432 | 0.9430 | 0.9434 | 0.9434 | 0.9430 |
| OP-2 | $d_E$ | 0.9612 | 0.9610 | 0.9614 | 0.9614 | 0.9610 |
| | $d_M$ | 0.9659 | 0.9660 | 0.9659 | 0.9659 | 0.9660 |

**Table 2.2:** Performance metrics

- *Sensitivity*. It is simply the true positive rate, i.e., the proportion of actual positives that are correctly identified as such:

$$ST = \frac{TP}{TP + FN} \tag{2.9}$$

- *Specificity*. Also known as the false positive rate, measures the proportion of actual negatives that are correctly identified as such:

$$SP = \frac{TN}{FP + TN} \tag{2.10}$$

- *Positive Predictive Value*. Also known as precision, measures the proportion of positive outcomes that are actually positive:

$$PPV = \frac{TP}{TP + FP} \tag{2.11}$$

- *Negative Predictive Value*. Measures the proportion of negative outcomes that are actually negative:

$$NPV = \frac{TN}{FN + TN} \tag{2.12}$$

In Table 2.2 the performance of the proposed system, in its four possible configurations, is summarized. In the next section we evaluate the proposed system from a biometric point of view and extract some conclusions about the performance and what is the most recommended configuration.

## 2.4 Discussion

In the above section we have shown the results of our proposal regarding its performance. Seven characteristics (including performance) are commonly demanded to a biometric system [178]: universality, uniqueness, permanence, collectability, acceptability, performance, and resistance to circumvention. In the following each of these characteristics is analyzed for our proposed system:

**Universality** The biometric characteristic must be universally applicable. In our case, we use the ECG signal, which can be collected from everyone who is alive. Normal values for a person at resting are in the rage of 60 to 100 beats per minute.

**Uniqueness** The biometric characteristic must be able to unequivocally identify the individuals within the target population. In this work we have proposed the use of the ECG. This signal has already been proved to be effective for biometrics purposes [43, 162]. In our case, we have checked whether features obtained from a compressed ECG signal (via Hadamard transform) can be used to identify individuals. As shown in Table 2.1, the number of misclassified samples is almost zero for all the configurations evaluated. This is a clear indicator about the effectiveness of the Walsh coefficients (lower ones) for the human identification task.

**Permanence** The biometric characteristic should be invariant over time. Nevertheless, physiological characteristics are not totally invariant during the entire life of an individual [44]. This means that the classifier model would have to be updated after five years since the model was generated. If we compare our system with other common solutions based on passwords [195], in which the user normally must update the password once per year, our proposed solutions is five times less demanding in terms of updating requirements.

**Collectability** The biometric characteristic should be quantitatively measurable. In our particular case, ECG signals can be easily gathered through a set of electrodes—3-lead or 12-lead system. Using these electrodes, the electrical activity of the heart can be recorded. More precisely, the ECG represents the potential differences between electrodes.

**Acceptability** It relates to how the user feels comfortable with the use of the biometric characteristic. We cannot do a strong presumption about this matter since we use a public dataset for our experimentation. Nevertheless, we can predict a high acceptability due to two main reasons: 1) the ECG signal is well-known to deal with heart diseases; and 2) the signal can be easily acquired—just three leads are sufficient for non-medical applications.

**Performance** Our proposed system offers a high accuracy level. The classification accuracy varies from 0.94 to 0.97 for the two configurations evaluated. Furthermore, and not less important, the identification system errors (i.e., false positive and false negative identification rates) are very low values: of the order of $10^{-2}$. In relation with the distance metric, the Manhattan distance seems to offer slightly better results for the 1-NN classifier than the Euclidean distance. From the computational point of view, OP-1 is less demanding since only Hadamard coefficients are necessary and the penalty in performance is small in comparison with OP-2.

We can compare our system with other proposals with similar results (see Table 2.3). Most other ECG-based biometric solutions achieve similar performance. Nevertheless, the main contribution of this proposal is the set of features used. Fiducial features has been proven to be effective but its calculation requires moderate computational capabilities [226]. In our case we use non-fiducial features trough the computation of the DWT. A matrix with ones and minus one values has to be stored in memory, in what is called the Waslh matrix. Note that the matrix size is fixed since the length of the ECG chunks does not vary. In our particular case, we have set this parameter through experimentation aiming at optimizing system performance. The Walsh Hadamard coefficients are obtained just by multiplying a vector (an ECG chunk of 256 samples) by the Walsh matrix (a matrix of size 256 x 256 with ones and minus ones). The complexity of this naive algorithm is $O(N^2)$ but this can be reduced to $(N \log N)$ using the Fast Walsh-Hadamard Transform.

**Resistance to Circumvention** This property is vital for an identification system. The biometric characteristic should prevent an attacker from impersonating an authorized user in the database. In our proposed system, this property is satisfied since the ECG signal (the complete wave) is characteristic of each person. Note that two persons can have identical heart rates but their ECG waves will be different. Previous studies have confirmed this matter and it is commonly assumed that the features of an ECG signal are mainly resistant against coun-

| System | Correctly Classified Instances |
|---|---|
| **Our system** | 94 % (OP-1) – 97 % (OP-2) % |
| ECG [160] | 86 % – 100 % (single day data acquisition) |
| EEG [211] | 72 % - 80 % (4-40 individuals) |
| EEG and ECG [183] | 97.9% (linear boundary) |
| Pulse-Response [178] | 88% –100% (small data set) |
| Finger-vein [237] | 98% (70 individuals) |
| Iris and Fingerprint [152] | 96% (small dataset) |
| Face & Iris [208] | 99% (UBIRIS v.2 and ORL) |

**Table 2.3:** Biosignal-based authentication proposals

terfeiting [216].

It is clear from all the above that the proposed system satisfied the characteristics required of a biometric system. Therefore the use of compressed ECG signals via Hadamard Transform is robust, effective, and efficient for human identification. We next provide reasoning about the implications of our proposal and extract some conclusions.

# 3

# Real-time Electrocardiogram Streams for Continuous Authentication

## 3.1 Introduction

Security applications are gaining momentum in modern societies. With the advent of information technologies, data and resources are available almost anytime, anywhere. One key aspect is to ensure that the access to these elements is provided for authorized users only. This need is usually referred to as access control [191].

As a prerequisite of access control, the identity of the user has to be established. This process is called *authentication* and is especially critical when sensitive data is at stake. For instance, access to medical records can be forbidden until being authenticated [124].

Authentication can be carried out by means of something the user *knows*, something the user *has* and/ or something the user *is* [110]. Among these three alternatives, the latter is receiving particular attention as a consequence of the evolution of *biometrical* systems, i.e., the acquisition of body-related variables called *biosignals* [200]. For example, entering a building after fingerprint recognition or accessing a smartphone application after facial scanning are two cases of these systems [135, 141]. Recent developments for medical devices open up the door to the access of biosignals in almost real-time. These devices can be placed over the skin (e.g., a external heart rate monitor), semi-implanted (e.g., an insulin pump) or within the body (e.g., a pacemaker or a neurostimulator). Implantable Medical Devices (IMDs) is the general term used to refer to electronic devices implanted within the body. IMDs are designed to provide a medical treatment, to monitor the patient's status, or to enable a particular capability in the patient [88].

Different biosignals have already been considered for authentication purposes, including the Electroencephalogram (EEG) [99], the Photoplethysmograph (PPG) [130] and the Electrocardiogram (ECG) [117]. Likewise, the continuous availability of biosignals enables performing an advanced form of authentication, called Continuous Authentication (CA). This variant is different from Non-Continuous Authentication (NCA). In NCA, the user is authenticated once at time $T$, for example when s/he is logged in a system with authentication checking. On the contrary, in a CA setting the user is authenticated every period of time $T_i$, thus ensuring the continued presence of the user.

Biosignal-based CA approaches have a direct benefit: users cannot transfer their privileges to other parties, since it must be the very same user who is periodically authenticated. Despite this benefit, one drawback is that biosignals evolve over time and may be slightly different from time to time. As a consequence, the authentication mechanism should be continuously enhanced and not static as time goes by [165].

Such a continuous enhancement and the adaptation to changes makes Artificial Intelligence (AI) techniques particularly suitable. In particular, as the process requires telling apart the legitimate user from other subjects, it can be considered a classification problem. This problem has been frequently solved through AI and, particularly, data mining and machine learning techniques. Machine learning focuses on the design of algorithms to make predictions after the identification of structural patterns in data . In general a model is created, trained with part of the existing dataset and evaluated with the remaining part of the dataset [231].

Since biosignals can be retrieved in real-time, this can be taken as an advantage to permanently refine the authentication mechanism. To this end, beyond machine learning, *data stream mining* can be applied. Data Stream Mining (DSM) is a recent IA technique that leverages data streams to adapt the classification model when a change is detected [15]. This is especially beneficial for the case of biosignals, due to their aforementioned evolution over time. Interested readers are urged to consult [68, 120, 174] for a detailed introduction to DSM and related concepts.

Despite the potential of DSM for biosignal-based CA, this approach has not been previously explored. To this end, this chapter presents the use of DSM for a particular type of cardiac signal, namely ECG data. The proposed solution allows the use of ECG streams in real-time applications in which the credentials of the users are validated in a continuous-fashion. For the generation of the ECG streams in the CA setting different approaches have been assessed. For completeness, the NCA scenario has been also evaluated.

## 3.2 Motivation

The use of biometrics is widespread nowadays, from the use of the touchscreen in smart devices [202] to a more common approach like fingerprint-based identification [169]. Biological signals are currently taking an important role in the authentication field [180] and they are considered useful biometric traits. Multiple physiological signals are used in this context, such as the EEG signal [99], the PPG signal [130], or the ECG signal [117].

Though many existing works deal with classical authentication using assorted biometric traits [229], continuous authentication systems and applications have been also extensively developed. For instance, Niinuma et al. [158] use the facial skin and color clothes to authenticate users. In the context of mobile devices, facial [138] and touch screen recognition [63] have been applied. Signal processing has also been used in this field, particularly PPG and ECG signals. Although some proposals work with PPG [17, 71], here we focus on those related to ECG signals since electrocardiograms are a richer signal from the information point of views—PPG signals only provide beats and average heart rate.

Focusing on ECG signal authentication, several works are devised. In [126] the QRS complex, the most stable component of ECG signal, is applied in the continuous authentication process. After preprocessing the QRS complex and extracting the cross-correlation of QRS signals between a pair of templates, the matching score is computed through different strategies, including using the mean, median, percentile and maximum values. Experiments attempt to analyze the permanence and stability of the biometric features extracted from the QRS complex in ECG signals on a time period of a day. Guennoun et al. [78] use several ECG features to perform continuous authentication. The Mahalanobis distance is then calculated between a heartbeat and a previously stored one such that results depend on a threshold when the process has been repeated for 35 heartbeats. The main limitation of this interesting proposal is that each ECG record only lasts 15 minutes and an experiment consumes around 30 seconds. In [145], the Autocorrelation / Linear Discriminant Analysis (AC/LDA) algorithm is applied for the design of the biometric features extracted from the ECG signal. Each time an authentication is performed the signal is preprocessed and the result is matched with the stored one. The proposal was tested with a population of 10 individuals and the leghth of each ECG record is only 5 minutes. A different approach is proposed in [42], the ECG signal is converted into strings to be later classified. This proposal shows promoting results but, as in previous works, the used ECG signals were recorded during a short time interval ($<$ 15 minutes).

Nonetheless, despite the use of ECG signals for continuous authentication, existing works are evaluated over cardiac signals of a few minutes length at maximum. The variability of the signals and, hence, that of the model, is not considered. An authentication model, created from an observed set of samples does not have to be always the same and it may evolve. Data streams are a useful way to manage this issue. In fact, they are already used in the context of data outsourcing [165] but, to the best of our knowledge, this contribution is the first time continuous authentication with ECG streams is applied.

## 3.3 Data analysis: Data mining vs. data stream mining

Data mining refers to the set of technologies to handle larger datasets to find patterns, trends or rules and explain data behavior [231]. These technologies have consolidated due to the huge amount of data which is everyday collected and handled. Indeed, this is a trend which continues growing at a fast pace in different areas, for instance in the healthcare context [154].

However, given the amount of data often available the question is: What if data cannot be fitted in memory? In this case smaller training sets are demanding such that algorithms process subsets of data at a time. Then, the goal is the development of a learning process linear in the number of samples. In other words, the problem is that while data mining handles too much data, it does not consider the continuous supply of data. Models cannot be updated when new information arrives and the complete training process has to be repeated. Furthermore the length of the data feed is hugely larger—for instance, imagine a cardiac signal monitored during the entire life of an individual.

Opposed to traditional data mining, Data Stream Mining (DSM) has emerged as a paradigm to address the continuous data problem. The core assumption is that training samples arrive very fast and should be processed and discarded to make room for new samples, thus being processed one time only. More specifically, DSM presents a set of different requirements [14]:

- **Uniqueness:** Each sample must be processed before a new one arrives and it has to be done only once, without being possible the retrieval of any previous samples.
- **Limitation of resources:** Use a limited amount of memory and work in a limited amount of time. Concerning the limitation of memory, this is one of the main motivations of using data streams because memory can be overloaded when too much data is stored in it. This restriction is physical and though it can be addressed using external storage, algorithms should scale linearly in the number of samples to work in a limited amount of time.
- **Immediacy:** An algorithm should be ready to produce the best model at any time regardless of the number of processed samples.

Concerning data mining algorithms, lazy and eager algorithms are key types to be distinguished [231]. In the former type no action is performed during the training phase, such that training data is stored and it waits until testing starts. By contrast, in eager algorithms a model is constructed from training data to apply testing on its regard. In the context of DSM, lazy and eager approaches are available but existing algorithms must be adapted to the data stream setting [15].

Other noteworthy types are parametric and non-parametric algorithms [187]. Parametric algorithms are those in which parameters are of fixed size and the model does not change regardless of the amount of data. Despite their simplicity, speed and less data required in the training phase, they are appropriate for simple problems and they are not well-fitted. Some examples are logistic regression [77] or naive Bayes [146]. On the contrary, non-parametric algorithms are useful for learning when there is too much data and no prior knowledge. Flexibility to fit to a number of functional forms and high performance are some of their main advantages, while they require a substantial amount of training data, they are slower in the training phase and the training data could be overfitted if not carefully performed. Some common algorithms in this class are support vector machines [11], decision trees [147] and the K-Nearest Neighbour (K-NN) [231]. Among all these algorithms, K-NN is often used for its simplicity and efficiency [3]. Basically, the problem K-NN solves is to identify the point in a dataset closer to a set of given ones. In the training phase,

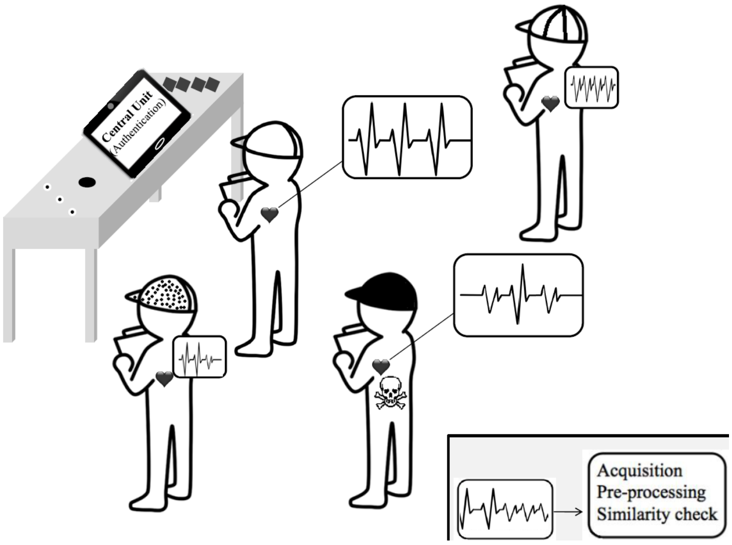

**Figure 3.1:** Example of an scenario for continuous ECG-based stream authentication.

any assumptions about the classification of samples is performed. In the classification, $K$ samples belonging to the training set that are closest to the sample are used as a good indicator to determine an unknown class, generally using a majority voting.

## 3.4 System overview

An application of our system is depicted in Figure 3.1. Imagine an air traffic control tower where there are controllers who should be permanently monitoring planes and, thus, verifying that everything runs smoothly. In this situation we have to consider that: 1) an intruder may enter into the tower trying to cause some damage; 2) a controller may try to do the work of another; and 3) physiological indices such as the heart rate is not constant and may vary according to each situation (e.g., too many plains about to take off or landing). In this regard, an authentication system requires the continuous authentication of each controller verifying that no impersonation attacks are performed and that each controller is in the work place no matter ECG fluctuations.

According to the example in Figure 3.1, the system works in the following way assuming that captured ECG signals are sent to a central unit (e.g., a smartphone or a nearby computer). Firstly, in the set-up phase, the ECG signal of each controller is observed (collected) for some time (i.e., 30 min) and, once cleaned and pre-processed, a reference model is constructed (similarity module). Secondly, in the operating phase, the system is prepared to start the authentication process, in this case verifying that each controller is in the tower throughout the office hours and feels well. In a first step, ECG records are cleaned, features are extracted and each ECG sample passes (or is discarded) by the similarity module. After that, the observed ECG signal of each controller is compared against the reference model (learner), also using part of the signal for learning and adjusting the model accordingly. Note that in case a change is produced, e.g., due to stress caused by 10 planes landing in a sort period of time, the ECG signal may change. However, as the model adapts dynamically to the situation, no alarm will be activated. In contrast, in case a big change in the ECG signal is detected, due to someone impersonating a controller or s/he feeling suddenly very dizzy, the authentication fails and an alarm is activated. The steps followed during the set-up and operation phases are summarized in Algorithm 1.

**Algorithm 1** ECG-based Authentication

```
procedure SET-UP PHASE
    1. Capture ECG records
    2. Pre-process & Extract features
    3. Build reference model for each user
end procedure
procedure OPERATION PHASE
    1. Capture ECG records
    2. Pre-process & Extract features
    3. Pass/discard ECG samples (similarity module)
    4. Authenticate samples (learner module)
    5. Update the learner (if necessary).
end procedure
```

## 3.5 System description

The general architecture of an ECG-based authentication system is displayed in Figure 3.2. Firstly, in this proposal we assume that the cardiac signal is acquired by an IMD (e.g., a pacemaker or an implantable cardioverter defibrillator), or perhaps by external sensors attached to the body of the individual. Once ECG signals are recorded, they need to be preprocessed before feature extraction. To do this the ECG signal is split into

**Figure 3.2:** General structure of an ECG-based authentication system.

time windows and, for each window, a set of numerical features are extracted. Then, the similarity module filters samples discarding those that do not seem to come from the user. Finally, the samples are classified using a classifier such as a decision tree, a support vector machine, a nearest neighbor algorithm, and so on. In fact, nearly all data mining algorithms can be tailored for coping with the data stream problem.

More details about each component of the ECG-based authentication system (see Figure 3.2) are explained in the following sections.

## 3.5.1 Dataset and pre-processing

Pacemakers and implantable cardioverter defibrillators are the most extended class of implantable devices (the first pacemakers date from the early 1950s [5]). An electrocardiogram (ECG) represents the electrical activity of the heart during a period of time. In particular, an ECG chart is composed of a set of waves: P, Q, R, S and T [126].

The ECG records are cleaned as the first step. For that, the zero-frequency component (DC bias) is eliminated and then the records are passed through a pass-band filter. An entire-raw ECG signal from a user $U_j$ is divided into windows of $L$ seconds:

$$ECG^{\mathcal{U}_j} = \{ECG^{\mathcal{U}_j}_{w(1)}, ECG^{\mathcal{U}_j}_{w(2)}, \ldots, ECG^{\mathcal{U}_j}_{w(N)}\} \tag{3.1}$$

where $N >> 1$. Subsequently each $ECG^{\mathcal{U}_j}_{w(i)}$ is the input of the feature extraction module.

## 3.5.2 Feature extraction and similarity module

Fiducial and non-fiducial approaches can be followed for extracting features of a physiological signal. Fiducial-based approaches are those in which characteristic points (e.g., Q, R and S peaks in ECG signals [69, 126]) in the time-domain can be used for the feature extraction. On the other hand, non-fiducial approaches obtain features from a transformed domain (e.g., Fourier or wavelet [2, 92]). In terms of performance, fiducial-based and transform-based approaches achieve similar results as reported in the comparative analysis by Odinaka et al. [160]. In our particular case, we opt for avoiding any sort of manipulation of the ECG signal in the time domain. The efficiency and simplicity of the system are the main goals that justify to work in a transform domain.

In particular, in the time-domain the ECG signal is only segmented into windows —this is the minimal possible manipulation of the signal. After that, a transform (TF) is applied over each ECG window and a set of $M$ coefficients is obtained:

$$F_{w(i)}^{\mathcal{U}_j} = TF(ECG_{w(i)}^{\mathcal{U}_j}) = \{f_{w(i)}^{\mathcal{U}_j}(0), f_{w(i)}^{\mathcal{U}_j}(1), \dots f_{w(i)}^{\mathcal{U}_j}(M)\} \tag{3.2}$$

Each of these feature vectors $F_{w(i)}^{\mathcal{U}_j}$ is passed through the similarity module which discards bad ECG samples, that is, samples which are considered to be too far from the reference model (outliers). The reasoning behind this is that the learner only analyzes "good" feature vectors and there is a benefit in the performance of the system in comparison with the obtained without this filtering.

Each user is responsible for the similarity checking module. For that, the user computes a reference matrix, which is called the reference module. To do so, in the set-up phase, the ECG signal is observed during a $T_R$ time interval (e.g., half an hour of continuous cardiac signal motorization). A matrix of $N$ average vectors is computed. Each of these vectors ($Y_i$ where $i = \{1, \dots, N\}$) represents an average value of ECG windows ($L$ seconds) in the Hadamard domain:

$$\overline{Y_i} = \frac{1}{T_R/(L \times N)} \cdot \sum_{q=1}^{\frac{T_R}{L \times N} \times i} F_{w(q + \frac{T_R}{L \times N} \times (i-1))}^{\mathcal{U}_j} \tag{3.3}$$

Then, the similarity module works in the following way. A set of $N$ new observed ECG windows ($F_{w(i^*)}^{\mathcal{U}_j}$, where $i^* = \{1, \dots, N\}$) are discarded or not depending on their similarity to the user's reference model. We use the Pearson's linear correlation coefficient ($corr$) to measure similarity between two matrices. Other similarity metrics could be used but we chose this due to its invariant behaviour to scale and shift changes. Mathematically, the module is described by the following equation:

$$\begin{cases} \text{Discard ECG samples} & \text{If } \left| corr\left( \begin{bmatrix} \overline{Y_1} \\ \overline{Y_2} \\ \cdots \\ \overline{Y_N} \end{bmatrix}, \begin{bmatrix} F^{\mathcal{U}_j}_{w(1^*)} \\ F^{\mathcal{U}_j}_{w(2^*)} \\ \cdots \\ F^{\mathcal{U}_j}_{w(N^*)} \end{bmatrix} \right) \right| < \delta \\ \text{Transmit ECG samples} & \text{Otherwise} \\ \left(\begin{bmatrix} F^{\mathcal{U}_j}_{w(1^*)} & F^{\mathcal{U}_j}_{w(2^*)} & \cdots & F^{\mathcal{U}_j}_{w(N^*)} \end{bmatrix}\right) & \\ \text{to the learner} & \end{cases} \tag{3.4}$$

where the parameter $\delta$ is tuned through experimentation.

Finally, the ECG streams are sent to the learner. We have evaluated two approaches: 1) buffered solution; 2) unbuffered solution. In the former, each ECG stream represents an average value of feature vectors (Hadamard domain) during an observation period $T_O$:

$$\overline{F}^{\mathcal{U}_j}_{w(i)} = \frac{1}{T_O/L} \cdot \sum_{q=1}^{\frac{T_O}{L} \times i} F^{\mathcal{U}_j}_{w(q + \frac{T_O}{L} \times (i-1))} \tag{3.5}$$

The unbuffered approach is the most demanding scenario as each ECG stream represents the feature vector of an ECG window, i.e., $F^{\mathcal{U}_j}_{w(i)}$.

Figure 3.3 depicts the creation of ECG streams for both proposed approaches. Also, considering that samples of different users can be received at different time and thus no order is expected, an illustrative example of several samples of a data stream of two users ($\{j, j^*\}$) is shown below:

$$\begin{cases} \begin{bmatrix} \overline{F}^{\mathcal{U}_j}_{w(i)} \overline{F}^{\mathcal{U}_j^*}_{w(i)} \overline{F}^{\mathcal{U}_j^*}_{w(i+1)} \cdots \overline{F}^{\mathcal{U}_j}_{w(N)} \overline{F}^{\mathcal{U}_j}_{w(N)} \end{bmatrix} & \text{Buffered approach} \\ \begin{bmatrix} F^{\mathcal{U}_j}_{w(i)}, F^{\mathcal{U}_j}_{w(i+1)}, F^{\mathcal{U}_j^*}_{w(i)} \cdots F^{\mathcal{U}_j}_{w(N)}, F^{\mathcal{U}_j}_{w(N)} \end{bmatrix} & \text{Unbuffered approach} \end{cases} \tag{3.6}$$

### 3.5.3 Learner

As introduced in Section 3.3, a wide set of methods (e.g., decision trees, bayesian methods, lazy, ensemble, etc.) can be used for the classification problem. Regardless of the used algorithm, a relevant aspect is how data is treated. Two approaches have been considered depending on whether the data is acquired in a continuous way or not and thus used in real-time or no real-time applications. In a real-time application in which cardiac records arrive continuously in a non-predefined order, an on-line analysis is used

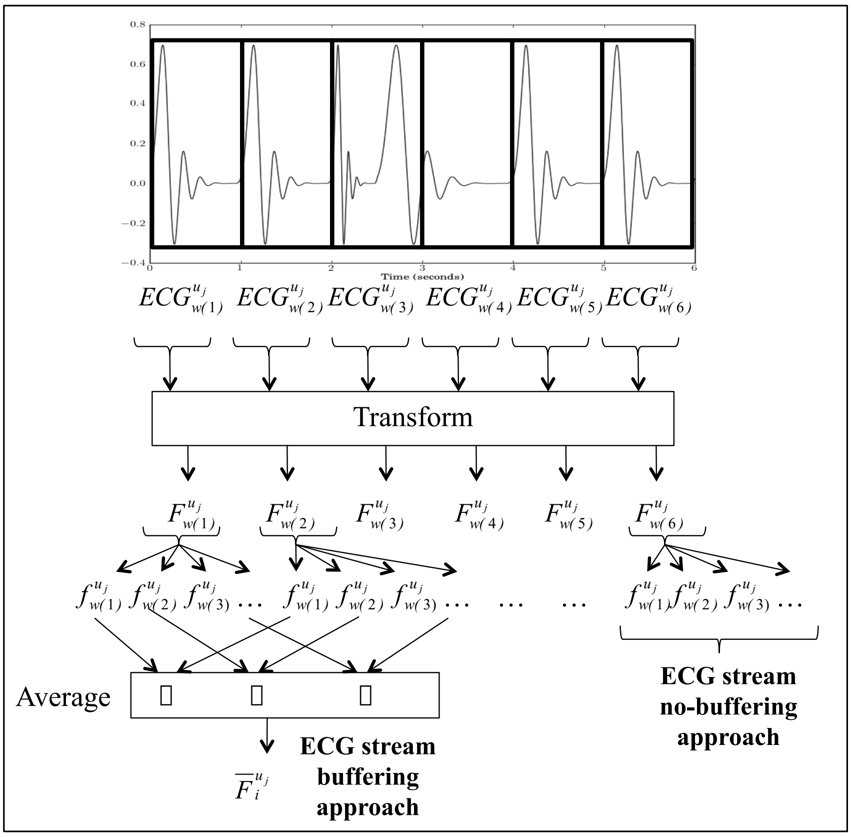
Time (seconds)
$ECG^{u_j}_{w(1)}$ $ECG^{u_j}_{w(2)}$ $ECG^{u_j}_{w(3)}$ $ECG^{u_j}_{w(4)}$ $ECG^{u_j}_{w(5)}$ $ECG^{u_j}_{w(6)}$
Transform
$F^{u_j}_{w(1)}$ $F^{u_j}_{w(2)}$ $F^{u_j}_{w(3)}$ $F^{u_j}_{w(4)}$ $F^{u_j}_{w(5)}$ $F^{u_j}_{w(6)}$
$f^{u_j}_{w(1)}$ $f^{u_j}_{w(2)}$ $f^{u_j}_{w(3)}$ …
Average
$\overline{F}^{u_j}_i$
ECG stream buffering approach
ECG stream no-buffering approach

**Figure 3.3:** ECG stream samples of buffered and unbuffered approaches.

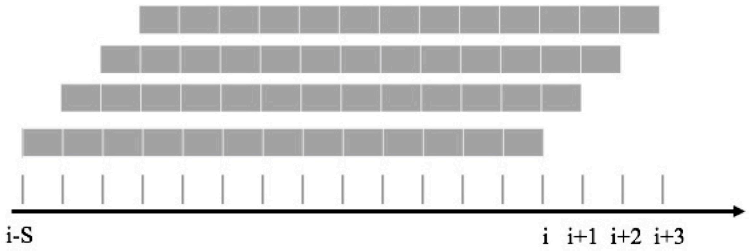

**Figure 3.4:** Example of a sliding window strategy.

and ECG streams are evaluated by interleaving testing and training (i.e., prequential evaluation) and following a sliding window strategy in which the size of the window is fixed and the buffer keeps the newest instances. Similar to the first-in, first-out data structures [70], and illustrated in Fig. 3.4, whenever a new instance is inserted into the window, another instance $j - S$ is forgotten —$S$ represents the window size. In particular, each new instance is used to test the model prior being used for training. Regarding the generation of data streams, buffered and unbuffered approaches are considered. Nonetheless, in a non real-time application, referred to as batch setting, the dataset can be split intro training and testing and there is not memory restrictions.

In terms of security, non real-time applications correspond to a NCA system. On the other hand, the analysis in real-time of the ECG data streams (user credentials) conforms with the requirements of a CA system.

# 3.6 Experimental validation

Established parameters and results achieved after experimentation are presented in the following sections.

## 3.6.1 Experimental settings

Table 3.1 provides the experimental setting used to validate our approach. The experiments were performed using the recordings of 10 individuals from the MIT-BIH Normal Sinus Rhythm Database [72]. The individuals under study do not show any relevant medical problem and were observed during a long time period. Besides, Table 3.1 provides a brief motivation for each choice of values.

The Walsh-Hadamard transform (HT) is the chosen transformation in our system. The HT performs a projection of a signal onto a set of square

| **Parameters** | **Value** | **Justification** |
|---|---|---|
| Pass-band filter | 0.67 Hz and 45 Hz | Using this filter, the main sources of noise such as the respiration noise or the power line noise are canceled |
| $N$ | 3 | As a trade-off between efficiency and be able to check the variability of the ECG signal |
| $\delta$ | $10^{-1}$ | A relative-low correlation threshold |
| $L$ | 2 seconds | To guarantee the observation of 2 or 3 heart beats depending on how fast each individual is beating |
| $M$ | 256 | Number of coefficients needed to keep the ECG information at a low level (e.g., PQRST waves) |
| $T_O$ | 3 minutes | Observation period to guarantee the stability of the ECG signal |
| $T_R$ | 30 minutes | Time-interval needed to characterize the “common” samples of a user. |
| K-NN | $K = 1$ | Greater values of parameter $K$ do not offer higher performances and increase system complexity |
| Max. Num. instances in memory in CA approach ($S$) | 10% of the tested dataset | Trade-off between memory efficiency and system performance |

**Table 3.1:** Established parameters

waves, called Walsh functions (WAL). Mathematically, the forward and inverse HT of a signal $x(t)$ of length $W$ are defined as

$$y_n = \frac{1}{W}\sum_{i=0}^{M-1} x_i WAL(n,i), n = 1,2,\ldots,W-1 \tag{3.7}$$

$$x_i = \frac{1}{W}\sum_{i=0}^{M-1} y_i WAL(n,i), n = 1,2,\ldots,W-1 \tag{3.8}$$

Using the HT is justified by two main reasons. On the one hand, this transform is computationally efficient as it just consists of a matrix multiplication (that of the signal and the Walsh matrix). On the other hand, it has the ability of compressing the input signal, with the majority of the signal information being kept on the lower coefficients in the transformed domain. Therefore HT is efficient is terms of computation and memory requirements. Note here that the usage of other transforms (e.g., Fourier or Wavelet) were evaluated and discarded, mainly, due to their complexities in terms of computational requirements.

As for the learner, the K-Nearest Neighbour (K-NN) is the algorithm used. In the NCA setting, the dataset is divided into training and testing—60/40 and 80/20 are the splits commonly used in this area [231]. In the CA analysis, a tailored K-NN is employed as the learner, which uses a buffer memory to keep a small portion of the instances (training ones). For updating, this buffer follows a sliding window strategy with a First-In-First-Out (FIFO) rule. We refer the reader to [70] for a detailed introduction to data streams and drift concept.

The reasoning of using this learner is twofold: 1) efficiency; and 2) simplicity [3]. Regarding efficiency, a K-NN often outperforms more complex learners [231]. In relation to simplicity, a K-NN does not require complex computations. In detail, new samples are classified taking into account the class to which a set of training samples (N nearest instances) belong. Although more complex learners could have been used, we consider the K-NN the most appropriated since it offers a high performance and its simplicity facilitates the implementation of the system in portable devices with constrained resources.

### 3.6.2 Results

The main goal of the system is to achieve a high performance in the identification of the users enrolled in the system. Two approaches have been evaluated depending whether the data is sent o not in a continuous fashion to the learner (i.e., core of the CA system):

**Non-Continuous Authentication (NCA)** A data mining approach makes sense when we deal with non-real time applications and there is not severe memory restrictions. During a first phase (training), data of all the enrolled users is recorded and stored in memory. The classifier is then trained using these samples. After that, credentials (ECG streams) of the users are checked (testing phase)—for instance, user credentials are verified each time she/he attempts to unlock the touchscreen of her/his smartphone. Note that, motivated by the need to achieve high performance, the buffered approach is applied in our experimentation.

In connection to the application scenario described in Sect. 3.4, the credentials of each controller would be checked when s/he logs on the system (e.g., whenever her/his computer is turned on).

**Continuous authentication (CA)** Classical approaches are not feasible when data is provided in a continuous way and memory restrictions exist. The use of an on-line analysis approach is much more suitable for processing data streams. Tools like Massive Online Analysis (MOA) [14], VFML [100] and RapidMiner [95] are commonly used for mining data streams.

Following the scenario introduced in Sect. 3.4, the credentials of each controller would be verified at regular time intervals. In the unbuffered approach there is only a distance of few seconds between intervals, and this distance considerably increases to hundreds of seconds in the buffered approach. Accordingly, we have evaluated both approaches in Sections 3.6.4.1 (buffered approach) and 3.6.4.2 (unbuffered approach).

### 3.6.3 Non-Continuous Authentication (NCA)

We have a population of individuals and average feature vectors have been acquired at regular intervals. For simplicity, we use regular intervals in our experiments. This is not a limitation and non-regular intervals might be also employed. In our experimentation, the population size is set to 10 and each individual was sensed during a period of 11 hours. To assess the impact of the training dataset, the performance of the system has been evaluated for different training sizes. Figure 3.5 shows the accuracy

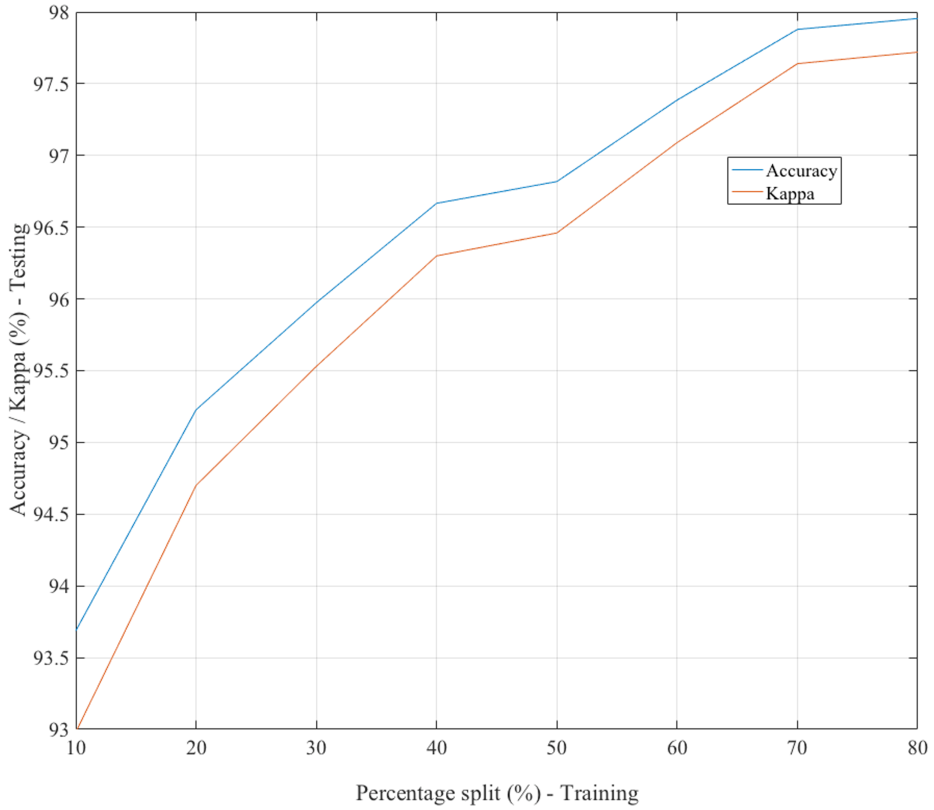

**Figure 3.5:** Assessment of the training size for a small population (NCA).

(correctly classified instances) and kappa statistic for several values of the training set. For the hard case (i.e., 10% of the whole dataset or, in other words, each individual is observed during around 1 hour) the accuracy is over 93.5%. Therefore, the system performs well even when the training phase is set to minimal values. When we use common values (60% or 80% [231]) for the training set, the performance is almost perfect (97.4% and 97.9%). The value of kappa, which is greater than 0.81 for all the training sizes, shows a perfect agreement [128]—that is, the influence of "random guessing" is minimal.

To guarantee the robustness of our results, we have also tested the classifier using a 10-fold cross-validation. The accuracy and kappa statistic are 97.90% and 97.68%, respectively. Apart from showing a significantly high True Positive Rate (97.9%), as expected from a good identification system, the weighted average False Positive Rate is extremely low (0.2%). The detailed accuracy by class and the confusion matrix are summarized in Tables 3.2 and 3.3, respectively. All in all, the metrics indicate that the system is very close to an ideal identification system.

| **Class** | **TP Rate** | **FP Rate** | **Precision** | **Recall** | **F-Measure** | **ROC Area** |
|---|---|---|---|---|---|---|
| 1 | 0.995 | 0.002 | 0.986 | 0.995 | 0.991 | 0.997 |
| 2 | 0.982 | 0.001 | 0.995 | 0.982 | 0.989 | 0.991 |
| 3 | 0.936 | 0.007 | 0.936 | 0.936 | 0.936 | 0.965 |
| 4 | 0.968 | 0.004 | 0.964 | 0.968 | 0.966 | 0.982 |
| 5 | 0.955 | 0.005 | 0.955 | 0.955 | 0.955 | 0.975 |
| 6 | 1.000 | 0.000 | 1.000 | 1.000 | 1.000 | 1.000 |
| 7 | 0.982 | 0.001 | 0.995 | 0.982 | 0.989 | 0.991 |
| 8 | 0.991 | 0.001 | 0.991 | 0.991 | 0.991 | 0.995 |
| 9 | 0.986 | 0.002 | 0.982 | 0.986 | 0.984 | 0.992 |
| 10 | 0.995 | 0.002 | 0.986 | 0.995 | 0.991 | 0.997 |
| **Weighted Avg.** | **0.979** | **0.002** | **0.979** | **0.979** | **0.979** | **0.988** |

**Table 3.2:** Accuracy by class —NCA setting.

|  | S1 | S2 | S3 | S4 | S5 | S6 | S7 | S8 | S9 | S10 |
|---|---|---|---|---|---|---|---|---|---|---|
| S1 | 219 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| S2 | 0 | 216 | 0 | 2 | 0 | 0 | 0 | 0 | 0 | 2 |
| S3 | 2 | 0 | 206 | 3 | 6 | 0 | 0 | 2 | 1 | 0 |
| S4 | 0 | 0 | 4 | 213 | 3 | 0 | 0 | 0 | 0 | 0 |
| S5 | 0 | 1 | 6 | 3 | 210 | 0 | 0 | 0 | 0 | 0 |
| S6 | 0 | 0 | 0 | 0 | 0 | 220 | 0 | 0 | 0 | 0 |
| S7 | 1 | 0 | 0 | 0 | 1 | 0 | 216 | 0 | 2 | 0 |
| S8 | 0 | 0 | 2 | 0 | 0 | 0 | 0 | 218 | 0 | 0 |
| S9 | 0 | 0 | 1 | 0 | 0 | 0 | 1 | 0 | 217 | 1 |
| S10 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 219 |

**Table 3.3:** Confusion matrix —NCA setting

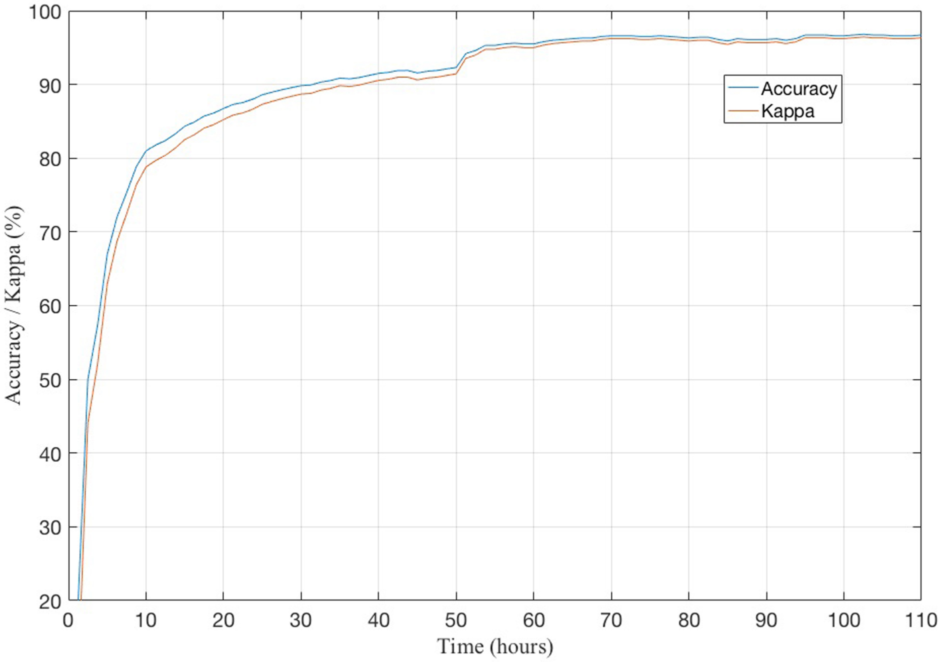

**Figure 3.6:** System performance: CA (buffered approach).

### 3.6.4 Continuous Authentication (CA)

The buffered and unbuffered approaches have been tested —see Section 3.5 for a detailed explanation in the generation of the used ECG streams.

#### 3.6.4.1 CA: buffered approach

We have tested the performance of the system for a population of individuals. According to the proposed use case (Section 3.4), assume that credentials (ECG streams) of controllers working in the same room of the tower are checked by a central unit in a continuous fashion at long-separated intervals. In this context, each subject generates an ECG data stream in a continuous way during a long period of time (i.e., 11 hours per individual in our experiments). Each sample of the stream represents an average value in the Hadamard domain, as explained in Section 3.5.2 (see Equation 3.5 for details). The population size has been set to 10 as in the NCA setting.

In Figure 3.6 we can see the evolution of the accuracy over the time using a prequential evaluation. The learner employed is a nearest neighbor (i.e., $K$-$NN$ with $K$ =1), as in the NCA setting but with a sliding window (maximum number of instances stored in memory) of reduced dimensions. In our experiments, 10% of the total instances are kept in memory. Having overcome the penalty of the first instances, the system exceeds the threshold of a 90% and swiftly stabilizes around 96% of correctly classified instances. Similarly, the kappa statistic rapidly exceeds the 80% threshold, approaching an almost perfect classifier performance.

#### 3.6.4.2 CA: unbuffered approach

We have assessed the system when the ECG streams are generated in a continuous way and at a high speed rate. Again, based on the proposed use case (Section 3.4), imagine a controller which has a cardiac problem and the authentication should be permanently to ensure the proper functioning of the system.

In the previously described scenario, we have a unique legitimate individual. Therefore, the system has to distinguish between two classes. That is, the data streams belong to the legitimate user (IMD holder in our example) or to any other unauthorized /fraudulent user (the attacker, in general terms). We have tested this setting for each one of the 10 individuals of the buffered approach. Therefore, in each experiment there is a legitimate user and the other ones are categorized within a unique class (fraudulent user). The performance (accuracy) for each of these aforementioned experiments in summarized in Figure 3.7. The accuracy is on average around 80%, with a standard deviation of 3.7%. Therefore, the system works well —in most cases, authorized users are correctly distinguished from intruders— and is particular well-suited to cope with the slight changes in the ECG data streams.

To assess the influence of whether the ECG streams are buffered or not, we have tested this setting using both approaches. In Tables 3.4 and 3.5 we show the obtained results. In terms of accuracy, the buffered approach offers a benefit of around 15% in comparison with the unbuffered approach. The Kappa statistic points out how the performance of the system switches from “substantial” to “almost perfect” accuracy when we move from the unbuffered to the buffered approach. Apart from performance metrics, the use of one approach or the other depends on the requirements demanded by the real-time application in question. The determining factor is the rate at which the ECG streams are examined. In the unbuffered approach, the ECG streams are provided almost instantly (i.e., intervals of two seconds). In contrast, only 20 examples/hour is the sample rate used in the buffered approach (i.e., intervals of 3 minutes). Therefore, the particular application of the system will driven the used option.

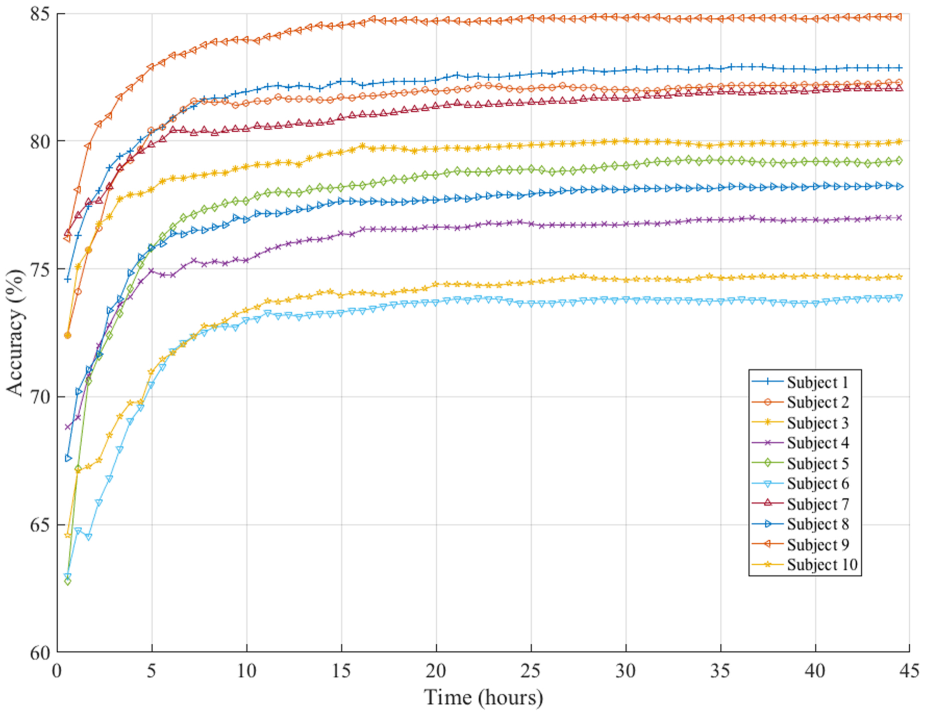

**Figure 3.7:** System performance: CA (unbuffered approach).

| Subject | Average Accuracy | Average Kappa |
|---|---|---|
| S1 | 81.98 | 63.95 |
| S2 | 81.39 | 62.79 |
| S3 | 79.28 | 58.57 |
| S4 | 76.00 | 52.05 |
| S5 | 77.81 | 55.58 |
| S6 | 72.75 | 45.49 |
| S7 | 81.02 | 62.04 |
| S8 | 77.17 | 54.35 |
| S9 | 84.12 | 68.18 |
| S10 | 73.51 | 47.12 |
| Average | 78.50 | 57.01 |

**Table 3.4:** CA: Unbuffered Approach (two classes).

| Subject | Average Accuracy | Average Kappa |
|---|---|---|
| S1 | 96.80 | 93.59 |
| S2 | 95.10 | 90.19 |
| S3 | 91.32 | 82.50 |
| S4 | 97.34 | 94.68 |
| S5 | 94.06 | 88.13 |
| S6 | 92.09 | 84.09 |
| S7 | 94.80 | 89.62 |
| S8 | 94.35 | 88.69 |
| S9 | 97.95 | 95.90 |
| S10 | 94.04 | 88.05 |
| Average | 94.79 | 89.54 |

**Table 3.5:** CA: Buffered Approach (two classes).

## 3.7 Discussion

Although some authors have already explored the problem of continuous authentication with cardiac signals (e.g., ECG [78] and PPG [17], the used datasets are made up of records with length of only a a few minutes), this is the first time that ECG records are interpreted and processed as data streams. In our opinion, a data stream approach fits perfectly the problem of CA, particularly in the case of ECG signals —and, more generally, physiological signals with a slight variability and a theoretical infinite length. We have considered the typical assumptions for classification in the DSM setting [15]: 1) each sample has a fixed number of attributes that are less than several hundreds; 2) the number of classes is limited and small (in our experiments, ten classes are considered at maximum); 3) we assume that the learner has a small memory; the size of the training dataset is larger than the available memory; and finally, 4) the speed rate of processing each sample is moderate high (the precise value is conditioned to the device that supports on-board the learner).

Data stream algorithms have the potential to deal with potential infinite amount of data. Regarding physiological signals, as far as we know, recordings are taken during a maximum period of 24 hours in the best case [171]. The execution time of the algorithm used scales linearly with the number of examples. In our experimental setting, the learner consumes several tens of milliseconds per sample using a Quad Core 2.7 GHz Intel Core i5 with 16GB of RAM. Using this value (or the equivalent if different equipment is used), an upper bound of the time necessary for processing an arbitrary number of examples may be computed.

Although important variations on ECG streams only occur after 5 years observation period [28], we can find slight variations from time to time — that is, data is not stationary. This is often referred as concept drift. To dealt with this, old instances should become irrelevant to characterize the current state of the system and this information would have to be forgotten by the learner. The interested reader can consult [70] for a detailed explanation of the main existing approaches in the literature. In our particular case, as explained in Sect. 3.5.3, we keep only the most recent samples in memory and the memory size is fixed —sliding window strategy.

Aside from using a limited memory, we can benefit from drift detection mechanisms that reset the learner model and trigger the learning of a new one when a significant change is detected. We have tested two well-known methods: Drift Detection Method (DDM) and Early Drift Dection Method

| Approach | Accuracy (average value) |
|---|---|
| KNN | 84.1200 ± 1.5095 |
| KNN-DDM | 84.8000 ±1.6290 |
| KNN-EDDM | 84.7300 ±1.7921 |
| KNN-DDM (with artificial noise) | 78.2500 ± 1.7060 |
| KNN-EDDM (with artificial noise) | 78.1600 ± 1.9831 |

**Table 3.6:** Average performance: CA (unbuffered approach) with drift detection

(EDDM) [76]. In a nutshell, DDM is based on monitoring the number of errors produced by the learner during prediction —errors are modelled by a binomial distribution. DDM performs well to detect abrupt changes and not very slow gradual changes. EEDM was proposed with the aim of improving the detection of gradual changes and keeping a good performance with abrupt changes. Instead of considering only the number of errors as in DDM, it also takes into account the distance (number of examples) between two classification errors.

The performance of the two aforesaid methods has been evaluated with one of the subjects of the CA (unbuffered approach) setting which is our more demanding scenario. The subject 9 has been selected for this experimentation without prejudice to the generality in the results. More precisely, DDM and EDDM algorithms are used as a wrapper on the KNN learner. We have tested two scenarios: 1) the original data stream; 2) artificial noise has been added to the original data —10% and 5% are the fractions of attributes values and class labels that have been disturbed, respectively. Figure 3.8 displays the obtained results and Table 3.6 summarizes the average values. In both cases, DDM and EDDM converge to the same accuracy values which points out that the gradual changes in the ECG records are not very slow. In terms of performance, the KNN with drift detection marginally improves our previous results of only using a KNN with sliding window. In addition, drift detection methods work well even when the data streams are quite noisy —the performance only suffers a brief dip. Note that we have overstated the used example since the noise remains during the whole data stream and often it is intermittent.

Finally, a key-aspect in the processing of cardiac signals is the time period during which the ECG is observed. This aspect is examined at the end of Section 3.6.4.2—see Tables 3.4 and 3.5 for details. In the buffered approach, each stream is linked with the observation of the ECG during a moderate long time period with the extra benefit of achieving a very high performance. In the unbuffered approach, the sending of the examples to

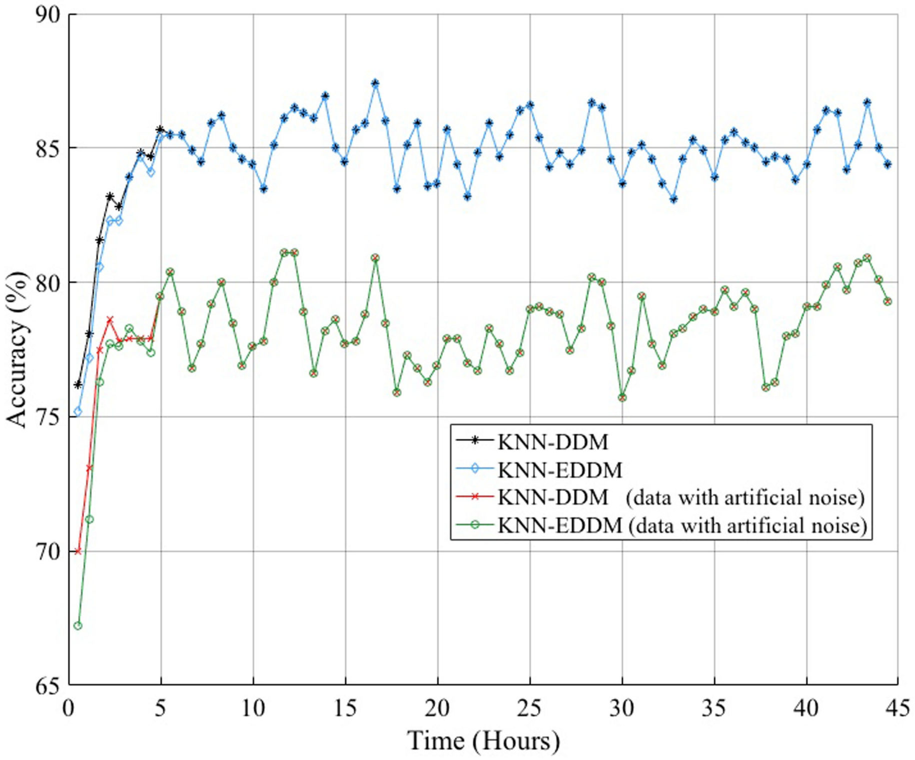

**Figure 3.8:** System performance: CA (unbuffered approach) with drift detection.

the learner is almost instantaneous with the penalty of a slightly degradation of the performance in comparison with the buffered approach. The choice of one approach or another would be conditioned by the processing speed rate demanded by the learner. In our particular case (a CA system), we have the possibility to check the credentials of an individual almost instantaneously (each two seconds) or just remain patient and proceed with the verification once every three minutes.

# 4

# ACIMD: A biometric distance bounding protocol

## 4.1 Introduction

Significant advances have been made in the healthcare domain over the past years. In particular, providing new communication capabilities to medical systems and devices benefits all actors [189, 190, 49]. Users can monitor their health status without interfering with their daily activities, the medical staff has fast remote access to medical data and can also quickly re-program these devices remotely. These new communication capabilities also reduce the global costs of healthcare operations [29].

Implantable medical devices (IMDs) is such an example of device having remote communication capabilities, including the access to telemetry data [143]. IMDs are electronic devices implanted within the body to treat a medical condition, to monitor a physiological organ and to actuate when necessary [86]. IMDs can be categorised in four main classes [29]: cardiac implanted devices (pacemakers and implantable cardioverter defibrillators), neurostimulators, drug delivery systems and biosensors. An illustrative example of a generic neurostimulator and an insulin pump is displayed in Fig. 4.1.

These monitoring and actuating operations are usually relayed by a nearby device communicating directly over the radio channel with the IMD and called Programmer. Two properties have to be achieved to enable this communication between IMD and Programmer: (1) we must ensure that the Programmer is authorized to interact with the implant: *access control*, (2) the data at stake must only be accessible to the two entities communicating: *confidentiality*. In this paper, we focus on the provision of (1): *access control*. Interested readers may refer to [29] for a survey on confidentiality issues.

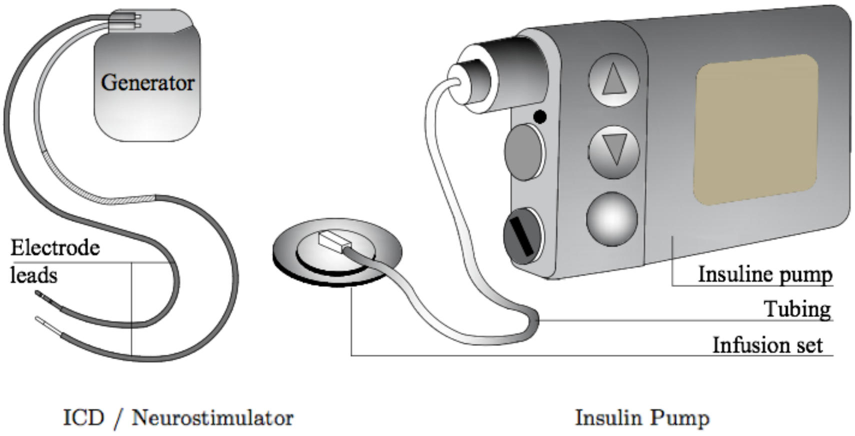

**Figure 4.1:** Example of IMDs

### 4.1.1 Access control

Access control mechanisms guarantee that the requester has the necessary privileges to execute a particular action. The existing solutions are very diverse, including those based on access control lists and certificates [64] or the ones based on biometrics [90, 186]. Other authors leave the access control responsibility to an internal [86] or external device [51, 73].

A particular branch in access control for IMDs is based on measuring the distance between this device and the Programmer. This technique is referred to as *distance bounding protocols* and require the following three definitions [8]:

**Definition 4.1.1** (Authentication)**.** One party is assured of both the identity of a second party and her presence at the time of the protocol execution.

**Definition 4.1.2** (Distance checking)**.** One party ($P$) is assured of the distance (or a property derived from this) to a second party ($V$) at some point of the protocol execution. The area in which $P$ is considered to be close enough to $V$ is called Neighbourhood Area (NA).

**Definition 4.1.3** (Distance bounding)**.** It combines identity verification (authentication) and distance checking. Regarding the distance between $P$ and $V$, an upper-bound limit is often used.

Distance bounding protocols were proposed by Brands and Chaum [20]. They were intended to cope with *mafia fraud* attacks, which are based on the relaying of messages between dishonest entities [52]. In particular, the attack consists in a man-in-the-middle attack between a honest verifier ($V$; e.g, IMD) and a legitimate prover ($P$; e.g., Programmer). The adversary is made up of two entities: a rogue prover ($\overline{P}$) and a rogue verifier ($\overline{V}$). $\overline{V}$ interacts with $P$ and $\overline{P}$ communicates with $V$, respectively. In addition, rogue entities (i.e., $\{\overline{P}, \overline{V}\}$) forward the messages received from the legitimate entities (i.e., $\{V, P\}$) between each other.

Distance bounding protocols guarantee to the IMD that the connected Programmer is in its Neighbourhood Area (NA) and is not a distant third party. For this purpose, these protocols check the delay between sending a bit and receiving its corresponding response bit (fast-rapid-bit exchange, see Fig. 4.2). This can be achieved through different means. The distance can be calculated based on the received signal energy/power (RSS) [60]. This sort of solutions are not reliable whether the adversary can increase the power of the emitted signals. A distance bounding protocol is one of the most used solutions. Since the mushrooming of radio frequency identification (RFID) devices, a large numbers of distance bounding protocols have appeared in the literature [111, 161, 9]. In the context of IMDs, for instance, Rasmussen *et al.* proposed a distance bounding protocol, which uses ultrasound signals to delimit the distance [177].

Among these techniques, a promising approach is the use of a key derived from an internal (measured by the IMD) and external (recorded by the Programmer/reader) physiological signal [114]. Thus, if the same key (or two ones with only few different bits) is obtained, the proximity between both entities is assumed. In the context of cardiac IMDs, Inter-Pulse Timing (IPI) is the common solution [4, 239] within this category.

### 4.1.2 Motivation and contribution

The motivation of this paper is twofold. On the one hand, current IPI-based solutions rely upon a fiducial point (i.e., R peaks in Electrocardiogram (ECG) or Photoplethysmograph (PPG) signals). They are tied to the assumption that this feature cannot be inferred from a distant place, which has been proven to not hold [24]. This allows a malicious party to illegally gain access to the IMD from a remote location.

On the other hand, access control for IMDs must accommodate with the two operation modes of IMDs, namely *normal* and *emergency* modes. The normal operation mode is the usual one that operates while no anomaly regarding the health of the patient is detected. In contrast, the emergency mode is triggered when the user suffers from a serious medical problem (e.g., a heart attack, a hypoglycemic episode or an epileptic attack) that endangers her life. Thus, access control mechanisms must meet a trade-off between level of security and speed of the authentication such that a programmer can have fast enough access to an IMD to quickly deal with

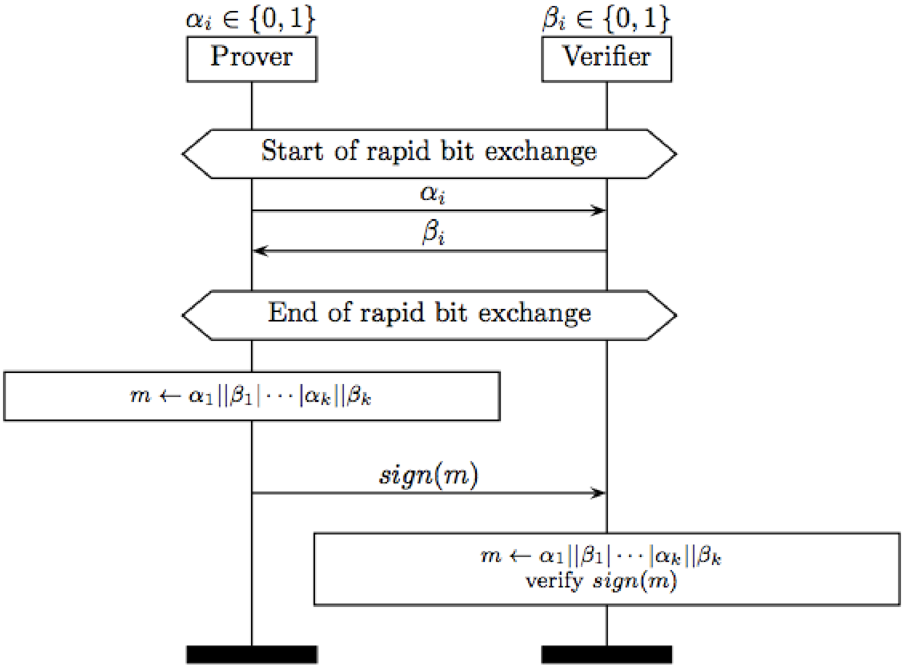
$\alpha_i \in \{0,1\}$
Prover
$\beta_i \in \{0,1\}$
Verifier
Start of rapid bit exchange
$\alpha_i$
$\beta_i$
End of rapid bit exchange
$m \leftarrow \alpha_1||\beta_1|\cdots|\alpha_k||\beta_k$
$sign(m)$
$m \leftarrow \alpha_1||\beta_1|\cdots|\alpha_k||\beta_k$
verify $sign(m)$

**Figure 4.2:** Brands and Chaum Distance Bounding Protocol [20]

emergency situations and actuate the IMD appropriately.

To address these issues, in this paper a novel distance bounding protocol (referred to as ACIMD) is proposed. ACIMD leverages the entire signal (i.e., several QRS complexes of an ECG record), thus limiting the attacker capabilities for remote acquisition. Particularly, ACIMD tests the proximity between the IMD and the Programmer by measuring the similarity between an internal and external physiological signals. Interestingly, ACIMD can work under the normal and emergency operation modes, which is beneficial for its real-world usage. ACIMD keeps computation and communication to a minimum to save battery and ease on-chip implementation. ACIMD has been tested with real ECG signals of 199 users who were recorded during a 24-hours period.

## 4.2 Methods and Materials

ACIMD is a distance bounding mechanism that ensures (1) that only an authorized Programmer is entitled to communicate with a given IMD and (2) that the Programmer is in the Neighbouring Area (NA) of the IMD. For this purpose, two main steps are carried out, namely authenticating the Programmer (Section 4.2.1) and checking its distance to the IMD (Section 4.2.2). The way in which ACIMD deals with normal and emergency

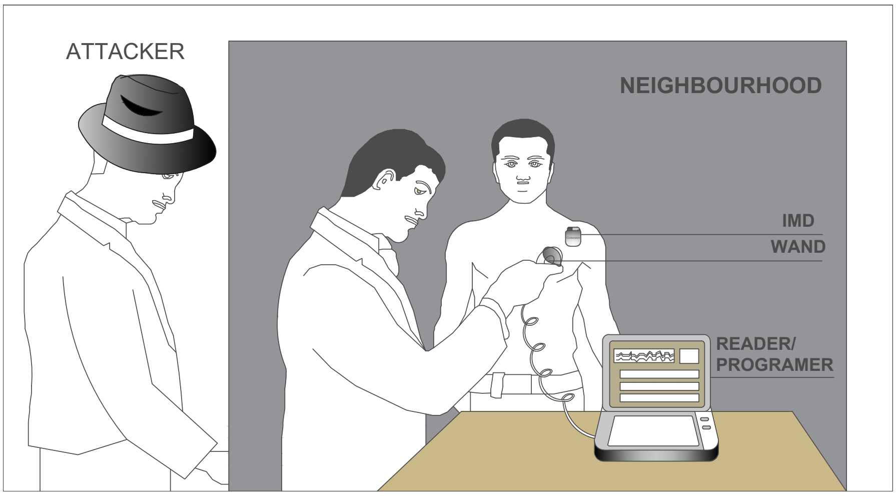

**Figure 4.3:** A typical scenario: IMD, Programmer and adversary

operation modes of IMDs is described in Section 4.2.3.

Fig. 4.3 illustrates a typical authentication scenario to facilitate the understanding of the interactions between the main entities. The IMD records an internal signal and the Programmer externally reads the same signal through a wand. The proximity between both devices is verified using our distance checking scheme. If both signals present enough similarity, the Programmer is considered to be within the NA of the implant. In contrast, any adversary is supposed to be out of NA and thus could not successfully complete the authentication.

### 4.2.1 Authentication

In order to both parties be sure on the identity of the other one, a key agreement scheme is applied [196]. Three main alternatives can be identified. On the one hand, we can assume that a pre-established key is shared between both entities. This approach raises the risk of endangering future communications if the key gets compromised or leaked to an adversary. Alternatively, as suggested in [115] or [37], a fuzzy extractor can be employed for the key generation. Nevertheless, solutions based on fiducial points like R-peaks in ECG or PPG signals are not secure from adversaries who can infer that peaks from a long distance [24].

Due to the drawbacks of the aforementioned solutions, we propose the use of a short-range and relatively secure channel for the transmission of a session key. In particular as suggested in [114] the use of photobiomodulation seems an interesting approach due to its resistance against eavesdroppers —it allows short-range communications and needs line-of-sight between the transmitter and the receiver. Photobiomodulation (or also known as Low-Level Light Therapy, LLLT) consists in the emission of light by a diode or laser in the spectral range of 600–1000 nm and at a low-power (<500 nW) [142]. Therefore, the first step of our proposed scheme consist in the exchange of the session key using LLLT. Let $ID_R$ and $ID_I$ be the identifiers of the Programmer and the IMD, $\{\cdot\}_{K_x}$ an authentication token using key $K_x$ and "||" the concatenation operation, the exchange of messages is as follows:

**Step 1:** The Programmer sends a "Wake-up" message and its identifier $ID_R$ to the IMD.

**Step 2:** The IMD replies three values: a session key $K_s$, its identifier $ID_I$ and finally the starting time $t_s$ for recording the physiological signal. This means that the first recorded-window $ECG^{(i)}_{I/R}$ starts at that particular time.

**Step 3:** During the signal acquisition phase, IMD and Programmer record ECG signals and compute $\delta$ and $\beta$, respectively (see Section 4.2.2 and Equation 4.3 for details).

**Step 4:** IMD sends to Programmer a random number $N_I$.

**Step 5:** Programmer generates a nonce $N_R$ and computes an authentication token. The authentication token is computed using $K_s$ and four input values: the nonces $\{N_R, N_I\}$, identifier $ID_R$ of the IMD, and finally $\beta$. Finally, Programmer sends $m_1$ message to IMD ($m1 = \{N_R||N_I||ID_I||\beta\}_{K_s}$).

**Step 6:** The IMD checks the correctness of the authentication token. In detail, it confirms the addressee of the message by checking the received identifier $ID_I$ and also verifies the validity of nonces $\{N_R, N_I\}$.

### 4.2.2 Distance checking

In ACIMD, the proximity of the Programmer and the IMD is assessed by comparing the ECG signals recorded by each device. The extraction of the features used for comparison rests on the wavelet transform and its coefficients. We give an introduction to the wavelet transform computation in Section 4.2.2.1, the reader is referred to [1, 23, 139] for a more detailed description. Afterwards, Section 4.2.2.2 describes ACIMD's distance checking mechanism.

#### 4.2.2.1 Wavelet Transform

The Continuous Wavelet Transform (CWT) is defined as:

$$X_f(a,b) = \frac{1}{\sqrt{a}} \int_{-\infty}^{+\infty} f(t) \cdot \psi^*(\frac{t-b}{a})dt \tag{4.1}$$

The results of the CWT are many wavelet coefficients, which are function of scale ($a$) and position ($b$), where $\{a, b\} \in \mathbb{R}$. The scale can be viewed as a compression factor and it is linked to the frequency. Low scale $a$ values (compressed wavelet) correspond to high frequency and high scale $a$ values (stretched wavelet) are equivalent to low frequency. The position factor represents the delay of the signal. $\psi$ is the basic wavelet function, the so-called mother wavelet (e.g., daubechies, biorthogonal, symlets, etc.), and the asterisk $*$ represents the operation of complex conjugate. The wavelet transform decompresses the signal into different scales with different levels of resolution by stretching (or compressing) the mother signal.

The calculation of wavelet coefficients at any value of scale ($a$) and position ($b$) is often redundant and requires a high amount of work. The analysis can be done more efficiently if scales and positions are power of two (dyadic scales and positions), which is called Discrete Wavelet Transform (DWT). In CWT, if $a = 2^m$ and $b = n \cdot 2^m$ and $\{m, n\} \in \mathbb{Z}$ we obtain the following equation:

$$X_{m,n} = \int_{-\infty}^{+\infty} f(t)[2^{-m/2} \cdot \psi(2^{-m} \cdot (t-n))]dt = \int_{-\infty}^{+\infty} f(t) \cdot \psi_{m,n}(t)dt \tag{4.2}$$

Wavelets can be calculated by iteration of filters with rescaling as described below. Two sets of coefficients are generated at each stage: approximations coefficients $Y_{j,k}$ and detailed coefficients $X_{j,k}$. In detail, these vectors are obtained by convolving $f(t)$ with a low-pass filter $h_0$ (LPF) and a high pass filter $h_1$ (HPF), followed by dyadic decimation – the signal is represented by only half the number of samples. The process at each step is summarized in Algorithm 2.

**Algorithm 2** Decomposition Algorithm

```
procedure AT LEVEL-K
    High-pass filter generates detailed coefficients: $X_{j,k}$
    Low-pass filter generates approximation coefficients : $Y_{j,k}$.
    Signals are down-sampling.
end procedure
```

#### 4.2.2.2 Mechanism description

The steps for distance checking in ACIMD are depicted in Fig. 4.4 and detailed as follows:

**Step 1:** Electrocardiogram signals are obtained from both the IMD ($ECG_I$) and the Programmer ($ECG_R$).

**Step 2:** The noise of the signals is eliminated and then ECG records are split into windows of $L_w$ seconds as further described in detail in Section 4.2.4. The $i$-th window of length $L_w$ is represented by $ECG_I^{(i)}$ when it comes from the IMD (or $ECG_R^{(i)}$ for those from the Programmer).

**Step 3:** A multi-level analysis of each window is performed using wavelet decomposition and then coefficients at level $D$ are extracted. In particular, "Daubechies-3" is the mother wavelet used in our experiments and $D$ is set to 3 —these parameters have been tuned through experimentation and taking into account the rule-of-thumb that the number of levels must be less than $log_2(L_w)$. The detailed coefficients at level 3 of the i-th $ECG_{I/R}^{(i)}$ window are represented by $X^{ECG_{I/R}^{(i)}}$. To deal with these coefficients, the values have been quantized using a dynamic quantizer with $2^8$ levels as in [186].

**Step 4:** A set of $N$ windows from the external and internal signals are used in the similarity checking module. The correlation coefficient has been the metric used for the comparison of the coefficients. The $N$ parameter is set in order to optimize the performance of the system and to minimize observation period of the signal, i.e. the time interval required for recording the internal and external signals. Mathematically,

$$\begin{aligned} S(\delta,\beta) &= S\left(X^{ECG_I^{(i)}}, X^{ECG_R^{(i)}}\right) = \\ &= corr\left(\begin{bmatrix} X^{ECG_I^{(i)}} \\ X^{ECG_I^{(i+1)}} \\ \ldots \\ X^{ECG_I^{(i+(N-1))}} \end{bmatrix}, \begin{bmatrix} X^{ECG_R^{(i)}} \\ X^{ECG_R^{(i+1)}} \\ \ldots \\ X^{ECG_R^{(i+(N-1))}} \end{bmatrix}\right) \end{aligned} \tag{4.3}$$

where $corr$ represents the correlation operation.

**Step 5:** A decision is taken based on the similarity of the signals. If both signals are considered *sufficiently close*, it means that the IMD and the Programmer are within the neighbourhood area. The proximity

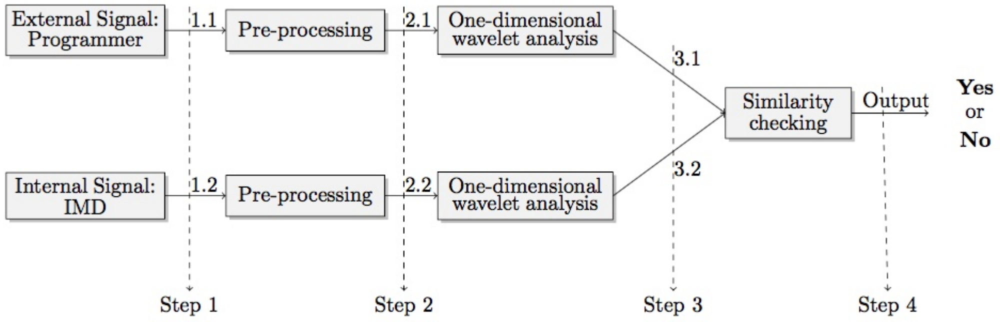

**Figure 4.4:** Distance checking mechanism

implies that the Programmer is able to record the ECG signal with all its fruitful components —the entire QRS complex is used. A threshold $\alpha$ is defined for this comparison, see Eq. 4.4.

$$\begin{cases} |S\left(X^{ECG_I^{(i)}}, X^{ECG_R^{(i)}}\right)| < \alpha & \text{Inside neighbourhood area} \\ Otherwise & \text{Outside neighbourhood area} \end{cases} \quad (4.4)$$

### 4.2.3 Normal and emergency modes of operation

ACIMD operates under normal and emergency scenarios. In a normal setting, the user keeps doing her daily routines and no restrictions of time and computation apply, apart from those intrinsic to IMDs. Thus the authentication procedure can be time-consuming to ensure the maximum level of security. On the contrary, in the emergency mode, keeping the IMD holder alive is the priority. The access to the implant should not be delayed by heavy security mechanisms. Thus, a lightweight authentication mechanism being less secure but faster can be considered.

In order to cope with these two scenarios, two modes of ACIMD are proposed.

#### 4.2.3.1 ACIMD in normal mode

In this mode, ACIMD performs the key agreement steps for authentication (recall Section 4.2.1). Moreover, distance checking procedure (recall Section 4.2.2.2) is also carried out. The scheme is depicted in Fig. 4.5.

The access control decision is based on the result of both procedures. In particular, after having the ECG signal of the holder and receiving message

$m_1$ from the Programmer, the IMD computes its answer $m_2$ as follows:

$$\begin{cases} m_2 = \{N_I||N_R||h(\beta||1)\}_{K_s} & \text{If } \perp auth \text{ and } |S(\delta,\beta)| < \alpha \\ m_2 = \{N_I||N_R||h(\beta||0)\}_{K_s} & \text{If } \perp auth \text{ and } |S(\delta,\beta)| \geq \alpha \\ m_2 = random\_value & \text{Otherwise} \end{cases}$$

where $h$ symbolizes a one-way hash function. This answer is sent to the Programmer, which verifies its correctness. In particular, if $m_2$ is valid with $h(\beta||1)$, it means that IMD and Programmer are (1) mutually authenticated and (2) within the Neigbourhood Area (NA) of IMD. If it is not the case, but $m_2$ is valid with $h(\beta||0)$, this means that the mutual authentication is successful but the reader is out of NA. Otherwise, none of these conditions are fulfilled.

#### 4.2.3.2 ACIMD in emergency mode

In emergency mode, we cannot assume that IMD and Programmer are under a controlled environment. For instance, this can be the case when the holder of the implant is in a foreign country or, for example, she is not in her corresponding referral hospital.

Therefore, in emergency mode only distance checking (Section 4.2.2.2) is applied. In Fig. 4.6 we sketch this mode of operation. Essentially, Programmer sends the ECG signal in clear to the IMD. This entity then computes the similarity with its internal signal and takes a decision following Equations 4.3 and 4.4, respectively.

Note that the security requirements are relaxed since the primary requirement becomes the speed and success of the process in order to keep the holder of the implant alive. The proposed solution is a trade-off between safety of the IMD holder and security of the system.

### 4.2.4 Dataset and Pre-processing

ACIMD has been evaluated using real physiological signals. Since ICDs and pacemakers are the most extended IMDs, electrocardiogram (ECG) signals have been used in our experimentation. In particular, cardiac signals from E-HOL-03-202-003 dataset (Telemetric and Holter ECG Warehouse of University of Rochester), are the ECG recordings used in our experiments [41]. In detail, this dataset was acquired using the SpaceLab-Burdick digital Holter recorder (SpaceLab-Burdick, Inc., Deerfield, WI) and a pseudo-orthogonal lead configuration with three electrodes $\{X, Y, Z\}$

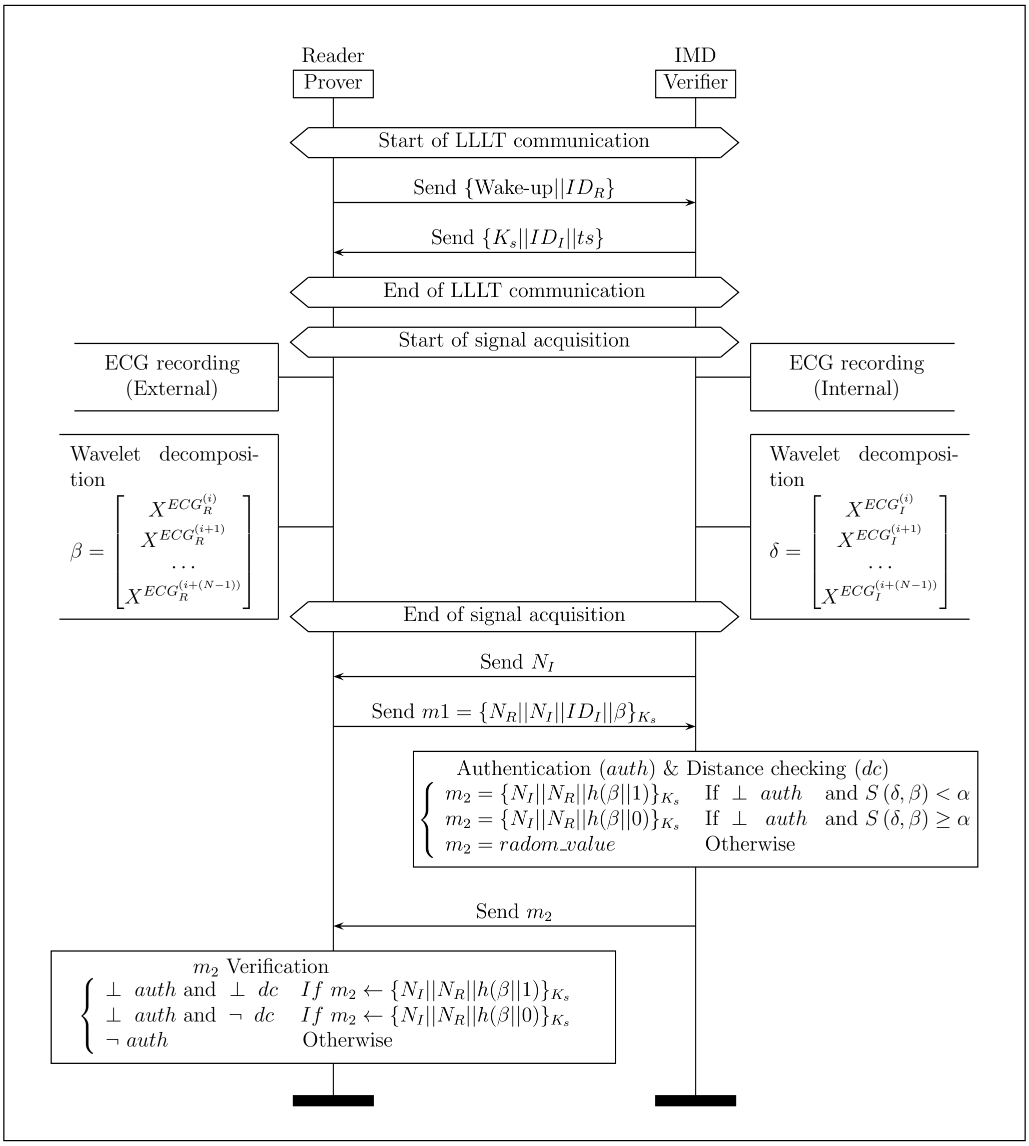
Reader
Prover
IMD
Verifier
Start of LLLT communication
Send {Wake-up$||ID_R$}
Send $\{K_s||ID_I||ts\}$
End of LLLT communication
Start of signal acquisition
ECG recording (External)
ECG recording (Internal)
Wavelet decomposition
$\beta = \begin{bmatrix} X^{ECG_R^{(i)}} \\ X^{ECG_R^{(i+1)}} \\ \dots \\ X^{ECG_R^{(i+(N-1))}} \end{bmatrix}$
Wavelet decomposition
$\delta = \begin{bmatrix} X^{ECG_I^{(i)}} \\ X^{ECG_I^{(i+1)}} \\ \dots \\ X^{ECG_I^{(i+(N-1))}} \end{bmatrix}$
End of signal acquisition
Send $N_I$
Send $m1 = \{N_R||N_I||ID_I||\beta\}_{K_s}$
Authentication $(auth)$ & Distance checking $(dc)$
$\begin{cases} m_2 = \{N_I||N_R||h(\beta||1)\}_{K_s} & \text{If } \perp auth \text{ and } S(\delta,\beta) < \alpha \\ m_2 = \{N_I||N_R||h(\beta||0)\}_{K_s} & \text{If } \perp auth \text{ and } S(\delta,\beta) \geq \alpha \\ m_2 = radom\_value & \text{Otherwise} \end{cases}$
Send $m_2$
$m_2$ Verification
$\begin{cases} \perp auth \text{ and } \perp dc & If\ m_2 \leftarrow \{N_I||N_R||h(\beta||1)\}_{K_s} \\ \perp auth \text{ and } \neg dc & If\ m_2 \leftarrow \{N_I||N_R||h(\beta||0)\}_{K_s} \\ \neg auth & \text{Otherwise} \end{cases}$

**Figure 4.5:** ACIMD in normal mode

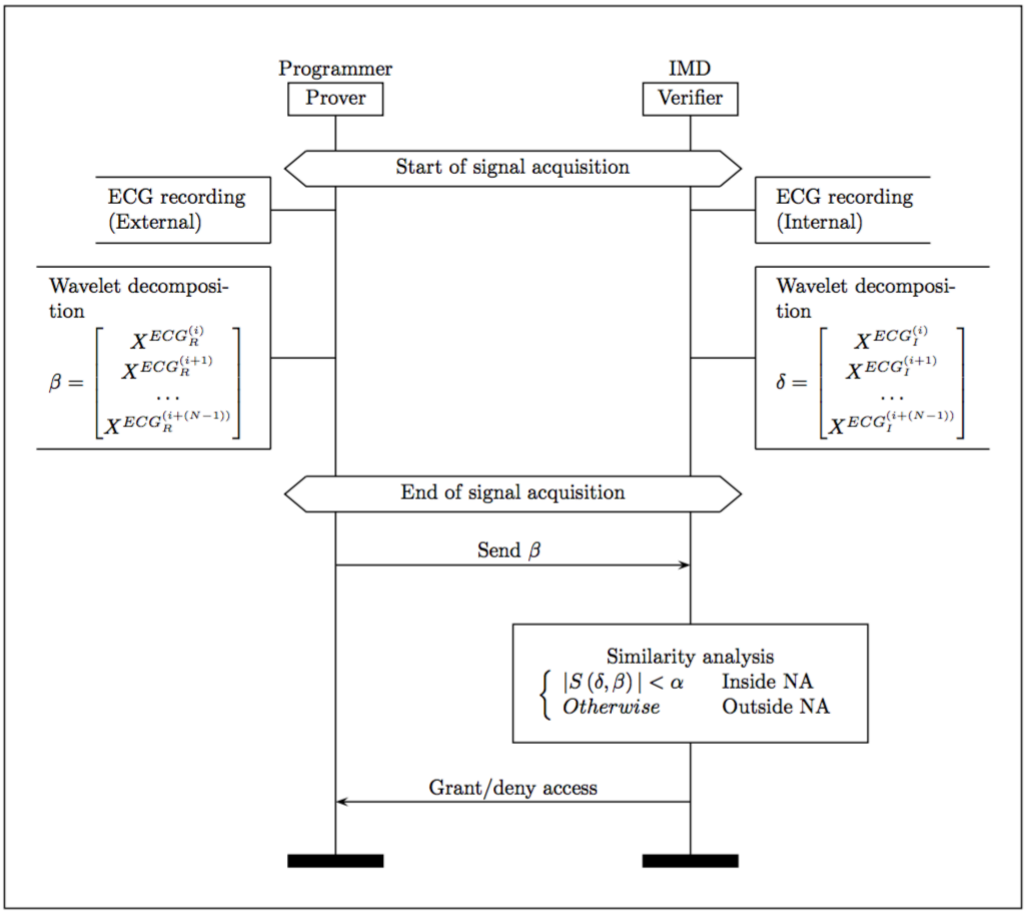
Programmer
Prover
IMD
Verifier
Start of signal acquisition
ECG recording (External)
ECG recording (Internal)
Wavelet decomposition
$\beta = \begin{bmatrix} X^{ECG_R^{(i)}} \\ X^{ECG_R^{(i+1)}} \\ \dots \\ X^{ECG_R^{(i+(N-1))}} \end{bmatrix}$
Wavelet decomposition
$\delta = \begin{bmatrix} X^{ECG_I^{(i)}} \\ X^{ECG_I^{(i+1)}} \\ \dots \\ X^{ECG_I^{(i+(N-1))}} \end{bmatrix}$
End of signal acquisition
Send $\beta$
Similarity analysis
$\begin{cases} |S(\delta,\beta)| < \alpha & \text{Inside NA} \\ Otherwise & \text{Outside NA} \end{cases}$
Grant/deny access

**Figure 4.6:** ACIMD in emergency mode

was used. The results shown in this proposal correspond to the pair of leads $\{Y, Z\}$. Thus, the lead $X$ is taken as $ECG_I$ and $Y$ is taken as $ECG_R$.

The rationale of using this dataset is four-fold. Firstly, as mentioned, cardiac implants are currently the most widespread IMD in the healthcare sector. Therefore, ECG records seem an interesting signal for our study. Secondly the dataset has a high number of individuals – 199 out of 203 have been employed since 4 had an insufficient file size. Thirdly, the recordings were taken during a long period of 24 hours. Last but not least, we can assume that the population is homogeneous (without any bias) since no important cardiac problems were detected over the subjects under study.

Before any other processing, ECG signals must be cleaned. We follow the procedure described below. We start eliminating the DC component. After that, a filter is used aiming to eliminate the respiration and the power-line source of noises. More precisely, ECG signals are passed over a pass-band filter with 0.67 Hz (lower-cut-off-frequency) and 0.45 Hz (upper-cut-off-frequency). The respiration noise is eliminated though the lower stop-band. On the other hand, the pass-band pursues to keep as much information as possible while the upper-stop band is related with the elimination of the power line noise.

Once cleaned, ECG records are split into windows of $L_w = 2$ seconds. Since a healthy individual beats between 60 and 100 times per minute, it entails that each window contains 2 or 3 heart beats. The usage of this window size is inspired on previous works in ECG identification with high accuracy rate [28, 160].

## 4.3 Results and Discussion

ACIMD has been assessed considering its three major elements – authentication, distance checking and its ability to operate on emergency scenarios. Each issue is addressed separately.

### 4.3.1 Authentication

The security of the entity authentication scheme is guaranteed by its compliance with ISO/IEC 9798-2 [19]. Assuming the use of a secure primitive and $L$ the length of session key $K_s$, the security of the protocol is upper bounded $\frac{1}{2^{2 \cdot L}}$ (cf. Section 4.2.1). Under the assumptions regarding the security of LLLT communications [114], which is used to transmit the session key, the above security upper bound holds true.

### 4.3.2 Distance checking

Considering an IMD and a Programmer being in its neighbourhood area (NA), we evaluate the accuracy of our system by computing the percentage of ECG signals recorded by each device that succeed the distance checking (Section 4.2.2). In addition, we evaluate the success rate of three considered adversary strategies:

**Definition 4.3.1** (Replay attack)**.** We define as $\mathcal{A}_R$ the advantage of an adversary to overpass the system by using signals of a previous sessions of the same subject. Mathematically,

$$p(\mathcal{A}_R) = p(|S(X^{ECG_I^{(i)}}, X^{ECG_I^{(j)}})| < \alpha) \qquad \text{where } i < j \tag{4.5}$$

**Definition 4.3.2** (Impersonation Attack)**.** We define as $\mathcal{A}_I$ the advantage of an adversary to overpass the system by using a signal captured from another subject than the holder of the IMD performing the authentication. There is no correspondence between the internal signal ($I$) recorded by the IMD and the external signal ($R'$) played by the attacker. The comparison is performed between the internal signal of a subject and the external signal of a different subject. It can be expressed as:

$$p(\mathcal{A}_I) = p(|S(X^{ECG_I^{(i)}}, X^{ECG_{R'}^{(i)}})| < \alpha) \qquad \text{where } I \not\Leftrightarrow R' \tag{4.6}$$

**Definition 4.3.3** (Random guessing)**.** We define as $\mathcal{A}_G$ the advantage of an adversary to overpass the system by random guessing. Mathematically,

$$p(\mathcal{A}_G) = p(|S(random, X^{ECG_I^{(i)}})| < \alpha) \tag{4.7}$$

Table 4.1 summarizes the accuracy of ACIMD for a normal authentication between two authorized devices (**Accuracy**) according to different $\alpha$ values. In addition, the success rate of an attack considering the three adversary advantages is provided. Four rows corresponding to four considered *configurations* are highlighted in the table. The choice of one or the other is conditioned by specific design goals.

Configuration-A is the one with highest accuracy (87.07%) but the adversary chances are a bit high (29.5% on average). Fortunately, configuration-B offers a similar accuracy (81.2%) while the success probability for the

| $\alpha$ | **Accuracy (%)** | $\mathcal{A}_R(\%)$ | $\mathcal{A}_I(\%)$ | $\mathcal{A}_G(\%)$ | |
|---|---|---|---|---|---|
| 0.050 | 87.069 | 36.083 | 35.902 | 16.574 | **Configuration-A** |
| 0.055 | 85.869 | 26.426 | 31.738 | 12.873 | |
| 0.060 | 84.637 | 25.222 | 28.007 | 9.441 | |
| 0.065 | 83.509 | 24.676 | 24.591 | 7.226 | |
| 0.070 | 82.404 | 17.083 | 21.432 | 5.246 | |
| 0.075 | 81.264 | 20.324 | 18.658 | 3.756 | **Configuration-B** |
| 0.080 | 80.126 | 14.398 | 16.185 | 2.696 | |
| 0.085 | 79.071 | 15.417 | 13.979 | 1.785 | |
| 0.090 | 78.022 | 10.907 | 12.078 | 1.227 | |
| 0.095 | 77.058 | 13.287 | 10.498 | 0.828 | |
| 0.100 | 76.064 | 9.426 | 9.022 | 0.528 | **Configuration-C** |
| 0.105 | 75.081 | 7.481 | 7.773 | 0.362 | |
| 0.110 | 74.086 | 5.157 | 6.626 | 0.211 | |
| 0.115 | 73.181 | 5.759 | 5.683 | 0.149 | |
| 0.120 | 72.275 | 4.602 | 4.882 | 0.092 | |
| 0.125 | 71.375 | 3.093 | 4.156 | 0.046 | **Configuration-D** |

**Table 4.1:** ACIMD performance: Accuracy and Adversary Advantages ($N = 3$ and $L_w = 2$ s.)

adversary reduces by around a half (14.23%). Configuration-C represents the case in which the accuracy is slightly lowered (76.06%) but the adversary chances are quite low (6.33%). The degeneration of configuration-3 is configuration-4 with a negligible probability of success for the adversary (2.39%) and a success rate for legitimate users of 71.38%. From a practical perspective, Configuration-C (or D) seems to be the most appropriate since it offers an acceptable accuracy while mitigating the three adversary advantages.

It is worth noting that the strategy to use signals of previous sessions or signals from other users achieve similar success rate. On the other hand, the success rate of an attacker using a random guessing approach is very low and decreases rapidly to zero when the parameter $\alpha$ increases.

Regarding distance checking, parameter $\alpha$ can be set as in "Configuration-C/D" in order to minimize the success rate of an adversary. The protocol might be executed several times in case the distance checking fails and the legitimate reader is, in reality, within the neighbourhood area.

### 4.3.3 Emergency mode

Considering that our authentication scheme must apply to emergency situations, the access to the implant must be guaranteed and fast in such scenarios. To ensure the access to the implant, the parameter $\alpha$ may be set to 0.05 (configuration-A) or a lower value in order to increase the success rate of authentication (>87%).

One key-point in emergency situations is the duration of signal monitoring. We experimentally tuned this parameter and selected an optimal value of 6 $s$. (considering $N = 3$ and $L_w = 2$ s.). The time consumed for computing the three wavelet transforms must be added to these 6 seconds. An upper bound on the time required for these operations is of few milliseconds, considering the performance of implementing a wavelet transform in a constrained device like a Field Programmable Gate Array (FPGA) [7, 113]. Therefore, only a few seconds are consumed in the emergency mode, which is reasonable to check proximity and to deal with the critical condition of an individual.

# 5

# Conclusions

## 5.1 Implantable Medical Devices

Implantable Medical Devices improve the quality of life of patients and, in some cases, play an important role in keeping them alive. The new generation of IMDs are increasingly incorporating more computing and communication capabilities. In this dissertation, we argued that advances on novel and smarter IMD designs must incorporate security solutions by design in order to provide the user with both safety and security guarantees. We have provided a comprehensive overview of the main security problems associated to the newest IMDs and have discussed how, in some cases, the patient's health can be seriously threatened by a malicious adversary. It is therefore evident that security mechanisms have to be incorporated into these devices. Further cooperation among researchers coming from manufacturing technologies, bioengineering, and computer security are necessary to guarantee both the patient's safety and the privacy and security of the data and communications.

Given the tensions among the different security objectives and the solutions proposed so far, it is unclear what the optimal choice would be. The question still remains an open problem. Many proposals provide a reasonably high security level but require too many resources (e.g., memory or computation), which is infeasible taking into consideration the need to save battery life. Alternatively, lightweight solutions are often vulnerable to attacks as a consequence of their weak designs.

Apart from purely engineering solutions, the procedures that both the medical personnel and the patients follow when interacting with the implants have to be considered, and existing regulations and standards should be also reviewed. However, nowadays these aspects are essentially ignored [218]. Devices must be used responsibly, and users must know various details about its functioning and the possible threats in order to raise security awareness.

Although some existing security solutions can be effective from a theoretical point of view, patients are very likely to reject them. The IMD is a computer system that is embedded into the human body. This is nowadays a special and delicate situation and the user opinion should be taken into account as far as possible. Interested readers can find in [50] some guidelines for designing security systems for IMDs considering the patient's point of view.

Looking even further ahead, medical implants open the door to other types of devices to improve human abilities, such as memory or perception, or even integrate our physiology with silicon-based components to improve our body. This looks certainly far ahead nowadays, but perhaps pacemakers and neurostimulators were also considered a remote possibility before they were introduced [228]. The field of computer security has to be ready to adapt and incorporate solutions for this new setting at the design phases, avoiding the develop-then-patch approach that has provided disastrous results for the Internet.

### 5.1.1 Compressed ECG signals

The integration of smart homes and telecare services aims to improve quality life and the possibilities for independent living through the use of new technologies and services. Smart devices at home pursue to increase comfort, energy efficiency, and security. On the other hand telecare services allow people to stay in their homes without prejudicing the quality of the health care services they are getting. The proper identification of the users is crucial to secure the systems. This task can be done through the features extracted from vital signals. Since the ECG signal is often monitored for medical purposes, we can take advantage of this and use also this vital signal for security purposes (e.g., identification or key generation). In our proposal we show how compressed ECG signals are robust and effective to unequivocally identify individuals.

In Chapter 2 we have evaluated the seven characteristics commonly demanded to biometrics systems. Apart from this, we would like to stress several additional characteristics of the proposed system. On the one hand, the use of compressed signals saves memory space, which could be critical in constrained devices like an implantable medical device such as a pacemaker or a holter monitor. Regarding the computational load, the penalty is very small since a matrix multiplication is only required to obtain the Hadamard coefficients. Furthermore no additional computations are required to extract the signal features—contrary to what occurs in systems

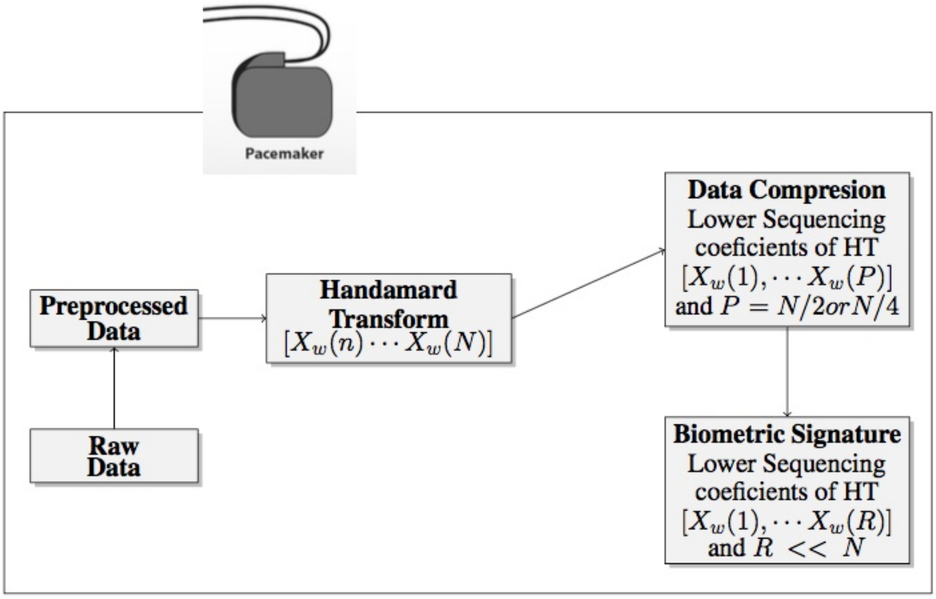

**Figure 5.1:** Pacemaker with data compression and biometrics signature modules

based on fiducial features. On the other hand, since only a small fraction of the coefficients (the lower ones) are employed, even if the attacker would acquired these coefficients, she could not reconstruct the original signal. In conclusion, the proposed system is privacy preserving and works with a highly compressed version of the signal. As illustration of how the proposed system might be integrated in implantable medical devices is sketched in Figure 5.1, showing how our proposal could contribute to the design of more secure medical applications and devices. For instance, a patient holding this sort of pacemaker could be remotely monitored once she is identified in a secure way using features extracted from her own heart signal.

As a future work, there are several research lines to continue with the ideas presented in this article. The proposal has been only tested with a database (MIT-BIH Normal Sinus Rhythm Database) of healthy individuals. Other databases, which include patients with a heart disease (e.g., MIT-BIH Arrhythmia Database or MIT-BIH Long-Term ST Database) or patients under stress conditions (e.g., MIT-BIH ST Change Database), could be employed to assess the use of compressed ECG signals. In line with this, in our proposal the Hadamard Transform is the core of our system for human identification. It would be interesting to perform a comparative study using a wide set of transforms (e.g., Fourier, Wavelet, Hadamard, etc.). Last but not least, the proposal could be extended to other vital signals like

EEG or EMG.

### 5.1.2 ECG streams

We are currently in an era in which our surrounding devices generate and transmit data in a continuous way. An example of these devices are those belonging to the Internet-of-Things (IoT) or the new generations of Implantable Medical Devices (IMDs) with wireless connectivity. These devices receive data continuously and very frequently in a non-orderly fashion. One use of such data is user authentication. In particular, the use of biological signals has been previously studied for authentication purposes. Cardiac signals (PPG or ECG) and brain signals (EEG) collected from IMDs or body sensors, are widely used for authentication and some authors have applied them to the CA scenario [17, 42]. However, given the continuous nature of the authentication process, the system has to be adapted to changes; for example, ECG signal may slightly change over time. Thus, Data Stream Mining (DSM) emerges as a promising technique to face this sort of problems. To the best of our knowledge, none of the existing solutions use ECG signals as data streams.

We exploit the full potential of DSM for designing a CA system using ECG streams. The proposed real-time system has been evaluated using records of 10 individuals monitored during approximately half a day. Our results show the potential of ECG streams for security purposes. In fact, the behaviour of the classifier, which is the core of the CA system, is almost perfect. The CA approach achieves an accuracy as high as the NCA approach but with the benefit of using a limited memory and being able to process data streams. Moreover, we have tested the buffered and unbuffered approaches in the CA setting to show how the use of one or another is driven by the requirements of the real time application (e.g., credentials/second that must be checked by the CA system). Finally, we have studied how drift detection techniques (e.g., DDM or EDDM) may help to deal with the existing changes in the ECG data streams —a wrapper approach has been tested. The results clearly indicate that drift detection techniques are effective to build robust CA schemes even under very noisy conditions.

As a future work, we plan to check whether the concept of ECG streams can be extended to other physiological signals. We hope this contribution can serve as seed to many other works that explore the use of biological signals for continuous authentication.

### 5.1.3 Emergency conditions

There is an agreed consensus about the benefits of incorporating telemetry into the new generation of IMDs. In particular, it improves the patients' quality of life, facilitates its management by the medical personnel and reduces costs. Unfortunately, security protection mechanisms are still missing in comercial IMDs. Among the security requirements, controlling which devices can read from or send commands to IMDs is paramount. For this purpose, in this paper an access control protocol called ACIMD has been introduced. ACIMD implements a distance bounding mechanism based on physiological signals, particularly electrocardiograms. In particular, our proposed scheme allows to verify the proximity between an IMD and a Programmer (distance checking) and also each entity is assured of both the identity of the other party and her presence during the protocol execution (mutual authentication).

ACIMD outperforms previous approaches since it considers the whole ECG signal, which is difficult to acquire remotely. Moreover, it can work under normal and emergency operation modes typical for IMDs. The feasibility of the proposal has been evaluated with real ECG data.

Future work will be focused on applying lightweight cryptographic primitives to reduce the cost of the different operations at stake. In addition, regarding the similarity analysis, other alternatives will be tested.

## 5.2 Publications

During this PhD, the research work done has resulted in some scientific papers which has been published in scientific journals [27, 28, 29, 31, 75, 170] and international conferences [30, 108]. This section describes the published works and the impact or ranking of each of the journals where they have been published.

| | |
|---|---|
| Title: | Security and privacy issues in implantable medical devices: A comprehensive survey |
| Authors: | Carmen Camara, Pedro Peris-Lopez, Juan E. Tapiador |
| Abstract: | Bioengineering is a field in expansion. New technologies are appearing to provide a more efficient treatment of diseases or human deficiencies. Implantable Medical Devices (IMDs) constitute one example, these being devices with more computing, decision making and communication capabilities. Several research works in the computer security field have identified serious security and privacy risks in IMDs that could compromise the implant and even the health of the patient who carries it. This article surveys the main security goals for the next generation of IMDs and analyzes the most relevant protection mechanisms proposed so far. On the one hand, the security proposals must have into consideration the inherent constraints of these small and implanted devices: energy, storage and computing power. On the other hand, proposed solutions must achieve an adequate balance between the safety of the patient and the security level offered, with the battery lifetime being another critical parameter in the design phase. |
| Keywords: | Implantable medical devices; Security; Privacy; m-Health; Survey |
| Journal: | Journal of Biomedical Informatics |
| DOI: | `https://doi.org/10.1016/j.jbi.2015.04.007` |
| Impact: | 2.447 of Impact Factor (Q1) |
| Year: | 2015 |

| | |
|---|---|
| Title: | Real-time electrocardiogram streams for continuous authentication |
| Authors: | Carmen Camara, Pedro Peris-Lopez, Lorena Gonzalez-Manzano, Juan Tapiador |
| Abstract: | Security issues are becoming critical in modern smart systems. Particularly, ensuring that only legitimate users get access to them is essential. New access control systems must rely on continuous authentication (CA) to provide higher security level. To achieve this, recent research has shown how biological signals, such as electroencephalograms (EEGs) or electrocardiograms (ECGs), can be useful for this purpose. In this paper, we introduce a new CA scheme that, contrarily to previous works in this area, considers ECG signals as continuous data streams. The data stream paradigm is suitable for this scenario since algorithms tailored for data streams can cope with continuous data of a theoretical infinite length and with a certain variability. The proposed ECG-based CA system is intended for real-time applications and is able to offer an accuracy up to 96%, with an almost perfect system performance (kappa statistic >80%). |
| Keywords: | Datastreams; Healthcare; Identification; Electrocardiogram |
| Journal: | Applied Soft Computing |
| DOI: | `https://doi.org/10.1016/j.asoc.2017.07.032` |
| Impact: | 3.541 of Impact Factor (Q1) |
| Year: | 2017 |

| | |
|---|---|
| Title: | Human Identification Using Compressed ECG Signals |
| Authors: | Carmen Camara, Pedro Peris-Lopez, Juan E. Tapiador |
| Abstract: | As a result of the increased demand for improved life styles and the increment of senior citizens over the age of 65, new home care services are demanded. Simultaneously, the medical sector is increasingly becoming the new target of cybercriminals due the potential value of users' medical information. The use of biometrics seems an effective tool as a deterrent for many of such attacks. In this paper, we propose the use of electrocardiograms (ECGs) for the identification of individuals. For instance, for a telecare service, a user could be authenticated using the information extracted from her ECG signal. The majority of ECG-based biometrics systems extract information (fiducial features) from the characteristics points of an ECG wave. In this article, we propose the use of non-fiducial features via the Hadamard Transform (HT). We show how the use of highly compressed signals (only 24 coefficients of HT) is enough to unequivocally identify individuals with a high performance (classification accuracy of 0.97 and with identification system errors in the order of $10^{-2}$). |
| Keywords: | Healthcare; Biometrics; Human Identification; ECG |
| Journal: | Journal of Medical Systems |
| DOI: | `https://doi.org/10.1007/s10916-015-0323-2` |
| Impact: | 0.705 of SJR (Q2) |
| Year: | 2015 |

| | |
|---|---|
| Title: | Effect of attacker characterization in ECG-based continuous authentication mechanisms for Internet of Things |
| Authors: | Pedro Peris-Lopez, Lorena González-Manzano, Carmen Camara, José María de Fuentes |
| Abstract: | Wearable devices enable retrieving data from their porting user, among other applications. When combining them with the Internet of Things (IoT) paradigm, a plethora of services can be devised. Thanks to IoT, several approaches have been proposed to apply user data, and particularly ElectroCardioGram (ECG) signals, for biometric authentication. One step further is achieving Continuous Authentication (CA), i.e., ensuring that the user remains the same during a certain period. The hardness of this task varies with the attacker characterization, that is, the amount of information about the attacker that is available to the authentication system. In this vein, we explore different ECG-based CA mechanisms for *known*, *blind-modelled* and *unknown* attacker settings. Our results show that, under certain configuration, 99.5 % of true positive rate can be achieved for a blind-modelled attacker, 93.5 % for a known set of attackers and 91.8 % for unknown ones. |
| Keywords: | Internet of Things; Electrocardiogram; Continuous authentication; Attacker model |
| Journal: | Future Generation Computer Systems |
| DOI: | `https://doi.org/10.1016/j.future.2017.11.037` |
| Impact: | 3.997 of Impact Factor (Q1). |
| Year: | 2017 |

| | |
|---|---|
| Title: | Encryption by Heart (EbH) – Using ECG for time-invariant symmetric key generation |
| Authors: | Lorena González-Manzano, José María de Fuentes, Pedro Peris-Lopez, Carmen Camara |
| Abstract: | Wearable devices are a part of Internet-of-Things (IoT) that may offer valuable data of their porting user. This paper explores the use of ElectroCardioGram (ECG) records to encrypt user data. Previous attempts have shown that ECG can be taken as a basis for key generation. However, these approaches do not consider time-invariant keys. This feature enables using these so-created keys for symmetrically encrypting data (e.g. smartphone pictures), enabling their decryption using the key derived from the current ECG readings. This paper addresses this challenge by proposing EbH, a mechanism for persistent key generation based on ECG. EbH produces seeds from which encryption keys are generated. Experimental results over 24 h for 199 users show that EbH, under certain settings, can produce permanent seeds (thus time-invariant keys) computed on-the-fly and different for each user—up to 95.97% of users produce unique keys. In addition, EbH can be tuned to produce seeds of different length (up to 300 bits) and with variable min-entropy (up to 93.51). All this supports the workability of EbH in a real setting. |
| Keywords: | ECG; Symmetric encryption; Time-invariant keys |
| Journal: | Future Generation Computer Systems |
| DOI: | `https://doi.org/10.1016/j.future.2017.07.018` |
| Impact: | 3.997 of Impact Factor (Q1). |
| Year: | 2017 |

| | |
|---|---|
| Title: | Non-invasive Multi-modal Human Identification System Combining ECG, GSR, and Airflow Biosignals |
| Authors: | Carmen Camara, Pedro Peris-Lopez, Juan E. Tapiador, Guillermo Suarez-Tangil |
| Abstract: | A huge amount of data can be collected through a wide variety of sensor technologies. Data mining techniques are often useful for the analysis of gathered data. This paper studies the use of three wearable sensors that monitor the electrocardiogram, airflow, and galvanic skin response of a subject with the purpose of designing an efficient multi-modal human identification system. The proposed system, based on the rotation forest ensemble algorithm, offers a high accuracy (99.6 % true acceptance rate and just 0.1 % false positive rate). For its evaluation, the proposed system was testing against the characteristics commonly demanded in a biometric system, including universality, uniqueness, permanence, and acceptance. Finally, a proof-of-concept implementation of the system is demonstrated on a smartphone and its performance is evaluated in terms of processing speed and power consumption. The identification of a sample is extremely efficient, taking around 200 ms and consuming just a few millijoules. It is thus feasible to use the proposed system on a regular smartphone for user identification. |
| Keywords: | Sensor data Bioinformatics; Human identification; Data mining; Ensemble classification |
| Journal: | Journal of Medical and Biological Engineering |
| DOI: | https://doi.org/10.1007/s40846-015-0089-5 |
| Impact: | 1.018 of Impact Factor (Q3) |
| Year: | 2015 |

| | |
|---|---|
| Title: | Beyond Security on Implantable Medical Devices (Poster) |
| Authors: | Carmen Camara, Pedro Peris-Lopez, Juan E. Tapiador |
| Conference: | EIT DIGITAL SYMPOSIUM — in conjunction with European Cyber week |
| Year: | 2016 |

| | |
|---|---|
| Title: | Dial 258x: play your app on your speakers to get the malware family (Poster) |
| Authors: | Guillermo Izquierdo-Moreno, Pedro Peris-Lopez, Carmen Camara |
| Conference: | USENIX Security, |
| Year: | 2017 |

# Bibliography

[1] P. S. Addison. *The Illustrated Wavelet Transform Handbook: Introductory Theory and Applications in Science, Engineering, Medicine and Finance*. Taylor & Francis, 1st edition, 2002.

[2] F. Agrafioti and D. Hatzinakos. ECG based recognition using second order statistics. In *6th Annual Conference on Communication Networks and Services Research (CNSR)*, pages 82–87, 2008.

[3] M. S. Aldayel. K-nearest neighbor classification for glass identification problem. In *nternational Conference on Computer Systems and Industrial Informatics*, pages 1–5, 2012.

[4] D. K. Altop, A. Levi, and V. Tuzcu. Towards using physiological signals as cryptographic keys in body area networks. In *9th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, PervasiveHealth '15, pages 92–99. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2015.

[5] O. Aquilina. A brief history of cardiac pacing. *Paediatric Cardiology*, 8(2):17–81, 2008.

[6] D. Arney, K. K. Venkatasubramanian, O. Sokolsky, and I. Lee. Biomedical devices and systems security. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 2376–2379, 2011.

[7] C. S. Avinash and J. S. R. Alex. Fpga implementation of discrete wavelet transform using distributed arithmetic architecture. In *International Conference on Smart Technologies and Management for Computing, Communication, Controls, Energy and Materials (IC-STM)*, pages 326–330, May 2015.

[8] G. Avoine, M. A. Bingöl, S. Kardaş, C. Lauradoux, and B. Martin. A Framework for Analyzing RFID Distance Bounding Protocols. *Journal of Computer Security – Special Issue on RFID System Security*, 19(2):289–317, March 2011.

[9] G. Avoine, S. Mauw, and R. Trujillo-Rasua. Comparing distance bounding protocols: A critical mission supported by decision theory. *Computer Communications*, 67:92 – 102, 2015.

[10] S.-D. Bao, C. C. Y. Poon, Z. Yuan-Ting, and L.-F. Shen. Using the timing information of heartbeats as an entity identifier to secure body sensor network. *IEEE Transactions on Information Technology in Biomedicine*, 12(6):772–779, 2008.

[11] A. Ben-Hur, D. Horn, H. T. Siegelmann, and V. Vapnik. Support vector clustering. *Journal of machine learning research*, 2(Dec):125–137, 2001.

[12] S. Bergamasco, M. Bon, and P. Inchingolo. Medical data protection with a new generation of hardware authentication tokens. 2001.

[13] T. P. Berger, J. D'Hayer, K. Marquet, M. Minier, and G. Thomas. The gluon family: A lightweight hash function family based on fcsrs. In *Progress in Cryptology - AFRICACRYPT 2012*, volume 7374 of *Lecture Notes in Computer Science*, pages 306–323. Springer Berlin Heidelberg, 2012.

[14] A. Bifet, G. Holmes, R. Kirkby, and B. Pfahringer. Moa: Massive online analysis. *J. Mach. Learn. Res.*, 11:1601–1604, Aug. 2010.

[15] G. Bifet, A. Holmes, R. Kirkby, and B. Pfahringer. Data stream mining: A practical approach. Technical report. University of Waikato. `http://www.cs.waikato.ac.nz/~abifet/MOA/StreamMining.pdf`, 2012.

[16] A. Bogdanov, M. Knezevic, G. Leander, D. Toz, K. Varıcı, and I. Verbauwhede. spongent: A lightweight hash function. In *Proc. of Cryptographic Hardware and Embedded Systems*, volume 6917 of *Lecture Notes in Computer Science*, pages 312–325. Springer Berlin Heidelberg, 2011.

[17] A. Bonissi, R. D. Labati, L. Perico, R. Sassi, F. Scotti, and L. Sparagino. A preliminary study on continuous authentication methods for photoplethysmographic biometrics. In *IEEE Workshop on Biometric Measurements and Systems for Security and Medical Applications (BioMS)*, pages 28–33, 2013.

[18] C. J. Borleffs, J. Thijssen, M. K. de Bie, J. B. van Rees, G. H. van Welsenes, L. van Erven, J. J. Bax, S. C. Cannegieter, and M. J. Schalij. Recurrent implantable cardioverter-defibrillator replacement is associated with an increasing risk of pocket-related complications. *Pacing and Clinical Electrophysiology*, 33(8):1013–1019, 2010.

[19] C. Boyd and A. Mathuria. *Protocols Using Shared Key Cryptography*, pages 73–106. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.

[20] S. Brands and D. Chaum. *Distance-Bounding Protocols*, pages 344–359. Springer Berlin Heidelberg, 1994.

[21] A. D. Brucker and H. Petritsch. Extending access control models with break-glass. In *Proceedings of the 14th ACM Symposium on Access Control Models and Technologies*, pages 197–206. ACM, 2009.

[22] I. Buhan, E. Kelkboom, and K. Simoens. A survey of the security and privacy measures for anonymous biometric authentication systems. In *Sixth International Conference on Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP)*, pages 346–351, Oct 2010.

[23] C. S. Burrus, R. A. Gopinath, and H. Guo. *Introduction to Wavelets and Wavelet Transforms: A Primer*. Prentice Hall, 1st edition, 1997.

[24] A. Calleja, P. Peris-Lopez, and J. E. Tapiador. *Electrical Heart Signals can be Monitored from the Moon: Security Implications for IPI-Based Protocols*, pages 36–51. Springer International Publishing, 2015.

[25] M. A. Callejon, D. Naranjo-Hernandez, J. Reina-Tosina, and L. M. Roa. A comprehensive study into intrabody communication measurements. *IEEE Transactions on Instrumentation and Measurement,*, 62(9):2446–2455, 2013.

[26] M. A. Callejon, L. M. Roa, J. Reina-Tosina, and D. Naranjo-Hernandez. Study of attenuation and dispersion through the skin in intrabody communications systems. *IEEE Transactions on Information Technology in Biomedicine*, 16(1):159–165, 2012.

[27] C. Camara, P. Peris-Lopez, L. Gonzalez-Manzano, and J. Tapiador. Real-time electrocardiogram streams for continuous authentication. *Applied Soft Computing*, 2017.

[28] C. Camara, P. Peris-Lopez, and J. E. Tapiador. Human identification using compressed ecg signals. *Journal of Medical Systems*, 39(11):1–10, 2015.

[29] C. Camara, P. Peris-Lopez, and J. E. Tapiador. Security and privacy issues in implantable medical devices: A comprehensive survey. *Journal of Biomedical Informatics*, 55:272 – 289, 2015.

[30] C. Camara, P. Peris-Lopez, and J. E. Tapiador. Beyond security on

implantable medical devices (poster). EIT DIGITAL SYMPOSIUM — in conjunction with European Cyber week, 2016.

[31] C. Camara, P. Peris-Lopez, J. E. Tapiador, and G. Suarez-Tangil. Non-invasive multi-modal human identification system combining ecg, gsr, and airflow biosignals. *Journal of Medical and Biological Engineering*, 35(6):735–748, 2015.

[32] I. Canada. Medical devices operating in the 401-406 mhz frequency band. `http://www.ic.gc.ca/eic/site/smt-gst.nsf/vwapj/rss243.pdf/$FILE/rss243.pdf`, 2010.

[33] M. Cempirek and J. Stastny. The optimization of the EEG-based biometric classification. *Applied Electronics*, pages 25–28, 2007.

[34] I. T. R. Center. Data breach report. Technical report, December 2014.

[35] P. Chadwick. Regulations and standards for wireless applications in ehealth. In *29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS)*, pages 6170–6173, 2007.

[36] A. D. C. Chan, M. M. Hamdy, A. Badre, and V. Badee. Wavelet distance measure for person identification using electrocardiograms. *IEEE Transactions on Instrumentation and Measurement*, 57(2):248–253, Feb 2008.

[37] S. Cherukuri, K. K. Venkatasubramanian, and S. K. S. Gupta. Biosec: a biometric based approach for securing communication in wireless networks of biosensors implanted in the human body. In *Proc. of International Conference on Parallel Processing Workshops*, pages 432–439, 2003.

[38] K. Cho and D. Lee. Biometric based secure communications without pre-deployed key for biosensor implanted in body sensor networks. In *Information Security Applications*, volume 7115 of *Lecture Notes in Computer Science*, pages 203–218. Springer Berlin Heidelberg, 2012.

[39] D. D. Clark and D. R. Wilson. A comparison of commercial and military computer security policies. In *IEEE Symposium on Security and Privacy*, pages 184–195. IEEE Computer Society, 1987.

[40] F. C. Commission. About medical device radiocommunications service. `http://wireless.fcc.gov/services/index.htm?job=service_bandplan&id=medical_implant`, 2009.

[41] J.-P. Couderc. The telemetric and holter {ECG} warehouse (thew):

The first three years of development and research. *Journal of Electrocardiology*, 45(6):677 – 683, 2012.

[42] D. P. Coutinho, A. L. N. Fred, and M. A. T. Figueiredo. Ecg-based continuous authentication system using adaptive string matching. In *Biosignals*, pages 354–359, 2011.

[43] E. J. da S. Luz, D. Menotti, and W. R. Schwartz. Evaluating the use of {ECG} signal in low frequencies as a biometry. *Expert Systems with Applications*, 41(5):2309 – 2315, 2014.

[44] H. P. Da Silva, A. Fred, A. Lourenço, and A. Jain. Finger ecg signal for user authentication: Usability and performance. In *IEEE Sixth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–8. IEEE, 2013.

[45] K. Daniluk and E. Niewiadomska-Szynkiewicz. Energy-efficient security in implantable medical devices. In *Computer Science and Information Systems (FedCSIS), 2012 Federated Conference on*, pages 773–778, 2012.

[46] M. Darji and B. Trivedi. Detection of active attacks on wireless imds using proxy device and localization information. In *Security in Computing and Communications*, volume 467 of *Communications in Computer and Information Science*, pages 353–362. Springer Berlin Heidelberg, 2014.

[47] M. de Sousa, G. Klein, T. Korte, and M. Niehaus. Electromagnetic interference in patients with implanted cardioverter-defibrillators and implantable loop recorders. *Indian Pacing Electrophysiol Journal*, 2(3):79–84, 2002.

[48] B. Defend, M. Salajegheh, K. Fu, and S. Inoue. Protecting global medical telemetry infrastructure. Technical report, Institute of Information Infrastructure Protection (I3P), 2008.

[49] F. Delmastro. Pervasive communications in healthcare. *Computer Communications*, 35(11):1284 – 1295, 2012.

[50] T. Denning, A. Borning, B. Friedman, B. T. Gill, T. Kohno, and W. H. Maisel. Patients, pacemakers, and implantable defibrillators: Human values and security for wireless implantable medical devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 917–926. ACM, 2010.

[51] T. Denning, K. Fu, and T. Kohno. Absence makes the heart grow fonder: New directions for implantable medical device security. In *Proceedings of the 3rd Conference on Hot Topics in Security*, HOTSEC'08, pages 5:1–5:7. USENIX Association, 2008.

[52] Y. Desmedt, C. Goutier, and S. Bengio. Special uses and abuses of the fiat-shamir passport protocol. In *Advances in Cryptology - CRYPTO '87*, volume 293 of *Lecture Notes in Computer Science*, pages 21–39. Springer, 1987.

[53] Dexcom. Seven plus cgm system. `http://www.dexcom.com/seven-plus`, Consulted on December 2017.

[54] T. Drew and M. Gini. Implantable medical devices as agents and part of multiagent systems. In *Proc. of the fifth international joint conference on Autonomous agents and multiagent systems*, AAMAS '06, pages 1534–1541. ACM, 2006.

[55] M. H. Eldefrawy, M. K. Khan, and K. Alghathbar. A key agreement algorithm with rekeying for wireless sensor networks using public key cryptography. In *2010 International Conference on Anti-Counterfeiting Security and Identification in Communication (ASID)*, pages 1–6, 2010.

[56] C. Falcon. Wireless medical devices: Satisfying radio requirements. *Medical Device & Diagnostic Industry*, page 80, 2004.

[57] FDA. Radio frequency wireless technology in medical devices - guidance for industry and food and drug administration staff. `http://www.fda.gov/downloads/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm077272.pdf`, 2013.

[58] FDA. Content of premarket submissions for management of cybersecurity in medical devices. `http://www.fda.gov/downloads/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/UCM356190.pdf`, 2014.

[59] A. Ferreira, R. Cruz-Correia, L. Antunes, P. Farinha, E. Oliveira-Palhares, D. W. Chadwick, and A. Costa-Pereira. How to break access control in a controlled manner. In *Proceedings of the 19th IEEE Symposium on Computer-Based Medical Systems*, pages 847–854. IEEE Computer Society, 2006.

[60] K. Fishkin and S. Roy. Enhancing rfid privacy via antenna energy analysis., 2003.

[61] U. Food and D. A. (FDA). Medical device safety. `http://wireless.fcc.gov/services/index.htm?job=service_bandplan&id=medical_implant`, Consulted on December 2017.

[62] K. Fotopoulou and B. Flynn. Optimum antenna coil structure for

inductive powering of passive rfid tags. In *IEEE International Conference on RFID*, pages 71–77, 2007.

[63] M. Frank, R. Biedert, E. Ma, I. Martinovic, and D. Song. Touchalytics: On the applicability of touchscreen input as a behavioral biometric for continuous authentication. *IEEE Transactions on Information Forensics and Security*, 8(1):136–148, 2013.

[64] E. Freudenthal, R. Spring, and L. Estevez. Practical techniques for limiting disclosure of rf-equipped medical devices. In *IEEE Engineering in Medicine and Biology Workshop,*, pages 82–85, 2007.

[65] K. Fu. Inside risks: Reducing risks of implantable medical devices. *ACM Communications*, 52(6):25–27, 2009.

[66] K. Fu. Trustworthy medical device software. In *Public Health Effectiveness of the FDA 510(k) Clearance Process: Measuring Postmarket Performance and Other Select Topics: Workshop Report*, 2011.

[67] F. Furbass and J. Wolkerstorfer. Ecc processor with low die size for rfid applications. In *IEEE International Symposium on Circuits and Systems*, pages 1835–1838, 2007.

[68] M. M. Gaber, J. Gama, S. Krishnaswamy, J. B. Gomes, and F. Stahl. Data stream mining in ubiquitous environments: state-of-the-art and current directions. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 4(2):116–138, 2014.

[69] Y. Gahi, M. Lamrani, A. Zoglat, M. Guennoun, B. Kapralos, and K. El-Khatib. Biometric identification system based on electrocardiogram data. In *Int. Conference on new technologies, mobility and security (NTMS)*, pages 1–5, 2008.

[70] J. Gama. *Knowledge Discovery from Data Streams*. Chapman and Hall/CRC, 2010.

[71] H. Gascon, S. Uellenbeck, C. Wolf, and K. Rieck. Continuous authentication on mobile devices by analysis of typing motion behavior. In *Sicherheit*, pages 1–12. Citeseer, 2014.

[72] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley. PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. *Circulation*, 101(23):e215–e220, 2000 (June 13). Circulation Electronic Pages: http://circ.ahajournals.org/cgi/content/full/101/23/e215 PMID:1085218; doi: 10.1161/01.CIR.101.23.e215.

[73] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, and K. Fu.

They can hear your heartbeats: Non-invasive security for implantable medical devices. *SIGCOMM Comput. Commun. Rev.*, 41(4):2–13, Aug. 2011.

[74] S. Gollakota, H. A. Hassanieh, B. Ransford, D. Katabi, and K. Fu. Imd shield: Securing implantable medical devices (poster). USENIX Association, 2011.

[75] L. González-Manzano, J. M. de Fuentes, P. Peris-Lopez, and C. Camara. Encryption by heart (EbH) – using ecg for time-invariant symmetric key generation. *Future Generation Computer Systems*, 77(Supplement C):136 – 148, 2017.

[76] P. M. Gonçalves, S. G. de Carvalho Santos, R. S. Barros, and D. C. L. Vieira. A comparative study on concept drift detectors. *Expert Systems with Applications*, 41(18):8144 – 8156, 2014.

[77] R. Greiner, X. Su, B. Shen, and W. Zhou. Structural extension to logistic regression: Discriminative parameter learning of belief net classifiers. *Machine Learning*, 59(3):297–322, 2005.

[78] M. Guennoun, N. Abbad, J. Talom, S. M. M. Rahman, and K. El-Khatib. Continuous authentication by electrocardiogram data. In *Science and Technology for Humanity (TIC-STH), 2009 IEEE Toronto international conference*, pages 40–42. IEEE, 2009.

[79] E. E. Gul and M. Kayrak. Common pacemaker problems: Lead and pocket complications. In M. R. Das, editor, *Modern Pacemakers - Present and Future*. InTech, 2011.

[80] J. Guo, T. Peyrin, and A. Poschmann. The photon family of lightweight hash functions. In *Advances in Cryptology – CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 222–239. Springer Berlin Heidelberg, 2011.

[81] S. K. S. Gupta, T. Mukherjee, and K. Venkatasubramanian. Criticality aware access control model for pervasive applications. In *Fourth Annual IEEE International Conference on Pervasive Computing and Communications*, pages 257–261, March 2006.

[82] D. Halperin, T. S. Heydt-Benjamin, K. Fu, T. Kohno, and W. H. Maisel. Security and privacy for implantable medical devices. *IEEE Pervasive Computing*, 7(1):30–39, 2008.

[83] D. Halperin, T. S. Heydt-Benjamin, B. Ransford, S. S. Clark, B. Defend, W. Morgan, K. Fu, T. Kohno, and W. H. Maisel. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. In *Proc. of the 29th Annual IEEE Symposium on Security and Privacy*, pages 129–142, 2008.

[84] W. Han, J. Zhang, and X. Yao. Context-sensitive access control model and implementation. In *The Fifth International Conference on Computer and Information Technology*, pages 757–763, 2005.

[85] S. Hanna. Regulations and standards for wireless medical applications. In *Third International Symposium on Medical Information & Communication Technology*, pages 1–5, 2009.

[86] J. A. Hansen and N. M. Hansen. A taxonomy of vulnerabilities in implantable medical devices. In *Proc. of the second annual workshop on Security and privacy in medical and home-care systems*, SPIMACS '10, pages 13–20, New York, USA, 2010. ACM.

[87] M. Hanson, H. Powell, A. Barth, K. Ringgenberg, B. Calhoun, J. Aylor, and J. Lach. Body area sensor networks: Challenges and opportunities. *Computer*, 42(1):58–65, Jan 2009.

[88] G. Haubrich, L. Twetan, and G. Rosar. Multiple band communications for an implantable medical device, July 27 2006. WO Patent App. PCT/US2006/000,961.

[89] D. He, S. Chan, and S. Tang. A novel and lightweight system to secure wireless medical sensor networks. *IEEE Journal of Biomedical and Health Informatics*, 18(1):316–326, Jan 2014.

[90] X. Hei and X. Du. Biometric-based two-level secure access control for implantable medical devices during emergencies. In *Proceedings IEEE INFOCOM*, pages 346–350, 2011.

[91] X. Hei, X. Du, J. Wu, and F. Hu. Defending resource depletion attacks on implantable medical devices. In *Proc. of IEEE Global Telecommunications Conference (GLOBECOM)*, pages 1–5, 2010.

[92] M. Hejazi, S. A. R. Al-Haddad, Y. P. Singh, S. J. Hashim, and A. F. A. Aziz. Ecg biometric authentication based on non-fiducial approach using kernel methods. *Digital Signal Processing*, 52:72–86, 2016.

[93] N. Henry, N. Paul, and N. McFarlane. Using bowel sounds to create a forensically-aware insulin pump system. In *Workshop on Health Information Technologies, HealthTech, USENIX*, pages 1–10, 2013.

[94] HIPPA. Security standards: Technical safeguards. 2(4):1–17, 2007.

[95] M. Hofmann and R. Klinkenberg. *RapidMiner: Data Mining Use Cases and Business Analytics Applications*. Chapman & Hall/CRC, 2013.

[96] S. Hosseini-Khayat. A lightweight security protocol for ultra-low power asic implementation for wireless implantable medical de-

vices. In *5th International Symposium on Medical Information Communication Technology (ISMICT)*, pages 6–9, March 2011.

[97] F. Hu, Q. Hao, M. Lukowiak, Q. Sun, K. Wilhelm, S. Radziszowski, and Y. Wu. Trustworthy data collection from implantable medical devices via high-speed security implementation based on ieee 1363. *IEEE Transactions on Information Technology in Biomedicine*, 14(6):1397–1404, 2010.

[98] F. Hu, Q. Sun, Y. Wu, M. Guo, J. Lu, J. Li, D. J. Gay, J. K. Garner, and A. L. Poellnitz. Implantable medical devices: Architectureand design. In *Telehealthcare Computing and Engineering: Principles and Design*, chapter 14, pages 359–406. 1 edition, 2013.

[99] J. Hu and Z. Mu. Eeg authentication system based on auto-regression coefficients. In *10th International Conference on Intelligent Systems and Control (ISCO)*, pages 1–5, 2016.

[100] G. Hulten and P. Domingos. VFML – a toolkit for mining high-speed time-changing data streams. 2003.

[101] T. Instruments. Msp430f156, 16-bit ultra-low-power mcu. `http://www.ti.com/lit/ds/symlink/msp430f156.pdf`.

[102] K. Inthavisas and D. Lopresti. Secure speech biometric templates for user authentication. *IET Biometrics*, 1(1):46–54, 2012.

[103] ISO. Information technology – security techniques – entity authentication – part 2: Mechanisms using symmetric encipherment algorithms, iso/iec 9798-2:2008. International Standard, 2nd ed., 1999.

[104] T. Isobe and K. Shibutani. Security analysis of the lightweight block ciphers xtea, led and piccolo. In *Information Security and Privacy*, volume 7372 of *Lecture Notes in Computer Science*, pages 71–86. Springer Berlin Heidelberg, 2012.

[105] C. W. Israel and S. S. Barold. Pacemaker systems as implantable cardiac rhythm monitors. *The American Journal of Cardiology*, 88(4):442 – 445, 2001.

[106] S. A. Israel, J. M. Irvine, A. Cheng, M. D. Wiederhold, and B. K. Wiederhold. Ecg to identify individuals. *Pattern Recognition*, 38(1):133 – 142, 2005.

[107] ITU-T. E-health standards and interoperability. `http://www.itu.int/dms_pub/itu-t/oth/23/01/T23010000170001PDFE.pdf`, 2012.

[108] G. Izquierdo-Moreno, P. Peris-Lopez, and C. Camara. Dial 258x: play your app on your speakers to get the malware family (poster). USENIX Security, 2017.

[109] A. K. Jain, A. Ross, and S. Pankanti. Biometrics: a tool for information security. *IEEE Transactions on Information Forensics and Security*, 1(2):125–143, 2006.

[110] A. K. Jain, A. Ross, and S. Prabhakar. An introduction to biometric recognition. *IEEE Transactions on circuits and systems for video technology*, 14(1):4–20, 2004.

[111] H. Jannati. Analysis of relay, terrorist fraud and distance fraud attacks on RFID systems. *International Journal of Critical Infrastructure Protection*, August 2015.

[112] H. Jiang, J. Zhang, D. Lan, K. K. Chao, S. Liou, H. Shahnasser, R. Fechter, S. Hirose, M. Harrison, and S. Roy. A low-frequency versatile wireless power transfer technology for biomedical implants. *Biomedical Circuits and Systems, IEEE Transactions on*, 7(4):526–535, 2013.

[113] R. M. Jiang and D. Crookes. Fpga implementation of 3d discrete wavelet transform for real-time medical imaging. In *18th European Conference on Circuit Theory and Design*, pages 519–522, Aug 2007.

[114] A. Juels and D. Bailey. Access control for implanted medical devices, 2013.

[115] A. Juels and M. Wattenberg. A fuzzy commitment scheme. In *Proceedings of the 6th ACM Conference on Computer and Communications Security*, CCS '99, pages 28–36. ACM, 1999.

[116] A. Kaadan and H. Refai. Securing wireless medical devices. In *IEEE Global Communications Conference (GLOBECOM)*, pages 942–948, 2012.

[117] S. J. Kang, S. Y. Lee, H. I. Cho, and H. Park. Ecg authentication system design based on signal analysis in mobile and wearable devices. *IEEE Signal Processing Letters*, 23(6), 2016.

[118] P. Karger, G. S. Kc, and D. Toll. Privacy is essential for secure mobile devices. *IBM Journal of Research and Development*, 53(2):5:1–5:17, 2009.

[119] W. Khalifa, A. Salem, M. Roushdy, and K. Revett. A survey of eeg based user authentication schemes. In *8th International Conference on Informatics and Systems*, pages BIO–55–BIO–60, May 2012.

[120] I. Khamassi, M. Sayed-Mouchaweh, M. Hammami, and K. Ghédira. Discussion and review on evolving data streams and concept drift adapting. *Evolving Systems*, pages 1–23, 2016.

[121] S. Kiranyaz, T. Ince, J. Pulkkinen, and M. Gabbouj. Personalized

long-term ecg classification: A systematic approach. *Expert Systems with Applications*, 38(4):3220–3226, 2011.

[122] P. Kitsos, N. Sklavos, M. Parousi, and A. N. Skodras. A comparative study of hardware architectures for lightweight block ciphers. *Computers & Electrical Engineering*, 38(1):148 – 160, 2012.

[123] L. Knudsen, G. Leander, A. Poschmann, and M. Robshaw. Printcipher: A block cipher for ic-printing. In *Cryptographic Hardware and Embedded Systems*, volume 6225 of *Lecture Notes in Computer Science*, pages 16–32. Springer Berlin Heidelberg, 2010.

[124] S. Krawczyk and A. K. Jain. Securing electronic medical records using biometric authentication. In *International Conference on Audio- and Video-Based Biometric Person Authentication*, pages 1110–1119. Springer, 2005.

[125] P. Kumari and A. Vaish. Brainwave based user identification system: A pilot study in robotics environment. *Robotics and Autonomous Systems*, 65(0):15 – 23, 2015.

[126] R. D. Labati, R. Sassi, and F. Scotti. Ecg biometric recognition: Permanence analysis of qrs signals for 24 hours continuous authenticationsafie2011electrocardiogram. In *IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 31–36. IEEE, 2013.

[127] T. J. Lamer. Treatment of cancer-related pain: When orally administered medications fail. *Mayo Clinic Proceedings*, 69(5):473–480, 1994.

[128] J. R. Landis and G. G. Koch. The measurement of observer agreement for categorical data. *Biometrics*, 33:159–174, 1977.

[129] N. Leavitt. Mobile phones: the next frontier for hackers? *Computer*, 38(4):20–23, 2005.

[130] A. Lee and Y. Kim. Photoplethysmography as a form of biometric authentication. In *IEEE Sensors*, pages 1–2. IEEE, 2015.

[131] S. Lee, K. Fu, T. Kohno, B. Ransford, and W. H. Maisel. Clinically significant magnetic interference of implanted cardiac devices by portable headphones. *Heart rhythm*, 6(10):1432–1436, 2009.

[132] Y. K. Lee, K. Sakiyama, L. Batina, and I. Verbauwhede. Elliptic-curve-based security processor for rfid. *IEEE Transactions on Computers*, 57(11):1514–1527, Nov 2008.

[133] C. Li, A. Raghunathan, and N. Jha. Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system. In *13th*

*IEEE International Conference on e-Health Networking Applications and Services (Healthcom)*, pages 150–156, June 2011.

[134] C. Li, A. Raghunathan, and N. K. Jha. Improving the trustworthiness of medical device software with formal verification methods. *IEEE Embedded Systems Letters*, 5(3):50–53, 2013.

[135] Q. Li, C. Jin, W. Kim, J. Kim, S. Li, and H. Kim. Multi-feature based score fusion method for fingerprint recognition accuracy boosting. In *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA)*, pages 1–4, 2016.

[136] J. Lindqvist, S. Liimatainen, and T. Katajamaki. Secure pairing architecture for wireless mobile devices. In *IEEE 63rd Vehicular Technology Conference*, volume 2, pages 823–827, May 2006.

[137] R. T. Lukins, S. Tisch, and B. Jonker. The latest evidence on target selection in deep brain stimulation for parkinson's disease. *Journal of Clinical Neuroscience*, 21(1):22 – 27, 2014.

[138] U. Mahbub, V. M. Patel, D. Chandra, B. Barbello, and R. Chellappa. Partial face detection for continuous authentication. *arXiv preprint arXiv:1603.09364*, 2016.

[139] S. Mallat. *A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way*. Academic Press, 3rd edition, 2008.

[140] V. S. Mallela, V. Ilankumaran, and N. S. Rao. Trends in cardiac pacemaker batteries. *Journal Indian Pacing Electrophysiol*, 4(4):201–212, 2004.

[141] A. A. Mandavkar and R. V. Agawane. Mobile based facial recognition using otp verification for voting system. In *IEEE International Advance Computing Conference (IACC)*, pages 644–649, 2015.

[142] A. Mandel and M. Hamblin. A renaissance in low-level laser (light) therapy – lllt. *Photonics & Lasers in Medicine*, 1(4):231 – 234, 2012.

[143] E. Marin, D. Singelée, F. D. Garcia, T. Chothia, R. Willems, and B. Preneel. On the (in)security of the latest generation implantable cardiac defibrillators and how to secure them. In *32nd Annual Computer Security Applications Conference (ACSAC 2016)*, page to appear. ACM, 2016.

[144] D. Marohn. Biometrics in healthcare. *Biometric Technology Today*, 14(9):9 –11, 2006.

[145] R. Matta, J. K. H. Lau, F. Agrafioti, and D. Hatzinakos. Realtime continuous identification system using ecg signals. In *24th*

*Canadian Conference on Electrical and Computer Engineering (CCECE)*, pages 001313–001316. IEEE, 2011.

[146] A. McCallum, K. Nigam, et al. A comparison of event models for naive bayes text classification. In *AAAI-98 workshop on learning for text categorization*, volume 752, pages 41–48. Citeseer, 1998.

[147] D. K. McIver and M. A. Friedl. Estimating pixel-scale land cover classification confidence using nonparametric machine learning methods. *IEEE Transactions on Geoscience and Remote Sensing*, 39(9):1959–1968, 2001.

[148] S. J. Medical. Cardiac rhythm management products. `http://professional-intl.sjm.com/products/crm/pacemakers/dual-and-single-chamber`, Consulted on December 2017.

[149] Medtronic. Implantable pacemaker and defibrillator information. *Patient Services*, 1(800):551–5544, x41835, 2006.

[150] Medtronic. Cardiac rhythm products - pacemakers. `http://www.medtronic.com/for-healthcare-professionals/products-therapies/cardiac-rhythm/pacemakers/index.htm`, Consulted on December 2017.

[151] Medtronic. Parkison's disease. `http://www.medtronic.eu/your-health/parkinsons-disease/device/our-dbs-therapy-products/activaRC/index.htm`, Consulted on December 2017.

[152] H. Mehrotra, A. Rattani, and P. Gupta. Fusion of iris and fingerprint biometric for recognition. In *Proceedings of the International Conference on Signal and Image Processing*, pages 1–6, 2006.

[153] A. J. Menezes, S. A. Vanstone, and P. C. V. Oorschot. *Handbook of Applied Cryptography*. CRC Press, Inc., 1st edition, 1996.

[154] D. Miljkovic, D. Aleksovski, V. Podpečan, N. Lavrač, B. Malle, and A. Holzinger. Machine learning and data mining methods for managing parkinson's disease. In *Machine Learning for Health Informatics*, pages 209–220. Springer, 2016.

[155] B. Miller. Vital signs of identity [biometrics]. *Spectrum, IEEE*, 31(2):22–30, Feb 1994.

[156] C. Miller. Mobile attacks and defense. *IEEE Security Privacy*, 9(4):68–70, 2011.

[157] A. Nait-Ali. Beyond classical biometrics: When using hidden biometrics to identify individuals. In *3rd European Workshop on Visual Information Processing (EUVIP)*, pages 241–246, 2011.

[158] K. Niinuma, U. Park, and A. K. Jain. Soft biometric traits for continuous user authentication. *IEEE Transactions on information forensics and security*, 5(4):771–780, 2010.

[159] N. Noiseux, P. Khairy, A. Fournier, and S. J. Vobecky. Thirty years of experience with epicardial pacing in children. *Cardiology in the Young*, 14:512–519, 2004.

[160] I. Odinaka, L. Po-Hsiang, A. D. Kaplan, J. A. O'Sullivan, E. J. Sirevaag, and J. W. Rohrbaugh. Ecg biometric recognition: A comparative analysis. *IEEE Transactions on Information Forensics and Security*, 7(6):1812–1824, Dec 2012.

[161] E. Pagnin, G. Hancke, and A. Mitrokotsa. Using distance-bounding protocols to securely verify the proximity of two-hop neighbours. *Communications Letters, IEEE*, PP(99), May 2015.

[162] S. Pal and M. Mitra. Increasing the accuracy of {ECG} based biometric analysis by data modelling. *Measurement*, 45(7):1927 – 1932, 2012.

[163] R. Palaniappan. Multiple mental thought parametric classification: A new approach for individual identification. *Int. Journal of Information and Communication Engineering*, 2(4), 2006.

[164] D. Panescu. Emerging technologies [wireless communication systems for implantable medical devices]. *IEEE Engineering in Medicine and Biology Magazine*, 27(2):96–101, 2008.

[165] S. Papadopoulos, Y. Yang, and D. Papadias. Cads: Continuous authentication on data streams. In *Proceedings of the 33rd international conference on Very large data bases*, pages 135–146. VLDB Endowment, 2007.

[166] C.-S. Park. Mechanism based on hospital authentication server for secure application of implantable medical device. *BioMed Research International*, 2014:1–14, 2014.

[167] S. Park, Y. Kim, B. Urgaonkar, J. Lee, and E. Seo. A comprehensive study of energy efficiency and performance of flash-based {SSD}. *Journal of Systems Architecture*, 57(4):354 – 365, 2011.

[168] S. Patel, K. Lorincz, R. Hughes, N. Huggins, J. Growdon, M. Welsh, and P. Bonato. Analysis of feature space for monitoring persons with parkinson's disease with application to a wireless wearable sensor system. In *29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS)*, pages 6290–6293, 2007.

[169] D. Peralta, S. García, J. M. Benitez, and F. Herrera. Minutiae-

based fingerprint matching decomposition: Methodology for big data frameworks. *Information Sciences*, 408:198 – 212, 2017.

[170] P. Peris-Lopez, L. González-Manzano, C. Camara, and J. M. de Fuentes. Effect of attacker characterization in ecg-based continuous authentication mechanisms for internet of things. *Future Generation Computer Systems*, pages –, 2017.

[171] PhysioNet. Physiobank. National Institute of General Medical Sciences (NIGMS) and the National Institute of Biomedical Imaging and Bioengineering (NIBIB). `https://physionet.org/physiobank/`, 2017.

[172] D. Povey. Optimistic security: A new access control paradigm. In *Proceedings of the 1999 Workshop on New Security Paradigms*, pages 40–45. ACM, 2000.

[173] J. Radcliffe. Hacking medical devices for fun and insulin: Breaking the human. scada system. In *Black Hat. Technical Security Conference*, 2011.

[174] S. Ramírez-Gallego, B. Krawczyk, S. García, M. Woźniak, and F. Herrera. A survey on data preprocessing for data stream mining: Current status and future directions. *Neurocomputing*, 239:39 – 57, 2017.

[175] A. Rana and L. Sportiello. Implementation of security and privacy in epassports and the extended access control infrastructure. *International Journal of Critical Infrastructure Protection*, 7(4):233 – 243, 2014.

[176] S. Rane, Y. Wang, S. C. Draper, and P. Ishwar. Secure biometrics: Concepts, authentication architectures, and challenges. *IEEE Signal Processing Magazine*, 30(5):51–64, 2013.

[177] K. B. Rasmussen, C. Castelluccia, T. S. Heydt-Benjamin, and S. Capkun. Proximity-based access control for implantable medical devices. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, pages 410–419. ACM, 2009.

[178] K. B. Rasmussen, M. Roeschlin, I. Martinovic, and G. Tsudik. Authentication using pulse-response biometrics. In *The Network and Distributed System Security Symposium (NDSS)*, 2014.

[179] R. Raymond, D and S. F. Midkiff. Denial-of-service in wireless sensor networks: Attacks and defenses. *IEEE Pervasive Computing*, 7(1):74–81, 2008.

[180] K. Revett, F. Deravi, and K. Sirlantzis. Biosignals for user authentication-towards cognitive biometrics? In *International Con-*

*ference on Emerging Security Technologies (EST)*, pages 71–76. IEEE, 2010.

[181] M. R. Rieback, B. Crispo, and A. S. Tanenbaum. Rfid guardian: A battery-powered mobile device for rfid privacy management. In *Information Security and Privacy*, volume 3574 of *Lecture Notes in Computer Science*, pages 184–194. Springer Berlin Heidelberg, 2005.

[182] M. R. Rieback, B. Crispo, and A. S. Tanenbaum. Is your cat infected with a computer virus? In *Fourth Annual IEEE International Conference on Pervasive Computing and Communications*, pages 179–189, March 2006.

[183] A. Riera, S. Dunne, I. Cester, and G. Ruffini. STARFAST: a wireless wearable eeg/ecg biometric system based on the ENOBIO sensor. In *International Workshop on Wearable Micro and Nanosystems for Personalised Health*, pages 1–4, 2008.

[184] E. Rissanen, B. Firozabadi, and M. Sergot. Towards a mechanism for discretionary overriding of access control. In *Security Protocols*, volume 3957 of *Lecture Notes in Computer Science*, pages 312–319. Springer Berlin Heidelberg, 2006.

[185] M. Rostami, W. Burleson, F. Koushanfar, and A. Juels. Balancing security and utility in medical devices? In *Proceedings of the 50th Annual Design Automation Conference*, DAC '13, pages 13:1–13:6. ACM, 2013.

[186] M. Rostami, A. Juels, and F. Koushanfar. Heart-to-heart (h2h): authentication for implanted medical devices. In *ACM Conference on Computer and Communications Security*, pages 1099–1112. ACM, 2013.

[187] S. Russell and P. Norving. *Artificial Intelligence: A Modern Approach (3rd Edition)*. Pearson Education Limited, 2014.

[188] S. Saechia, J. Koseeyaporn, and P. Wardkein. Human identification system based ECG signal. In *IEEE TENCON*, pages 1–4, 2005.

[189] M. Salajegheh, A. Molina, and K. Fu. Privacy of home telemedicine: Encryption is not enough. *Journal of Medical Devices*, 3(2), 2009.

[190] H. A. Salam and B. M. Khan. Use of wireless system in healthcare for developing countries. *Digital Communications and Networks*, 2(1):35 – 46, 2016.

[191] R. S. Sandhu and P. Samarati. Access control: principle and practice. *IEEE communications magazine*, 32(9):40–48, 1994.

[192] C. Sandner and R. Amirtharajah. Power management. In *IEEE Custom Integrated Circuits Conference*, pages 1–1, Sept 2013.

[193] H. Savci, A. Sula, Z. Wang, N. Dogan, and E. Arvas. Mics transceivers: regulatory standards and applications [medical implant communications service]. In *IEEE Proceedings SoutheastCon*, pages 179–182, April 2005.

[194] S. Schechter. Security that is meant to be skin deep: Using ultraviolet micropigmentation to store emergency-access keys for implantable medical devices. `http://research.microsoft.com/apps/pubs/default.aspx?id=12213`, 2004.

[195] B. Schneier. Changing passwords. `https://www.schneier.com/blog/archives/2010/11/changing_passwo.html`, November 2010.

[196] B. Schneier. *Applied Cryptography: Protocols, Algorithms and Source Code in C / Edition 1*. John Wiley & Sons, Inc., 2015.

[197] J. Schulman. Stimulating and sensing network inside the human body. In *International Workshop on Wearable and Implantable Body Sensor Networks*, pages 95–98, April 2006.

[198] N. Security and P. C. (SPC). Break-glass: An approach to granting emergency access to healthcare systems, 2004.

[199] L. Seltzer. Securing your private keys as best practice for code signing certificates. `https://www.symantec.com/content/en/us/enterprise/white_papers/b-securing-your-private-keys-csc-wp.pdf`, 2013.

[200] J. L. Semmlow and B. Griffel. *Biosignal and medical image processing*. CRC press, 2014.

[201] M. Seyedi, B. Kibret, D. T. H. Lai, and M. Faulkner. A survey on intrabody communications for body area network applications. *IEEE Transactions on Biomedical Engineering*, 60(8):2067–2079, 2013.

[202] D.-H. Shih, C.-M. Lu, and M.-H. Shih. A flick biometric authentication mechanism on mobile devices. In *International Conference on Informative and Cybernetics for Computational Social Systems (ICCSS)*, pages 31–33. IEEE, 2015.

[203] L. Shin. How biometrics could improve health security. *Fortune*, 2015.

[204] S. Shivshankar and K. Summerhayes. *Challenges of Conducting Medical Device Studies*. Institute of Clinical Research, 2007.

[205] V. Shnayder, B.-r. Chen, K. Lorincz, T. R. F. Fulford Jones, and

M. Welsh. Sensor networks for medical care. In *Proceedings of the 3rd international conference on Embedded networked sensor systems*, SenSys '05, pages 314–314. ACM, 2005.

[206] K. M. Silay, C. Dehollain, and M. Declercq. A closed-loop remote powering link for wireless cortical implants. *IEEE Sensors Journal*, 13(9):3226–3235, 2013.

[207] H. Silva, H. Gamboa, and A. Fred. One lead ecg based personal identification with feature subspace ensembles. In *Machine Learning and Data Mining in Pattern Recognition*, volume 4571 of *Lecture Notes in Computer Science*, pages 770–783. Springer Berlin Heidelberg, 2007.

[208] H. M. Sim, H. Asmuni, R. Hassan, and R. M. Othman. Multimodal biometrics: Weighted score level fusion based on non-ideal iris and face images. *Expert Systems with Applications*, 41(11):5390 – 5404, 2014.

[209] R. N. Simons, D. G. Hall, and F. A. Miranda. Rf telheidu2011emetry system for an implantable bio-mems sensor. In *IEEE MTT-S International Microwave Symposium Digest*, volume 3, pages 1433–1436, 2004.

[210] K. Singh and V. Muthukkumarasamy. Authenticated key establishment protocols for a home health care system. In *3rd International Conference on Intelligent Sensors, Sensor Networks and Information*, pages 353–358, 2007.

[211] Y. N. a. Singh, S. K. Singh, and A. K. Ray. Bioelectrical signals as emerging biometrics: Issues and challenges. *ISRN Signal Processing*, 2012:1–13, 2012.

[212] C. Strydis, R. M. Seepers, P. Peris-Lopez, D. Siskos, and I. Sourdis. A system architecture, processor, and communication protocol for secure implants. *ACM Trans. Archit. Code Optim.*, 10(4):57:1–57:23, 2013.

[213] R. Sullivan and A. Ferriter. Prevent life-threatening communication breakdowns. *Nursing*, 38(2):17, 2008.

[214] S. Sun. Multitask learning for eeg-based biometrics. In *19th International Conference on Pattern Recognition (ICPR)*, pages 1–4, 2008.

[215] M. Suresh, P. G. Krishnamohan, and M. Holi. GMM modeling of person information from EMG signals. In *IEEE Recent Advances in Intelligent Computational Systems (RAICS)*, pages 712–717, 2011.

[216] M. M. Tantawi, K. Revett, M. F. Tolba, and A. Salem. On the use

of the electrocardiogram for biometrie authentication. In *8th International Conference on Informatics and Systems*, pages BIO–48–BIO–54, May 2012.

[217] R. Tarricone and A. D. Tsouros. *Home Care in Europe: The Solid Facts*. WHO Regional Office Europe, 2008.

[218] B. N. Technology. Medical device hack attacks may kill, researchers warn. `http://www.bbc.co.uk/news/technology-17631838`, 2012.

[219] TheVerge. Dick cheney had the wireless disabled on his pacemaker to avoid risk of terrorist tampering. `http://www.theverge.com/2013/10/21/4863872/dick-cheney-pacemaker-wireless-disabled-2007`, 2013.

[220] P. Tresadern, T. F. Cootes, N. Poh, P. Matejka, A. Hadid, C. Levy, C. McCool, and S. Marcel. Mobile biometrics: Combined face and voice verification for a mobile platform. *IEEE Pervasive Computing*, 12(1):79–87, 2013.

[221] K. K. Venkatasubramanian, A. Banerjee, and S. K. S. Gupta. Pska: Usable and secure key agreement scheme for body area networks. *IEEE Transactions on Information Technology in Biomedicine*, 14(1):60–68, 2010.

[222] K. K. Venkatasubramanian and S. K. S. Gupta. Security for pervasive health monitoring sensor applications. In *In Proceedings of 4th International Conference on Intelligent Sensing and Information Processing (ICISIP)*, pages 197–202, 2006.

[223] J. A. Von Arx and K. Najafi. On-chip coils with integrated cores for remote inductive powering of integrated microsystems. In *International Conference on Solid State Sensors and Actuators*, volume 2, pages 999–1002 vol.2, Jun 1997.

[224] D. Vouyioukas and A. Karagiannis. *Telemedicine Techniques and Applications*. Intech, 2011.

[225] F. Walker, S. C. Siu, S. Woods, D. A. Cameron, G. D. Webb, and L. Harris. Long-term outcomes of cardiac pacing in adults with congenital heart disease. *Journal of the American College of Cardiology*, 43(10):1894 – 1901, 2004.

[226] Y. Wang, F. Agrafioti, D. Hatzinakos, and K. N. Plataniotis. Analysis of human electrocardiogram for biometric recognition. *EURASIP J. Adv. Signal Process*, 2008, Jan. 2008.

[227] Z. L. Wang and J. Song. Piezoelectric nanogenerators based on zinc oxide nanowire arrays. *Science*, 312(57771):242–246, 2006.

[228] K. Warwick. Upgrading humans via implants - why not? `http://www.19.bbk.ac.uk/index.php/19/article/view/488`, 2008.

[229] J. Wayman, A. Jain, D. Maltoni, and D. Maio. *An introduction to biometric authentication systems*. Springer, 2005.

[230] J. G. Webster. *Design of cardiac pacemakers*. IEEE Press, 1995.

[231] I. H. Witten, E. Frank, M. A. Hall, and C. J. Pal. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, 2016.

[232] D. Wu, K. Warwick, Z. Ma, M. Gasson, B. J. G., S. Pan, and T. Aziz. Prediction of parkinson's disease tremor onset using a radial basis function neural network based on particle swarm optimization. *International Journal of Neural Systems*, 20(02):109–116, 2010.

[233] S. Xiao, A. Dhamdhere, V. Sivaraman, and A. Burdett. Transmission power control in body area sensor networks for healthcare monitoring. *IEEE Journal on Selected Areas in Communications*, 27(1):37–48, 2009.

[234] F. Xu, Z. Qin, C. Tan, B. Wang, and Q. Li. Imdguard: Securing implantable medical devices with the external wearable guardian. In *Proceedings IEEE INFOCOM*, pages 1862–1870, 2011.

[235] R.-F. Xue, K.-W. Cheng, and M. Je. High-efficiency wireless power transfer for biomedical implants by optimal resonant load transformation. *Circuits and Systems I: Regular Papers, IEEE Transactions on*, 60(4):867–874, 2013.

[236] A. Yakovlev, S. Kim, and A. Poon. Implantable biomedical devices: Wireless powering and communication. *IEEE Communications Magazine*, 50(4):152–159, 2012.

[237] J. Yang, Y. Shi, and J. Yang. Personal identification based on finger-vein features. *Computers in Human Behavior*, 27(5):1565 – 1570, 2011.

[238] M. Zhang, A. Raghunathan, and N. K. Jha. Medmon: Securing medical devices through wireless monitoring and anomaly detection. *IEEE Transactions on Biomedical Circuits and Systems*, 7(6):871–881, 2013.

[239] G. Zheng, G. Fang, R. Shankaran, M. Orgun, J. Zhou, L. Qiao, and K. Saleem. Multiple ecg fiducial points based random binary se-

quence generation for securing wireless body area networks. *IEEE Journal of Biomedical and Health Informatics*, PP(99):1–1, 2016.

[240] H. Zhu, X. R., and Y. J. High speed intra-body communication for personal health care. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 709–712, Sept 2009.

[241] B. Ziaie, M. Nardin, J. Von Arx, and K. Najafi. A single channel implantable microstimulator for functional neuromuscular stimulation. In *Proceedings 7th International Conference on Solid State Sensors and Actuators*, pages 266–269, 1993.