Low Power Sensing and Processing in Wearable Biomedical Devices for Personalized Health Monitoring

Thèse N° 9721

Grégoire Casimir Joseph SURREL

2019

EPFL
ÉCOLE POLYTECHNIQUE
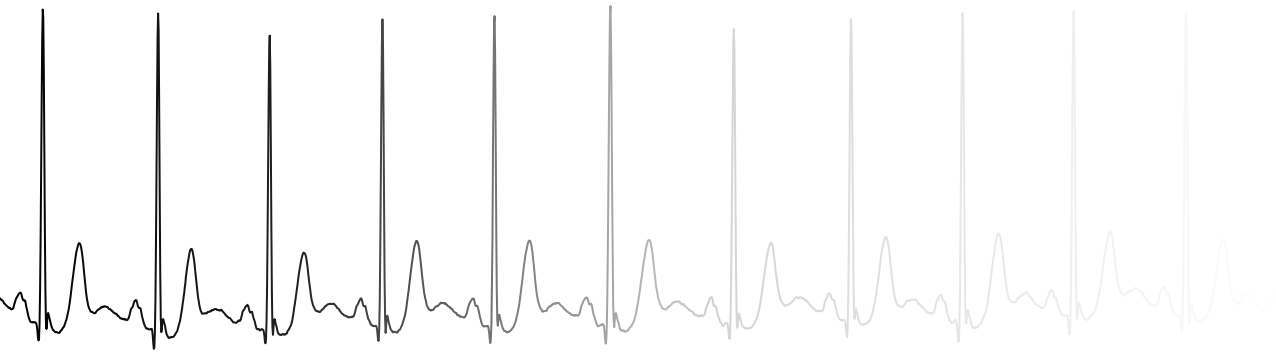FÉDÉRALE DE LAUSANNE

If the bees disappeared off the surface of the globe
then man would only have four years of life left
— Albert Einstein *(apocryphal quote)*

To you, my reader,

Acknowledgements

First of all, I need to thank all my jury members, Prof. Andreas Burg, Prof. Carlotta Guiducci, Dr Elisabetta Farella, and Dr Sani Nassif. They all kindly accepted this role, that implies dedicating a enough of time to evaluate this thesis and attend my private defense. I am very grateful and I appreciate it a lot.

Then, it is essential for me to thank Prof. David Atienza Alonso who is much more than my thesis director. It all started during my last year of master, with both an interesting course about programming applications for smartphones and tablets, as well as a semester project in his Embedded Systems Laboratory. This semester project was already more research than I noticed at first, as its purpose was to use a device originally intended to record electrocardio-grams as a means of tracking the muscle fatigue over time. This required the research and development of set of algorithms, enabling the live acquisition, processing and display of the signal. When I considered the possibility of applying for a PhD position, I certainly remembered the nice environment combined with interesting research projects ongoing in the laboratory. This is the visible result of the tremendous amount of work David puts into creat-ing a stimulating environment. It can be simply logistical help, or detailed administrative knowledge on how specific things should be done. David is also available to answer questions that can come up, and of course provide accurate guidance and helpful insights about the research being conducted. I believe there are many more things being done behind the scenes that let the PhD students in the laboratory focus on the research and teaching required.

How my PhD was not

Besides this very professional aspect, there is another essential humane interaction. David is a very understanding person and knows how to take care of mistakes or can provide help when more personal problems occurs. For all these reasons, my PhD was not the stereotypical hell that one might think of, but a very enriching experience. Of course there were challenges on the way, but it never turned into an unbearable duty to be done one day after another. For all of that, David, thank you very much, it was very much appreciated.

Then, I need to thank all the persons that followed me in my research, giving me pointers to go further. Some little implementation ideas grew into full research subjects. Other times, mundane questions turned out to uncover an interesting topic to explore. Overall, they provided me building blocks of my thesis, with knowledge in specific domains I did not have at the time. Finally, they allocated me some time to discuss my research, and, last but not least, proofread my writings. I really appreciated to work with all of them. Therefore, I need to thank all of you, in chronological order, Francisco Rincón, Srinivasan Murali, Amir Aminifar, and finally Tomás Teijeiro.

From the Embedded Systems Laboratory, everybody contributed to the nice atmosphere. Many came, others went. A full exhaustive list would be too long, but everyone had a role. Taking care of the administrative and technical tasks, Homeira Salimi, Francine Eglese and Rodolphe Buret were essential to get everything working smoothly, whether it is about planning trips for conferences or keeping all the equipment up and running. I had many good moments with my office-mates, Loris Duch, Yasir Qreshi and Alexandre Levisse. Among the other persons from the lab, I need to thank Dionisije Šopić for his liveliness and secrets, Fabio Dell'Agnola for being our trustworthy Superman, Wellington Silva De Souza for the insightful network-related brainstormings, Marina Zapater, Halima Najibi, Farnaz Forooghiar and Arman Iranfar for the nice teaching collaboration, and Miguel Peón for the chocolate. Ali Pahlevan enjoyed all the ducks I gave him, naming them one after another. There are many other reasons for acknowledgments but I will soon end up with an (almost) infinite list.

Not necessarily related to the workplace, many other people were important during my PhD. I want to thank all the people from the Rock Affinity dance association, my flatmates, my friends and my family for the support and

good times during these years. Sometimes, my PhD was also about new friendships. Pablo Garcia del Valle is one man of a kind, with whom I enjoyed music, rockets, hackerspaces, green tea, movies, etc. Also, among my friends, I can name Marie-Claire Ung for all the cultural outings, the trio made of Elisabetta De Giovanni, Артём Андреев (Tim), and Andrew Simon, with whom I spent countless valuable hours! Other times, encounters become much more than just friendship, and that role comes to Maude Serex. We found each other towards the end of my PhD research, with the joy of sleepless nights dedicated to writing the present manuscript. She certainly made these days more pleasant. To all of you, thanks for being part of my life.

Finally, nothing would have happened without my father, Yves Surrel, who whispered one day when I was finishing my master in electrical engineering:

"Have you considered doing a PhD?"

Lausanne, August 22, 2019   Grégoire Surrel
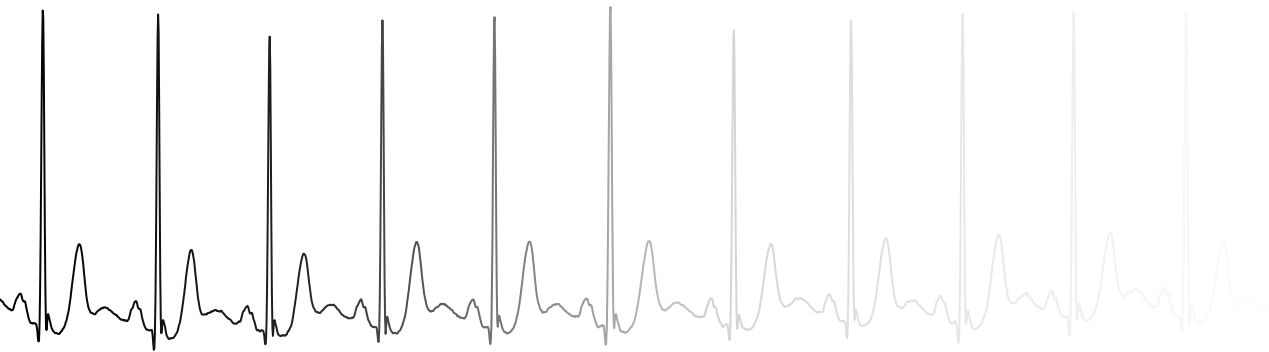
Abstract

WHETHER it is for personal use or for medical application, wearable sensors are becoming more and more widespread. This is the industry answer to two parallel trends. First, the public show a wish to collect data about their own lifestyle. This rather new effect appeared with the rise of smartphones, smartwatches and other heart-rate belts for athletes, with the promise of understanding and improving their health and performance. From a medical point-of-view, these commercial devices are not usable because their results come from unproven algorithms, far from clinical trials. However, health-care professionals see a strong benefit of having trustworthy wearable devices, which can be used by their patients for extended periods of time. Having *in situ* medical-grade data collection is able to give insights about the patient's health status and its development. There are however limitations to what can be performed.

Ideally, such a device should record as much data as possible, without requiring any set-up nor maintenance. Because of technical limitations, there are choices to be made. For example, which signals are captured, or what should be the minimal battery life? Any additional burden put on the patient hinders using the device, especially in the case of a non immediate life-threatening situation such as a daily screening.

State-of-the-art biomedical wearable devices deploy multiple strategies to match the specification's stringent requirements. First and foremost, from an energy point-of-view, it is expensive to sense signals. The more data

is collected, the smaller the battery lifetime. Therefore, extensive research is pursued beforehand to minimize the number of signals required. If a pathology can be reliably identified with less data, it is a meaningful saving. Additional energy savings are possible by minimizing the amount of data to transmit. Pushing this idea to the extreme, if the diagnostics algorithm can be efficiently run on the device itself, sending only the results is an excellent approach, which is being adopted by an increasing number of state-of-the-art sensor nodes.

In this thesis, I first propose event-driven approaches for sensing bio-signals, designed to take into account the signal's behavior. The rationale in changing the sampling strategy is based on the fact that oversampling is frequent, because this approach provides strong guarantees about the data quality. However, it is not taking into account the actual signal's temporal characteristics, where a constant value is digitized with the same rate as a high-frequency pattern. Taking into account the signal's evolution, several kinds of events are envisioned for triggering the measure, from a simple level-crossing strategy to a more refined knowledge-based one. This event-driven sampling paradigm minimizes the amount of data to process, leading to an increased battery lifetime.

Secondly, as a specific application case, I consider a widespread but underdiagnosed disease: obstructive sleep apnea. This disease is connected to increased risks of multiple cardiovascular diseases, motivating the need for early diagnosis and treatment. For improved applicability, the target device I consider is an existing wearable biomedical node. Following the optimization process described previously, I rely on a single biological signal processed directly on the sensor. Using machine learning, a limited number of features are extracted and filtered before being used for the obstructive sleep apnea assessment. Afterwards, the device only requires to send the results to a base station. This is a practical solution for a non-intrusive long-lasting obstructive sleep apnea screening.
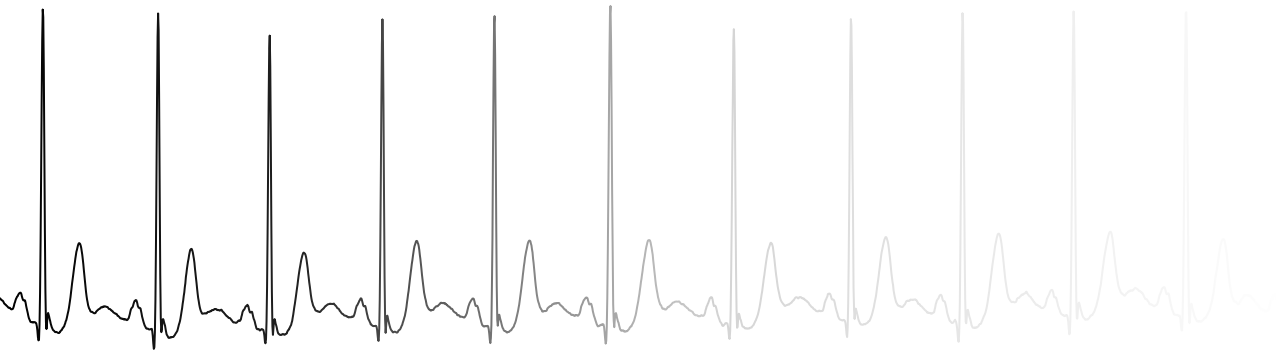
Résumé

QUE ce soit pour une utilisation personnelle ou pour une application médicale, les appareils mobiles sont de plus en plus répandus. C'est la réponse de l'industrie à deux tendances conjointes. Premièrement, le grand public montre le désir de collecter des données concernant son mode de vie. Cette tendance plutôt nouvelle est apparue avec l'adoption des téléphones multifonctions, montres intelligentes et autres ceintures cardiofréquence-mètres, avec la promesse de mieux comprendre et améliorer leur santé et performance. D'un point de vue médical, ces appareils commerciaux ne sont pas utilisables car leurs résultats proviennent d'algorithmes non éprouvés, loin d'essais cliniques. Néanmoins, les professionnels de la santé voient un important bénéfice à avoir des appareils portatifs fiables, qui peuvent être utilisés par leurs patients pendant des périodes temporelles étendues. Avoir une acquisition de données *in situ* approuvées pour une utilisation médicale est capable de donner des pistes concernant l'état de santé du patient ainsi que son développement. Il y a néanmoins des limitations à ce qui peut être fait.

Idéalement, un tel appareil devrait enregistrer autant de données que possible, sans nécessiter ni configuration ni maintenance. En raison de limitations techniques, il est nécessaire de faire des choix. Par exemple, quels sont les signaux mesurés, ou quelle devrait être l'autonomie minimale de la batterie? Chaque contrainte additionnelle ajoutée sur le patient limite l'utilisation de l'appareil, particulièrement dans le cas d'une situation sans danger de mort immédiat, tel qu'un suivi quotidien.

L'état de l'art des appareils biomédicaux portatifs déploie de multiple straté-gies pour satisfaire aux contraintes exigeantes du cahier des charges. Premiè-rement, d'un point de vue de la consommation énergétique, il est coûteux de capter des signaux. Plus il y a des données collectées, plus l'autonomie de la batterie est réduite. De fait, une recherche extensive est menée en amont pour minimiser le nombre de signaux requis. Si une pathologie peut être fidèle-ment détectée avec un ensemble réduit de données, il s'agit d'une économie sensée. Des économies supplémentaires sont possibles en minimisant la quantité de données à transmettre. En poussant cette idée jusqu'à l'extrême, si l'algorithme de diagnostic peut fonctionner efficacement sur l'appareil lui-même, alors n'envoyer que les résultats est une approche excellente, laquelle est adoptée par un nombre croissant de capteurs de l'état de l'art.

Dans cette thèse, je commence par proposer des approches événementielles pour la capture de signaux bio-médicaux, conçues pour prendre en compte le comportement du signal. La raison du changement de strategie d'échantillon-nage vient fait que le sur-échantillonnage est fréquent, car cette approche fournit de solides garanties concernant la qualité des données. Néanmoins, cela ne prend pas en compte les caractéristiques temporelles réelles du si-gnal, où une valeur constante est échantillonnée avec la même cadence qu'un motif à haute-fréquence. Prenant en compte l'évolution du signal, plu-sieurs types d'événements sont considérés pour déclencher la mesure, depuis une simple stratégie de franchissement de seuil à une plus avancée, basée sur la connaissance que l'on a du signal. Cet paradigme d'échantillonnage événementiel minimise la quantité de données traitées, ce qui mène à une augmentation de la durée de vie de la batterie.
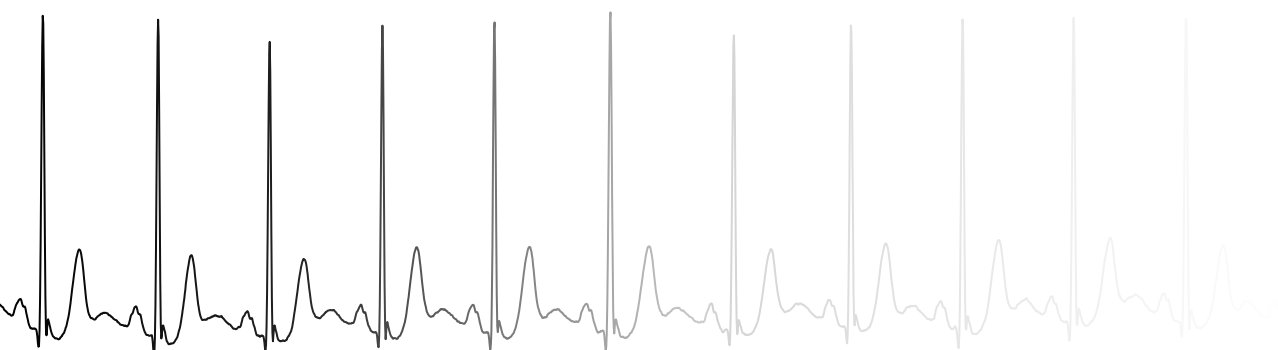
Deuxièmement, en tant qu'application concrète, je considère le cas parti-culier d'une maladie très répandue mais sous-diagnostiquée : le syndrome d'apnée obstructive du sommeil. Cette maladie est liée à des risques accus de développer de multiples maladies cardiovasculaires, ce qui motive la néces-sité d'un diagnostic et d'un traitement précoce. Pour une applicabilité accrue, l'appareil considéré est un capteur biomédical existant. Suivant le procédé d'optimisation décrit précédemment, je me base sur un signal biologique unique qui est traité directement sur le capteur. En utilisant l'apprentissage statistique, un nombre limité de motifs sont extraits et filtrés avant d'être

utilisés pour l'évaluation du syndrome d'apnée obstructive du sommeil. Ensuite, l'appareil a uniquement besoin d'envoyer les résultats à une station de collecte. C'est une solution pratique pour un suivi non-intrusif et de longue durée du syndrome d'apnée obstructive du sommeil.

Contents

Contents

List of Figures

List of Tables

Acronyms

Acronyms

Introduction

H EALTH-CARE has been morphing according to scientific, technological, and social changes, trying to answer to the patients' needs. With its progress, health threats changed, presenting new challenges to the scientific and medical community. Also, the population is aging, and causing the current health-care model to become unsustainable. Thankfully, technology has the means to provide the main solution for improving the situation.

1.1 The Way to Modern Health-Care

Coming from tribal beliefs and reaching evidence-base treatments, medicine has come a long way. The first known controlled experiment happened in 1747. At the time, scurvy was a widespread disease among sailors for centuries. During an expedition, James Lind had the idea to divide the sick into two separate groups observing two different diets. The one group having citrus fruits to eat saw a clear improvement of their health status [1]. This was a significant step towards a better understanding of the disease and how to prevent it.

Over the course of few centuries, the situation improved dramatically. Mankind has developed the knowledge and tools required to fight an increasing number of pathologies, with the direct result of a reduced infant mortality and lengthened life expectancy. Indeed, considering only the modern history between 1990 and 2013, the mortality rate before five years

per 1000 live births decreased from 90 to 46, while the global life expectancy increased by seven years, up to 71 [2]. There is however a rising criticism of the modern medical system, described as being too systematic, impersonal, and inadequate [3]. This criticism fuels some dangerous behaviors. One of them is turning towards *alternative medicines* or *complementary medicines*, where the treatments are commonly not better than placebo, and sometimes simply harmful. However, with practitioners usually dedicating more time and attention to patients, patients tend to prefer this more humane interaction.

This problem of disconnection between the medical community and the patients is well-known and acknowledged, up to recommending the integration of alternative medicine into the conventional framework [4]. With an improved human connection comes a better therapeutic relationship and thus a superior recovery. Nonetheless, it is highly discouraged to disregard conventional medicine, as it delays proper diagnostics and treatment. However, this is a curative approach of health-care, and still can be improved further.

Preventive health-care is more beneficial than curative care [5, 6]. However, because of the public healthcare system overload [7], an optimal curative care is difficult, let alone a preventive one which requires additional examinations. Without enough health-care professionals for in-depth routine checkups of an aging population, a technological approach can help in this regard. Tracking the most prominent risk factors along with the early signs of a developing disease, the diagnostic and treatment can start early. With an improved recovery, the therapy is lighter than when the condition is more developed and spread in the body because of a symptoms-based approach.

1.2   Shaping the Future of Health-Care

Worldwide, Non-Communicable Diseases (NCDs) became the first cause of death [8]. They reached an estimated 41 million death per year, with 15 million between the ages of 30 and 69 years. This is a group of diseases associated with sub-optimal environment and lifestyle, worsening the health status year after year. As shown in Figure 1.1, a sedentary lifestyle is a known

aggravating factor for NCDs, as it is linked with unbalanced diets and a lowered physical activity [9]. Part of this NCD category are heart-attacks, stroke, cancer, chronic obstructed pulmonary disease, asthma, and diabetes.

Modifiable behavioural risk factors, with their respective annual death toll

Tobacco   Unhealthy diet   Alcohol   Physical inactivity

7.2 million   4.1 million   3.3 million   1.6 million

Non-communicable diseases associated with behavioral risk factors

Cardiovascular diseases   Cancer   Other NCDs   Chronic respiratory diseases   Diabetes
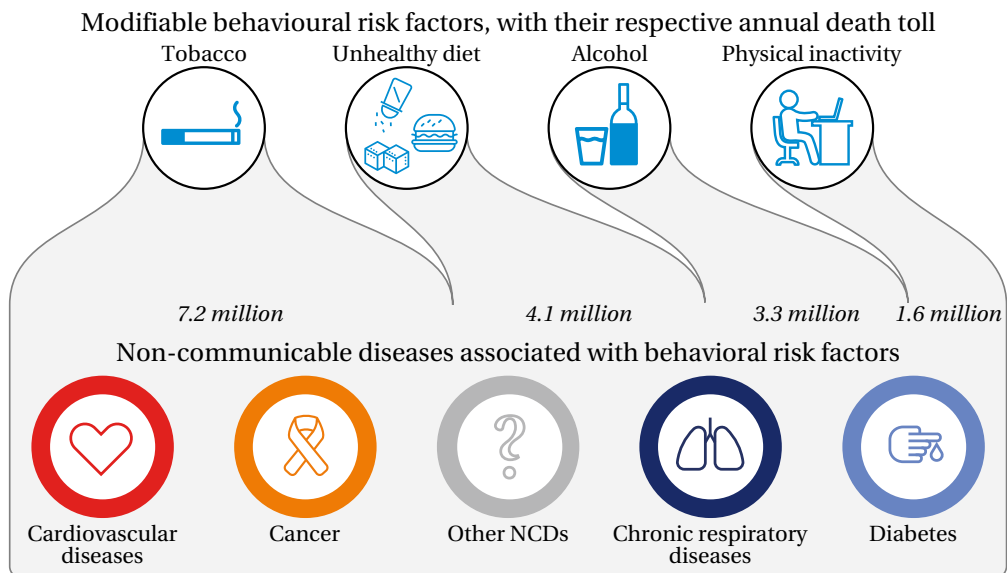
Figure 1.1 – Modifiable behavioral risk factors triggering non-communicable diseases, based on [10].

Among the population younger than 70, NCDs are responsible for 52% of deaths worldwide. From these NCDs, as depicted in Figure 1.2, 37% of them are cardiovascular diseases. With an aging population, because NCDs unfold with time, NCDs become a major concern. According to the United Nations, the number of people older than 60 increased by 48% between 2000 and 2015 [11], while the population older than 79 increased by 77%. Consequently, the incidence of NCDs saw a similar increase, with the World Health Organization estimating that cardiovascular diseases are the most prevalent diseases, causing more than 25% deaths worldwide [12].

As NCDs are partially preventable because of their link with behavioral risk factors [10], it is the perfect target for preventive care. Indeed, the current system is mostly focused on symptom-based medicine, with diagnostics performed on the visible effects of a disease. The most notable exception

Non-communicable diseases

Communicable maternal, perinatal and nutritional conditions

Injuries

| 52% | 34% | 14% |

| 37% | 27% | 23% | 8% | 5% |

Cardiovascular diseases

Malignant neoplasm

Other NCDs

Respiratory diseases

Diabetes mellitus
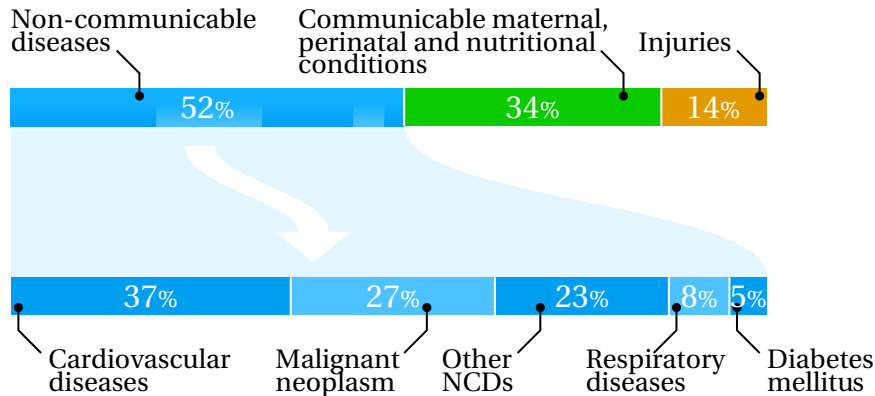
Figure 1.2 – Proportion of deaths worldwide for people under the age of 70, along with the proportion of each NCD. Most people are dying because of a NCD, and most likely because of a cardiovascular disease.

being vaccines, which are a brilliant example of preventive care in the context of transmissible diseases. Indeed, when considering only measles between 2000 and 2017, vaccination resulted in a 80% drop in measles deaths worldwide, with an estimated 21.1 millions deaths avoided [13]. In the context of NCDs, given their nature, one solution is careful long-term monitoring. However, with current medical systems, massive health-monitoring of an ever-increasing population with risks is unrealistic because of the resources required (human, facilities along their equipment). Therefore, it is necessary to find another solution for early diagnostics with a personalized approach, to provide an improved care.

1.2.1 Wireless Sensor Nodes

The *quantified self-tracking* trend, with people monitoring an increasing number of aspects of their daily life, is fueled by a multitude of devices and services [14]. This is connected to the development and spread of smartphones in the population on a global scale. Indeed, smartphones are handheld devices with a wireless connectivity that can be turned into a gateway to data collection. Optionally, with their increasing raw processing power, data processing can happen directly on the device, without the need to rely on any cloud computing. The activity trackers and sports trackers are the

most visible result of the research and development lead by the industry, with the promise of a better understanding and improvement of the user's lifestyle. The category of Wireless Sensor Nodes (WSNs) also includes devices dedicated to specific body-monitoring applications, such as the blood pressure or concentration of glucose in the blood in case of diabetes. Already widely available for the public, heart-rate monitors and pedometer are usable with sport-tracking applications. WSNs, together with a smartphone, form a Wireless Body Sensor Network (WBSN), as shown in Figure 1.3.
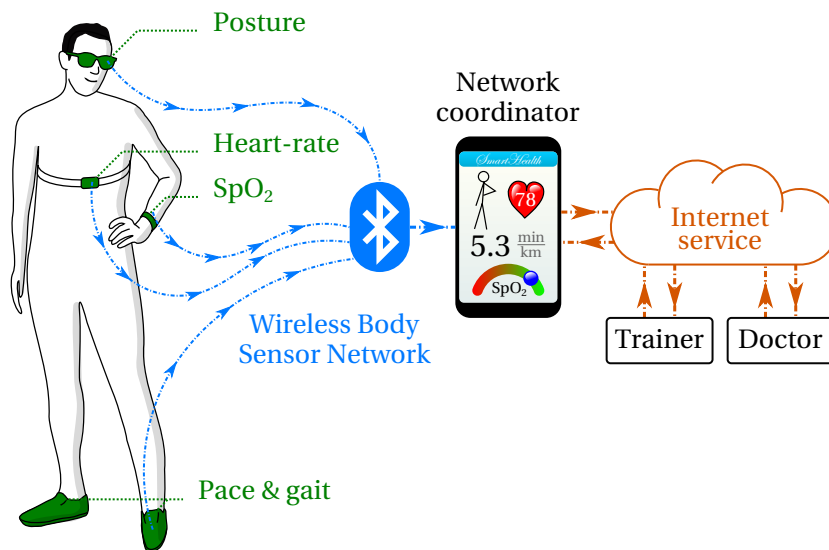
Figure 1.3 – Example of a WBSN for professional athlete training. Multiple sensors are used together to collect metrics about the person's performance. The network coordinator then uploads the data online, making it available to the trainer, or to the physician if needed.

Going further, medical-grade wearable devices are a key aspect of the future of health-care. Easy to use, easy to carry around, and more affordable than a full clinical setup, they are in the best position to collect relevant data for the physician. In the case of cardiovascular diseases, being able to collect and analyze the electrocardiogram (ECG) *in situ* is crucial as it provides meaningful data. When considering the immediate threat of stroke, foreseeing the attack greatly improve the safety of the patient [15–17].

Targeting human beings, such WSNs have multiple design constraints that can be summed up to a single one: the device ought to be usable, otherwise it will be dismissed. Therefore, the setup must be very accessible and easy, tolerant to mispositioning. The device must be unobtrusive, not preventing everyday life actions. The device must be relevant, yielding accurate data. Finally, the device must have a long battery-life as charging it is an additional burden for the user and ultimately data is lost if the battery is empty. In other words, WSNs must be convenient by being small and efficient. The required small size limits the battery size as well as the processing capabilities, therefore optimizing the energy-consumption is one essential for developing a viable and relevant product.

One solution for lowering the energy spent by the WSN is to process the data directly on the node. An application-agnostic device need to collect and transmit all the data it records, but any application-specific device can directly perform the required computation, transmitting only the final results. In the context of sport-tracking, cardiac monitoring devices only report the heart-rate. However, these devices need to measure the ECG in order to detect the beats. Transmitting the full ECG recording for remote processing is possible, but at the price of an increased energy-consumption from the additional burden put on the wireless link. However, as the savings in transmitting less data are significantly more important than the additional processing on the sensor [18], it is possible to provide non-intrusive sensors lasting for weeks rather than hours. This is a widespread approach that is used by a majority of application-specific WSNs.

Another solution for lowering the energy consumed by the sensor is to reduce how many data is collected. In the medical domain, for the ECG, the current medical standard is to have 12 channels for recording. Setting aside the long setup required and the increased design complexity, a 12 channels recording requires more energy for acquisition and signal processing. Keeping the same algorithmic performance for diagnostics with a lowered number of inputs is an optimization worth doing. Coming back to the cardiac monitor for sports, only a single channel is required to characterize the heart rate. The other 11 channels of a medical-grade recording are not required for this specific application. This approach is very efficient in reducing the overall

energy-consumption, but it is required to prove that the use of fewer features achieves at least a similar performance as the original solution. Any lowered performance can still be useful for performing diagnostics of a disease on a wider population, with the added risk of stress induced in case of false positive detection. However, as WSNs can be used for extended periods of time by the patients, the final diagnostic made by the algorithms can rely on more data, refining the estimate and improving the final diagnostic performance.

WSNs are therefore a promising solution for personalized health-care, provided they are developed to be as easy to use for the public while being medically significant. Doing so, it opens the way to a preventive patient-centric care. This approach brings benefits to both the patients, with early detection of diseases with a low impact on their daily lives, and to the medical system as a whole, easing the diagnostics and simplifying the cure thanks to an early treatment.

1.2.2 The Case of Obstructive Sleep Apnea

Obstructive Sleep Apnea (OSA) is a specific example of a widespread condition that can benefit from the WBSN technology. Indeed, OSA is a sleep disorder and an aggravating factor for high blood pressure [19], stroke [20], clinical depression [21], and lowered cognitive skills [22]. This translates to an increased risk of sudden death [23]. However, OSA must not be confused with the less prevalent Central Sleep Apnea (CSA), which is connected to a brain disorder preventing the normal breathing [24].

From 3% to 24% [25] of the population is affected by OSA, depending on the lifestyle. Worldwide, 5% of people are subject to OSA [25–27]. Even though OSA has a high prevalence in the population, an estimated 90% of the persons affected go undiagnosed [28]. With such prevalence and associated risks of developing a NCD, an early detection and long-term screening is highly desirable.

Physiologically, OSA is caused by an obstruction of the aerial airways. As seen in Figure 1.4, the path from the nose or mouth to the lungs is blocked. Depending on the situation, this condition is treatable, using a Continuous Positive Airway Pressure (CPAP) machine, which is an external equipment

forcing air in the lungs when an OSA episode is detected. In case of light to moderate OSA, a simpler mouthpiece or implant can be prescribed. A surgery is recommended when the patient has a severe OSA while not responding to the CPAP machine.
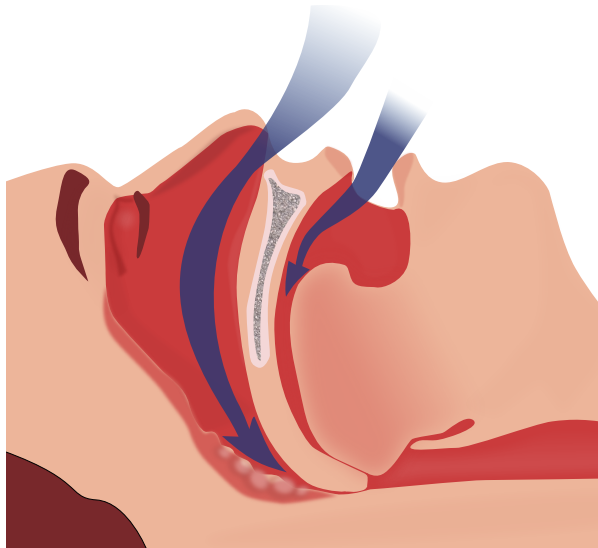
Figure 1.4 – The blocked airways in the throat prevent a normal breathing when sleeping. The air flow in blue is stopped before reaching the lungs. [Modified from Habib M'henni's work, Public Domain]

Diagnosing of OSA is difficult as it happens while asleep. Most of the time, the first consultation with a doctor is related to an excessive sudden snoring noticed by the partner. In the case of CSA, little to no snoring is reported. An OSA event consists of an absence of breathing for an extended period (generally longer than 20 seconds), often finishing with sudden gasping for air. Symptoms include daytime sleepiness, fatigue, decreased cognitive skills (attention, vigilance, concentration, motor skills, and memory), headaches, and lower libido. The most reliable solution for diagnosis is the polysomnography (PSG) for conducting a sleep study.

The PSG, while having the best performance in terms of signal quality, has some significant drawbacks. Because of the constraints, patients are reluctant to accept a PSG [28]. Even if the patient agrees, the sleep quality is impaired

by all the equipment attached to the body, and the change of environment is an aggravating factor [29]. Given the impracticality of PSG and because of the stakes of OSA, a database was released publicly to stimulate research on non-intrusive OSA detection, possibly in a home setup [30]. It contains 70 recordings of ECG with annotations of OSA events. With only a single ECG channel, a strong correlation was found between the heart's signal and OSA, while being less invasive and expensive than the traditional PSG. This is a way which can be used to an improved diagnostics, and, *in fine*, partially reduce the risks of NCDs.

1.3   Contributions

In this thesis, I bring contributions in two domains related to personalized health monitoring. This research is scoped as a multi-level optimization of WSN, at the sensing level on the one hand, by optimizing the amount of data captured by the biosensors, and on the application level on the other hand, by providing an improved system design for non-intrusive online OSA detection and screening.

1.3.1   Bio-signal Sensing

In the context of personalized medicine, wearable medical WSNs can become prevalent as they enable long-term non-intrusive health monitoring. Such devices must be both very efficient and not cumbersome to use, and therefore provide a long battery life. Since signal digitization is the main energy expenditure in optimized setups [31, 32], improving the sampling strategy is crucial for improving the patient's experience. Targeting ECG-based devices, I present in Chapter 2 multiple event-triggered solutions for wearable systems, and medical WSN in particular.

As an introduction, I present in Section 2.1 the widespread concept of uniform sampling, along with its problems of high energy consumption and data overload. Event-triggered sampling is another strategy which can improve the situation by taking into account the signal's temporal characteristics to only sample when required, thus lowering the energy spend by reducing the amount of data collected.

I propose three event-triggered strategies in Section 2.2. The first one is based on a simple level-crossing to trigger a measurement. My second proposal triggers a measurement as soon as the error between the linear interpolation of the previous measurements and the signal is greater than a threshold. Finally, my last method requires knowledge about the signal to sample. While the values conform to the expected outcome, it is possible to lower the sampling rate. However, when the signal varies more than the defined tolerance, the more data is collected. This knowledge-based method can be applied to ECG as it is a well-defined and regular signal. Compared to the level-crossing Time Analog-to-Digital Converter (T-ADC) approach, it has the main benefit of lowering the average sampling rate when the signal has a uniform variation, such as a ramp.

I present an evaluation of the performance of the three proposed approaches in Section 2.3. Each event-driven sampling strategy considered, and compared to compressed sensing as it is also used to reduce the amount of data handled by the system. Because reducing the data alone is not enough to prove the energy savings, I perform a detailed energy analysis for each part of the system.

Finally, in the context of wearable sensors, the data quality can be impaired by a very noisy signal. Depending on the application, it is worth to perform a signal correction before using it. In Section 2.4, I consider the context of the ECG in a sports tracker WSN and I explain the inner workings of a fast and lightweight heart-beat series correction.

1.3.2   Obstructive Sleep Apnea Screening

From the sensing optimization, I then proceed to the application level by focusing on the case of OSA for a multi-level optimization. Being an under-diagnosed but widespread risk factor for NCDs, OSA is a relevant candidate to target for non-intrusive diagnostics and screening. Existing solutions are impractical for multiple reasons such as inconvenience, cost, or constraints. Presented in Chapter 3, a non-intrusive wearable WSN reaches high performance for precisely detecting minutely OSA.

After an introduction to OSA and the existing algorithms for ECG-based detection of OSA in Section 3.1, I present the target device I consider in Section 3.2. Being a WSN of a WBSN, it has limited battery and processing resources. I, therefore, need to consider limitations during the entirety of the algorithm's design.

Next, I explain in Section 3.3 how the features are generated and evaluated offline, and I optimize the results in the online system. To accomplish this step, I select the best features to reach a high OSA diagnosis accuracy. The cost of extracting features from the ECG is variable, hence, I choose each features to be lightweight and optimized in order to save energy.

Once I selected the features, I process them as explained in Section 3.4, where I perform optimizations to enable the integration of the detection in the WSN considered for OSA detection. This is a multi-steps process, requiring me to pre-filter the ECG signal and post-filter the algorithm's results before yielding the final diagnostics.

Finally, I explain the experimental setup in Section 3.5, while I present the results in Section 3.6. Going further than the State-of-the-Art of OSA detection from single-lead ECG, I propose a specific tuning of the algorithm, taking into account each patient's physiological characteristics. This raises the quality of the analysis, with notable improvements for the few patients which were poorly diagnosed with the generic configuration.

Overall, my proposed OSA detection technique is compatible with existing online cardiac-monitoring solutions [33–37]. Adding a minimal Central Processing Unit (CPU) load, it can be combined with other algorithms, such as arrhythmia detection [38].

Sensing Bio-Signals

B IO-SIGNALS are signals that can be measured from a living being. Depending on the measurement, the signal can be of multiple natures. When measuring the electric potential difference (bio-potential) between two electrodes, it is an electrophysiological signal. The less invasive version uses electrodes on the surface of the skin, while more accurate ones require the setup of electrodes in the body. Bio-potential measurement is commonly used for the following applications and targets:

- electrocardiogram (ECG): heart muscle (the less frequent invasive version is called intracardiac electrogram),

- electromyogram (EMG): muscles, usually skeletal muscles,

- electroencephalogram (EEG): brain (the less frequent invasive version is called electrocorticography),

- electrooculogram (EOG): eyes muscles, for tracking eye movements,

- Other specific body parts can be measured but are less common (e.g. cochlea, stomach, olfactory epithelium, glottis, etc.).

It is possible to measure mechanical signals (e.g. body movements), magnetic ones (e.g. from the brain activity), acoustics sounds (e.g. knee joint friction).

Specific sensors are developed accordingly, providing an adequate ergonomic body interface and bio-compatibility.

In this thesis, I focus specifically on bio-potentials and especially on the ECG signal due to its practical importance. An example of a clean ECG is available in Figure 2.1. However, the general principle and techniques are applicable to any other signal.
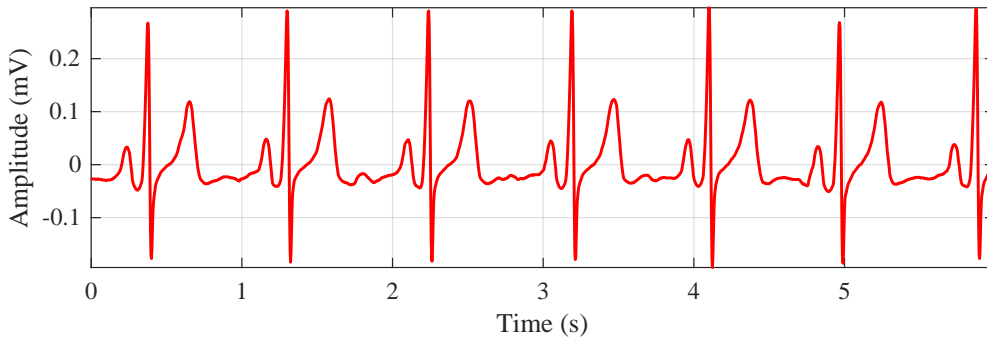
Figure 2.1 – Preview of 6 seconds of a real ECG. [Data source: Record A00848 from the PhysioNet Challenge 2017 database, between $t = 0s$ and $t = 6s$]

Before being able to perform a software analysis of a bio-signal, including the ECG, it needs to be converted from the analog domain to the digital domain. In this chapter, with the aim of lowering the energy consumed by this mandatory step of digitization, I explore the application of an unconventional event-driven approach to signal acquisition.

Additionally, ambulant sensing of bio-signals has an added complexity because of a degraded signal quality. I present an energy-efficient algorithm to quickly correct erroneous heart-beats detected from a noisy ECG signal.

2.1   Introduction

Sampling signals is the root of signal processing. Since the early days of digital signal processing, uniform sampling has been the most common approach. However, recent works have shown that it is not always the optimal solution and alternatives are possible.

2.1.1   The Data Overflow Problem

As global health improved, cardiovascular diseases became the most prevalent public health problem. In particular, the World Health Organization evaluated that coronary (or ischemic) heart disease combined with stroke account for 26.6% deaths worldwide [12], motivating the need for long-term heart monitoring.
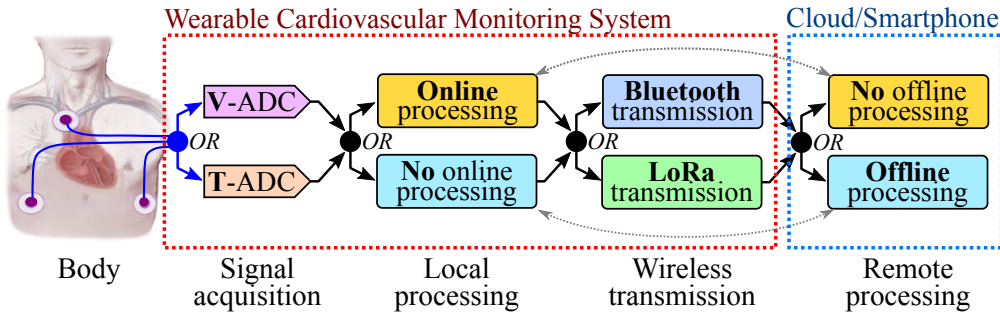
Figure 2.2 – Structure of the cardiovascular monitoring system, along with the final data sink on a smartphone or in the cloud. I consider in this work different possibilities of signal acquisition (uniform or event-triggered sampling), data processing (online in the sensor or offline on a remote device) and transmission (short-range Bluetooth Low-Energy (BLE) or mid-range Long Range Wide-area network (LoRa)). In total, there are eight different scenarios envisioned which can significantly impact the energy budget. The light-dotted arrows are constraints: the ECG processing needs to happen only once, whether it is online or offline.

Because of the costs of health-care professionals, their limited available time, and the requirements for facilities, such personalized healthcare is impossible without a technological solution. Lowering the constraints of cost and setup for the patient, non-intrusive Wireless Sensor Node (WSN) are a key solution to the problem. Thus, low-power cardiac monitoring wearable systems can have a major role in the future of global health. The functional structure of wearable WSN for cardiovascular monitoring in a possible Internet of Things (IoT) framework is shown in Figure 2.2, with different possibilities for signal acquisition, processing, and wireless data transmission.
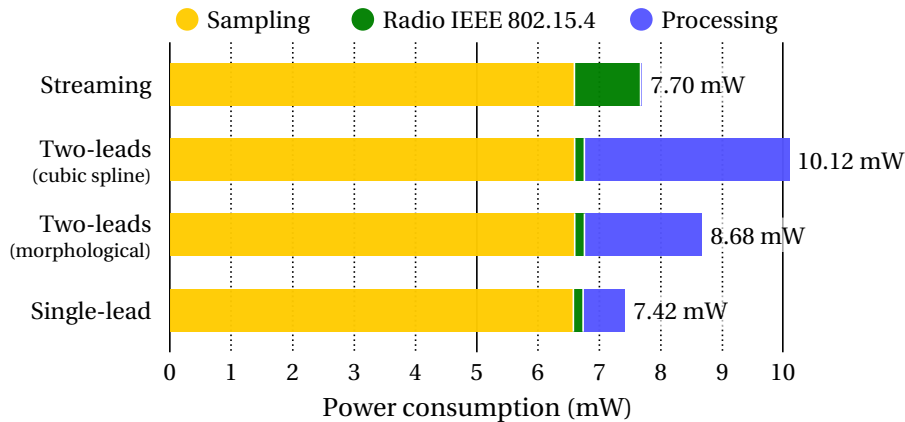
Figure 2.3 – Energy spent on a WSN for ECG-based cardiac monitoring per second. All the results are published in [31]. Four different strategies are considered for the application:

Streaming: in this scenario, the signal is sampled and streamed using the platform's low-power wireless link (IEEE 802.15.4). No signal processing or analysis is performed on the WSN itself.

Single-lead: in this scenario, a single lead of the ECG is considered. The amplitude and timing information of the nine fiducial points of the beat-beat are extracted, transmitting only the results on the wireless link. While bringing savings energy-wise, the improvement is marginal.

Two-leads: these two scenarios consider the analysis of the signal using two ECG leads. Using more channels improves the results' quality concerning the nine fiducial points. As in the single-lead scenario, only the results of the analysis are transmitted wirelessly. The two scenarios only differ by the pre-processing implemented, where the morphological filtering is an optimized approach compared to the cubic spline filtering.

As seen in Figure 2.3, prior art [31] shows that digitizing the ECG is the most expensive process for such optimized systems, whether the signal is processed on-node (65.2% of energy dedicated to sampling in the least optimized processing case) or streamed wirelessly (85.7%). Based on the fact that sampling the ECG systematically represents more than 50% of the energy budget, I

consider two very different approaches for signal acquisition, which will be presented hereafter.

Like in many industries, most bio-medical applications rely on collecting data according to the Nyquist-Shannon theorem. Indeed, this theorem guarantees that all the frequencies contained in the signal are perfectly acquired provided that the uniform sampling frequency for the Analog-to-Digital Converter (ADC) is at least twice the highest frequency in the signal. Conversely, if the sampling rate is insufficient, ghost frequencies appear in the sampled signal due to the spectral folding phenomenon.

However, even when respecting the theoretical constraints, uniform sampling has two main issues. First, digitization is always executed at constant rate that is at least twice higher than the highest frequency in the signal. Thus, this solution ignores the signal's temporal properties. Consequently, it leads to over-provisioned sampling as the signal does not constantly reach the maximum frequency allowed before spectral folding. Secondly, defining a high-enough sampling rate is not always possible.

The over-sampling effect of uniform sampling is critical in the context of resource-constrained wearable technologies for health monitoring [32, 39, 40] because more data samples means more energy required to measure, then process, then store or transmit the acquired data. One particular example of such resource-constrained wearable WSN is the *e-Glass* system [39] for real-time epilepsy monitoring, which has all the electronics hidden inside the glasses to reduce the negative impact of social stigma for the patients. Such wearable medical systems or IoT devices are extremely limited in terms of processing power, communication bandwidth, memory storage, and battery lifetime. Once the device's software is optimized to handle the amount of data coming from a uniform sampling with an efficient on-node processing, the signal digitization is the main energy expenditure [31, 32]. Changing the sampling paradigm is a solution for better energy efficiency and, *in fine*, longer battery lifetime for end-users.

2.1.2 An Old-New Hope for Efficient Data Sampling

From the *Animalia* kingdom, the brain is an event-triggered system, down to it's smallest component. Neurons react to action potentials (Figure 2.4) they receive, because this is an energy-efficient approach. However, event-triggering is not exclusive to biological systems. In the computing world, external interrupt signals on micro-controllers/processors mimick this approach.
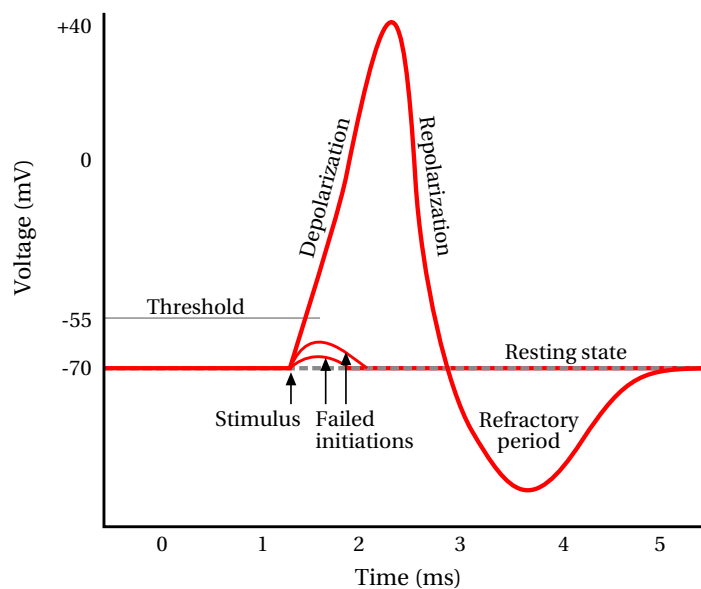
Figure 2.4 – Action potential generated by a neuron. An external trigger, called stimulus, will activate the neuron if it is higher than a threshold. In turn, the neuron will transmit the input signal to the following neurons by generating another action potential. For example, in the eye, a photon (the stimulus) will activate a rod cell (a photosensitive neuron) that will transmit the signal to the cortex. [Original by Chris 73, updated by Diberri. Modified for this thesis according to the CC-BY-SA 3.0 license.]

However, less has been done in the context of event-driven digital sensing. In particular, constant analog signals do not need new samples when considering that the absence of samples carry the information about the signal being invariant. Hence, new data points are required only if the signal is changing. Following this principle, the faster a signal changes, the faster the

signal is sampled. A Time Analog-to-Digital Converter (T-ADC), also known as Time-to-Digital Converter (TDC) in the literature, measure time between two events. Rather than sampling the signal value with regular time intervals (Figure 2.5a), a level-crossing T-ADC samples the time needed for the signal to change from one value to another (Figure 2.5b). The T-ADC levels can be separated by regular intervals, but they can also follow irregular intervals (i.e., exponential, squared, or even custom). When comparing a Voltage Analog-to-Digital Converter (V-ADC) and a T-ADC in Figure 2.5, there is a fundamental change in the process. It is however not easy to claim that one solution outperforms the other, as the performance of each sampling scheme is signal-dependent.

In the rest of this thesis, I will use the acronym V-ADC instead of ADC to prevent any ambiguity about the measurement done. Symmetrically, I will always refer to T-ADC for event-driven ADCs.

Section 2.2 presents three different event-driven approaches for sensing the ECG using a T-ADC. Besides a level-crossing solution, it is possible to consider different kind of events, based on the error between the sampled signal and the analog one. Going further, I show the possibility to dynamically update the error threshold depending on the knowledge of the ECG.

2.1.3   Data Correction

According to [41, 42], analyzing the Heart-Rate Variability (HRV) is key in cardiovascular analysis because the variation of the heart-rate contains indicators of current diseases or impending cardiac diseases. However, the algorithms need to work on long and accurate series of beats. With a patient in a hospital bed, the setup is ideal to record multiple hours of high quality ECG. It is nonetheless significantly more complex for people not hospitalized, as the screening must be minimally intrusive, therefore lowering the number of electrodes to use. Indeed, an ECG in a hospital usually relies on 12 leads, while a WSN should ideally use a single lead.

Whether the WSN relies on V-ADC or T-ADC sampling strategy, one major challenge is the real-time analysis of the single-lead ECG acquired in everyday-life in ambulatory conditions. Wearable devices experience more noise in the

(a) Voltage Analog-to-Digital conversion   *(b)* Time Analog-to-Digital conversion
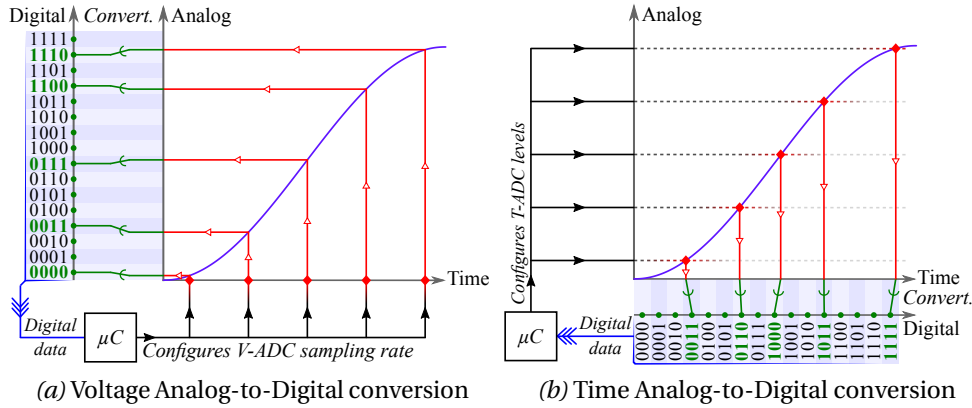
Figure 2.5 – Comparison of the working principle of the usual V-ADC (2.5a) and the event-triggered T-ADC (2.5b). • As a first step, the micro-controller μC configures the chip responsible for the analog-to-digital conversion, visible as a **black arrows** ➔.
• In the case of both V-ADC and T-ADC, the **red lozenges** ◆ shows the measurement trigger. In the V-ADC case, the trigger happens on regular time intervals, whereas in the T-ADC case, the signal itself triggers the sampling by crossing the configured thresholds.
• The measured value is visible following the **red arrows** ⇒ until they reach the axis.
• Whether it is voltage (V-ADC) or time (T-ADC), this analog value needs to be digitized, using a given number of bits, which establishes the resolution of the measure. This conversion is showed with the **green arrows** ⟶. It illustrates the origin of digital noise, making visible the rounding of each sample towards discrete digital values, which are represented as **green dots** ●.
• Finally, each sample's **binary value** (in green) is returned to the micro-controller μC following the **blue arrows** ⇛.

measured signal than a stationary acquisition from a medicalized bed, which motivates the need for robust signal acquisition and processing [43–47]. As the accuracy of the recorded signals is not optimal, this new generation of devices is more prone to algorithmic errors due to artifacts generated by, for

example, baseline wandering (illustrated in Figure 2.6) or muscular activity. In Section 2.4, I tackle the problem of erroneous heart-beats detected in a noisy signal by presenting a lightweight algorithm for on-line heart-beat classification and correction that relies on a probabilistic model to determine whether a heart-beat is likely to happen under certain timing conditions.
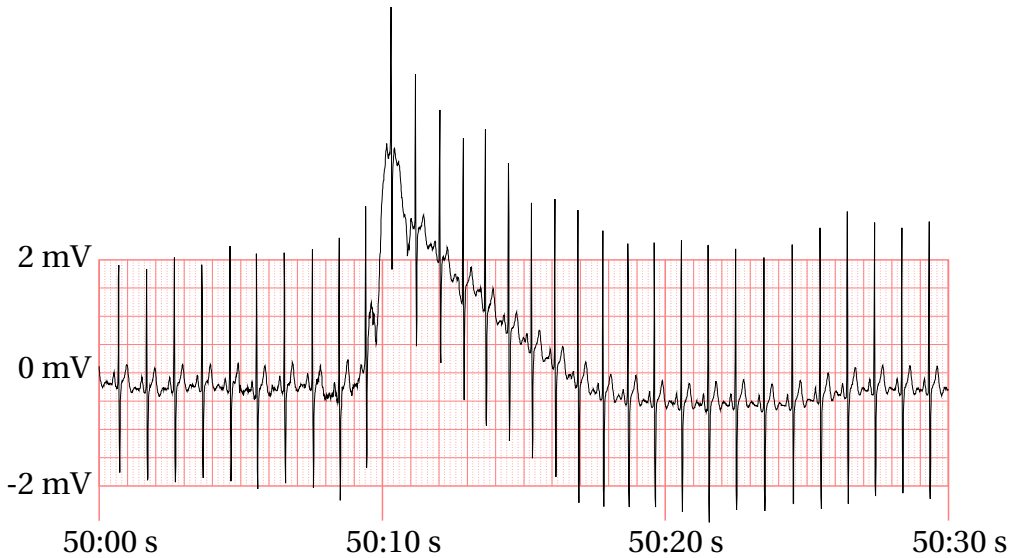
Figure 2.6 – Example of a 30 seconds extract from an ECG with baseline wandering. Measuring the ECG is subject to multiple sources of noise which impair signal processing. The case of baseline wandering is caused by a change in the recording resting potential originated from the patient's breath or a change in the electrode-skin impedance. [Data source: Record a05 of the Apnea-ECG Database, between $t = 3000s$ and $t = 3030s$]

2.2 Event-Triggered Paradigm in Sensing and Sampling for Electrocardiograms

Event-triggered sampling covers multiple approaches. The following subsections explain these different paradigms applied to the specific case of ECGs. However, in order to have an intuitive understanding of why the ECG is a suitable candidate for event-triggered sampling, a brief description is advised.

Each heart-beat in an ECG is observed as a sequence of three wave components (annotated in Figure 2.7), signals of the inner workings of the heat:

1. P wave: electrical activation of the atria (atrial depolarization), which results in atrial contraction, or atrial systole.

2. QRS complex: electrical activation of the ventricles (depolarization of the right and left ventricles) and contraction of the large ventricular muscles.

3. T wave: electrical recovery of the ventricles (repolarization).

The process of identifying each component from the ECG is called delineation.

ECG signals combine periods of high frequency and low frequency, because of the biological rhythm of the heart. This is visible on Figure 2.7, where each QRS complex is translated in the frequency domain with an increased power for frequencies between 25 Hz and 55 Hz. With specific frequencies in the range of 15 Hz to 25 Hz, the P waves and the T waves are also visible from the spectrogram. This is a scenario where T-ADCs should perform much better than V-ADCs, because the existence of high frequencies at the moment of the heart contraction (R peak) requires the V-ADC to sample at a high rate to ensure the quality of spectral analyses. Comparatively, between two beats, the ECG is composed of low frequencies.

2.2.1 Compressed-Sensing

Strategies similar to compressed sensing help lower the average sampling rate, even in the specific case of sampling the ECG [48]. As I am considering the same objective of reducing the amount of data, I am using the same implementation as a reference to compare the event-triggered solutions against.

There are however drawbacks to compressed sensing. First, as the decision is taken in advance whether to sample or not, they may statistically skip important data points. Understandably, this is a bet that can miss crucial events, namely irregular beats conveying signs of pathologies. Secondly, de-
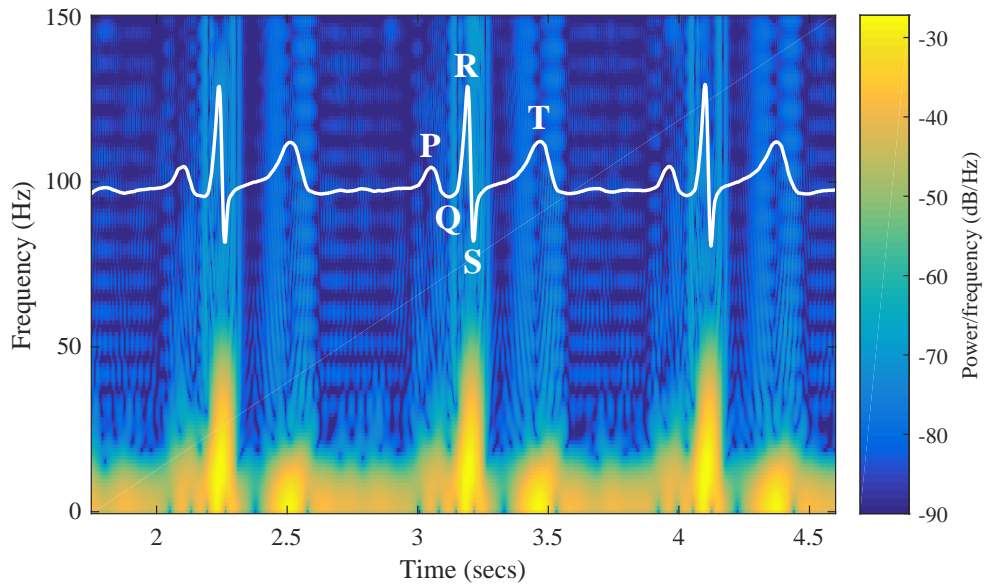
Figure 2.7 – The ECG, as a solid white line, is superimposed over its spectrogram.

Among the three heart-beats in the ECG, the middle one is annotated with each component: the P wave, the QRS complex, and then the T wave. The ECG's range is from -0.190 mV to +0.296 mV.

The evolution of the power for each frequency is illustrated in the background, where high power appear as bright yellow and low power in dark blue. [Data source: Record A00848 from the PhysioNet Challenge 2017 database, between $t = 1.75s$ and $t = 4.6s$, sampled at 300 Hz. Analysis settings: FFT window: 32 samples, window overlap: 31 samples, FFT size: 1024]

compressing the signal is computationally expensive as it requires finding the solutions to an under-determined system. With a costly signal reconstruction, compressed sensing cannot be performed on a low-power system. In comparison, an event-triggered sampling strategies also brings the benefits of reduced sampling rate, but without the speculative risk of choosing when to sample each value.

2.2.2 Level-Crossing Event Triggering

When comparing systems using different sampling strategies, I consider the distance between the linearly interpolated signal and the source signal as the error to minimize. In the two cases illustrated in Figure 2.8, I highlight the errors of both solutions, which are the results of the design constraints of the to approaches. On one hand, for the traditional V-ADC, a low sampling rate can miss the main R peak from the QRS complex. On the other hand, the T-ADC measures an accurate timing of the R peak, while accumulating error on the smaller T wave.

(a) Error for Voltage Analog-to-Digital conversion

(b) Error for Time Analog-to-Digital conversion
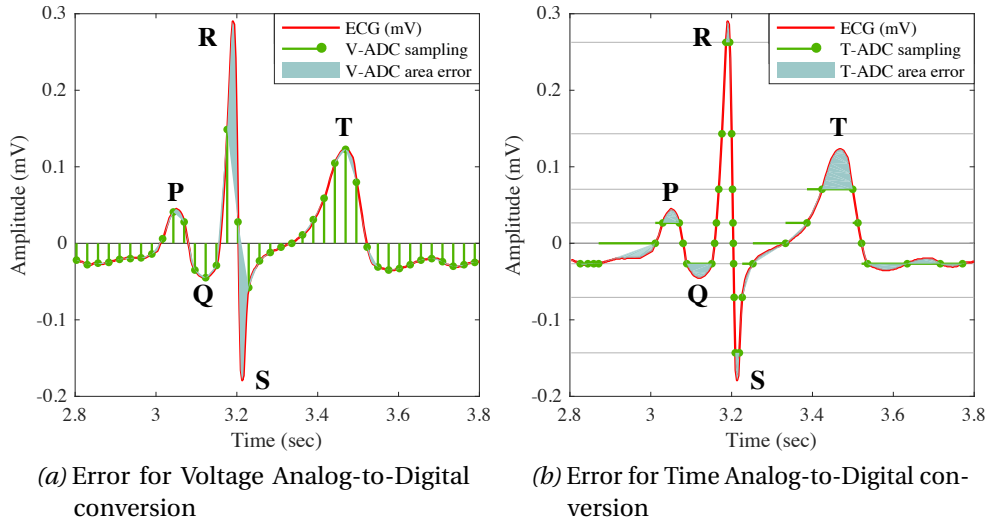
Figure 2.8 – Defining the reconstruction error between the classical V-ADC (2.8a) and the event-triggered T-ADC (2.8b).
Both the V-ADC and T-ADC scenarios show the same number of samples. Due to the low sampling frequency, the V-ADC signal is missing almost completely the R peak from the QRS complex, whereas the T-ADC, relying on an exponential scale of levels, captures the most important feature. [Data source: Record A00848 from the PhysioNet Challenge 2017 database, between $t = 2.8s$ and $t = 3.8s$]

Even though the ECG contains a lot of information, it is not always required to have the full information about the P, QRS, and T waves. Depending on the application, desired accuracy and use-case, partial data is sometimes enough

to compute the required results without reducing the quality of the analysis. For instance, in [49], the authors explain the methodology for embedded classification of heart-beats. The first part of the analysis is to classify if any heat-beat is normal or pathological, using a peak detector. Any suspicious beat detected triggers the activation of the three-leads recording mode, with full ECG delineation for further analysis.

This aforementioned Obstructive Sleep Apnea (OSA) detection system relies on a ECG-specific V-ADC ADS1191 from Texas Instruments. Two problems arise from lowering the sampling frequency for a coarser acquisition. First, it is impossible to significantly reduce the sampling frequency as the heart-beats need to be detected reliably. Secondly, the algorithm needs accurate timing between the heart-beats, otherwise the quality of the results decrease dramatically. Lowering the sampling frequency lowers the temporal resolution of the heart-beat detection. As a consequence, any energy saved by lowering the sampling frequency is paid with the price of a lowered performance.

However, switching to a T-ADC signal acquisition process is beneficial for two reasons. First, it becomes difficult to miss the heart-beat as it is the highest peak in the signal. It will therefore quickly cross multiple thresholds, clearly flagging its presence in the triggers received. Secondly, even with a coarse configuration (i.e. low number of thresholds), the precision is lost only in the peak height but the R-peak timing information is preserved.

The event generation from the level-crossing strategy starts with a simple demonstration concept in Figure 2.9. On the left part, a ladder of resistors acts as a voltage divider between *Vdd* and *Vss*. With *N* resistors, the ladder defines *N*-1 additional values, driving the sampling. While easy to understand, this simple design has several flaws. First, the levels are fixed at design time. A slightly improved version would replace *Vdd* and *Vss* by controllable voltages. This way, it becomes possible to change the value of all the levels at run-time. It is however impossible to dynamically change the individual spacing of each level: with constant-value resistors, a linear scale will always produce levels uniformly spread between *Vdd* and *Vss*. Last the series of resistors draws a constant current, along with the series of comparators. A more
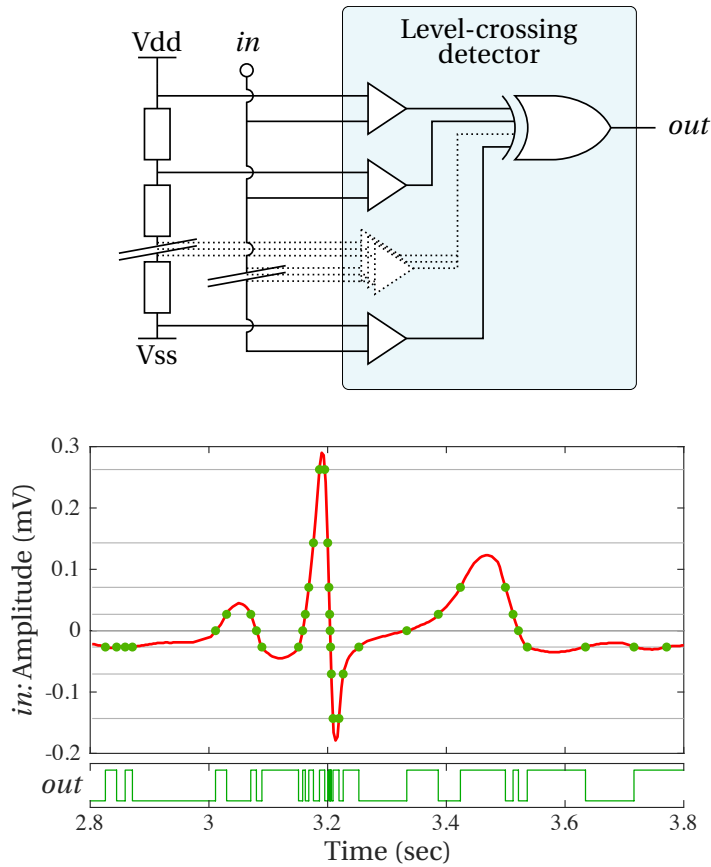
Figure 2.9 – Example of a simple analog front-end for a multi-level comparator generating the events required for time measurement. A feature this example front-end lacks is the capability to output which level is crossed by the signal, whether as an absolute value (e.g. "Signal is 3.27 V"), or as an evolution (e.g. "Signal rose up to the next level"). This trivial front-end can only be used to create edges to trigger time-measurements. The output signal for a small ECG input when resistances creating an exponential scale is shown in the lower half. [Data source: Record A00848 from the PhysioNet Challenge 2017 database, between $t = 2.8s$ and $t = 3.8s$]

clever design would rely on only two comparators with their voltage reference changed at run-time, tracking the signal level after level. An energy-efficient implementation exists [50], targeting specifically event-driven IoT sensor nodes.

2.2.3 Error-Based Event Triggering

Putting the focus on the signal reconstruction error, the event-triggered sampling task is a minimization problem. In the case of ECG acquisition, the objective is to obtain a digital representation of the analog signal that is sufficient for ECG processing, while reducing the global number of samples.

A family of methods well suited for this problem is polygonal approximation, also called piece-wise linear representation or linear path simplification [51]. These methods assume that the input signal can be represented as a sequence of linear segments, and they apply different techniques to obtain the minimum number of segments satisfying some error criterion.

Within this family, one method especially suitable for sampling time series is the Wall-Danielsson algorithm [52]. This method has linear complexity, works online, and only needs one signal sample in advance to estimate the approximation error. This method follows a bottom-up approach, in which points are merged to the current linear segment until an error threshold is reached, and then a new segment is created. Figure 2.10 illustrates the results with different thresholds, using the same ECG segment as in 2.8. In parallel of the ECG, each segment generated by the Wall-Danielsson algorithm (between black dots •) has its approximation error shown in the lower part. The error is measured as the signed integral between the original signal and the current segment, as illustrated in Figure 2.8.

Taking the case of the Figure 2.10a, the last segment (in purple, between plus signs +) is being created. Its error is increasing, and falling twice, because the small oscillations from the ECG cancel each other. At the precise point in time shown in the Figure (i.e. at the end of the plot), the error is close to 0 V.s, which means that the area from the dotted line and the ECG is the evenly spread above and below the between the dotted line.

Conversely, the Wall-Danielsson algorithm cannot guarantee optimality neither in the number of points nor in the selected samples. This is visible in Figure 2.10b, where the red asterisk symbols ∗ show samples taken on a straight line, and could be removed if the algorithm knew the future.

(a) Approximation threshold: 1 mV·ms *(b)* Approximation threshold: 0.4 mV·ms
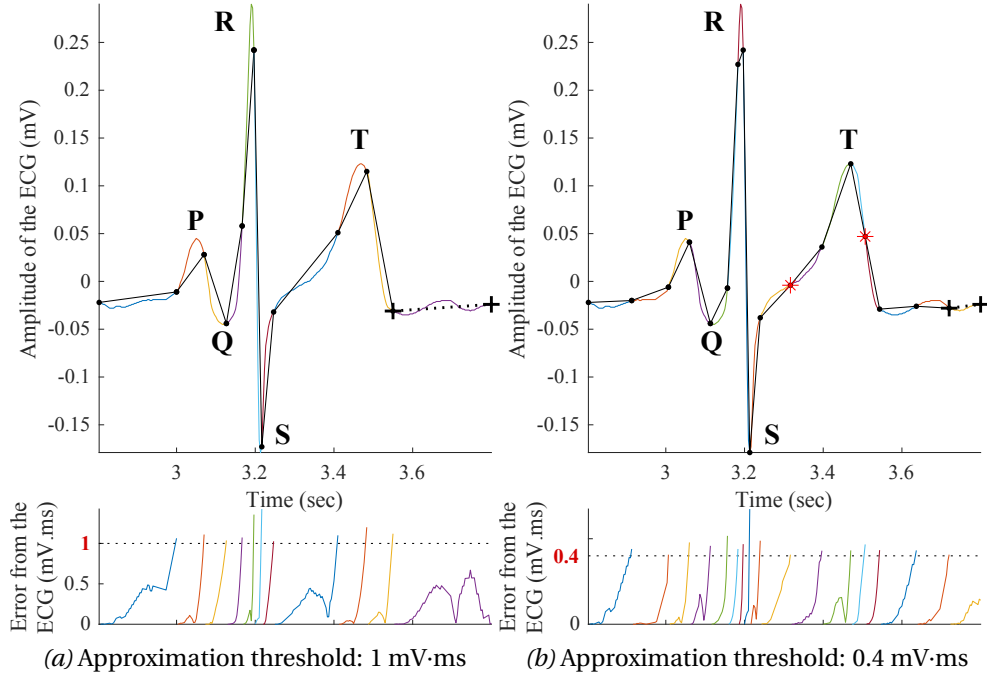
Figure 2.10 – Wall-Danielsson approximation (black line with round markers •) of the ECG with multiple thresholds, with a maximum allowed error of 1 mV·ms (sub-figure 2.10a) and of 0.4 mV·ms (sub-figure 2.10b).
Each colored ECG segment (upper part) is split by the algorithm when reaching the maximum error threshold. The corresponding error and the tolerance threshold is depicted in the lower part, with matching colors. Additionally, the dotted line between the plus signs + illustrate a segment being constructed by the algorithm, as its error is below the defined threshold. The red asterisk symbols ∗ are extra samples generated because of the non-optimality of the algorithm. [Data source: Record A00848 from the PhysioNet Challenge 2017 database, between $t = 2.8s$ and $t = 3.8s$]

This algorithm overcomes two main shortcomings of the classical level-crossing T-ADC method. First, in a level-crossing situation, an almost constant signal oscillating near a level generates more events than required. Second, fast linear changes generate numerous events. With polygonal approximation, the number of samples is not affected by constant displace-

ments of the signal level, and linear changes are always represented by just
two samples, no matter the slope value. I use this algorithm as the basis for
a novel adaptive-sampling method for ECG signals, which is explained in
Section 2.3.

The use of higher order approximation functions, such as quadratic or cubic
splines [53], is not suitable for this specific problem. The main reasons are the
absence of linear methods for building the approximation, the complexity of
the reconstruction procedure and the higher dependency on floating-point
precision.

2.2.4   Knowledge-Based Adaptive Sampling for ECG Signals

An important ECG feature that can be exploited to reduce the amount of cap-
tured data is the physiological regularity underlying the processes observed
in the signal trace. In particular, under a normal sinus rhythm situation, the
same heart-beat pattern is repeated at a frequency between 60 and 100 times
per minute. Thus, if this situation is detected on a signal fragment, from that
point onward it would be enough to capture just the information needed to
identify a change in the underlying rhythm.

This idea is illustrated in Figure 2.11, showing a 24 seconds ECG segment. As
long as I observe three regular P-QRS-T heart-beat patterns with a normal
distance between them, I can drastically reduce the detail of the signal just to
be able to check that the regularity is maintained. This results in a rougher
signal, but detailed enough to observe the regular heart-beats at the expected
time points. When an unexpected event breaks this regularity, the procedure
is able to increase the detail and support a more precise analysis of the new
situation.

The pseudo-code of the method for knowledge-based adaptive sampling is
shown in Algorithm 1. There are four main variables, sig', $llim$, $ulim$ and
$thres$, that correspond to the linearly interpolated sub-sampled signal, the
lower and upper temporal limits for the next expected QRS complex accord-
ing to the current rhythm, and the threshold for the polygonal approximation
algorithm. Two constants, mRR and MRR represent the minimum and maxi-
mum possible RR interval (distance between consecutive QRS complexes),

Figure 2.11 – Event-triggered adaptive sampling of an ECG fragment using polygonal approximation. Top: Original signal, sampled at 360 Hz. Bottom: Resulting signal of the adaptive sampling method. The detection of a regular rhythm enables a substantial reduction of the sampling frequency by getting a much coarse representation of the signal. After a rhythm change (second vertical line), the sampling frequency is increased to allow capturing the details of the abnormal area. [Data source: Record 119 of the MIT-BIH Arrhythmia DB, lead *MLII*, between 17:10 and 17:24]

and their values are obtained from the *gqrs* algorithm [54], which locates QRS complexes in an ECG signal.

Each cycle of the algorithm has four main steps:

1. obtain the signal up to $ulim$ (Line 8),

2. sub-sample it with the current threshold and, interpolate it (Lines 8-9)

3. look for the next QRS complex (Line 10),

4. update the approximation threshold and the time limits for the next QRS complex (Lines 12-32).

The sub-sampling function is called `WD()`, and it is a direct implementation of the Wall-Danielsson algorithm [52]. It receives as input a signal and an

Algorithm 1 Adaptive sampling algorithm. It takes as input an ECG signal sig and returns an event-triggered sampled version sig' obtained by polygonal approximation with dynamic thresholding. For compatibility with the QRS detector, sig' is linearly interpolated at the same frequency of sig.

1: **function** ADAPTIVE_SAMPLE(sig)
2:     **const** $mRR = 0.28s, MRR = 2.4s$
3:     **var** $sig' = []$ ▷ Resulting subsampled signal
4:     **var** $llim, ulim = 0, MRR$
5:     **var** $thres = 2mV \cdot ms$
6:     **var** $beats = []$ ▷ Dynamically detected QRSs
7:     **while** len(sig') < len(sig) **do**
8:         $x, y = $ WD(sig[len(sig') : $ulim$], $thres$)
9:         $segment = $ LINTERP(x, y)
10:       $qrs = $ GQRS($segment$)[0]
11:       $sig' = $ CONCATENATE($sig', segment$[: qrs])
12:       **if** qrs is null & LENGTH($beats$) > 1 **then**
13:           $thres = 2mV \cdot ms$
14:           $ulim = 2 \cdot ulim - beats[-1]$
15:       **else**
16:           $delineation = $ QRSDEL(qrs)
17:           $w = $ MAX($delineation.waves$)
18:           $nthres = 0.25 \cdot (w.amp \cdot w.dur/2)$
19:           **if** $nthres < thres$ **then**
20:               $thres = nthres$
21:           **else if** $qrs \geq llim$ **then**
22:               $thres = 0.9 \cdot thres + 0.1 \cdot nthres$
23:           **end if**
24:           $llim, ulim = qrs + mRR, qrs + MRR$
25:           APPEND($beats, qrs$)
26:           $rr_0 = beats[-1] - beats[-2]$
27:           $rr_1 = beats[-2] - beats[-3]$
28:           **if** $|rr_0 - rr_1| \leq 0.2s \lor 0.8 \leq \frac{rr_0}{rr_1} \leq 1.2$ **then**
29:               $llim = qrs + $ MIN($0.8 \cdot rr_0, rr_0 - 0.2s$)
30:               $ulim = qrs + $ MAX($1.2 \cdot rr_0, rr_0 + 0.2s$)
31:           **end if**
32:       **end if**
33:     **end while**
34:     **return** sig'
35: **end function**

error threshold, and returns the time and value of the selected points, that are subsequently interpolated by means of the `LINTERP()` function. The error threshold is updated as a fraction of the area of the largest wave of the observed QRS complex, which is estimated from its amplitude and duration (line 18). The delineation of the QRS complex and its constituent waves is made using the *qrsdel* algorithm [55] (line 16). After this, the time limits for the following QRS complex are updated considering the regularity of the observed rhythm (lines 24-30). The adopted definition of regularity has been taken from the *Construe* algorithm's implementation [56], which sets a maximum instant variation of the RR interval of a 20% or 200 ms (lines 28-30).

In order to avoid one-time increases that could lead to an excessive signal deformation, a conservative approach has been adopted for the updating of the polygonal approximation threshold. Thus, if the new threshold $nthres$ is lower than the previous one, the new one is immediately applied (line 20). However, if the new threshold is higher and the QRS was observed on time (between $llim$ and $ulim$), then the increase is applied in a 90%—10% fashion, requiring some time to get a large increase (line 22). The initial value selected for the threshold is 2 mV·ms. According to the experimental results shown below, a fraction of 0.25 for threshold update has shown a good compromise between the resulting sampling frequency and the QRS detection performance.

2.3   Experimental Evaluation

To demonstrate the potential of this knowledge-based adaptive sampling approach for ECG signals, I performed a series of experiments comparing the performance of a standard QRS detection algorithm using different sampling methods. The details of these experiments are explained next. The full non-aggregated results are available in Appendix A. , with more information on how to reproduce the results.

2.3.1   Dataset and Detection Algorithm

The database used in the experimental evaluation is the MIT-BIH Arrhythmia database available on PhysioNet [57], which is widely used in the literature to

evaluate QRS detection algorithms. This database contains 48 ECG records of 30 minutes duration sampled at 360 Hz. For my WSN-based ECG acquisition case I consider, the event-triggered sampling procedure is single lead. I decided to use the *MLII* lead, as it is the most frequent one in the database. Since records 102 and 104 do not have this lead available, they were not used for the experiments. Hence, the total time of the database is 23 hours.

The implementation of the *gqrs* algorithm used for QRS detection is from the WFDB software package from PhysioNet [58]. This enables me to use a openly available and trustworthy resource for unbiased performance assessment.

2.3.2   Compared Methods

Comparing each approach previously described requires a systematic process. As a reference, I consider the performance using the 360 Hz uniform sampling (US) from the database. I then applied the same methodology for the three approaches I considered, namely:

- Level-Crossing (L.C.) Sampling: to convert the database to a level-crossing scheme, I considered a regularly spaced set of levels. Whenever the signal crossed a level, both the time and level value were reported. Finally, I benchmarked the performance of a linear interpolation re-sampled at 360 Hz.

- Compressed Sensing (C.S.): I compressed the database using the same approach as defined in [48], with a compression ratio of 50%. After decompression, I obtained the initial sampling frequency of 360 Hz, not requiring additional processing.

- Knowledge-Based (K.B.) Sampling: as for the level-crossing sampling, I obtained series of times and values. This approach also required a linear interpolation at 360 Hz to run the performance benchmark.

2.3.3   Performance Validation

For assessing the influence of the event-triggered sampling strategy on the QRS detection algorithm's performance, Algorithm 1 was applied to

all records in the data-set that have the *MLII* lead available. Then, the *gqrs* algorithm was run both on the original signal and in the resultant sub-sampled and interpolated signal. The results of both runs were then compared with the manual beat annotations contained in the `.atr` files using the *bxb* application from the WFDB toolkit, using the F_1 score (harmonic average between the sensitivity and precision of the beat classification against the reference labels) as the performance metric.

Figure 2.12 shows how the overall F_1 and the average sampling frequency varies depending on the fraction constant I apply to update the approximation threshold (line 18 in Algorithm 1). As seen from the results, reducing the sampling frequency more than 10× (down to 34.10 Hz) has no noticeable performance penalty, and reducing it up to 37× (9.59 Hz) has less than 1% of F_1 decrease. From this curve, I select a constant fraction of 0.25 as a good compromise between a 26× reduction of the sampling frequency (13.60 Hz) and only a 0.06% of F_1 penalty. This is an acceptable performance decrease to qualify for medical certification [48], while minimizing the sampling frequency. It should be noted that in this scheme each sample has to be larger, since it has to encode both the time and the physical value of the signal.

The detailed sensitivity, positive predictivity, and F1 score for the 46 selected records in the database compared with the results in the original signal are available in Figure 2.13. All differences are highlighted, and correspond mostly to sensitivity decreases due to an excessive simplification of the signal. However, it is also possible to observe some increased positive predictability in the detection on the sub-sampled signals, probably due to noise removal during the approximation.

Then, Table 2.1 shows a performance comparison between the proposed method and the other sampling strategies, including ordinary uniform sampling, compressed sensing and event-triggered sensing. The considered performance metrics are specificity, positive predictivity and the combined F_1 score. The compressed sensing method has been applied as explained in [48], while the adopted level-crossing scheme is linear with a threshold every 200 μV. Results show that for a similar F_1 score, compressed-sensing halves the sampling frequency while level-crossing divides it by more than
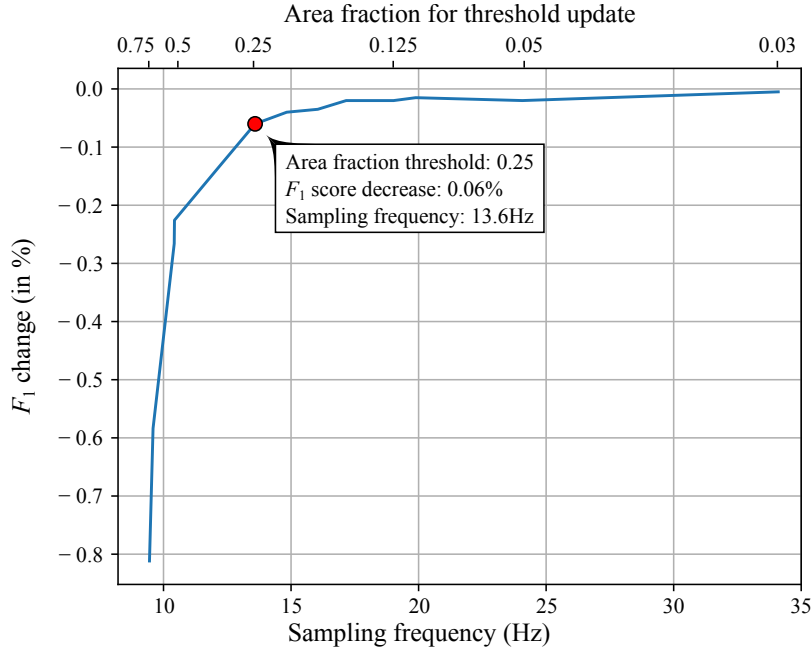
Figure 2.12 – Evolution of the QRS detection error with the average sampling frequency and the area fraction used for threshold update.

eight. The knowledge-based approach outperforms the two other strategies. In case a simpler hardware implementation, the level-crossing model should be considered as a first step.

2.3.4 Energy Consumption for Wearable Cardiac Monitoring Systems

To evaluate the energy consumption of different ADCs in the same conditions, I used a LaunchPad board from Texas Instruments [59], which is an affordable board hosting a MSP430 Micro-Controller Unit (MCU). It can accept plug-in modules called Boosterpacks [60]. The ADC chips were on BoosterPacks. Because the Launchpad and BoosterPacks do not provide any software nor hardware means to measure the energy consumed by the connected BoosterPack, I designed a custom printed-circuit board, seen in Figure 2.14. It consists of an intermediary board fitting between the Launchpad and its BoosterPack with selectable shunts resistors and test-points for easy access.
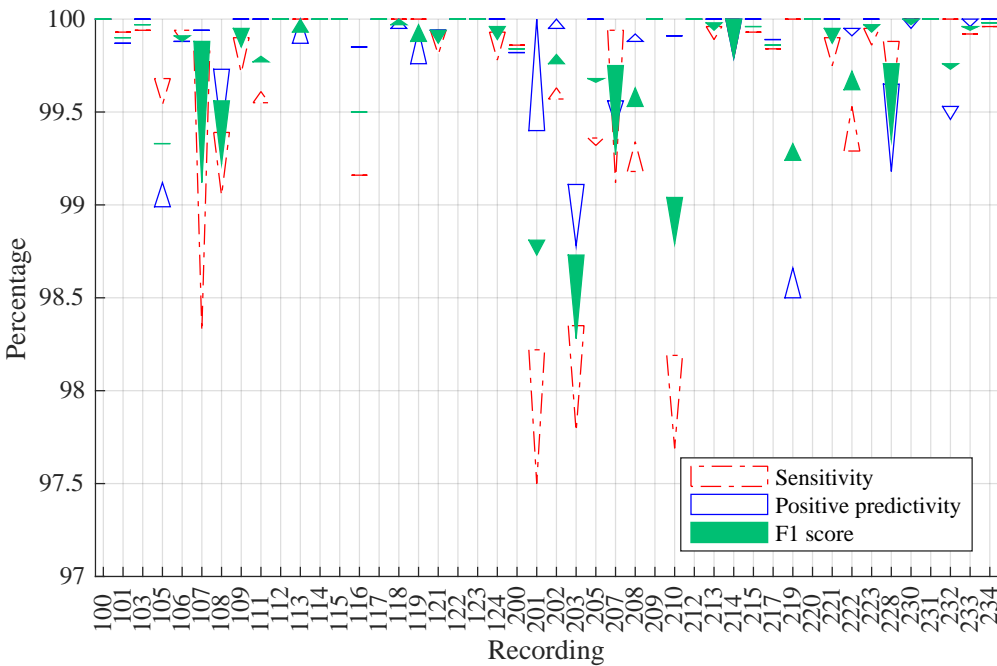
Figure 2.13 – Evolution of sensitivity (dashed red outline), positive predictivity (solid blue outline), and F1 score (filled green triangles) for each recording on the database. The base of each triangle is the result for uniform sampling, while the last vertex points to the result for a knowledge-based sampling with a constant ratio for dynamic threshold update of 0.25

I modeled the consumption of a full cardiac monitoring system active for 23 hours to match the database considered in Section 2.3.1. As illustrated in Figure 2.2, I studied multiple approaches to sampling, processing and transmitting data. All my results in Table 2.2 consider V-ADC or T-ADC (level-crossing or knowledge-based) sampling, embedded online processing or remote offline processing, and finally BLE [61] or LoRa [62] for wireless data transmission.

In all cases considered, the medium-range low-power LoRa wireless communication is higher-power than BLE, even though it becomes competitive with smaller payloads. This is therefore a solution to consider when no Bluetooth-enabled smartphone is available.

Sampling Strategy	Se	+P	F_1	fs (Hz)
Uniform Sampling (U.S.)	99.73	99.85	99.79	360.0
Compressed Sensing (C.S.)	99.64	99.82	99.73	180.0
Level-Crossing (L.C.)	99.66	99.83	99.74	43.7
Knowledge-Based (K.B.)	99.62	99.84	99.73	13.6

Table 2.1 – QRS detection performance comparison among different sampling strategies and resulting average sample rate for the 46 selected records from the MIT-BIH Arrhythmia DB.

Data Flow	Stream [J]		Online [J]	
Strategy	V-ADC	T-ADC	V-ADC	T-ADC
(from Table 2.1)	U.S.	L.C. / K.B.	U.S.	L.C. ≈ K.B.
V-ADC [50]	27.7	0.18	27.7	0.18
Sig. proc. [63]	-	-	23.9	23.9
BLE tx. [64]	56.7	11.8 / 9.94	7.06	7.06
LoRa tx. [65]	660	88.0 / 66.0	22.0	22.0
Total (BLE)	84.4	12.0 / **10.1**	58.7	31.2
Total (LoRa)	**687**	88.2 / 66.2	73.7	46.1

Table 2.2 – Model of the total energy consumption for the scenarios considered, with different sampling strategies (Uniform-Sampling V-ADC, Level-Crossing T-ADC, or Knowledge-Based T-ADC), signal processing (streamed to another device or local online analysis), and wireless communication protocols (BLE or LoRa). The values are reported for 23h of run-time.
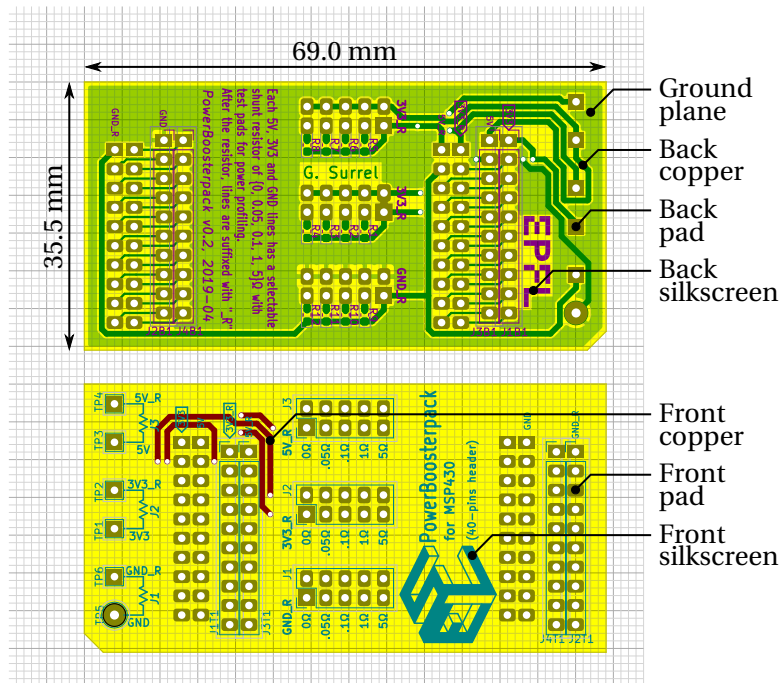
Figure 2.14 – Design of a Boosterpack for the TI Launchpad board. This custom board is designed to fit in-between a TI Launchpad board with the 40-pins Boosterpack header and a Boosterpack. Selecting the shunt resistor for each power rail (including the ground line), it is possible to precisely measure the current drawn by the Boosterpack.

Focusing on the data acquisition and processing, the most straightforward approach of streaming the V-ADC sampled ECG signal is the least efficient. My results match the conclusions in [31], where ECG streaming consumes more than on-device delineation. However, a paradigm shift occurs with state-of-the-art T-ADCs devices. Indeed, the much lower sampling frequency enables savings both at the sampling level and at the transmission level. Paying examining to the results, the level-crossing T-ADC is very close to the most advanced knowledge-based T-ADC, even though the sampling frequency is higher (43.7 Hz vs 13.6 Hz). This is because in the knowledge-based strategy, each sample has 10 bits for time and 12 bits for voltage, whereas the level-crossing strategy only requires one bit to encode the direction of the signal change as the levels are known. Therefore, a direct streaming of the T-ADC

values for remote processing is the most efficient solution. From these results, wearable WSN battery lifetime is multiplied by a factor of four to five.

2.4   Noise-Robust Correction of Beat Series

One important challenge of WSNs is to handle highly variable noise levels without affecting the final application performance. In that sense, there are some scenarios in which we can relax (or strengthen) our initial assumptions, increasing the robustness of the system to noise. In the context of sport trackers, tracking the quickly changing heart-rate is essential, even though the physical activity generates a noisy signal. Similarly, some games rely on the player's heart-rate [66] to change the game-play and therefore need to update in a matter of seconds. In this context, the assumption of a normal heart-rhythm is valid. That is to say, arrhythmia or ectopic (additional) beats are not to be expected. As a consequence, it becomes possible to take advantage of this information to build a fast and efficient heart-beat correction algorithm to integrate in a system that is robust to the presence of noise or artifacts in the signal.

For this use-case targeting healthy users, I designed an algorithm that can quickly decide if a beat is occurring at an expected time or if there is a problem in the heart-beat series (e.g., a skipped, an extra or a misplaced beat). When an error is detected, the series is repaired accordingly. The presence of error in the heart-beat series adversely affects the HRV [67], which is one metric used by sports tracking solutions when individual beat information is available. Therefore, any occurrence of erroneous beat should be corrected, regardless of the spectral estimation technique [67]. In the situation described in this section where the recordings are done in the field where motion artifacts are expected, correcting the series before its processing is a necessary step. An automated correction of the series is interesting for extracting the HRV indices live for an immediate feedback for the user.

Considering that simply removing the problematic parts of the heart-beat series skews the resulting HRV indices, there is a need for an more sophisticated approach. Using a probabilistic model of the heart-beats occurrences, the

proposed algorithm classifies and corrects each RR interval (see Figure 2.15) with respect to four known situations [68].
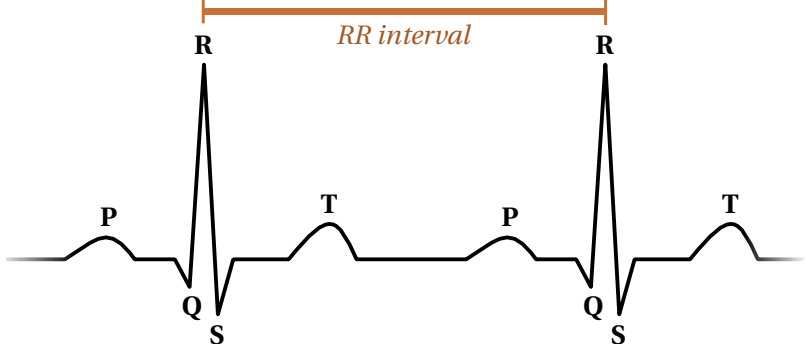
Figure 2.15 – Schematic diagram of the ECG of a normal sinus rhythm for a human heart with the P-wave, the QRS-complex and then the T-wave. [Image modified from the work of Agateller (Anthony Atkielski), Public Domain.]

2.4.1   Reference Algorithm

The reference algorithm is described in [68], where the RR intervals between heart-beat events are modeled following an Inverse-Gaussian (IG) probabilistic distribution

$$ig(x; \mu, \lambda) = \sqrt{\frac{\lambda}{2\pi x^3}} \times \exp\left(\frac{-\lambda(x - \mu)^2}{2\mu^2 x}\right),$$

where $\mu > 0$ is the mean and $\lambda > 0$ is the shape factor. The μ and λ parameters are estimated by maximizing the log-likelihood applying an exponential decay over the recent history of recorded events. They are computed using the following equations [69]:

$$\hat{\mu} = \frac{\sum_{i=1}^{n} w_i RR_i}{\sum_{i=1}^{n} w_i},$$

$$\frac{1}{\hat{\lambda}} = \frac{1}{n} \sum_{i=1}^{n} w_i \left(\frac{1}{RR_i} - \frac{1}{\hat{\mu}}\right),$$

where $w_i = w(\tau_i) = e^{-\alpha\tau}$ which is the weighting function and RR_i is the i^{th} time interval between two beats.

2.4.2 Beat Classification and Correction

In the current implementation, there are four different situations that can be identified and corrected if needed. The corrected beat is considered in the IG beat-classification model for the next beat. For each beat, the four different outcomes are evaluated and scored according to their respective likelihood of occurring. After a relative weighting of the individual scores, the highest one determines the final classification. Figure 2.16 shows the different cases, with the corresponding distributions. Given the current beat u_k, the next beat u_{k+1} is expected near the first and tallest IG (solid green line). If the next beat u_{k+1} happens before this time, it will have a low likelihood, linked to an extra beat. Finally, if the next beat u_{k+1} is measured near the maximum likelihood of the second IG distribution (dashed blue line), it is likely to reveal a missed beat.

Normal beat N: In the case of a normal beat, the newest RR interval will match closely the expected timing estimated from the recent history. Therefore, the score for beat i being a normal beat is $p_N = ig(RR_i; \mu_{i-1}, \lambda_{i-1}) - \eta_N$, where η_N is a constant relative weighting factor. In Figure 2.16 this occurs when the new beat u_{k+1} is near the maximum of the probability density of the IG function.

Extra beat e: If an extra beat u_{k+1} is in the series, the next beat u_{k+2} will be close to the expected timing (i.e., near the probability density peak). Thus, $p_e = ig(RR_i + RR_{i+1}; \mu_{i-1}, \lambda_{i-1}) - \eta_e$, where η_e is a constant relative weighting factor. This is the second case in Figure 2.16, where the new beat u_{k+1} is close to a zero-probability, that is to say very unlikely. An extra beat is corrected by removing it from the series, that is to say by considering the RR interval between the beat u_k and u_{k+2}.

Skipped beat s: A skipped beat is detected if the RR interval matches the expected time of two beats. In other words, its score is $p_s = ig(RR_i; 2\mu_{i-1}, 2\lambda_{i-1}) - \eta_s$, where η_s is a constant relative weighting

41
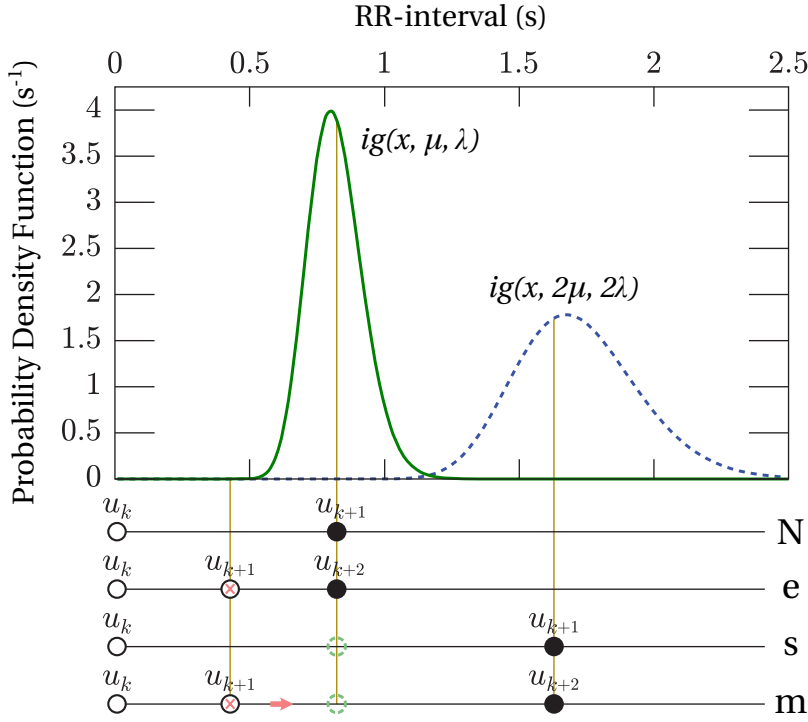
Figure 2.16 – Score for each heart-beat correction situation, with normal beats **N**, extra beats **e**, skipped beats **s**, or misplaced beats **m**.

factor. This is the third case in Figure 2.16, where the beat u_{k+1} is at the second maximum probability peak.

A skipped beat is corrected by inserting a beat by splitting the RR interval in two. The approach used to insert a beat is important for two reasons. First, the insertion should be as close to the real position as possible to keep a high quality HRV analysis. Second, as the inserted beat influences the future IG modeling, any error will also penalize the future beat classification and correction. Therefore, multiple strategies are considered:

1. splitting the RR interval in two equal parts so $RR_i = RR_{i+1}$,

2. setting the beat to the maximum probability computed from the previous beat so $RR_i = \frac{\sqrt{4\lambda^2\mu^2+9\mu^4}-3\mu^2}{2\lambda}$,

3. iterating its position using a Newton-Raphson approach to maximize the likelihoods of the missed and the next beats together,

4. have as a compromise an average between solutions 1 and 2.

As the skipped beat correction strategy influences the classification of the next beats, the performance for the classification and correction needs to be evaluated. The results of this assessment is visible in Figure 2.17 where the strategies 1 and 2 already offer a good performance (small distance between the real beat position and its reconstruction). Given the accuracy and low processing workload, the more demanding strategy 3 is not desirable for the considered use case. Regarding the last strategy, the results are never better than for strategy 2. Therefore, in this work, I chose the second solution, using the maximum of the probability density function, as it produces the best results.

Misplaced beat m: In the case a beat is misplaced (too early or too late), it needs to be re-positioned. This scenario is close to either a skipped beat followed by an extra beat, or the reverse situation, as depicted in Figure 2.16 in the case m. Indeed, the second beat happens near the maximum likelihood of the second IG, while the previous (misplaced) one has a low probability of being a legitimate beat.

2.4.3 Workflow

Different approaches were followed to have an efficient embedded algorithm with a lowered processing load. Each modification has been tested using the same process involving a training phase on one part of the Fantasia Database [70] and evaluating its performance on the remaining set of files (`f1o01, f1o05, f1o10, f1y01, f1y03, f1y08, f2o04, f2y03 and f2y06`).

The performance is assimilated to the classification error rate on corrupted series, where each 4^{th} beat is altered. The generated sequence of beats is `NNNeNNN(s)NNNm` with `N` being a normal beat, `e` an extra beat, `s` indicating a skipped beat (therefore not in the series) and `m` a misplaced beat. This way,
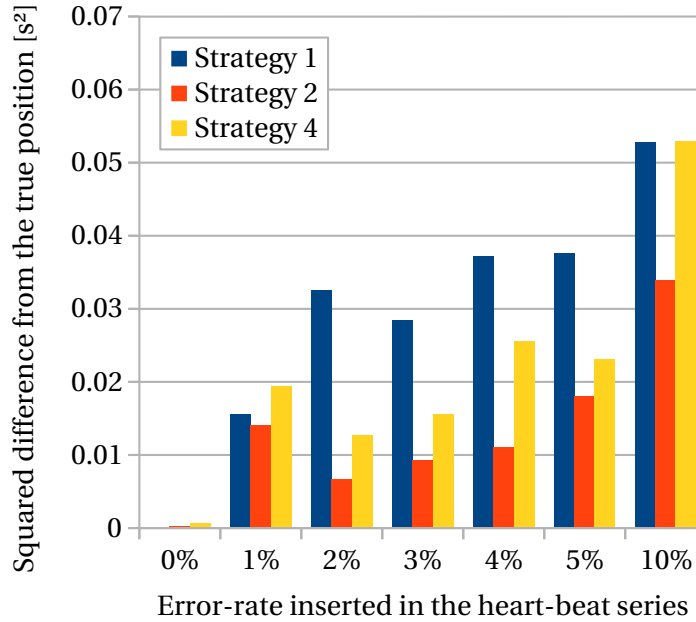
Figure 2.17 – Evaluation of the influence of the correction method on the global performance of the algorithm. Depending on the global error rate in the series, the correction to be applied is not always the same. If the series has more than around 15% of erroneous beats, then splitting the RR interval in halves is a good solution. If the series is cleaner, solution 2 is the best, that is to say to add the beat at the maximum likelihood timing. It turns out that solution 4, i.e. the average between solutions 1 and 2, never yields the best results.

the altered sequence is deterministic, with information on how each beat should be classified by the algorithm.

2.4.4 Coefficients Optimization

Before evaluating the performance of a given configuration, the relative coefficients for classification are optimized to maximize the classification performance. This is an offline training, that is only done once before evaluating the algorithm's performance. This computationally intensive training to tune the parameters can be afforded as it is not part of the embedded algorithm.

As the parameter design space is large, I develop an iterative algorithm (Algorithm 2). Starting on line 2 with a fixed set of values for the parameters η, the algorithm performs a straightforward iterative tuning (Lines 4–15): as long as the norm of vector of modifications (`deltas`) to apply is greater than a threshold (Line 4), each parameter η_N, η_e, η_s and η_m to optimize is slightly modified according to the aforementioned vector (Line 6). If the final classification rate is improved, the modification is saved and the next modification will be of a higher magnitude (Lines 8–9). In the other case (Line 11), the magnitude is lowered and turned in the opposite direction (as the sign is changed). On each iteration, the `deltas` vector is multiplied by 1.1 to increase the speed towards the optimal value of the coefficients (Line 13).

Algorithm 2 Parameter tuning algorithm for optimizing the coefficients of the heart-beat correction algorithm.

1: **function** OPTIMIZE_COEFFICIENTS
2:     **const** `coefs` $= \{0, 0, 0, 0\}$
3:     **const** `deltas` $= \{0.1, 0.1, 0.1, 0.1\}$
4:     **while** $||$`deltas`$|| > 0.05$ **do**
5:         **for** η'_i in `coefs` **do**
6:             $\eta'_i = \eta_i + $`delta`$_i$
7:             **if** ResultsImproved **then**
8:                 $\eta_i \leftarrow \eta'_i$
9:                 `delta`$_i = $`delta`$_i \times 1.5$
10:             **else**
11:                 `delta`$_i = $`delta`$_i / -2.5$
12:             **end if**
13:             `delta`$_i = $`delta`$_i \times 1.1$
14:         **end for**
15:     **end while**
16:     return `coefs`
17: **end function**

Once the parameters have been optimized, the final performance evaluation is performed on the files not used for the training, yielding the classification score for the considered variant of the classification algorithm.

2.4.5 Algorithm Tuning and Simplification

The algorithm has been iteratively modified to lower the computational load and/or improve the results. Each individual step it detailed hereafter.

2.4.5.1 Window-Size Reduction

The first step to lighten the processing is to reduce as much as possible the window size. Indeed, the longer the window, the more memory is used for storing the data and the more operations are needed to compute the $\hat{\mu}$ and $\hat{\lambda}$ estimations. As seen in Figure 2.18, progressively shrinking the window down to a few RR intervals is improving the accuracy, until a point where the error-rate rises extremely quickly. I settled on using the last six RR intervals as it has a very low resource usage as well as yielding the best results. Fewer samples are not enough to compute accurate data to model the series, while more samples tend to be too conservative when the heart rate changes in time, detecting the rhythm variation as an error.

2.4.5.2 Outlier Removal

Any error in the series used for computing $\hat{\mu}$ and $\hat{\lambda}$ is assimilated as a salt and pepper noise: the majority of RR intervals length are close to the normal length. However, when there are some extra or skipped beats in the samples used to model the IG, the corresponding RR intervals are respectively much smaller or bigger. A median filtering removes the extrema, but it favors over-conformant distributions, with μ and λ sharpening the curve around the median RR interval length. Overall, this additional filtering is prejudicial, leading to more classification error.

2.4.5.3 Weighting

An exponential decay is used to have more weight on the newest RR intervals than on the older ones. With a small window length, the weighting is not as relevant as with longer window lengths. Using a constant weighting instead of an exponential decay function gives a classification error-rate of 0.34% instead of 0.36%, which means the results are a insignificantly improved and uses no inverse exponential computation which is not adequate for a low-power MCU.
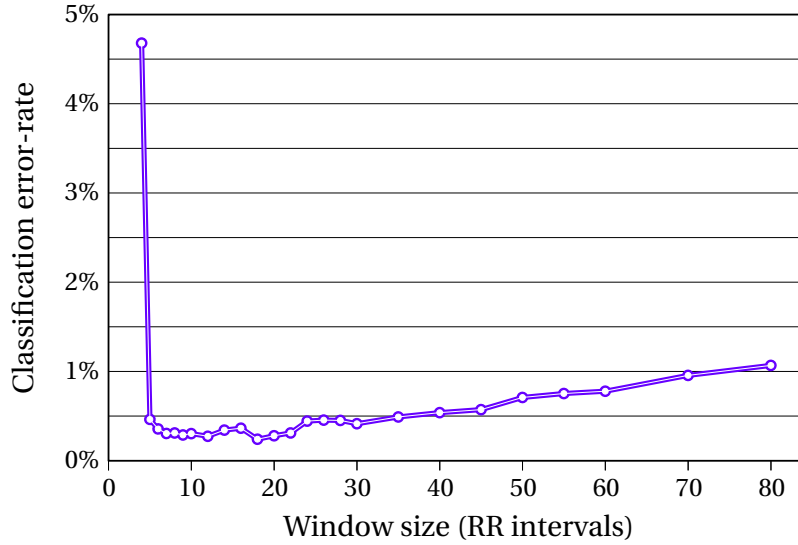
Figure 2.18 – Classification error-rate depending on the windows' length used by the heart-beat correction algorithm. With the window length shrinking, the results are slightly improving, until a lower limit is reached, making the classification totally unreliable, as there is not enough data anymore to compute meaningful $\hat{\mu}$ and $\hat{\lambda}$ for the IG model.

2.4.5.4 Different Estimator for λ

The standard deviation of the IG distribution is $\sigma = \sqrt{\mu^3 / \lambda}$. From this formula, it is possible to compute an estimate of λ. From the implementation, this $\hat{\lambda}$ estimator yields a classification error-rate of 0.30% instead of 0.34%. This is a small benefit balanced by a higher processing load noticed with a 69% time penalty.

2.4.5.5 Final Configuration

Targeting WSNs, the final algorithm runs the last six RR-intervals, without weighting and with the original estimators for μ and λ. The alternative configurations are rejected because detrimental to either the system's classification performance or battery-life.

The confusion matrix for the classification is reported in Table 2.3. The two antagonist cases of an extra-beat and a skipped beat are never confused one

Table 2.3 – Confusion matrix for the heart-beat classification in the final system configuration.

Prediction:	N	e	s	m
N	50642	1	83	30
e	54	7189	0	7
s	1	0	7242	3
m	61	0	34	7157

Table 2.4 – Classification error rate for the different files used for the performance analysis of the heart-beat correction algorithm.

File	$beats_{tot}$	Err_{tot}	Err_N	Err_e	Err_s	Err_m
f1o01	7951	0	0	0	0	0
f1o05	6353	2	0	0	1	1
f1o10	9144	2	1	0	1	0
f1y01	9748	**216**	111	3	47	55
f1y03	8479	8	3	0	1	4
f1y08	8071	**45**	18	2	5	20
f2o04	7654	1	1	0	0	0
f2y03	7549	1	1	0	0	0
f2y06	7603	6	3	0	2	1

another. Most errors are due to normal beats being classified as skipped beats. This happens when the subject is relaxing and the heart-rate is slowing down.

2.4.6 Database Analysis

Taking a closer look at the individual results exposed in Table 2.4, two of them stand out with a much higher classification error-rate: the files f1y01 and f1y08 have a 2.22% and 0.56% error rate whereas the other are between 0% and 0.09%.

Further investigation reveals that those two files have the highest HRV computed both using the Standard Deviation of Normal-to-Normal heart beat intervals (SDNN) or Root Mean Square of the Successive Differences (RMSSD). This is because a high HRV is linked to more irregularities in the heart-rate, with the body adapting quickly to the environment. A healthy individual

presents a high HRV, while a reduced HRV can often be linked with a degraded health state. When sorting all the 40 files of the database by increasing RMSSD (Figure 2.19), seven of the nine recordings used for assessing the algorithm performance are in the 10 files having the lowest HRV. This is not representative of the full database, skewing the model by training on extremely regular heart-rates.
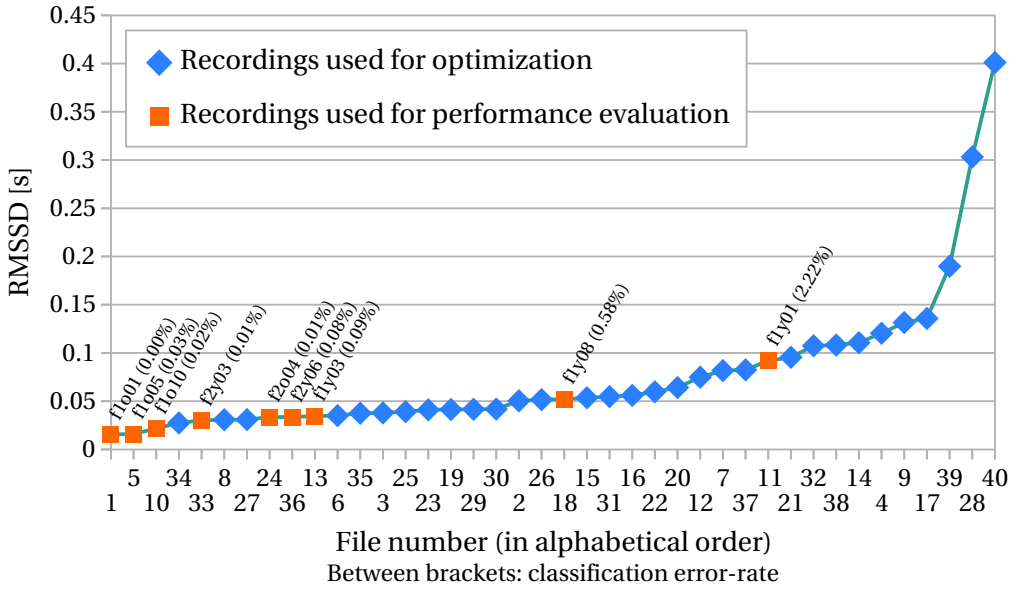
Figure 2.19 – The HRV of each recording of the database has been computed using the RMSSD. The diamonds are the files used for the automatic tuning of the algorithm whereas the highlighted squares are the files used for verifying the performance of each optimization. As seen in Table 2.4, the higher the RMSSD, the higher the classification error-rate, reaching more than 2% for the worst case.

Given the variability of the results depending on the individual files, I ran several additional analyses using both manual selection of files depending on the HRV or random shuffling of the files used in the database.

2.4.6.1 Manual File Selection

Four tests are performed to see the influence of the HRV on the classification. The size of the training database is maintained to 31 files, and the perfor-

mance evaluation on the remaining nine files. For each test, the files used for the performance evaluation are the files ranked 5 to 13 when sorting the database by RMSSD or SDNN, ordering from low to high or the opposite.

From Figure 2.20, we have the confirmation that when the files are chosen according to the HRV, the results are changing significantly. Training using the files with a high RMSSD gives the best classification, with a really low error-rate. On the other hand, as the performance evaluation is done on the remaining set of files (having a lower RMSSD), the classification is easier for the algorithm.
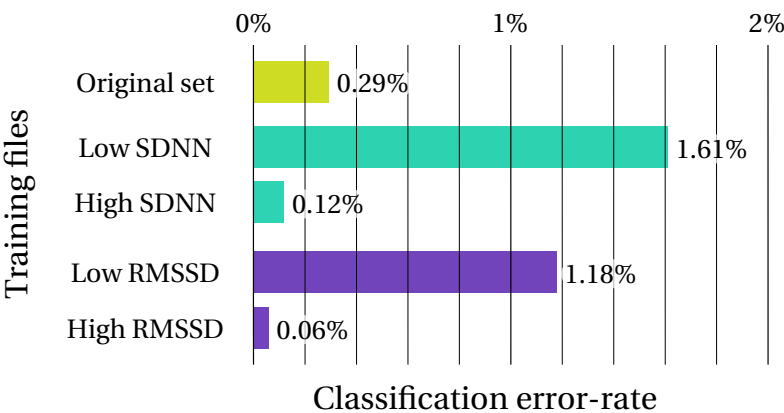
Figure 2.20 – Influence of the database used for training on the final classification. The recordings used are sorted by their HRV, either considering SDNN or RMSSD. When the training has been done using most of the high-HRV recordings, the classification is better.

2.4.6.2 Statistical Analysis

To further check the influence of the recordings used for training, as well as the size of the database, an automatic shuffling of the recordings used for training and performance evaluation has been performed. The number of training files range from 10% to 90% of the full database, evaluation being always performed on the remaining recordings.

The results, visible in Figure 2.21, indicate that the training dataset can be reduced to a few recordings, as long as they are representative of the total
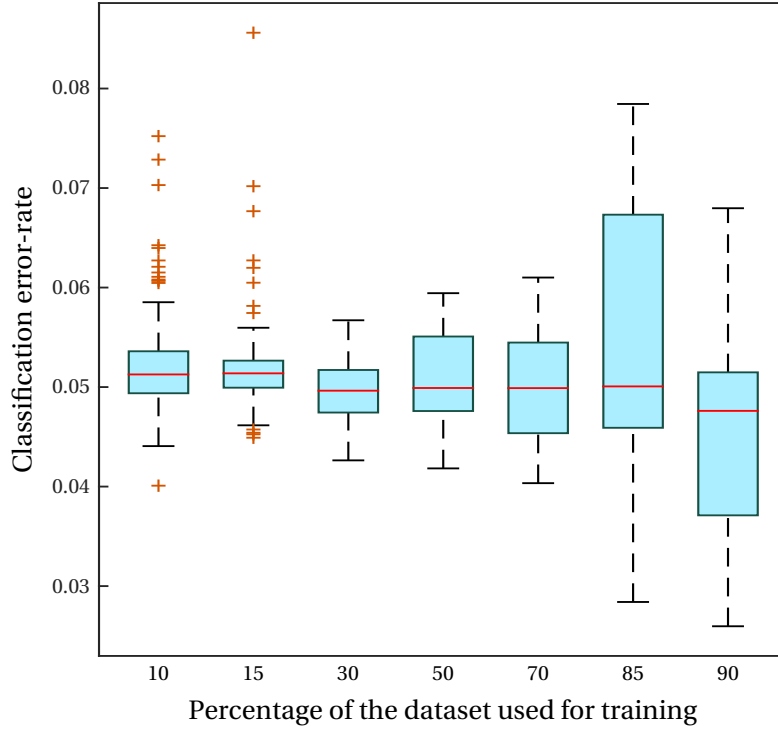
Figure 2.21 – Heart-beat classification error-rate when shuffling the recordings of the database, while varying the number of files used for training and performance evaluation. The number of trials in the groups from 10% to 90% are respectively 117, 100, 70, 45, 15, 15. Outliers are plotted individually beyond the whiskers, further than $\pm 2.7\sigma$.

dataset. From this analysis, only four representative recordings are enough to get a high classification performance. In the case we have no information about the individual recordings, using ten out of the 40 files is the best way to lower the variability of the result, and avoiding data dredging.

2.5 Summary and Concluding Remarks

In this chapter, using an event-triggered sampling strategy, I demonstrated that the sampling rate required to perform a widespread task in current wearable medical systems (heart-rate monitoring) can be reduced between one and two orders of magnitude while retaining the same performance as
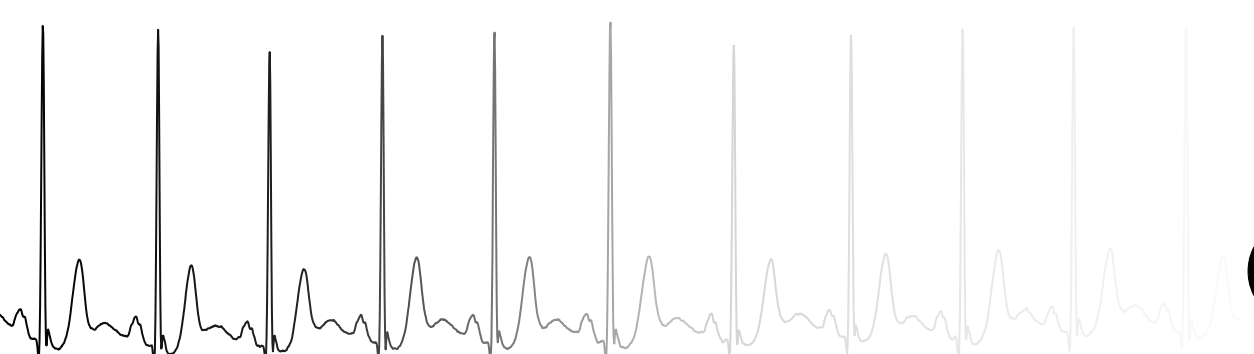
a compressed-sensing approach with 50% compression ratio. Hence, the observed magnitude of the difference allows for optimism, with high energy-savings for a hardware implementation, along with a lowered processing and storage for micro-controllers thanks to a lower amount of data to process. When considering full systems, such as wearable cardiac monitoring devices, further developing and the use of T-ADCs opens the way to a promising future of long-lasting solutions for a non-intrusive and personalized healthcare.

With the recorded ECG, for wearable WSNs systems requiring fast feedback to the user, I presented a low-power algorithm correcting possible mistakes in a series of heart-beats, applying a correction immediately when a new beat is detected. Starting form a reference algorithm, I explored different approaches to make it lightweight while keeping a high performance. It involved mainly reducing the number of loops when computing intermediary results and reducing the amount of data needed to get a high classification accuracy. Other modifications could improve slightly the classification results but with an extra processing cost. For the application considered, this slightly improved performance cannot justify the penalty on the battery life. Overall, the memory load has been divided by a factor 10 from the reference implementation. However, when targeting medical applications for online diagnostics, better algorithms exist. Indeed, without the need to immediately classify every new beat as soon as it is detected, there are fewer constraints on the design space. It nevertheless prevents the use of applications where immediate feedback is needed, such as training or gaming.

Non-Intrusive Screening of Obstructive Sleep Apnea

IN the previous chapter, I addressed the optimization of Wireless Sensor Nodes (WSNs) by targeting the fundamental building block of all systems: signal sensing. However, by completely turning around the perspective, and focusing on the specifics of a target application, we can also develop a set of tailored strategies to achieve further optimization. In this chapter, I describe a novel algorithm for WSNs targeting the monitoring of Obstructive Sleep Apnea (OSA).

3.1   Introduction and State of the Art

Nowadays, OSA is a common sleep disorder that involves partial or complete obstruction of the upper airway. In the U.S.A. alone, an estimated number of 3.8 million people between 30 and 60 years old are affected by this condition [71, 72]. Depending on the population lifestyle, the prevalence of OSA ranges from 3% to 24% [25], with an estimated 5% worldwide [25–27]. This disorder is an aggravating factor for multiple health diseases, including cardiovascular [73] (high blood pressure [19], stroke [20]) and clinical depression [21] as well as decreased memory and cognitive skills [22]. People with OSA present a higher rate of sudden deaths [23]. While this condition is treatable, 90% of the subjects go undiagnosed [28]. As shown in Figure 3.1, OSA affects 25% of the adult male population in the U.S.A., with a small minority of subjects being diagnosed. Hence, there is a need for accessible obstructive sleep apnea screening.

OSA   Unaffected population

| 24% | 76% |

| 90% | 10% |

Undiagnosed   Diagnosed

Figure 3.1 – Proportion of male adults in the U.S.A. being affected by OSA. In the upper bar, the green part on the right corresponds to people without OSA. The yellow part on the left is divided in two, with a zoom on the lower bar. Among this population of affected people, most of them (the diagonal lines) are undiagnosed. Over the total population considered, only 2.5% are diagnosed with OSA.

Despite major progress, there is still a need to develop a non-intrusive solution for home OSA monitoring for two reasons. First, there is low incentive for patients with low to moderate OSA to use external breathing equipment, such as Continuous Positive Airway Pressure (CPAP) [74]. Second, existing solutions are bulky, time-consuming, expensive and intrusive [75–77]. Because the population-wide capacity of performing full OSA diagnosis does not match the recommended capacity [78], OSA testing or screening are only available in dedicated facilities for the most severe cases: the patient is required to go to a sleep center or hospital where his or her sleep will be monitored extensively for two non-consecutive nights. A full polysomnography (PSG) will be done to record electrocardiogram (ECG), electromyogram (EMG), electroencephalogram (EEG), electrooculogram (EOG), and nasal airflow [79]. Altogether, it requires 22 electrodes plus a respiration mask. This is an intrusive setup that disturbs the patient's sleep quality. Moreover, the acquired data needs to be analyzed by a specialist. Given these constraints, more than 80% of patients are reluctant to undergo a PSG [28]. Additionally, both the risks associated with an external respiratory help and the impracticality of OSA screening leaves most of the affected population without any kind of monitoring [28], hence the opportunity for a simple yet efficient Internet of Things (IoT) solution usable for home screening.

In the context of non-intrusive OSA monitoring, previous studies have shown that it is possible to detect OSA automatically based on single-lead ECG recordings [30, 80]. The existing wearable devices for OSA detection, either commercially available or at the research state [81–98], focus only on signal acquisition. The processing and identification of OSA events comes afterwards, as an additional offline phase, after downloading the data to a more powerful platform. Nonetheless, due to the high rate of sudden death of people with OSA because of additional cardiac issues [23], there is a clear need for personal real-time systems. Devices which integrate in the bed cannot currently assess the problem of cardiac monitoring, and do not work well when two persons are in the same bed. Therefore, non-intrusive personal wearable systems need to be investigated [99, 100].

In terms of OSA detection techniques and classification accuracy, the best one from a single-lead ECG recording reported in the literature is 92.5% classification accuracy [89]. It is obtained by considering multiple features extracted from the frequency-domain as well as the ECG morphology. However, the main drawback of this solution is that the classification is manual and time-consuming. In fact, as a comparative figure, the accuracy of a manual classification from an expert using the full polysomnogram signal is 93% [101]. Among the fully automatic classification techniques, the approach proposed by [97] achieves the best results with an accuracy of 90.6%. When using exclusively time-domain features, non-linear statistics reach an accuracy of 85.6% along with 72.1% sensitivity and 91.2% specificity [94]. An "if-then" decision tree is considered by [99] with a reported accuracy of 93.2%, relying on features derived from heart-beats. However, it has the drawback of using an undefined subset of the database, so the results cannot be compared as-is to the previous numbers.

Finally, neural networks have also been considered for OSA classification and show a performance of 82,1%, a sensitivity of 88.4% and a specificity of 72,3% [98], which is on par with the algorithms mentioned previously. Two groups have considered a similar approach [98, 102] using a Long Short-Term Memory Recurrent Neural Network (LSTM-RNN) which report 82.1% up to 100% but they consider only a limited set of recordings which prevents replicating the results using the same performance assessment.

However, among all the devices and techniques available, none is a high-accuracy wearable system with real-time and long-term screening while satisfying the need of cardiac monitoring. In this chapter, I present a solution based on a WSN for automatic, energy-efficient, non-intrusive, yet accurate OSA detection and screening using a single-channel ECG recording. The main contributions are:

- an online OSA detection technique compatible with cardiac-monitoring [33–37], with a time-complexity of $O(n)$, which is the theoretical lower bound. This is achieved by developing an efficient outlier removal and through performing OSA assessment in time-domain, which removes the need for computationally expensive frequency-domain analyses (see Section 3.4).

- an autonomous and energy efficient OSA screening system with a battery lifetime for continuous monitoring during 46 days, measured experimentally on an actual hardware platform (see Section 3.6.5).

- a system, while energy efficient, which is also comparable with the state-of-the-art in terms of performance, reaching 88.2% accuracy, 80.0% sensitivity and 93.9% specificity ($F1 = 84.7\%$) thanks to adopting a patient-specific perspective, which takes into account the distinct profile of each patient (see Section 3.5.2).

The rest of this chapter is organized as follows. First, the targeted hardware and software platform is presented in Section 3.2. Next, Section 3.3 explains how features are generated and evaluated offline for optimizing the results in the final online system. Then, Section 3.4 describes the implementation and optimization of an energy-efficient OSA detection technique. The setup used for testing the system designed is defined in Section 3.5, while in Section 3.6, the proposed system is evaluated experimentally with respect to energy efficiency and classification accuracy, along with patient-specific configuration. Finally, Section 3.7 closes the current Chapter, stating that using the proposed patient-specific technique, it is possible to achieve high classification accuracy for OSA detection, while having a longer battery lifetime than the state-of-the-art systems.

3.2 Sleep-Apnea Monitoring System

In these sections, I first describe the wearable WSN, and then I present the software architecture.

3.2.1 Target Wearable Platform

For this work, I consider the SmartCardia INYU wearable sensor (Figure 3.2) as my target device for the final implementation. INYU is an energy-efficient wearable device providing a single-lead ECG recording with a 24-bit Analog-to-Digital Converter (ADC) operating at a frequency from 250 Hz to 16 kHz. The ECG is measured using silver-chloride electrodes by impedance pneumography [103]. The device is equipped with the STM32L151RDT6 [104], an ultra-low power 32-bit micro-controller, which can operate at a maximum frequency of 32 MHz. It features 48 kB of Random Access Memory (RAM) and 384 kB of flash storage, and it is powered using 710 mAh battery.
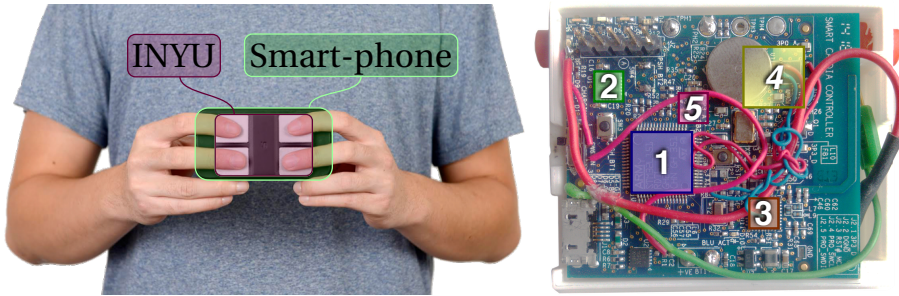
Figure 3.2 – INYU sensor and prototype.
Front: 1: STM32L151RDT6 (ARM Cortex-M3 Micro-Controller Unit (MCU), 384 kB Flash, 48 kB RAM), 2: MPU-6000 (6-axis I^{2}C motion sensor), 3: nRF8001 (Bluetooth Low-Energy (BLE) radio).
Back: 4: AFE4300 (Analog front-end for body impedance measurement), 5: ADS1191 (Analog front-end for ECG applications).

In the considered case of continuous monitoring, the device is held in place with a dedicated chest-strap shown in Figure 3.3.
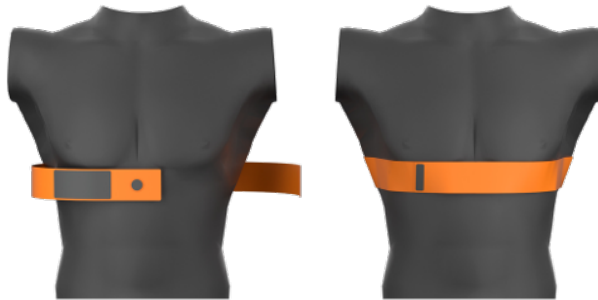
Figure 3.3 – Chest-strap for folding the INYU wearable WSN in place on the torso. This is applicable for a hand-free use, such as practicing for sports, or overnight screening.

Given the internal capabilities and connectivity possibilities, this device can work as a fully autonomous device for several days of continuous recording, uploading the recorded and processed data to a base station when one becomes available. For example, a BLE compliant smartphone can be used for this purpose, which can afterwards display the data on-screen or upload it online to a remote medical service to be manually checked by a physician.

Two questions arise from this target platform. The first one is about the signal sensing. In the previous chapter (2) of this thesis, I worked on improving the signal acquisition aspect, relying on event-triggered sensing applied to the specific case of ECG. The current chapter uses a hardware platform leveraging a commercially available Voltage Analog-to-Digital Converter (V-ADC) chip. This is a conscious choice, because developing and relying on custom chips greatly limits the applicability of a novel signal analysis. By using off-the-shelf components, it becomes possible to easily provide a hardware and software solution to the lack of OSA diagnostics and screening, even if there is room for further improvements using custom silicon. Considering how widespread and under-diagnosed OSA is, an affordable solution for diagnostics and screening OSA on a large scale is meaningful. The second question comes with the body-impedance measurement capability. This feature can be used to measure the impedance between the two arms. When breathing, the chest moves and changes the impedance accordingly. Therefore, it has the capability to measure breathing *via* the impedance change of the chest. While the breathing rate has a direct relation with OSA events, there

are however two major drawbacks. First, ambulant breathing monitoring has not been proven as a reliable solution for OSA detection. Second, while the considered device can measure the ECG by being in contact with the chest with an elastic chest-strap, the breathing measurement requires electrodes on the sides of the chest. Targeting an established and non-cumbersome setup, I decided to rely exclusively on the ECG.

3.2.2   Software Structure

The overall software structure and flow of the system I propose is shown in Figure 3.4. The ECG is first acquired using medical electrodes connected to the chest of the patient. Then, an initial noise filtering is performed to remove artifacts caused, for example, by power lines, electrode parasitic motion, or baseline drift [105, 106]. Towards this, *erosion* and *dilation* morphological filters are used [107, 108], which can be implemented efficiently in energy-constrained wearable systems. After filtering the signal, the ECG delineation (extraction of fiducial points in the signal related to the physiological behavior) is done based on the wavelet transform [109], relying on the fact that the different wave parts are made of different frequency components. The output of the delineation is then used to run two different automatic diagnostics. First, the OSA detection, which is the focus of this chapter and is discussed in Sections 3.3 and 3.4, and a cardiac monitoring for additional evaluation of the patient. In parallel, the raw data is compressed [48] and stored for further offline analysis.

3.3   Offline Feature Extraction and OSA Learning Phase

In this section, while keeping in mind the stringent energy-constraints of a wearable WSN, I first identify the relevant features in Subsection 3.3.1 and then, I train a classifier for OSA detection using the selected features in Subsection 3.3.2.
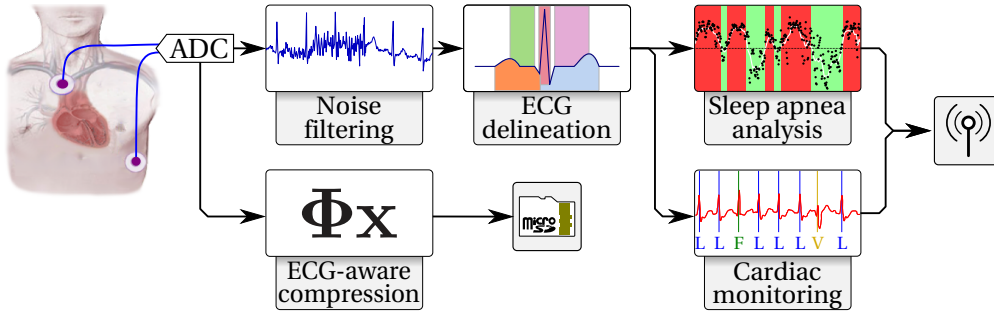
Figure 3.4 – Overview of the processing blocks integrated in the device used for the online OSA analysis in the proposed system [110]. The ECG is first filtered to remove the noise, then the fiducial points are extracted. Finally, the signal is analyzed to detect OSA and cardiac pathologies. Raw ECG data is stored compressed on a memory for further offline analysis by an expert.

3.3.1 Features Extraction

The first step towards designing a classifier is to extract relevant information from the source signal. This is known as feature extraction and is often done by detecting patterns or computing statistical metrics. Because of the stringent energy-constraints of wearable platforms and WSNs in particular, I carefully select and optimize the few most significant features to detect OSA from the ECG. This contributes to diminishing the Central Processing Unit (CPU) footprint down to the point where it is small enough for an online embedded OSA classification in a WSN. This is motivated by the fact that an autonomous device with online analysis has a longer battery life than streaming the uniformly sampled data to a remote device such as a smartphone, provided that the embedded computations required for signal processing and classification are sufficiently lightweight [38, 111–113].

The most informative features for detecting OSA are selected for a set of features derived from the data in multiple ways. First, I independently evaluate the ones considered in the literature [97]. In [97], 52 features are derived from the ECG morphology and an additional 36 features are derived from an ECG Derived Respiration (EDR) signal, i.e., a total of 88 features. Second, I generate

other features directly from the ECG, namely the EDR, the RR-interval series and the RS-amplitude series along with their respective spectra. Third, besides statistical features (*min, max, mean, std, rmssd, sdnn*), for the ECG, EDR, RR-intervals and RS-amplitudes, I also consider auto-regressive processes, i.e., $x_i = y_i - \frac{1}{4} \sum_{j=i-4}^{i-1} y_j$, as prior study [88] reports an improved performance. Finally, I generate several features by integrating or differentiating parts of the ECG signal.

Some features, while potentially very informative, can be energy-hungry to generate. This is the case for frequency-domain transforms, involving a lot of floating-point operations and memory allocations. Indeed, floating-point arithmetic is not available on the platform considered and therefore available only through software emulation, and the total amount of memory available is very limited. Therefore, a trade-off exists between the energy consumed for extracting the features and the classification performance. For designing an energy-efficient OSA detection algorithm for wearable WSNs, it is essential to keep the number of features extracted low but meaningful.

To identify which features are the most relevant for OSA classification, all the features were first generated for both the training and testing set. To select the features for the final system, a minute-by-minute classification was performed, where the classification accuracy identifies how relevant each feature is for OSA detection. By applying forward feature selection [114], the single feature providing the best accuracy was selected. Iterating, the next feature that, combined with the previously selected ones, provides the best classification accuracy was added.

In a comparison of different algorithms for apnea detection from ECG recordings, it appears that the most common features are generated from the time series of heart beats, the ECG morphology, and from the EDR signal [115]. In particular, during an OSA event, there is a shift of the signal's energy towards low frequencies for two distinct time-series: the series of time intervals between two heart beats (RR-intervals in Figure 3.5) and the series of R-peak amplitudes with respect to S-amplitudes (RS-amplitude in Figure 3.5). This shift of the signal's energy is illustrated in Figure 3.6 with the spectrogram of the RR-intervals series. It shows the minute-by-minute frequency-spectrum
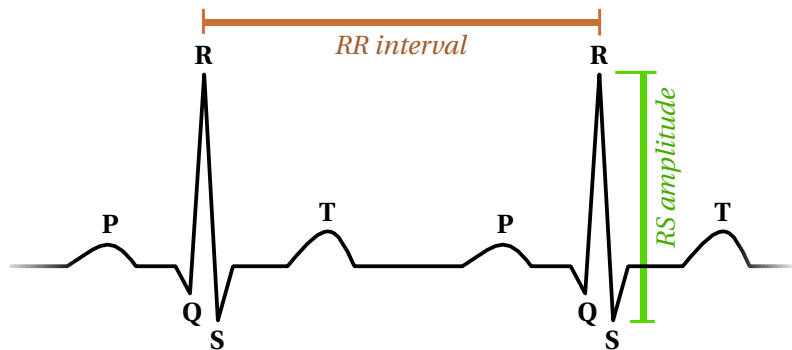
Figure 3.5 – ECG morphology with RR-interval and RS-amplitude labeled [116]. [Image modified from the work of Agateller (Anthony Atkielski), Public Domain.]

of the RR-intervals for a complete overnight recording. Frequencies of the signal are displayed on the *y*-axis versus the time on the *x*-axis. The dark red color indicates high energy for that frequency and minute, whereas the blue is linked to low energy. The manually-labeled OSA ground truth by the medical expert is shown below. In these annotations, the *HIGH* value is linked to OSA events whereas normal breathing is captured by the *LOW* value. As this figure shows, there is a clear correlation between the OSA events and the signal's energy in low frequencies. A very similar trend is observed for the RS-amplitude series. As a result, both RR-intervals spectrum and the RS-amplitude spectrum were selected as both are correlated with OSA.

Computing signal's spectrum is, however, energy-hungry. I, therefore, need to optimize this process. First, I isolate the most relevant frequency-band from each spectrum (RR-interval and RS-amplitude). Thus, a parameter sweep is used to select the optimal frequency-band bounds of the signal correlated with OSA events, with respect to classification accuracy. A 2D-map of accuracy (see Figure 3.7) is created from an exhaustive exploration of all possible frequency band bounds. It represents the OSA classification accuracy obtained for each pair of low and high bounds. The apnea frequency-band bounds are on the axes and the classification accuracy is linked to the graph color: the brighter the color, the higher the accuracy. This figure, which is computed for RR-intervals, shows a clear frequency interval regarding the lower bound of the band, ranging from 0.010 to 0.020, while the upper bound
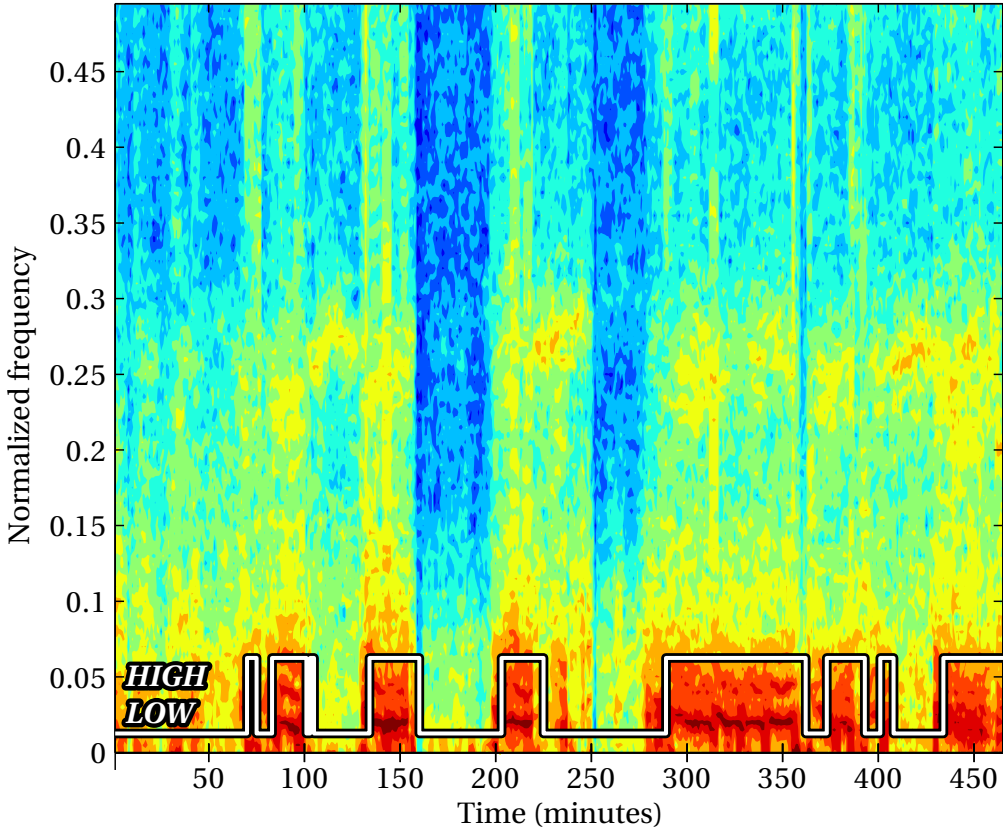
Figure 3.6 – Spectrogram of the log-power of the RR-intervals series from the recording x32 (see Section 3.5.1) with the labeled OSA logic signal on a same time axis.

has more tolerance, from 0.045 to 0.075, in terms of normalized frequency. This whole area provides the highest classification results, reaching 76% when using the raw RR-intervals without filtering. Therefore, only a single feature is needed from the spectrum of RR-intervals for classification, which can be computed in an energy-efficient way (cf. Section 3.4.2). Similarly, the best frequency-band bounds are found for the RS-amplitude time-series.

3.3.2 Features Combination and Learning Phase

Based on the discussion in the previous section, I only use the relative energy in a specific frequency band for both RR-intervals and RS-amplitudes time-
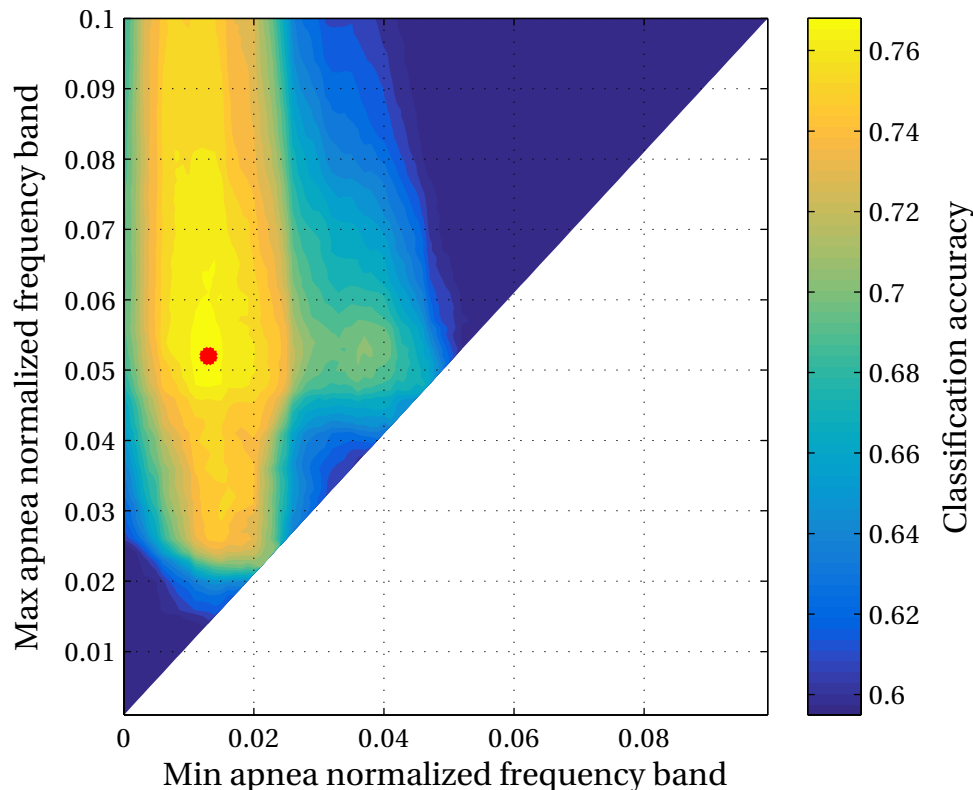
Figure 3.7 – Classification accuracy (normalized between 0 and 1) for OSA when varying frequency-band bounds considering the RR-intervals time series. The circular dot is placed at the position of the best normalized frequency band.

series. As they are two features not correlated to each other, it means they provide complimentary information for the classification. This is illustrated in Figure 3.8 where each sample data is positioned in the RR-RS relative energy plot. The sample data for OSA events are in yellow, and normal sleep points are in blue. The red line (linear combination of RR and RS apnea scores) separates the two classes by maximizing the margin between the two classes.

In the literature, several classifiers were considered and the classification accuracy obtained is in the range between 80% and 90% [117]. Additionally, Random Forest [118] has the highest performance in terms of classification
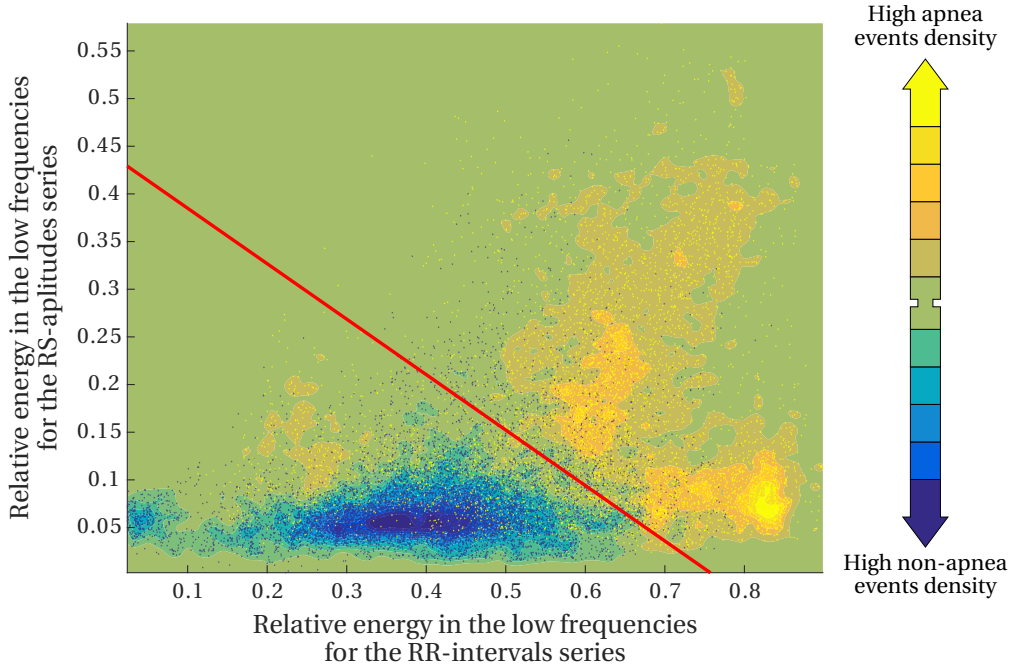
Figure 3.8 – Distribution of the *apnea* (yellow) and *non-apnea* (blue) events from the training set of recordings for the two features generated from the series of RR-intervals and RS-amplitudes. The two classes are separated based on a linear Support Vector Machine (SVM) classifier.

accuracy. In my setup to reproduce the reported results, the Random Forest classifier reaches a classification accuracy of 88.1% at most. On the other hand, when using a SVM classifier [119], I reach 90.1% classification accuracy, whether I use a linear, Gaussian or polynomial kernel. Therefore, in terms of classification accuracy, the results obtained using the SVM classifier are up to 2% better than Random Forest's results.

In terms of computational complexity, linear SVM can be implemented even with limited computational resources. As I design an energy-constrained system, I use linear SVM in my OSA detection technique, because of its high classification performance and high computational efficiency at run-time (see Section 3.4.4).

3.4   Online Sleep-Apnea Detection Technique

In this section, I propose an energy efficient OSA detection technique that can directly run on an energy-constrained wearable WSN. Previous work and devices focused on offline analysis on a computer or cloud system [84–98]. Conversely, I aim to provide an online ECG analysis running on a wearable system. Therefore, my main goal here is to lower the energy consumption of the OSA detection, while maintaining high classification accuracy on the wearable device.

The overall flow of the online OSA analysis, after the ECG noise filtering and ECG delineation, is shown in Figure 3.9. My OSA screening method allows different interval lengths but in particular in this thesis, I use a 60 seconds analysis as the database used provides only a minute-by-minute labeling. From the ECG, I generate RR-intervals and RS-amplitudes time-series, as they are the most relevant features (see Section 3.3). First, I use my own enhanced Thompson-Tau filter (Subsection 3.4.1) to remove the outliers from erroneous beats caused mostly by motion artifacts and muscle noise, inherent to ambulant systems. Then, I compute the power in the two apnea frequency bands of the spectrums in order to obtain the RR apnea score as well as the RS apnea score (Subsection 3.4.2). Afterwards, I apply a moving average filter to smooth the variability of both apnea scores (Subsection 3.4.3). Finally, according to the SVM classifier trained in Subsection 3.3.2, if the linear combination of the smoothed apnea scores is greater than a threshold, I label the corresponding minute as apnea.

3.4.1   Low-Complexity Outlier Removal

As heart-beat outliers have strong negative effects on frequency analysis [67], detecting and removing these outliers is critical. Therefore, I apply my own low-complexity version of the Thompson-Tau filter on the series of RR-intervals before apnea scoring. The original Thompson-Tau outlier removal algorithm is provided [120] and its average time complexity is $O(n\log(n))$, because of sorting the entire input array. However, I propose faster outlier removal by reducing the average time complexity down to $O(n\log(k))$ (see Algorithm 3). My algorithm first starts by computing the initial mean and
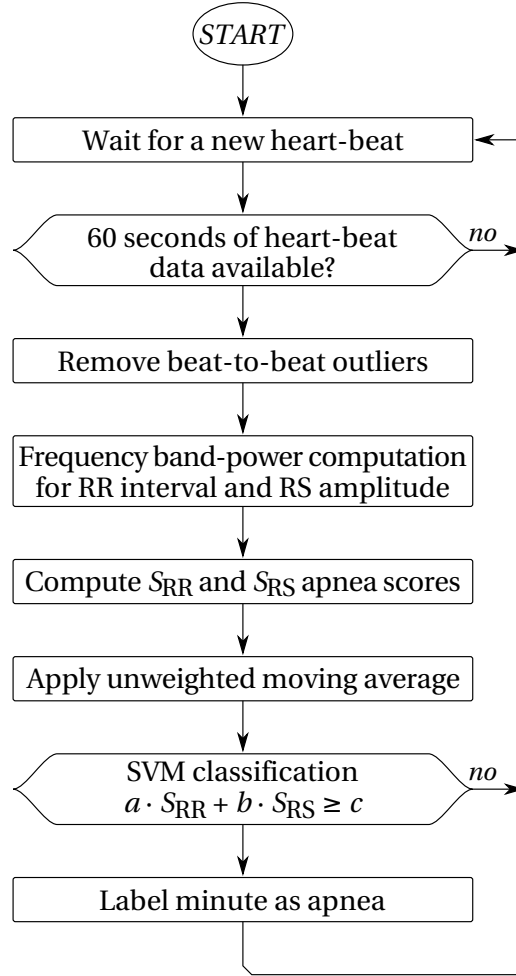
Figure 3.9 – Flowchart of the proposed online OSA screening technique. It is applied when considering both the RR-intervals and the RS-amplitudes of the beats.

standard-deviation of the series (Lines 3–4). In Lines 5–6, it finds the k-th largest and the k-th smallest values using the QuickSelect algorithm [121], which only sorts the beginning and ending of the input array. In the subsequent Lines 8 and 13, I test if the smallest or largest value is conforming to the Student's-t distribution [122] of the input. If not (Lines 9–12 and 14–17),

then I move the *start* or *end* indices of the series to exclude the new outlier and update both the mean and standard-deviation according to [123]:

$$
\begin{aligned}
\bar{x}_{n-1} &= \frac{n\bar{x}_n - x_n}{n-1}, \\
s_{n-1} &= \sqrt{\frac{(n-1)s_n^2 - (x_n - \bar{x}_n)(x_n - \bar{x}_{n-1})}{n-2}},
\end{aligned}
\tag{3.1}
$$

where $\bar{x}_i$ and s_i are the mean and standard deviation of the series up to the i-th value x_i, respectively. I repeat this process until both the smallest and largest values are conforming to the distribution (Lines 7–19).

The Student's-t distribution however has a τ parameter which depends on the filter strength, which is a fixed design parameter, and the length of the series. By removing outliers from the series, its value should be updated. However, recomputing it is expensive, especially for a CPU without support for floating-point operations. Indeed, in this case, a software emulation of floating-point needs to be included, providing a low-performance compatibility solution. Indeed, as it is derived from the inverse cumulative distribution function of the Student's distribution (Equation 3.2), no straightforward integer-based computation is reachable.

$$
\begin{cases}
-\sqrt{v}\cdot\sqrt{\dfrac{1}{I_{2x}^{-1}(0.5v,0.5)} - 1} & 0 < x < 1/2 \\[2mm]
0 & x = 1/2 \\[2mm]
\sqrt{v}\cdot\sqrt{\dfrac{1}{I_{2(1-x)}^{-1}(0.5v,0.5)} - 1} & 1/2 < x < 1 \\[2mm]
-\infty & x \le 0 \\[1mm]
\infty & x \ge 1
\end{cases}
\tag{3.2}
$$

With $I_x^{-1}(a,b)$ the inverse regularized beta function.

There is nonetheless a workaround for this complexity. Indeed, the Figure 3.10 shows the evolution of the τ factor with a weak filter of 1% strength, depending on the input vector length. We notice that past a vector length of 200, τ has a constant value of 2.54 ± 0.01. Because the algorithm works on five minutes of data at once, even with a very low heart-rate of 40 beats-per-minute, the
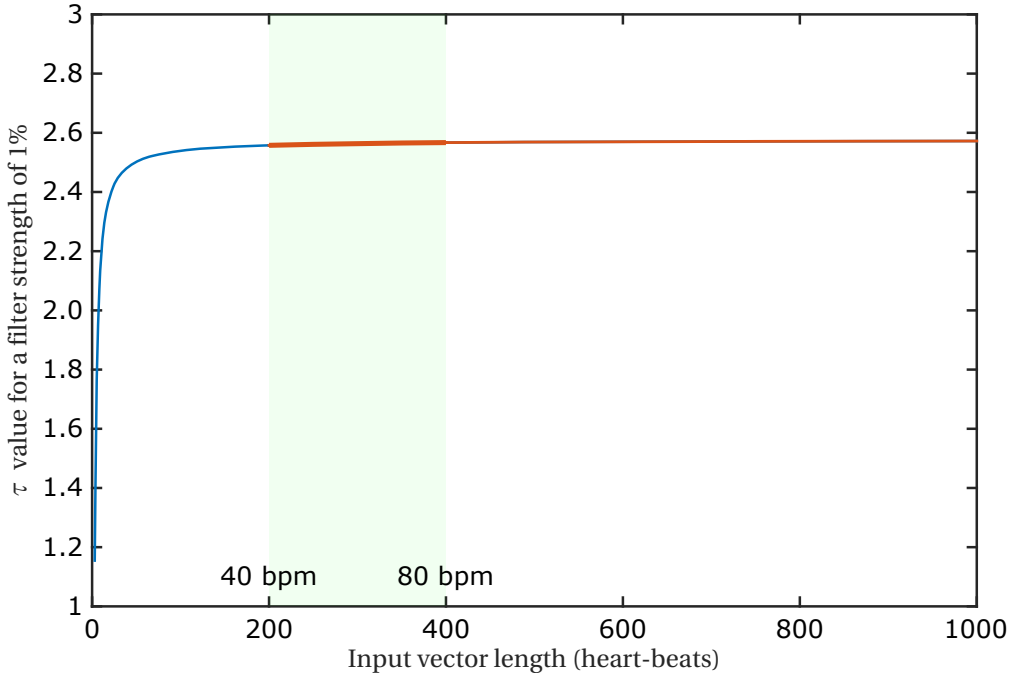
Figure 3.10 – Evolution of the τ factor modeling the Student's-t distribution with a filter strength of 1%. In orange, for a vector length greater than 200, the value of τ changes less than 1%. With the algorithm working with five-minutes segments, 200 values are for a heart-rate of 40 beats-per-minute, while 80 bpm creates 400 values, which is the range for someone asleep. The filter strength only affects the vertical scaling, for the vector length we are considering.

minimal vector length to consider is 200. Therefore, it is sensible to consider the τ value to be constant while the filtering is in progress.

The average complexity of Algorithm 3 is $O(n)$. In the impossible case where all values are outliers, the complexity is the same as in the original implementation. However, only a few values are outliers. Indeed, Figure 3.11 shows the classification accuracy while changing the filter's tolerance for outliers detection. As this figure shows, there is a 5% increase of the classification accuracy if a very tolerant filter is used, thus discarding the few biggest outliers. It means that removing a restricted number of the most non-conforming

---

Algorithm 3 Low-complexity outlier removal.

1: **function** FASTOUTLIERREMOVAL($\boldsymbol{x}, \tau, k$)
2:     $N \leftarrow$ LENGTH($\boldsymbol{x}$); $start \leftarrow 1$; $end \leftarrow N$;
3:     $\mu \leftarrow$ COMPUTEMEAN($\boldsymbol{x}$);
4:     $\sigma \leftarrow$ COMPUTESTD($\boldsymbol{x}$);
5:     $low \leftarrow$ QUICKSELECT($\boldsymbol{x}, k$); $\triangleright$ Sorts the beginning of the array
6:     $high \leftarrow$ QUICKSELECT($\boldsymbol{x}, N - k$); $\triangleright$ Sorts the ending of the array
7:     **while** $\boldsymbol{x}[start] < low$ or $\boldsymbol{x}[end] > high$ **do**
8:         **if** $low > \tau \cdot \sigma$ **then**
9:             $\mu \leftarrow$ UPDATEMEAN($\boldsymbol{x}[start], N, \mu$);
10:             $\sigma \leftarrow$ UPDATESTD($\boldsymbol{x}[start], N, \sigma, \mu$);
11:             $start \leftarrow start + 1$;
12:             $N \leftarrow N - 1$;
13:         **else if** $high < \tau \cdot \sigma$ **then**
14:             $\mu \leftarrow$ UPDATEMEAN($\boldsymbol{x}[end], N, \mu$);
15:             $\sigma \leftarrow$ UPDATESTD($\boldsymbol{x}[end], N, \sigma, \mu$);
16:             $end \leftarrow end - 1$;
17:             $N \leftarrow N - 1$;
18:         **end if**
19:     **end while**
20: **end function**

---

samples brings significant improvements. I can, therefore, abort sorting early, thus providing significant energy savings.
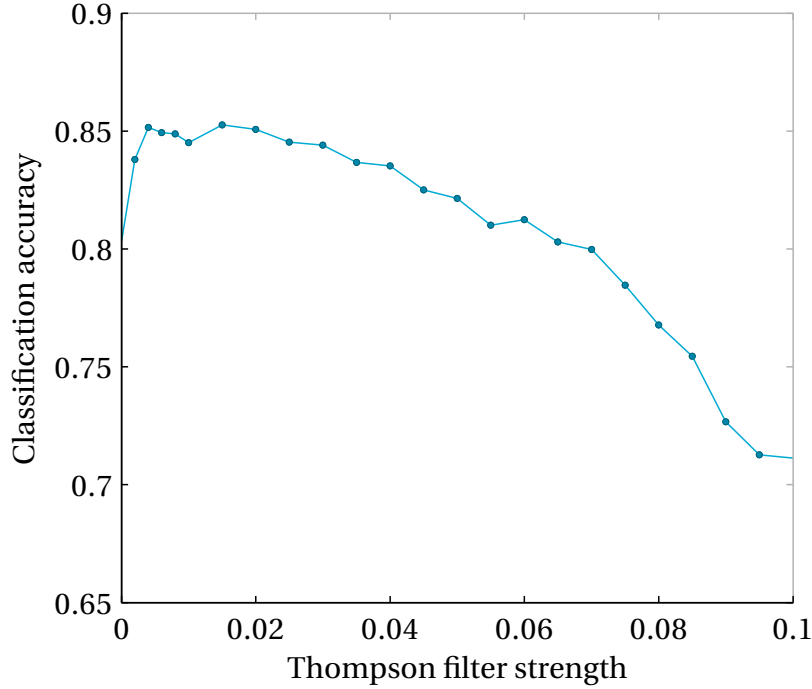
Figure 3.11 – Evolution of the accuracy depending on the strength of the Thompson filter. Even a very tolerant (weak) removal significantly improves the classification accuracy, which means that only removing the few biggest outliers is desirable.

3.4.2 Apnea Scoring

In this subsection, I define two apnea scores and discuss the time-domain analysis to compute both of them. They are the features used by the linear SVM for classifying OSA events.

Definition: The apnea score S is defined as the relative energy in the apnea band $E_{\text{apnea-band}}$ compared to the total signal energy E_{total}:

$$S = \frac{E_{\text{apnea-band}}}{E_{\text{total}}} \tag{3.3}$$

This enables us to consider the apnea score computed on the series of RR-intervals S_{RR} with the corresponding apnea band $E_{RR\text{-apnea-band}}$ as well as the apnea-score computed on the series of RS-amplitudes S_{RS} with the corresponding apnea band $E_{RS\text{-apnea-band}}$. With these definitions, the apnea score is necessarily bounded between zero and one.

To compute the energy in a given frequency band, we traditionally use a Fast Fourier Transform (FFT). This is however not possible because the data points (each beat) are irregularly spaced in time. In this case, a Lomb-Scargle normalized periodogram is typically used [124], but the drawback of this approach is the complexity associated with the generation of the whole frequency power spectrum at run time. Indeed, the frequency transform yields as many features as discrete frequencies, which involves energy-hungry computations and a large memory space. Contrarily, in my case, I only need to compute the energy of a signal in the corresponding apnea band $E_{\text{apnea-band}}$ as well as the total energy of the signal E_{total} to compute an apnea score S.

Thanks to Parseval's Theorem [125], I can compute directly the energy contained in the full signal, straight from the time-domain data. Therefore, because it is very lightweight, it is possible to substantially lower the memory and CPU usage and save energy without sacrificing the classification accuracy. Using this approach, I compute the total energy of the signal E_{total}.

$$\sum_{n=0}^{N-1} |\boldsymbol{x}(n)|^2 = \frac{1}{N} \sum_{k=0}^{N-1} |\boldsymbol{X}(k)|^2, \tag{3.4}$$

where $\boldsymbol{X}(\cdot)$ is the Discrete Fourier Transform (DFT) of signal $\boldsymbol{x}(\cdot)$, and $|x|$ is the norm of x.

For the energy of the signal in the apnea band $E_{\text{apnea-band}}$, it is also possible to avoid a frequency-domain transform. With a carefully designed time-domain band-pass filter, we remove the frequencies outside the apnea band. A first-order digital Butterworth band-pass filter [126] has been implemented to compute the coefficients a_i and b_i of the transfer function:

$$H(z) = \frac{\sum_{i=0}^{m} b_i z^{-i}}{\sum_{i=0}^{k} a_i z^{-i}}. \tag{3.5}$$

Using these coefficients, I obtain an energy-efficient time-domain filter. The time-domain Infinite Impulse Response digital filter is given as follows and is illustrated in Figure 3.12:

$$a_0 y(n) = \sum_{i=0}^{m} b_i x(n-i) - \sum_{i=1}^{k} a_i y(n-i). \tag{3.6}$$
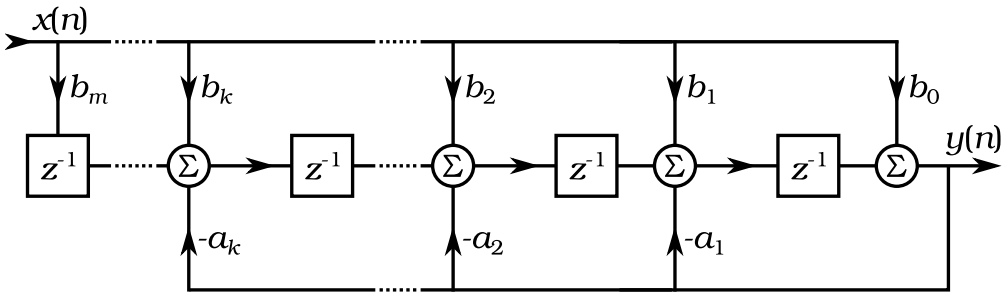
Figure 3.12 – Diagram of a rational transfer function of a filter where $a_0 = 1$ for normalization.

Having filtered the frequencies outside the apnea band, I can once again use Parseval's Theorem to obtain the apnea band energy $E_{\text{apnea-band}}$. Once this is done, the apnea score is computed according to the definition given in Equation 3.3.

Considering the time-domain energy computation, along with the time-domain band-pass signal filtering, the apnea-scoring is more efficient from the energy-consumption point of view because of the reduced algorithmic complexity, compared to computing the apnea-score using a frequency domain transform.

3.4.3 Apnea-Score Low-Pass Filtering

The reference minute-by-minute sleep-apnea labels reveal that OSA is a signal that changes relatively slowly. Intuitively, it is understandable as it is unlikely to have a single minute containing an apnea event in a long period of non-apnea sleep and vice-versa. Therefore, I consider the evolution of the apnea

score over several minutes. Hence, I use a simple unweighted moving average (see Appendix B.) to lower the raw apnea score variability, as follows:

$$\boldsymbol{x}(i) = \frac{1}{2m+1} \sum_{j=-m}^{m} \boldsymbol{x}(i+j). \tag{3.7}$$

I optimize the moving average window length by maximizing the OSA classification on the training set when changing the window length from zero minutes up to 31 minutes. The classification accuracy based on the RR apnea score increases by 3-4%, if I consider a window length of 13±4 minutes. Similarly, in the case of RS apnea score, the best improvement in accuracy observed for the database used is achieved for a five minutes window.

3.4.4 Online Classification of OSA

Once the ECG signal is measured and filtered, the features generated and filtered, the final step is to actually label the minute as apnea or not. I classify each minute of signal using a linear SVM using the parameters from the offline training done in Section 3.3. It is especially energy-efficient as it only requires computing the following condition:

$$a \cdot S_{\mathrm{RR}} + b \cdot S_{\mathrm{RS}} \geq c,$$

with the parameters a, b and c computed during an offline training (cf. Section 3.3.2). A visual representation and discussion of the algorithm's results working only with the RR-apnea score is available in Appendix C. .

3.5 Experimental Setup

This section defines the database used as well as the setup for testing the performance of the system.

3.5.1 Apnea-ECG Benchmark Recordings

The recordings used for training and testing are publicly available on PhysioNet as the apnea-ECG database [127]. They were made available for stim-

ulating research about non-intrusive OSA detection. I use this database to be able to compare my proposed system against prior studies, as I report my results using the same experimental methodology and metric based on the same signals. The 70 single-lead ECG recordings were sampled at a frequency of 100 Hz and manually labeled minute-by-minute by an expert for sleep apnea and hypopnea events, without distinction between both. The heart-beat timestamps used to retrieve the RR-intervals are provided by PhysioNet using an automatic delineation.

The recordings come from a set of 32 subjects, namely, healthy and with OSA. From those subjects, four subjects contributed to four recordings each, two subjects contributed to three recordings each, 22 subjects contributed to two recordings each and four patients contributed to a single recording. Then, the database is divided in two groups of 35 recordings, one for training and one for testing. In each group, the number of apnea events represent around 38% of the data. The total number of recorded minutes is 34313, and I include 17045 of them in the training set and the remaining 17268 in the testing set.

The duration of the recordings ranges from 6 h 41 min to 9 h 38 min, with an average duration of 8 h 12 min. The normal breathing time varies between 11 and 535 minutes, whereas for the problematic breathing, it ranges from 0 to 534 minutes. Overall, 62% of the minutes in the database are labeled as apnea. This means that a system classifying everything as apnea would only reach 62% accuracy. The amount of breathing-disordered minutes is used to classify the patients in three different groups: the apnea group A was defined as having 100 or more minutes of OSA and the control group C showed less than 5 minutes of disordered breathing. The remaining cases belong to group B, classified as borderline, i.e., with between 5 and 99 minutes with apnea during the recording.

3.5.2   OSA Classification Performance

To have a performance comparison with prior works, I use the overall classification accuracy when working on the testing set recordings. The proposed OSA monitoring system performs a minute-by-minute analysis, assigning either the *non-apnea minute* label, or the *apnea minute* one. I also provide the

specificity and sensitivity to fully characterize the system. These indicators are defined as follows:

Sensitivity (or True Positive Rate):

$$TPR = \frac{TP}{TP + FN} = \frac{TP}{RP}, \tag{3.8}$$

Specificity (or True Negative Rate):

$$TNR = \frac{TN}{FP + TN} = \frac{TN}{RN}, \tag{3.9}$$

Accuracy:

$$Acc = \frac{TP + TN}{RP + RN}, \tag{3.10}$$

where TP is the number of true positives (correctly classified minute as apnea), TN is the number of true negatives (correctly classified minute as non-apnea), FP is the number of false positives (misclassified minute as apnea), FN is the number of false negatives (misclassified minute as non-apnea), and RP and RN are, respectively, the number of real positives (apnea minutes from the ground truth) and real negatives to classify (non-apnea minutes from the ground truth).

3.6   Experimental Results

In this section, I evaluate my approach in terms of classification performance: The following four subsections report the classification accuracy under different assumptions.

First, Subsection 3.6.1 reports the performance using exactly the same setup as the PhysioNet Challenge. This enables an easy comparison with other work using the same database. The next three subsections rely on focusing of patients rather than individual recordings. To evaluate my patient-specific approach, the recordings are grouped by patient using the metadata provided along the database, using the reported age, sex, height and weight, without

Table 3.1 – Grouping of OSA recordings on a per-patient basis based on the database meta-data.

Patient	Recordings	Patient	Recordings
01	a01, a14	17	b05, x11
02	a02, x14	18	c01, x35
03	a03, x19	19	c02, c09
04	a04, a12	20	c03, x04
05	a05, a10, a20, x07	21	c04, x29
06	a06, x15	22	c05, x33
07	a07, a16, x01, x30	23	c06
08	a08, a13, x20	24	c07, x34
09	a09, a18	25	c10, x18
10	a11	26	x02
11	a15, x27, x28	27	x06, x24
12	a17, x12	28	x09, x23
13	a19, x05, x08, x25	29	x10
14	b01, x03	30	x13, x26
15	b02, b03, x16, x21	31	x17, x22
16	b04, c08	32	x31, x32

ambiguity. The results of this grouping are shown in Table 3.1, sorted in the same order as the original files. As some patients (patients 10, 23, 26 and 29) contributed to a single recording, they could not be used for the patient-specific study and, therefore, have been excluded. This is because I require one recording for training, and at least another recording for accessing the classification performance. My results from the patient-specific approach are reported in Table 3.2.

3.6.1 PhysioNet Challenge Classification Accuracy

In this work, I consider the exact same setup defined for the PhysioNet Challenge [128]. After the optimization and porting to embedded C, the accuracy reached is 85.7% while the sensitivity is 81.4% and specificity is 88.4% ($F1 = 81.3\%$). When only considering the RR-intervals series, the accuracy is 82.2%, with 73.3% sensitivity and 87.6% specificity ($F1 = 75.7\%$). These results are better by 0.1% of the best time-domain classifiers [94], and 3 to

Table 3.2 – Classification accuracy and relative improvements reported (in %) on a patient-specific basis. This table shows the change of the accuracy when changing from a patient-agnostic (P.A.) classifier to a patient-specific (P.S.) one and then to a ideal patient-specific (P.I.) classifier.

Patient	1	2	3	4	5	6	7	8	9	11	12	13	14	15	16
P.A.	64	54	90	72	74	66	80	83	86	92	80	85	95	70	98
$\frac{P.S.-P.A.}{P.A.}$	17	67	1	-2	6	11	-5	-1	5	-7	8	5	2	30	2
P.S.	75	90	91	70	78	73	76	82	90	85	86	89	97	91	100
$\frac{P.I.-P.S.}{P.S.}$	1	0	1	32	4	14	5	4	0	9	8	4	0	0	0
P.I.	76	90	92	93	81	84	79	86	90	93	92	93	97	91	100

Patient	17	18	19	20	21	22	24	25	27	28	30	31	32	**Avg.**	**Std.**
P.A.	98	100	100	100	98	99	98	99	100	98	69	100	87	**84**	**12**
$\frac{P.S.-P.A.}{P.A.}$	0	0	0	0	2	1	1	0	0	-3	1	0	11	**4**	**13**
P.S.	97	100	100	100	100	99	99	100	100	95	70	100	96	**88**	**10**
$\frac{P.I.-P.S.}{P.S.}$	0	0	0	0	0	0	0	0	0	3	20	0	0	**3**	**7**
P.I.	97	100	100	100	100	99	99	100	100	98	83	100	96	**91**	**7**

5% lower to the absolute best algorithms [97, 98], with the added benefit of online energy-efficient analysis on a wearable device.

3.6.2 Patient-Ideal (P.I.) Classification Accuracy

In order to quantify how much improvement can be expected from the patient-specific approach, for each patient, training and testing is done on the same recordings. The overall classification accuracy achievable with my technique in this ideal case, where I have full knowledge of the recordings, is 91.3% (specificity = 92.7%, sensitivity = 89.2%). This is the best achievable performance in terms of classification accuracy with my OSA detection technique.

3.6.3 Patient-Agnostic (P.A.) Classification Accuracy

To compare with a patient-specific approach, I group the recordings in Table 3.1 by patients. I use the first recording of each patient for training and the remaining recordings for testing. I leave out four patients as they contributed

to only a single recording. Using this approach, I reach a classification accuracy of OSA events of 84.5% (specificity = 88.1%, sensitivity = 79.2%).

3.6.4 Patient-Specific (P.S.) Classification Accuracy

Even though the overall classification accuracy of OSA events is high (more than 85%), few patients present significantly lower classification results (below 70%). After grouping the recordings per patient, I observe that they have a distinct apnea profile for the energy distribution in the RR and RS apnea bands. I, therefore, propose a patient-specific SVM classifier training in my OSA screening system. The recordings were done without using a CPAP machine, helping the patients to breath during their sleep. Such a machine can alter the features extracted from the ECG by the proposed solution, and therefore change the classification outcome.

In the patient-agnostic situation, I use 35 recordings coming from different patients for training, as per the PhysioNet Challenge rules, to have comparable results with the state of the art. In the patient-specific situation, I use the first overnight recording from the patient to train the SVM classifier.

With an exhaustive analysis across all the patients, I observe that, on the one hand, this patient-specific approach is overall 4.4% better than the patient-agnostic setup (specificity = 93.9%, sensitivity = 80.0%). On the other hand, my patient-specific OSA detection technique is 3% less accurate than the patient-ideal approach. Even though in a very few cases the performance is slightly degraded, this tuning brings significant improvements for the majority of patients with a very low accuracy in a patient-agnostic setting (see Table 3.2 and Figure 3.13). This observation opens the possibility of a long-term patient screening for a wider range of patients.

On the other hand, let me consider, as an example, Patient 15 who has contributed to four recordings. Without any patient-specific training, all four recordings receive a poor classification results, as the initial patient-agnostic accuracy is 69.8%. However, when training using one recording from the patient, the classification accuracy increases to 88.2% (I avoid data-dredging by excluding the recording used training). These results are similar when using any of the recordings for training, and testing against the other recordings.
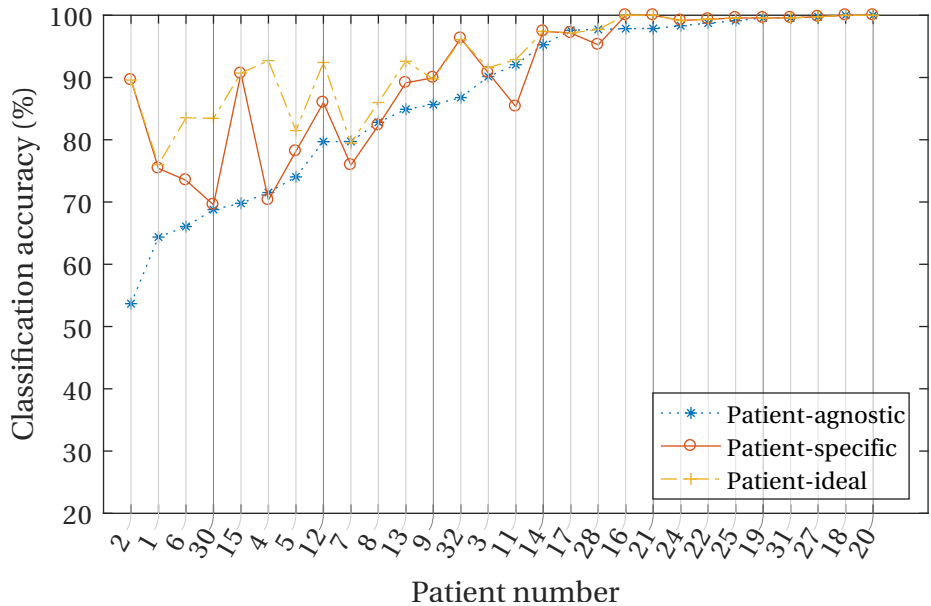
Figure 3.13 – Evolution of the OSA classification accuracy in the context of patient-agnostic, patient-specific and patient-ideal classification. The patients have been sorted according to the patient-agnostic classification accuracy to better show the classification accuracy changes with the patient-specific and patient-ideal cases.

Let me now consider Patient 1. The classification accuracy is poor compared to the majority of patients (below 70% for both the patient-specific and patient-agnostic cases). To find the best performance achievable in theory, I both train and test my technique on the recordings from the patient. This ideal classification can achieve optimal accuracy as it is based on data-dredging. The ideal accuracy, under my assumptions, is 76.0%, which is comparable with the classification accuracy obtained by my patient-specific training (75.4%).

To give more insight about the performance of my system, I plot the number of used features against the classification accuracy. As only few publications report the number of used features, many publications are not in the chart. Apart from my system, the best performing one with the lowest number of features is [97], with 87.7% accuracy using 9.7 features (non-integer value

because it is the average number of features used in multiple data splits from a cross-validation process). That is to say, it has 3% better accuracy at the cost of more than four times as many features.

It is interesting to compare this work with its different configurations (Patient-Agnostic, and Patient-Specific) against the existing publications, where the number of features if linked to the classification accuracy. However, most of the publications referenced earlier in Section 3.1 do not mention even an approximate number of features used. Thankfully, most of the best performing ones did so (accuracy greater than 80%), and are therefore included in Figure 3.14. The work I present in the current chapter, while not providing the absolute best result, is systematically using the lowest number of features. It enables an energy-efficient integration in a WSN used for cardiac monitoring.
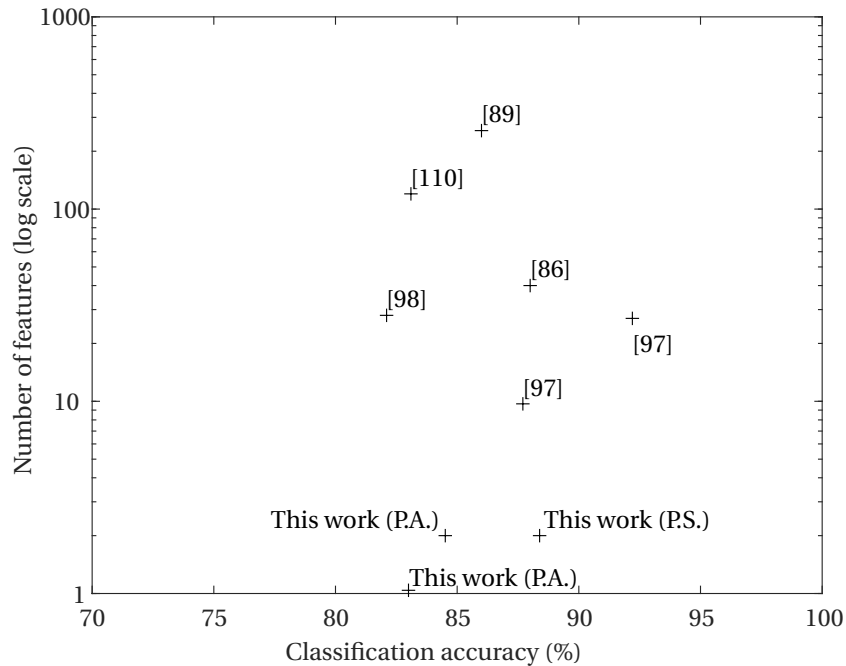
Figure 3.14 – Comparison between the number of features used against the classification accuracy among the published papers.

3.6.5   Energy Consumption Characterization

Targeting an autonomous hardware platform INYU, in addition to the classification performance, it is important to take energy consumption into consideration as a design goal. Therefore, in this Subsection, I evaluate the energy efficiency of my system experimentally. I use the commercially available Gecko EFM32 development board (with the same ARM Cortex-M3 core as in the INYU) with the provided Simplicity Studio software as a full energy profiler is integrated.

The detailed energy analysis of the OSA detection algorithm (Figure 3.15) has been performed on a recorded set of data, spanning over 12.3 hours, with an average heart-rate of 87 beats per minute. The total active time is 38.78 seconds, which represents a duty cycle of 0.085%. The average current drawn by the micro-controller while active is 10.5 mA.
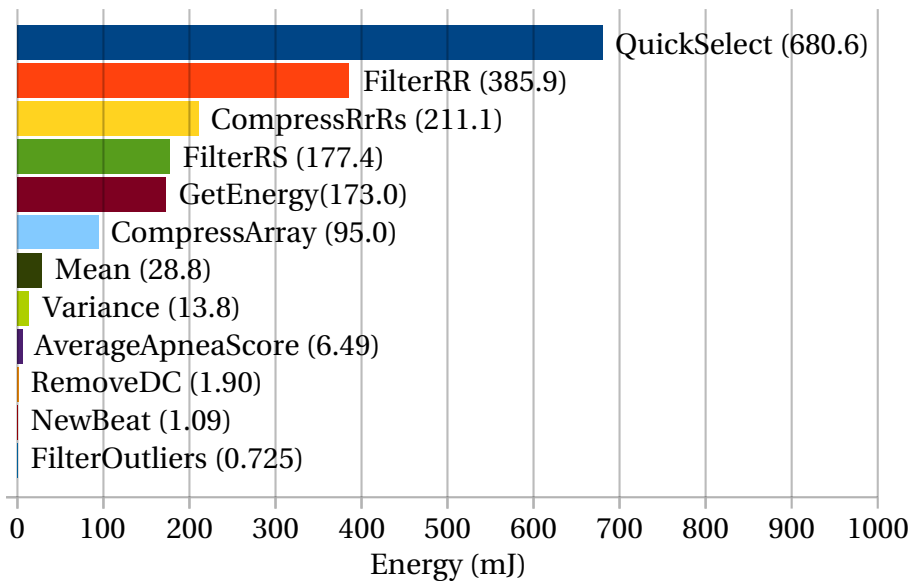
Figure 3.15 – Energy consumption of the OSA detection algorithm. Each bar represents the energy consumed by a single function in the C code running on the energy-profiling board.

In the case where the device is only used for OSA detection, the energy consumption results are reported in Table 3.3. The ECG measurement is active 100% of the time. Concerning the software, two main parts are required: the

Table 3.3 – Current used for OSA detection on the target device. The currents drawn are based both on the manufacturer's datasheets and are confirmed experimentally with measurements. As the duty cycle is data dependent, I determine an average active time for various ECGs from PhysioNet recordings, thus reflecting the variability observed for real signals.

Operation	Current (mA)	Duty cycle (%)	Avg. current (mA)
ADS1191 [129]	0.427	100	0.427
MPU-6000 [130]	0.005	100	0.005
Signal acquisition subsystem			**0.432**
ECG delineation	10.5	1.667	0.175
Apnea processing	10.5	0.085	0.009
Idle time	0.018	98.25	0.018
STM32 [104] data processing subsystem			**0.202**
nRF8001 [131]	11	0.0007	0.008
Wireless subsystem			**0.008**
Total			**0.637**

ECG delineation to detect the heart-beat and the OSA detection. In both cases, the micro-controller is in its active state, drawing 10.5 mA. When idle, the CPU is in an energy saving mode, drawing 0.018 mA. By performing an analysis of the different individual consumptions, we get an average current of 0.636 mA. As the battery is rated at 710 mAh, the total lifetime is approximately 1115 hours (46.5 days). Thus, the OSA detection technique in this thesis chapter is one order of magnitude more energy-efficient than a previous design [110], while having better classification results whether we consider a patient-agnostic or patient-specific approach. This is mainly due to the fast outliers removal and time-domain apnea score computation.

In the hypothetical case where a WSN designer changes the sampling strategy from a traditional V-ADC to a novel Time Analog-to-Digital Converter (T-ADC) as explained in Chapter 2, the energy-consumption results would be very different. With a signal acquisition subsystem consuming a current as low as 2.2 μA [50], the sampling becomes two orders of magnitude smaller, with the signal processing becoming the most significant. As we see from Table 3.3,

considering the objective of providing an autonomous system and therefore process the data on the WSN, the ECG delineation would become the major energy consumer. While the delineation algorithm can likely be improved for lower power processing, it is not sure the savings would be enough. A partial improvement can be in the form of a hardware-accelerated bio-signal processing acceleration [132]. Therefore, the system would benefit from spending more time in a power-saving sleep mode. Another improvement would be the use of a full-custom hardware design for ECG delineation and OSA detection. This is however a complex, application specific development. In all cases, the system I designed in this chapter has the benefit of relying on existing components available for use. This is opposed to using a T-ADC based sampling or specific hardware processing developments that are at research state at best, which nonetheless are promising opportunities for the future WSN.

3.7   Summary and Concluding Remarks

OSA is an aggravating factor for different health conditions, including cardiovascular diseases. Despite the high rate of obstructive sleep apnea, only a small fraction of the population is diagnosed and monitored. Therefore, in this chapter, I designed an online ultra-low power wearable OSA monitoring system. The performance and energy efficiency of my system are evaluated experimentally, in a patient-specific setting. My system has a classification accuracy of 88.2%, for a minute-by-minute classification, with a battery lifetime of 46.8 days. Thanks to its Bluetooth link, this wearable sensor can upload its analysis to an online web-service for a continuous monitoring, tracking the evolution of the disease.

Conclusion

C ONCLUDING this thesis, this chapter highlights the main contributions of my research targeting low-power biomedical devices from complementary optimization perspectives. As such, these wearable Wireless Sensor Node (WSN) have strong constraints in terms of battery life. This is the main limitation, because it limits the amount of operations that can be performed, whether it is for sampling bio-signals like the electrocardiogram (ECG), processing data, or transmitting the results using a wireless connection. This is the framework in which I pursued both lowering the energy spent for signal acquisition and signal processing in the case of Obstructive Sleep Apnea (OSA).

4.1   Signal Sensing

Sampling signals is a necessity for signal processing. However, it turns out to be a major consumer of energy. I pursued research towards lower power sampling, by exploring the use of non-Nyquist sampling for the case of cardiovascular monitoring systems. My contributions are the following:

- I did set up a methodology to compare the performance of event-triggered sampling systems in comparison of a classical uniform sampling approach. This is applied to the context of cardiovascular applications, where the quality of the signal analysis is crucial. Using this framework, it is possible to directly compare the performance results of different approaches.

- I provided three event-driven sampling schemes, with different levels of complexity. Starting from an intuitive level-crossing strategy, I highlight its benefits and drawbacks. Building on this initial approach, I propose an alternative that triggers measurements whenever a modeled error is reaching a design threshold. This enables the system's user to specify the admissible distance between the sampled signal and the original signal. The last sampling strategy relies on the knowledge we have of the measured signal. In case of an ECG, it is a periodic signal, with a heart-rate constrained between known bounds, and with a specific variability. With this knowledge, it is possible to dynamically change the acceptable sampling error depending on evolving requirements, lowering further the energy consumption of the WSN. By reducing the average sampling rate from 360 Hz down to 13.6 Hz, it becomes possible to increase the battery lifetime by 4× with a marginal impact on the accuracy of heart-rate analysis.

- Finally, I proposed a simple probabilistic heart-beat correction algorithm for real-time applications. This is a research topic tackling a problem where noisy ECGs are the cause of erroneous delineations for healthy users. From four different scenarios (normal beat, skipped beat, extra beat, and misplaced beat), my algorithm models likelihood of each situation, and immediately classifies and corrects each new

detected beat accordingly. This approach is therefore relevant for any system with a direct interaction with a user, such as a sport tracking smart-phone app, or a video-game that adapts the game-play depending on the player's heart-rate. Overall, the proposed algorithm achieves 99.5% sensitivity in the detection and correction of erroneous beats. In addition, it features a fast response time when the activity level of the user changes, thus enabling its use in situations where the heart rate rises quickly.

List of publications:
- Event-Triggered Sensing for Low-Power Cardiovascular Monitoring Systems, *publication pending in IEEE Council on Electronic Design Automation*
- Real-Time Probabilistic Heart Beat Classification and Correction for Embedded Systems, *Computing in Cardiology Conference (CinC)*, 2015

4.2   Screening Obstructive Sleep Apnea

After a hardware approach where I targeted the first component of a WSN in the full signal acquisition and processing chain, I presented a wearable, accurate, and energy efficient system for monitoring OSA on a long-term basis. This system reduces the gap between home health- care and professional supervision. The main contributions are the following:

- From a multitude of features extracted from the ECG, I selected only two of them, combined in a way that maximizes the classification accuracy. Compared to the State-of-the-Art, I reached a high performance rate while relying on a significantly smaller set of features. This is a design constraint imposed by the nature of wearable WSNs, as the low processing resources limit the number of features that can be extracted and processed on the device, not even considering the lowered battery-life.

- For the C implementation of the algorithm, many algorithmic optimizations were applied. First, a Thompson-Tau outlier removal is thoroughly optimized, reducing its complexity to the theoretical lower bound. Second, the spectral power band linked to OSA is extracted from the ECG thanks to a time-domain analysis, rather than using an energy-hungry frequency-domain transform. This approach has the advantage of improving the processing speed, lowering the memory used, and thus reducing the overall energy-consumption while keeping the same classification accuracy.

- Finally, I went further than the State-of-the-Art by performing a specific training of the algorithm for each patient. For this, I performed a grouping of the recordings on a per-patient basis. This method increases the performance of the system, notably for patients for which the default configuration is not optimal.

List of publications:

- Design of Ultra-Low-Power Smart Wearable Systems, *Proceedings of the IEEE 16th Latin-American Test Symposium (LATS)*, 2015
- Low-Power Wearable System for Real-Time Screening of Obstructive Sleep Apnea, *IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, 2016
- Online Obstructive Sleep Apnea Detection on Medical Wearable Sensors, *IEEE Transactions on Biomedical Circuits and Systems (TBioCaS)*, 2018
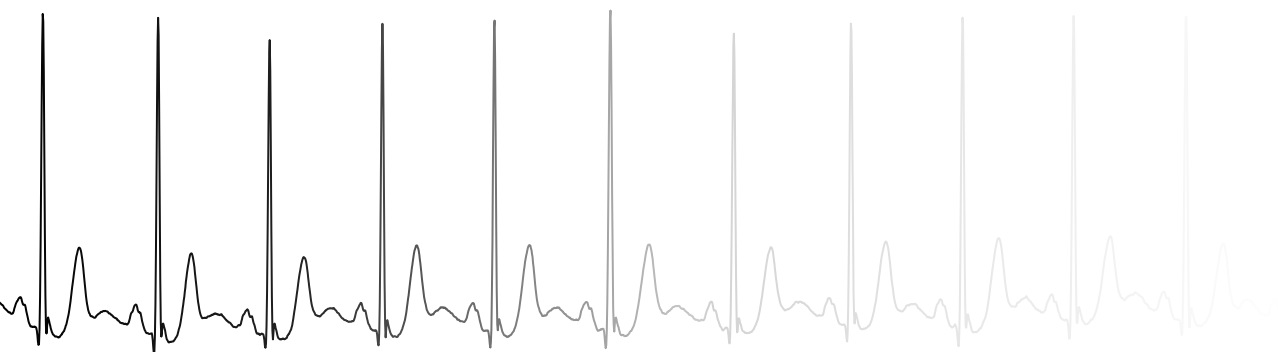
Future Work

T HE work presented in this thesis opens multiple directions for future
research and development projects, depending on the objective con-
sidered. The Time Analog-to-Digital Converter (T-ADC) hardware, as it exists
currently [50], is promising for a number of applications, requiring the de-
velopment of the right tools to make most of it. It can however be pushed
further by considering strategies similar to the error-based sampling scheme.
As for Obstructive Sleep Apnea (OSA), is worth considering neural network
detection as a recent work shows promising results [133]. Also, because of the
high OSA prevalence, it is worth investigating solutions for a broader reach of
the population.

5.1 The Signal Sensing Case

Considering the potential reduction in the volume of acquired, processed
and transmitted data that the proposed approach (Chapter 2) may achieve in
the light of the results of this thesis, I envision different future lines of work at
the system and hardware levels.

5.1.1 System Level Work

At the system level, the efforts should be clearly focused on signal processing
and algorithms optimization. Because of the ubiquity of uniform sampling,
almost all the available tools assume this scenario. One major example is

the Fourier Transform. The Lomb-Scargle periodogram is an equivalent for irregularly sampled signals, but has shortcomings for asymptotic analysis. Therefore, building strong mathematical models and tools is essential for such event-triggered sampled signals. In this sense, the development of efficient filtering methods is particularly relevant.

On the other hand, successful domain-specific algorithms should be adapted to the non-uniform sampling scenario, in order to avoid interpolated reconstructions that unnecessarily increase the processing requirements. As an example, when using a T-ADC sensing for the electrocardiogram (ECG), it becomes highly relevant to start working on adapted versions of algorithms used for the different tasks that are usually performed on the raw signal: heart beat detection and classification, wave delineation, etc. The quantification of the performance and energy efficiency of both a T-ADC and a Voltage Analog-to-Digital Converter (V-ADC) approach is all the more relevant, along with the estimated potential benefits at the application level.

5.1.2   Hardware Level Work

Event-triggered systems show clear benefits in terms of energy savings in the context of complex biological signals. An extensive study between V-ADC and T-ADC would highlight the specific strength of each solution, across signals shapes and frequencies. Given the difference of measurement domains (voltage against time), a proper evaluation requires comparing the energy-consumption of two approaches with a constant signal accuracy.

The following list can be considered as potential research ideas in the domain of T-ADC:

- Design a Wireless Sensor Node (WSN) prototype using a T-ADC chip that is able to continuously measure and process bio-signals, keeping the power consumption at a minimum. An extension of the existing capabilities of the existing T-ADC [50] is worth considering as waking up an external micro-controller of each new value has a non-negligible overhead. The capacity of sending an interrupt `DATA_READY` after N triggers or T milliseconds opens the way to batch processing of the

collected data. This is of particular interest in novel multi-core Internet of Things (IoT) processors, such as the PULP Platform [134].

- Design a fully configurable level-crossing front-end for analog signals. Even though the principle design is straightforward (multiple voltage comparators), an integrated design must pay close attention to the global energy consumption. Such a front-end has an analog input and digital outputs, which makes it suitable to have configurable levels. With such reconfigurability, it is possible to envision using only two comparators working in tandem, tracking exclusively the levels above and below the current voltage of the signal. As the levels are externally defined by software, there is no need to transmit the absolute voltage measured, but only one `UP/DOWN` bit is enough.

- Design an error-overflowing front-end for analog signals. Indeed, a level-crossing strategy has shortcomings. A constant signal between two thresholds would have a constant error. A ramp will regularly trigger the T-ADC, providing no meaningful information. More generally, any signal with a constant derivative can be fully characterized with only two points in time and voltage. Based on this, an analog front-end integrating the error between the signal and the linear interpolation of the last two measurements points would be extremely efficient for a wide class of common signals. In this sense, we consider a promising line of work the hardware implementation of the Wall-Danielsson method [52], with the objective of obtaining a new analog-to-digital conversion device that:

  1. directly provides the sequence of linear segments that best approximate the input signal,

  2. can be dynamically reconfigured to adapt the error threshold to the target task (in the case of ECG processing, this could be heart rate analysis, wave delineation, etc.).

The implementation could be done in the digital domain, after a classical Analog-to-Digital Converter (ADC) stage, or directly in the ana-

log domain. The advantages and drawbacks of these two approaches should be studied in depth.

5.2 The Obstructive Sleep Apnea Case

When considering the case of OSA, two different directions of research and development are possible, whether we consider the need of an additional device or not.

5.2.1 With an Additional Specialized Device

The case where we consider the need for a specific device is the context considered for this thesis. Building on this work, multiple ways of research can be envisioned:

- Use a T-ADC sensing approach to lower the energy consumption of the device. Because of the paradigm shift induced, the sampling and wireless transmission of the raw signal is not anymore unaffordable energy-wise. Processing the data on-device therefore requires additional optimizations of the algorithms used. Indeed, I showed that, by changing the signal sensing from V-ADC to T-ADC, the main current drain is the ECG delineation algorithm. Another approach in delineation and implementation is therefore required to keep the online signal processing relevant compared to the wireless streaming.

- While the OSA detection algorithm presented in this thesis is accurate and energy-efficient, it is not the most accurate one. Other approaches (using multiple other features extracted from the signal, or relying on a trained neural network) reported a patient-agnostic classification above 90%. It is interesting to consider porting a more accurate algorithm to a similar WSN, improving the diagnostics for the patients. This is to be combined with efforts towards embedded and energy-efficient neural network for edge-processing [135, 136].

- OSA is connected to cardiac problems. The results of an online cardiac analysis can possibly be used to refine the OSA analysis. For example, by modifying an existing arrhythmia-detection algorithm, it is possi-

ble to give a *pathological score* to each beat individually, using this information for improving the OSA classification.

5.2.2 Without Any Specific Equipment

Requiring a specific device for using the OSA screening algorithm is a drawback, even if it were easily available and inexpensive. Ideally, using existing devices should be considered as the ideal situation. In that mindset, multiple possibilities are envisioned:

- Personal activity trackers (fitness trackers and smartwatches) are getting widespread, and the majority of them feature a pulse oximeter, yielding a photoplethmography (PPG) and therefore an indication of the heart-rate. Hence, it would be interesting to evaluate the performance of the OSA detection algorithm relying only on the RR-intervals measured with the PPG, thus using devices already available on the mass market. In case multiple signals are available, sensor fusion could be the way for improving the OSA classification accuracy.

- Smartphones are becoming ubiquitous. Relying on the sensors they provide, it is likely that they can be used to detect OSA. Not going as far as to extract the heart-beats using the microphone, analyzing the snoring and breathing patters is a promising approach. In all cases, these techniques should ideally be able to discriminate each individual in the room.
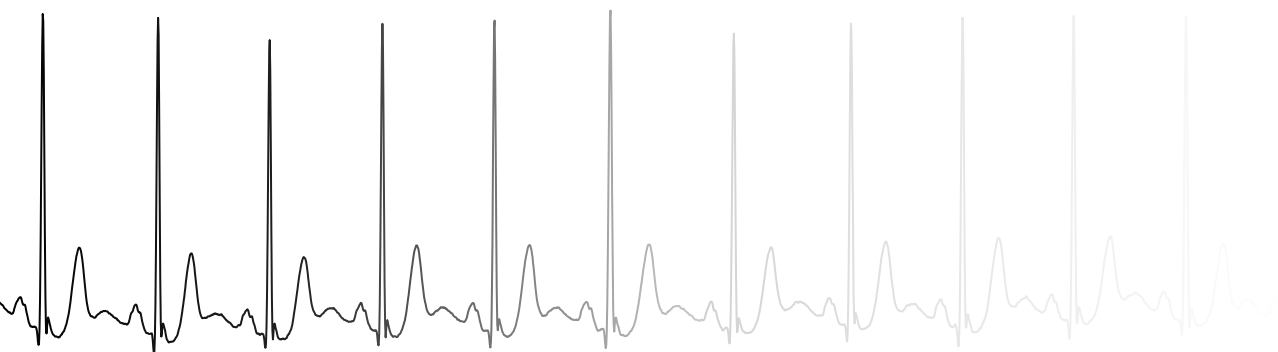
Appendix

Following, you can find the appendices and supplementary material to this thesis. The data is also available in a processable format, accessible using the link provided from each appendix.

A. Signal Sensing Experimental Results

With the aim of supporting reproducible research, the full source code of the algorithms to reproduce the experiments presented has been published under an Open Source License, along with the individual results: https://c4science.ch/diffusion/8502

The table below shows the dis-aggregated sensitivity (Se), positive predictivity (+P), F1 score and average sampling frequency (fs) for each of the compared methods considered in Chapter 2, with each evaluated record from the MIT-BIH Arrhythmia Database. The analysis of the results is available in Section 2.3.

Record	Periodic Sampling (360 Hz)			Compressed Sensing (180 Hz)			Level Crossing				Knowledge Based			
	Se	+P	F1	Se	+P	F1	Se	+P	F1	fs	Se	+P	F1	fs
100	100.00	100.00	100.00	99.89	100.00	99.94	100.00	100.00	100.00	32.80	100.00	100.00	100.00	12.56
101	99.93	99.87	99.90	99.87	99.87	99.87	99.93	99.87	99.90	33.18	99.93	99.87	99.90	9.52
103	99.94	100.00	99.97	99.88	100.00	99.94	99.94	100.00	99.97	43.90	99.94	100.00	99.97	11.44
105	99.68	98.99	99.33	99.49	99.08	99.28	99.44	98.66	99.05	44.07	99.54	99.12	99.33	11.94
106	99.94	99.88	99.91	99.82	99.94	99.88	99.82	99.88	99.85	45.59	99.88	99.88	99.88	13.57
107	99.83	99.94	99.88	99.72	100.00	99.86	99.78	99.94	99.86	76.21	98.32	99.94	99.12	11.35
108	99.39	99.73	99.56	99.12	99.53	99.32	98.58	98.58	98.58	30.38	99.05	99.39	99.22	23.32
109	99.90	100.00	99.95	99.81	100.00	99.90	99.90	100.00	99.95	49.23	99.71	100.00	99.85	8.57
111	99.55	100.00	99.77	99.44	100.00	99.72	99.44	100.00	99.72	32.74	99.61	100.00	99.80	14.31
112	100.00	100.00	100.00	99.95	100.00	99.97	100.00	100.00	100.00	28.82	100.00	100.00	100.00	16.14
113	100.00	99.87	99.93	99.93	99.80	99.86	100.00	99.87	99.93	50.07	100.00	100.00	100.00	11.20
114	100.00	100.00	100.00	99.94	100.00	99.97	100.00	100.00	100.00	40.77	100.00	100.00	100.00	8.43
115	100.00	100.00	100.00	99.94	100.00	99.97	100.00	100.00	100.00	47.44	100.00	100.00	100.00	11.15
116	99.16	99.85	99.50	99.01	99.85	99.43	98.96	99.90	99.43	70.13	99.16	99.85	99.50	12.36
117	100.00	100.00	100.00	99.92	100.00	99.96	100.00	100.00	100.00	29.38	100.00	100.00	100.00	15.48
118	100.00	99.95	99.97	99.95	100.00	99.97	100.00	99.95	99.97	57.95	100.00	100.00	100.00	19.49
119	100.00	99.76	99.88	99.94	99.64	99.79	100.00	99.70	99.85	50.18	100.00	99.94	99.97	11.26
121	99.94	99.94	99.94	99.87	100.00	99.93	99.74	99.94	99.84	20.72	99.81	99.94	99.87	7.73
122	100.00	100.00	100.00	99.95	100.00	99.97	100.00	100.00	100.00	45.89	100.00	100.00	100.00	9.74
123	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	41.28	100.00	100.00	100.00	11.34
124	99.93	100.00	99.96	99.85	100.00	99.92	99.85	100.00	99.92	37.24	99.78	100.00	99.89	6.92
200	99.86	99.82	99.84	99.82	99.91	99.86	99.82	99.82	99.82	46.14	99.86	99.82	99.84	16.88
201	98.22	99.40	98.81	97.90	99.27	98.58	97.04	99.93	98.46	19.11	97.50	100.00	98.73	12.78
202	99.57	99.95	99.76	99.52	99.95	99.73	99.52	100.00	99.76	35.30	99.63	100.00	99.81	10.44

Record	Periodic Sampling (360 Hz)			Compressed Sensing (180 Hz)			Level Crossing				Knowledge Based			
	Se	+P	F1	Se	+P	F1	Se	+P	F1	fs (Hz)	Se	+P	F1	fs (Hz)
203	98.35	99.11	98.73	97.94	99.31	98.62	98.19	98.86	98.52	55.19	97.78	98.78	98.28	20.31
205	99.36	100.00	99.68	99.27	100.00	99.63	99.32	100.00	99.66	31.20	99.32	100.00	99.66	13.13
207	99.94	99.56	99.75	99.94	99.81	99.87	99.62	99.37	99.49	30.24	99.12	99.43	99.27	9.30
208	99.18	99.88	99.53	99.18	99.92	99.55	99.43	99.88	99.65	51.96	99.34	99.92	99.63	17.18
209	100.00	100.00	100.00	99.96	100.00	99.98	100.00	100.00	100.00	52.97	100.00	100.00	100.00	18.03
210	98.19	99.91	99.04	98.00	99.91	98.95	97.96	99.95	98.94	31.01	97.69	99.91	98.79	12.16
212	100.00	100.00	100.00	99.91	100.00	99.95	100.00	100.00	100.00	57.20	100.00	100.00	100.00	15.64
213	99.96	100.00	99.98	99.85	100.00	99.92	99.93	100.00	99.96	86.82	99.89	100.00	99.94	15.94
214	100.00	100.00	100.00	99.95	100.00	99.97	100.00	99.95	99.97	45.92	99.79	99.79	99.79	9.11
215	99.93	100.00	99.96	99.86	100.00	99.93	99.96	100.00	99.98	56.22	99.93	100.00	99.96	21.80
217	99.84	99.89	99.86	99.78	99.89	99.83	99.84	99.89	99.86	56.15	99.84	99.89	99.86	13.12
219	100.00	98.50	99.24	99.94	96.57	98.23	100.00	98.99	99.49	44.30	100.00	98.66	99.33	13.28
220	100.00	100.00	100.00	99.94	100.00	99.97	100.00	100.00	100.00	48.38	100.00	100.00	100.00	11.85
221	99.90	100.00	99.95	99.85	100.00	99.92	99.95	100.00	99.97	39.17	99.75	99.91	99.87	14.24
222	99.29	99.95	99.62	98.82	100.00	99.41	98.87	100.00	99.43	30.70	99.53	99.91	99.72	18.76
223	99.95	100.00	99.97	99.91	100.00	99.95	99.95	100.00	99.97	42.16	99.86	100.00	99.93	12.62
228	99.88	99.65	99.76	99.88	99.77	99.82	99.82	99.65	99.73	35.40	99.53	99.18	99.35	16.74
230	100.00	100.00	100.00	99.95	100.00	99.97	100.00	100.00	100.00	49.53	100.00	99.95	99.97	12.80
231	100.00	100.00	100.00	99.92	100.00	99.96	100.00	100.00	100.00	35.79	100.00	100.00	100.00	11.21
232	100.00	99.53	99.76	100.00	99.66	99.83	100.00	99.40	99.70	23.27	100.00	99.46	99.73	17.46
233	99.92	100.00	99.96	99.88	99.96	99.92	99.92	99.96	99.94	56.61	99.92	99.96	99.94	22.06
234	99.96	100.00	99.98	99.91	100.00	99.95	99.96	100.00	99.98	41.71	99.96	100.00	99.98	10.80
Gross	99.73	99.85	99.79	99.64	99.82	99.73	99.66	99.83	99.74	43.70	99.62	99.84	99.73	13.62

B. Weighted Moving Average for Apnea-Score Filtering

For further investigation to optimize the length of the classification window, a low-pass filter is applied to the apnea-score, as explained in Subsection 3.4.3. While implemented as a simple unweighted moving average, the weighting of each sample can have an influence on the classification accuracy. Therefore, I considered different kind of weighted moving average filters and window lengths, always symmetrical with respect to the central value of the window:

1. a **square** weighting, corresponding to the unweighted scheme,

2. a **triangular** weighting, with linear decrease from the central value to zero towards the end of the considered window,

3. three **exponential** decreases of varying strengths,

4. a **square-root** falloff.

From Figure B. .1, we notice the absence of any significant benefit to change the window shape, as the evolution and peak value of each one is similar. Indeed, the accuracy is improved by 1.5% for a window length between 11 and 23 minutes, depending on the shape. When targeting a low-power implementation, the computational simplicity saves energy, hence the choice of an unweighted moving average (i.e. the square weighting in Figure B. .1).

Figure B. .1 – Classification accuracy of OSA with different weighted average configurations. The line style is connected to the weighting shape.

C.   RR-Apnea Score and Obstructive Sleep Apnea Evaluation

As explained in Chapter 3, the OSA is correlated with features that are extracted from the ECG. Presented in Figure 3.6, there is a visible link between the low frequencies contained in the series of durations between two heartbeats (RR-intervals) and OSA.

A frequency spectrum can however give wrong impressions when seeing it. One possible reason is the difficulty to represent 2D data with a color map, where a possible linearity can be visually distorted because of a non-linear brightness evolution. For this reason, I represented in Figure C. .1 the evolution of the sleep apnea score defined in Subsection 3.4.2 for one recording, along with the corresponding OSA labels provided with the recording.

Provided as additional information, I completed the figure with the intermediate and final results from my OSA detection algorithm: the unfiltered and filtered apnea estimation.

While the moving average filtering applied (see Subsection 3.4.3) provides no obvious improvement for this specific recording, it showed being useful when evaluating its effects over the whole database. Indeed, before the moving average filtering, there are 421 minutes correctly classified out of 501 minutes. The filtering rises the number of correctly classified minutes to 432. A detrimental effect of the filtering is visible for example between the minutes 230 and 270. There are four minutes originally correctly classified as non-OSA that, after filtering, are erroneously classified as OSA. On the other hand, between minutes 325 and 335, six minutes are properly corrected from non-OSA to OSA.
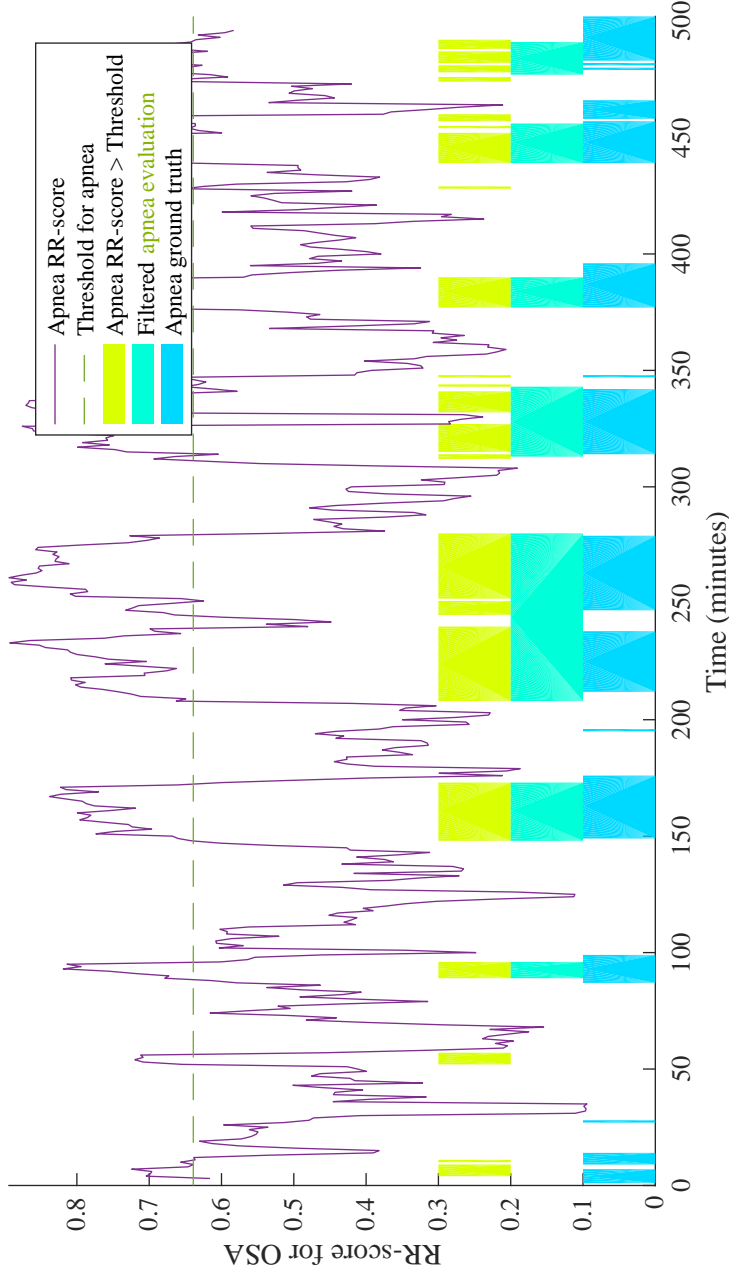
Figure C. .1 – RR-score for OSA computed from the ECG, along with the chosen threshold. On the lower part of the graph, three representation of Boolean signals:

• from .2 to .3: the direct result of the OSA evaluation, from the apnea RR-score and the threshold,

• from .1 to .2: the unweighted moving average of the direct apnea evaluation, with a window length of 13 minutes,

• from 0 to .1: the OSA labels provided by the database, annotated by an expert.

[Data source: Record a19 from the PhysioNet Challenge 2000 ECG-Apnea database]
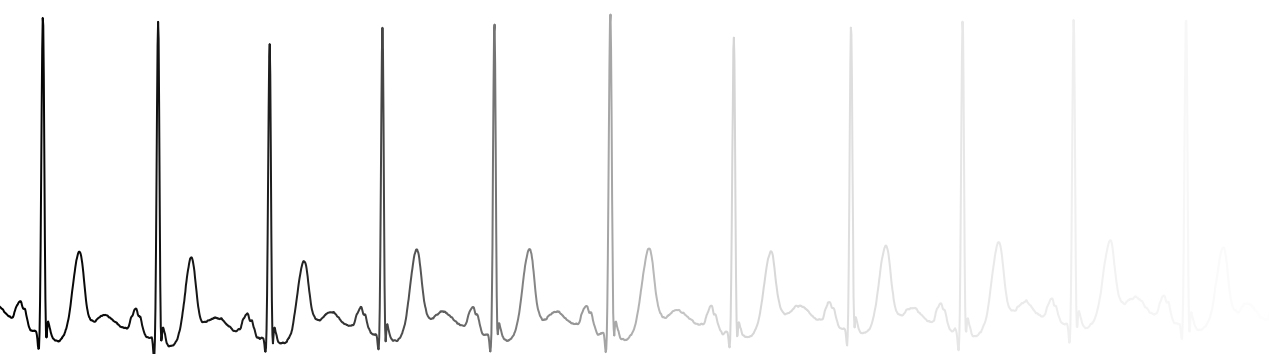
Bibliography

[1] Kumaravel Rajakumar et al. Infantile scurvy: a historical perspective. *Pediatrics*, 108(4):E76, 2001.

[2] World Health Organization. *World health statistics 2015*. World Health Organization, 2015.

[3] Jacques Mézard. *Dérives thérapeutiques et dérives sectaires : la santé en danger*. French Senate: Commission d'enquête Mouvements à caractère sectaire, 2013.

[4] Victoria Maizes, David Rakel, and Catherine Niemiec. Integrative medicine and patient-centered care. *Explore: The Journal of Science and Healing*, 5(5):277–289, 2009.

[5] HR Leavel. E gurney clark as „the science and art of preventing disease, prolonging life, and promoting physical and mental health and efficiency. *Preventive Medicine for the Doctor in his Community (3rd ed.). Huntington, NY: Robert E. Krieger Publishing Company*, 1979.

[6] Bert Arnrich, Oscar Mayora, Jakob Bardram, and Gerhard Tröster. Pervasive healthcare. *Methods of information in medicine*, 49(01):67–73, 2010.

[7] Stephen M Petterson, Winston R Liaw, Robert L Phillips, David L Rabin, David S Meyers, and Andrew W Bazemore. Projecting us primary care

physician workforce needs: 2010-2025. *The Annals of Family Medicine*, 10(6):503–509, 2012.

[8] World Health Organization. Noncommunicable diseases. https://www.who.int/en/news-room/fact-sheets/detail/noncommunicable-diseases. [Online; accessed 2019-06-03].

[9] Mohammad H Forouzanfar, Ashkan Afshin, Lily T Alexander, H Ross Anderson, Zulfiqar A Bhutta, Stan Biryukov, Michael Brauer, Richard Burnett, Kelly Cercy, Fiona J Charlson, et al. Global, regional, and national comparative risk assessment of 79 behavioural, environmental and occupational, and metabolic risks or clusters of risks, 1990–2015: a systematic analysis for the global burden of disease study 2015. *The Lancet*, 388(10053):1659–1724, 2016.

[10] World Health Organization. Noncommunicable diseases (ncds) and mentha health: challenges and solutions. https://www.who.int/nmh/publications/ncd-infographic-2014.pdf. [Online; accessed 2019-06-03].

[11] Department of Economic and Social Affairs Population Division. *World Population Ageing*. United Nations, 2015.

[12] *World health statistics 2016: monitoring health for the SDGs sustainable development goals*. World Health Organization, 2016.

[13] World Health Organization. Measles. https://www.who.int/en/news-room/fact-sheets/detail/measles. [Online; accessed 2019-06-03].

[14] Melanie Swan. Emerging patient-driven health care models: an examination of health social networks, consumer personalized medicine and quantified self-tracking. *International journal of environmental research and public health*, 6(2):492–525, 2009.

[15] John R Marler, BC Tilley, M Lu, Thomas G Brott, PC Lyden, JC Grotta, JP Broderick, SR Levine, MP Frankel, SH Horowitz, et al. Early stroke treatment associated with better outcome: the ninds rt-pa stroke study. *Neurology*, 55(11):1649–1655, 2000.

[16] The ATLANTIS et al. Association of outcome with early stroke treatment: pooled analysis of atlantis, ecass, and ninds rt-pa stroke trials. *The Lancet*, 363(9411):768–774, 2004.

[17] Edward C Jauch, Jeffrey L Saver, Harold P Adams Jr, Askiel Bruno, JJ Connors, Bart M Demaerschalk, Pooja Khatri, Paul W McMullan Jr, Adnan I Qureshi, Kenneth Rosenfield, et al. Guidelines for the early management of patients with acute ischemic stroke: a guideline for healthcare professionals from the american heart association/american stroke association. *Stroke*, 44(3):870–947, 2013.

[18] Rubén Braojos, Hossein Mamaghanian, Alair Dias Junior, Giovanni Ansaloni, David Atienza, Francisco J Rincón, and Srinivasan Murali. Ultra-low power design of wearable cardiac monitoring systems. In *Proceedings of the 51st Annual Design Automation Conference*, pages 1–6. ACM, 2014.

[19] Paul E. Peppard, Terry Young, Mari Palta, and James Skatrud. Prospective Study of the Association between Sleep-Disordered Breathing and Hypertension. *New England Journal of Medicine*, 342(19):1378–1384, 2000.

[20] H. Klar Yaggi, John Concato, Walter N. Kernan, Judith H. Lichtman, Lawrence M. Brass, and Vahid Mohsenin. Obstructive Sleep Apnea as a Risk Factor for Stroke and Death. *New England Journal of Medicine*, 353(19):2034–2041, 2005.

[21] CM Schröder and R O'Hara. Depression and Obstructive Sleep Apnea (OSA). *Ann Gen Psychiatry*, 4:13, 2005.

[22] Jeffrey S Durmer and David F Dinges. Neurocognitive consequences of sleep deprivation. In *Seminars in neurology*, volume 25, pages 117–129. Copyright© 2005 by Thieme Medical Publishers, Inc., 333 Seventh Avenue, New York, NY 10001, USA., 2005.

[23] Apoor S Gami, Daniel E Howard, Eric J Olson, and Virend K Somers. Day–night pattern of sudden death in obstructive sleep apnea. *New*

England Journal of Medicine, 352(12):1206–1214, 2005.

[24] U.S. Department of Health and Human Services. Sleep apnea. https://www.nhlbi.nih.gov/health-topics/sleep-apnea. [Online; accessed 2016-08-16].

[25] Terry Young, Paul E Peppard, and Daniel J Gottlieb. Epidemiology of obstructive sleep apnea: a population health perspective. *American journal of respiratory and critical care medicine*, 165(9):1217–1239, 2002.

[26] JinKwan Kim, KwangHo In, JeHyeong Kim, SeHwa You, KyungHo Kang, JaeJeong Shim, SangYeub Lee, JungBok Lee, SeungGwan Lee, Chan Park, et al. Prevalence of sleep-disordered breathing in middle-aged korean men and women. *American journal of respiratory and critical care medicine*, 170(10):1108–1113, 2004.

[27] Jamie CM Lam and Mary SM Ip. An update on obstructive sleep apnea and the metabolic syndrome. *Current opinion in pulmonary medicine*, 13(6):484–489, 2007.

[28] Terry Young, Linda Evans, Laurel Finn, Mari Palta, et al. Estimation of the clinically diagnosed proportion of sleep apnea syndrome in middle-aged men and women. *Sleep*, 20(9):705–706, 1997.

[29] HW Agnew Jr, Wilse B Webb, and Robert L Williams. The first night effect: An eeg studyof sleep. *Psychophysiology*, 2(3):263–266, 1966.

[30] Thomas Penzel, George B Moody, Roger G Mark, Ary L Goldberger, and J Hermann Peter. The apnea-ecg database. In *Computers in cardiology 2000*, pages 255–258. IEEE, 2000.

[31] Francisco Rincón, Joaquin Recas, Nadia Khaled, and David Atienza. Development and evaluation of multilead wavelet-based ecg delineation algorithms for embedded wireless sensor nodes. *IEEE Transactions on Information Technology in Biomedicine*, 15(6):854–863, 2011.

[32] Grégoire Surrel, Amir Aminifar, Francisco Rincón, Srinivasan Murali,

and David Atienza. Online obstructive sleep apnea detection on medical wearable sensors. *IEEE Transactions on Biomedical Circuits and Systems*, 12(4):762–773, 2018.

[33] D. Sopic, A. Aminifar, A. Aminifar, and D. Atienza. Real-time classification technique for early detection and prevention of myocardial infarction on wearable devices. In *IEEE Biomedical Circuits and Systems Conference*, 2017.

[34] D. Sopic, E. De Giovanni, A. Aminifar, and D. Atienza. A hierarchical cardiac rhythm classification methodology based on electrocardiogram fiducial points. In *Computing in Cardiology*, 2017.

[35] E. De Giovanni, A. Aminifar, A. Luca, S. Yazdani, Jean-Marc Vesin, and D. Atienza. A patient-specific methodology for prediction of paroxysmal atrial fibrillation onset. In *Computing in Cardiology*, 2017.

[36] Victoriano Montesinos Canovas, Fabio Isidoro Tiberio Dell'Agnola, Adriana Arza Valdes, Amir Aminifar, and David Atienza Alonso. Multimodal acute stress recognition using off-the-shelf wearable devices. Technical report, 2019.

[37] Farnaz Forooghifar, Amir Aminifar, and David Atienza Alonso. Self-aware wearable systems in epileptic seizure detection. In *Euromicro Conference on Digital System Design*, pages 426–432. IEEE, 2018.

[38] Francisco J Rincón, Laura Gutiérrez, Mónica Jiménez, Víctor Díaz, Nadia Khaled, David Atienza, Marcos Sánchez-Elez, Joaquín Recas, and Giovanni De Micheli. Implementation of an automated ecg-based diagnosis for a wireless body sensor platform. In *Proceedings of the International Conference on Biomedical Electronics and Devices (BIODEVICES 2009)*, pages 88–96, 2009.

[39] Dionisije Sopic, Amir Aminifar, and David Atienza. e-glass: A wearable system for real-time detection of epileptic seizures. In *IEEE International Symposium on Circuits & Systems*, pages 1–5. IEEE, 2018.

[40] Dionisije Sopic, Amin Aminifar, Amir Aminifar, and David Atienza.

Real-time event-driven classification technique for early detection and prevention of myocardial infarction on wearable systems. *IEEE Transactions on Biomedical Circuits and Systems*, (99):1–11, 2018.

[41] J.T. Bigger G. Breithardt S. Cerutti R.J. Cohen P. Coumel E.L. Fallen H.L. Kennedy R.E. Kleiger F. Lombardi A. Malliani A.J. Moss J.N. Rottman G. Schmidt P.J. Schwartz D.H. Singer A.J. Camm, M. Malik. Heart rate variability : Standards of measurement, physiological interpretation, and clinical use. task force of the european society of cardiology and the north american society of pacing and electrophysiology. *Circulation*, 93(5):1043–1065.

[42] U Rajendra Acharya, K Paul Joseph, Natarajan Kannathal, Choo Min Lim, and Jasjit S Suri. Heart rate variability: a review. *Medical and biological engineering and computing*, 44(12):1031–1051, 2006.

[43] Physionet/computing in cardiology challenge 2011. https://physionet.org/challenge/2011/. [Online; accessed 2019-06-10].

[44] Physionet/computing in cardiology challenge 2014. https://physionet.org/challenge/2014/. [Online; accessed 2019-06-10].

[45] Physionet/computing in cardiology challenge 2017. https://physionet.org/challenge/2017/. [Online; accessed 2019-06-10].

[46] Roberta Colloca. Implementation and testing of atrial fibrillation detectors for a mobile phone application. 2013.

[47] Julien Oster and Gari D Clifford. Impact of the presence of noise on RR interval-based atrial fibrillation detection. *Journal of electrocardiology*, 48(6):947–951, 2015.

[48] Hossein Mamaghanian, Nadia Khaled, David Atienza, and Pierre Vandergheynst. Compressed sensing for real-time energy-efficient ecg compression on wireless body sensor nodes. *IEEE Transactions on Biomedical Engineering*, 58(9):2456–2466, 2011.

[49] Rubén Braojos, Giovanni Ansaloni, and David Atienza. A methodology

for embedded classification of heartbeats using random projections. In *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2013*, pages 899–904. IEEE, 2013.

[50] Giovanni Rovere, Schekeb Fateh, and Luca Benini. A 2.2-μw cognitive always-on wake-up circuit for event-driven duty-cycling of iot sensor nodes. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 8(3):543–554, 2018.

[51] Eamonn Keogh, Selina Chu, David Hart, and Michael Pazzani. Segmenting time series: A survey and novel approach. In *Data mining in time series databases*, pages 1–21. World Scientific, 2004.

[52] Karin Wall and P. E. Danielsson. A fast sequential method for polygonal approximation of digitized curves. *Computer Vision, Graphics, & Image Processing*, 28(2):220–227, 1984.

[53] Marta Karczewicz and Moncef Gabbouj. ECG data compression by spline approximation. *Signal Processing*, 59(1):43–59, 1997.

[54] PhysioNet. Gqrs - wfdb application guide. https://physionet.org/physiotools/wag/gqrs-1.htm. [Online; accessed 2018-03-22].

[55] T. Teijeiro, P. Félix, and J. Presedo. A Noise Robust QRS Delineation Method Based on Path Simplification. *Proceedings of the Computing in Cardiology Conference*, 42:209–212, 2015.

[56] Tomas Teijeiro, Paulo Felix, Jesus Presedo, and Daniel Castro. Heartbeat Classification Using Abstract Features from the Abductive Interpretation of the ECG. *IEEE Journal of Biomedical and Health Informatics*, 22(2):409–420, 2018.

[57] G. B. Moody and R. G. Mark. The impact of the MIT-BIH arrhythmia database. *IEEE Engineering in Medicine and Biology Magazine*, 20(3):45–50, 2001.

[58] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. Ch. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley. Phys-

ioBank, PhysioToolkit, and PhysioNet : Components of a New Research Resource for Complex Physiologic Signals. *Circulation*, 101(23):e215–e220, 2000.

[59] Ti launchpad development kits - overview - ti.com. https://www.ti.com/tools-software/launchpads/overview.html. [Online; accessed 2018-03-04].

[60] About boosterpack plug-in modules - ti.com. https://www.ti.com/tools-software/launchpads/boosterpacks/about.html. [Online; accessed 2018-03-04].

[61] Bluetooth technology website. https://www.bluetooth.com/. [Online; accessed 2019-06-10].

[62] Home page | lora alliance. https://lora-alliance.org/. [Online; accessed 2019-06-10].

[63] Srinivasan Murali, Francisco Rincon, and David Atienza. A wearable device for physical and emotional health monitoring. In *Computing in Cardiology Conference*, pages 121–124. IEEE, 2015.

[64] Online Power Profiler - Nordic DevZone. https://devzone.nordicsemi.com/power/. [Online; accessed 2019-03-17].

[65] Moises Nunez Ochoa, Arturo Guizar, Mickael Maman, and Andrzej Duda. Evaluating lora energy efficiency for adaptive networks: From star to mesh topologies. In *IEEE Wireless and Mobile Computing, Networking and Communications*, pages 1–8, 2017.

[66] Skip a beat heart rate game - play to the beat of your heart rate! http://skipabeatgame.com/. [Online; accessed 2019-06-11].

[67] G.D. Clifford and L. Tarassenko. Quantifying Errors in Spectral Estimates of HRV Due to Beat Replacement and Resampling. *IEEE Transactions on Biomedical Engineering*, 52(4):630–638, 2005.

[68] L. Citi, E.N. Brown, and R. Barbieri. A real-time automated point-

process method for the detection and correction of erroneous and ectopic heartbeats. *IEEE Transactions on Biomedical Engineering*, 59(10):2828–2837, 2012.

[69] M. C. K. Tweedie. Statistical properties of inverse gaussian distributions. i. *The Annals of Mathematical Statistics*, 28(2):367, 1957.

[70] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. Ch. 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.

[71] Terry Young, Mari Palta, Jerome Dempsey, James Skatrud, Steven Weber, and Safwan Badr. The occurrence of sleep-disordered breathing among middle-aged adults. *New England Journal of Medicine*, 328(17):1230–1235, 1993.

[72] Vishesh Kapur, Kingman P Strohl, Susan Redline, Conrad Iber, George O'connor, and Javier Nieto. Underdiagnosis of sleep apnea syndrome in us communities. *Sleep and Breathing*, 6(02):049–054, 2002.

[73] Thomas Gaisl, Daniel J. Bratton, and Malcolm Kohler. The impact of obstructive sleep apnoea on the aorta. *European Respiratory Society Journal*, 46(2):532–544, 2015.

[74] Adaptive servo ventilation. https://www.sleepassociation.org/adaptive-servo-ventilation/. [Online; accessed 2017-02-16].

[75] Redmond B Shouldice, Louise M O'brien, Ciara O'brien, Philip de Chazal, David Gozal, and Conor Heneghan. Detection of obstructive sleep apnea in pediatric subjects using surface lead electrocardiogram features. *Sleep*, 27(4):784–792, 2004.

[76] J Victor Marcos, Roberto Hornero, Ian T Nabney, Daniel Álvarez, and Félix Del Campo. Analysis of nocturnal oxygen saturation recordings using kernel entropy to assist in sleep apnea-hypopnea diagnosis. In *IEEE International Conference of the Engineering in Medicine and Biol-*

ogy Society, pages 1745–1748. IEEE, 2011.

[77] Bijoy Laxmi Koley and Debangshu Dey. Real-time adaptive apnea and hypopnea event detection methodology for portable sleep apnea monitoring devices. *IEEE Transactions on Biomedical Engineering*, 60(12):3354–3363, 2013.

[78] W Ward Flemons, Neil J Douglas, Samuel T Kuna, Daniel O Rodenstein, and John Wheatley. Access to diagnosis and treatment of patients with suspected sleep apnea. *American journal of respiratory and critical care medicine*, 169(6):668–672, 2004.

[79] American Academy of Sleep Medicine, Conrad Iber, et al. *The AASM manual for the scoring of sleep and associated events: rules, terminology and technical specifications*. American Academy of Sleep Medicine, 2007.

[80] Christian Guilleminault, Robert Riley, and Nelson Powell. Obstructive sleep apnea and abnormal cephalometric measurements: implications for treatment. *Chest*, 86(5):793–794, 1984.

[81] N. Oliver and F. Flores-Mangas. HealthGear: a real-time wearable system for monitoring and analyzing physiological signals. In *International Workshop on Wearable and Implantable Body Sensor Networks (BSN'06)*, pages 4–64, 2006.

[82] SleepImage. http://www.sleepimage.com/. [Online; accessed 2016-05-03].

[83] Sleeptracker. http://www.thesleepwatch.com/. [Online; accessed 2016-05-03].

[84] Majdi Bsoul, Hlaing Minn, and Lakshman Tamil. Apnea MedAssist: real-time sleep apnea monitor using single-lead ECG. *IEEE Transactions on Information Technology in Biomedicine*, 15(3):416–427, 2011.

[85] Jessica M Kelly, Robert E Strecker, and Matt T Bianchi. Recent developments in home sleep-monitoring devices. *ISRN neurology*, 2012,

2012.

[86] MR Jarvis and PP Mitra. Apnea patients characterized by 0.02 Hz peak in the multitaper spectrogram of electrocardiogram signals. In *Computing in Cardiology 2000*, pages 769–772. IEEE, 2000.

[87] B Raymond, RM Cayton, RA Bates, and MJ Chappell. Screening for obstructive sleep apnoea based on the electrocardiogram-the Computing in Cardiology challenge. In *Computing in Cardiology 2000*, pages 267–270. IEEE, 2000.

[88] P De Chazal, C Heneghan, E Sheridan, R Reilly, P Nolan, and MO Malley. Automatic classification of sleep apnea epochs using the electrocardiogram. In *Computing in Cardiology 2000*, pages 745–748. IEEE, 2000.

[89] JN McNames and AM Fraser. Obstructive sleep apnea classification based on spectrogram patterns in the electrocardiogram. In *Computing in Cardiology 2000*, pages 749–752. IEEE, 2000.

[90] PK Stein and PP Domitovich. Detecting OSAHS from patterns seen on heart-rate tachograms. In *Computing in Cardiology 2000*, pages 271–274. IEEE, 2000.

[91] JE Mietus, CK Peng, P Ch Ivanov, and Ary Louis Goldberger. Detection of obstructive sleep apnea from cardiac interbeat interval time series. In *Computing in Cardiology 2000*, pages 753–756. IEEE, 2000.

[92] Zvi Shinar, A Baharav, and S Akselrod. Obstructive sleep apnea detection based on electrocardiogram analysis. In *Computing in Cardiology 2000*, pages 757–760. IEEE, 2000.

[93] MJ Drinnan, J Allen, P Langley, and A Murray. Detection of sleep apnoea from frequency analysis of heart rate variability. In *Computing in Cardiology 2000*, pages 259–262. IEEE, 2000.

[94] C Maier, M Bauch, and H Dickhaus. Recognition and quantification of sleep apnea by analysis of heart rate variability parameters. In *Computing in Cardiology 2000*, pages 741–744. IEEE, 2000.

[95] M Schrader, C Zywietz, V Von Einem, B Widiger, and G Joseph. Detection of sleep apnea in single channel ECGs from the PhysioNet data base. In *Computing in Cardiology 2000*, pages 263–266. IEEE, 2000.

[96] F Ng, I Garcia, P Gomis, A La Cruz, G Passariello, and F Mora. Bayesian hierarchical model with wavelet transform coefficients of the ECG in obstructive sleep apnea screening. In *Computing in Cardiology 2000*, pages 275–278. IEEE, 2000.

[97] Philip De Chazal, Conor Heneghan, Elaine Sheridan, Richard Reilly, Philip Nolan, and Mark O'Malley. Automated processing of the single-lead electrocardiogram for the detection of obstructive sleep apnoea. *IEEE Transactions on Biomedical Engineering*, 50(6):686–696, 2003.

[98] André Miguel da Silva Pinho, Nuno Pombo, and Nuno M Garcia. Sleep apnea detection using a feed-forward neural network on ecg signal. In *IEEE e-Health Networking, Applications and Services*, pages 1–6. IEEE, 2016.

[99] Shu-Han Fan, Chia-Ching Chou, Wei-Chen Chen, and Wai-Chi Fang. Real-time obstructive sleep apnea detection from frequency analysis of edr and hrv using lomb periodogram. In *IEEE International Conference of Engineering in Medicine and Biology Society*, pages 5989–5992. IEEE, 2015.

[100] J Dhara Prathap, Ekanath Rangan, and Rahul Krishnan Pathinarupothi. Real-time and offline techniques for identifying obstructive sleep apnea patients. In *IEEE Computational Intelligence and Computing Research*, pages 1–4. IEEE, 2016.

[101] Philip De Chazal, Conor Heneghan, et al. Apparatus for detecting sleep apnea using electrocardiogram signals, 2006. US Patent 7,025,729.

[102] Rahul K Pathinarupothi, R Vinaykumar, Ekanath Rangan, E Gopalakrishnan, and KP Soman. Instantaneous heart rate as a robust feature for sleep apnea severity detection using deep learning. In *IEEE EMBS International Conference on Biomedical & Health Informatics*, pages

293–296. IEEE, 2017.

[103] ECG electrode for sensitive skin | ambu BlueSensor VLC. http://www.ambu.com/corp/products/patient_monitoring_and_ diagnostics/product/ambu%C2%AE_bluesensor_vlc-prod837.aspx. [Online; accessed 2017-02-16].

[104] STM32L151RD - ultra-low-power ARM cortex-m3 MCU with 384 kbytes flash, 32 MHz CPU, USB, 3xop-amp - STMicroelectronics. http: //www.st.com/en/microcontrollers/stm32l151rd.html. [Online; accessed 2017-02-16].

[105] John G Webster. Reducing motion artifacts and interference in biopotential recording. *IEEE Transactions on Biomedical Engineering,* pages 823–826, 1984.

[106] James C Huhta and John G Webster. 60-hz interference in electrocardiography. *IEEE Transactions on Biomedical Engineering,* pages 91–101, 1973.

[107] Yan Sun, Kap Luk Chan, and Shankar Muthu Krishnan. Ecg signal conditioning by morphological filtering. *Computers in biology and medicine,* 32(6):465–479, 2002.

[108] Rubén Braojos, Giovanni Ansaloni, David Atienza, and Francisco J Rincón. Embedded real-time ecg delineation methods: A comparative evaluation. In *IEEE International Conference on Bioinformatics & Bioengineering,* pages 99–104. IEEE, 2012.

[109] Nicolas Boichat, Nadia Khaled, Francisco Rincon, and David Atienza. Wavelet-based ecg delineation on a wearable embedded sensor platform. In *International Workshop on Wearable and Implantable Body Sensor Networks,* pages 256–261. IEEE, 2009.

[110] Grégoire Surrel, Francisco Rincón, Srinivasan Murali, and David Atienza. Low-power wearable system for real-time screening of obstructive sleep apnea. In *IEEE Computer Society Annual Symposium on VLSI,* pages 230–235. IEEE, 2016.

Bibliography

[111] Cc2640r2f (active). http://www.ti.com/product/cc2640r2f/datasheet.

[112] Marco Crepaldi, Chen Li, Keith Dronson, Jorge Fernandes, and Peter Kinget. An ultra-low-power interference-robust ir-uwb transceiver chipset using self-synchronizing ook modulation. In *Solid-State Circuits Conference Digest of Technical Papers, IEEE International*, pages 226–227. IEEE, 2010.

[113] Soumya Subhra Basu, Loris Gérard Duch, Ruben Braojos Lopez, Giovanni Ansaloni, Laura Pozzi, and David Atienza Alonso. An inexact ultra-low power bio-signal processing architecture with lightweight error recovery. In *CODES+ ISSS: International Conference on Hardware/Software Codesign and System Synthesis*, number EPFL-CONF-229969, 2017.

[114] Jiliang Tang, Salem Alelyani, and Huan Liu. Feature selection for classification: A review. *Data Classification: Algorithms and Applications*, page 37, 2014.

[115] T Penzel, J McNames, P De Chazal, B Raymond, A Murray, and G Moody. Systematic comparison of different algorithms for apnoea detection based on electrocardiogram recordings. *Medical and Biological Engineering and Computing*, 40(4):402–407, 2002.

[116] Euan A Ashley and Josef Niebauer. *Cardiology explained*. Remedica, 2004.

[117] Baile Xie and Hlaing Minn. Real-time sleep apnea detection by classifier combination. *IEEE Transactions on Information Technology in Biomedicine*, 16(3):469–477, 2012.

[118] Tin Kam Ho. Random decision forests. In *Document Analysis and Recognition, 1995., Proceedings of the Third International Conference on*, volume 1, pages 278–282. IEEE, 1995.

[119] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.

[120] Michele Rienzner. Find Outliers with Thompson Tau. https: //www.mathworks.com/matlabcentral/fileexchange/27553-find-outliers-with-thompson-tau/content/find_outliers_Thompson.m, 2010. [Online; accessed 2016-04-04].

[121] Charles Anthony Richard Hoare. Algorithm 65: find. *Communications of the ACM*, 4(7):321–322, 1961.

[122] Student. The probable error of a mean. *Biometrika*, pages 1–25, 1908.

[123] BP Welford. Note on a method for calculating corrected sums of squares and products. *Technometrics*, 4(3):419–420, 1962.

[124] Georgios Karakonstantis, Alamelu Sankaranarayanan, Mohamed M Sabry, David Atienza, and Andreas Burg. A quality-scalable and energy-efficient approach for spectral analysis of heart rate variability. In *Design, Automation and Test in Europe Conference and Exhibition (DATE), 2014*, pages 1–6. IEEE, 2014.

[125] Marc-Antoine Parseval. Mémoire sur les séries et sur l'intégration complète d'une équation aux différences partielles linéaires du second ordre, à coefficients constants. *Mém. prés. par divers savants, Acad. des Sciences, Paris,(1)*, 1:638–648, 1806.

[126] Stephen Butterworth. On the theory of filter amplifiers. *Wireless Engineer*, 7(6):536–541, 1930.

[127] Penzel, Thomas, Moody, George B, Mark, Roger G, Goldberger, Ary L, and Peter, J Hermann. Apnea-ECG Database. [Online; accessed 2016-05-02].

[128] GB Moody, RG Mark, A Goldberger, and T Penzel. Stimulating rapid research advances via focused competition: the computers in cardiology challenge 2000. In *Computers in Cardiology 2000*, pages 207–210. IEEE, 2000.

[129] Texas Instruments. Complete low power integrated analog front end for ecg applications. http://www.ti.com/product/ADS1191. [Online;

accessed 2018-02-02].

[130] TDK InvenSense. Mpu-6000 motion sensor. https://store.invensense. com/ProductDetail/MPU6000-InvenSense-Inc/420595/. [Online; accessed 2018-02-02].

[131] nRF8001 Bluetooth low energy Connectivity IC. https://www. nordicsemi.com/eng/Products/Bluetooth-low-energy/nRF8001. [Online; accessed 2019-03-17].

[132] Loris Duch, Soumya Basu, Rubén Braojos, Giovanni Ansaloni, Laura Pozzi, and David Atienza. Heal-wear: An ultra-low power heterogeneous system for bio-signal analysis. *IEEE Transactions on Circuits and Systems*, 64(9):2448–2461, 2017.

[133] Lei Wang, Youfang Lin, and Jing Wang. A RR interval based automated apnea detection approach using residual network. *Computer Methods and Programs in Biomedicine*, 176:93–104, 2019.

[134] ETH Zürich and University of Bologna. PULP-Platform. https://www. pulp-platform.org/. [Online; accessed 2019-04-28].

[135] En Li, Zhi Zhou, and Xu Chen. Edge intelligence: On-demand deep learning model co-inference with device-edge synergy. In *Proceedings of the 2018 Workshop on Mobile Edge Communications*, pages 31–36. ACM, 2018.

[136] Swarnava Dey, Jayeeta Mondal, and Arijit Mukherjee. Offloaded execution of deep learning inference at edge: Challenges and insights. In *IEEE International Conference on Pervasive Computing and Communications Workshops*, pages 855–861. IEEE, 2019.

Grégoire Surrel

Engineer and Doctor in Electrical Engineering & Computer Science

CONTACT INFO

Rue Neuve 3, 1020 Renens, Switzerland
+41 76 604 13 06 | +33 6 28 25 38 89
gregoire@surrel.org

ABOUT ME

Nationality: French
Date of birth: 17 Sept. 1989 (29 y.o.)
Driving license: yes (B)

KEYWORDS

Computer Science	Android App development	Java
Data mining	Data visualization	Web
Electronic design	Research & development	Infosec
Privacy	Methodology	Security

LANGUAGES

- French: native speaker
- English: C2 (Cambridge CAE 2013)
- Spanish: B1 (intermediate)
- Esperanto: A2 (elementary)

HIGHLIGHTS

Award for responsible disclosure of a vulnerability   May 2018

Two EPFL websites presented cross-site scripting (XSS) vulnerabilities. I reported them using the internal bug bounty program, with a minimal proof-of-concept

Published open-source app on the Android Play Store: Notifications pour Facebook 🔗

Status (July 2019): more than 10'000 installs, evaluation of 4.1/5 ★★★★★
The app mirrors the notifications from the site to the phone, either for privacy oriented users or size/performance constraints. Sources on Github 🔗

Data visualization project for the World Terrorism database 🔗   Fall 2017

Data crunching for building an interactive browser-based D3.js visualization about terrorism. The repository with a process book is available on Github 🔗

EXPERIENCE

Doctorate in EPFL, Embedded Systems Laboratory   Oct. 2014
                                                                            5 years

Title: Low Power Sensing and Processing in Wearable Biomedical Devices for Personalized Health Monitoring

Teaching: · Android apps for tablets, smartphones and smartwatches 🔗
            · Internet-of-Things for Industry 🔗

Papers: · Real-Time Probabilistic Heart Beat Classification and Correction for Embedded Systems, CinC, 2015 🔗
            · Design of Ultra-Low-Power Smart Wearable Systems, IEEE LATS, 2015 🔗
            · Low-Power Wearable System for Real-Time Screening of Obstructive Sleep Apnea, IEEE ISVLSI, 2016 🔗
            · Online Obstructive Sleep Apnea Detection on Medical Wearable Sensors, IEEE TbioCaS, 2018 🔗
            · Event-Triggered Sensing for High-Quality and Low-Power Cardiovascular Monitoring Systems, *pending*

Webmaster for Rock Affinity   Sept. 2016
                                                                            3 years

Mission: Install and configure the association's website, with a public part showing relevant information for the visitors, and a private one used for internal documentation. The upcoming events are automatically synchronized with the association's Facebook page

Master project in EPFL, MultiMedia group

Feb. 2014
6 months

Mission: Design the hardware part of a motherboard dedicated to Ultra-HD (4K) video encoding and decoding for live TV broadcast, from building the architecture and choosing the components to doing the hardware layout and routing

Internship at Hilti, Research & Development department

July 2012
13 months

Mission 1: Design an electronic platform for an embedded remote tracking system

Mission 2: Build the Python server back-end and the client front-end for interacting with the data through an intuitive yet powerful web-interface

Presidency at Galiléo, association competing for the Eurobot robotic cup

Sept. 2010
2 years

Mission 1: Lead the association and its members to reach the annual goal

Mission 2: Take part in the development of the robot, both on the hardware and software parts. It involves parts selection, system architecture set-up and development of the robot's strategy

Internship at ST-Microelectronics, quality control operator in production plant

July 2011
2 months

Mission 1: Guarantee the optical quality of the produced silicon wafers in a clean-room

Mission 2: Wrap and label the wafers for shipping

Scientific Education — M.Sc. degree in Electrical Engineering

Master Erasmus exchange at EPFL, Electrical & Electronics Engineering

Sept. 2013
6 months

Courses: Lab on app development for tablets and smartphones 🔗, Analog circuits design 🔗, Digital humanities 🔗, Embedded systems 🔗, Flexible bio-electronics 🔗, Integrated circuits technology 🔗

Semester project: Bio-signals processing using Matlab for estimating muscle fatigue using a wearable EMG-monitoring device

Bachelor & Master CPE Lyon, in Electronics, Computer Science & Telecommunications

2007
5 years

Course blocks:
- Computer Science
- Electronics
- Mathematics, Signal and Images
- Physical sciences

Software

OS:	GNU/Linux, Mac OS X, Microsoft Windows
Development:	git (GitHub username: gsurrel)
Desktop:	Java (Android), C/C++ (Qt framework), Python
Webdev:	PHP, HTML5, CSS3, JavaScript, Leaflet, D3
Science:	Matlab/Octave, Pspice, LabView, ModelSim
Graphics	Inkscape, Gimp, Blender, SolidWorks

Interests

Contributions: Ethical Software, OpenStreetMap, Wikipedia

Sports: Rock-climbing, cycling, dancing, winter sports