CHARACTERIZING THE NOISE ASSOCIATED WITH SENSOR PLACEMENT AND MOTION ARTIFACTS AND OVERCOMING ITS EFFECTS FOR BODY-WORN PHYSIOLOGICAL SENSORS

A Dissertation

by

VISWAM NATHAN

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	Roozbeh Jafari
Committee Members,	Nancy Amato
	Gerard L. Coté
	Ricardo Gutierrez-Osuna
Head of Department,	Dilma Da Silva

May 2018

Major Subject: Computer Engineering

ABSTRACT

Wearable sensors for continuous physiological monitoring have the potential to change the paradigm for healthcare by providing information in scenarios not covered by the existing clinical model. One key challenge for wearable physiological sensors is that their signal-to-noise ratios are low compared to those of their medical grade counterparts in hospitals. Two primary sources of noise are the sensor-skin contact interface and motion artifacts due to the user's daily activities. These are challenging problems because the initial sensor placement by the user may not be ideal, the skin conditions can change over time, and the nature of motion artifacts is not predictable. The objective of this research is twofold. The first is to design sensors with reconfigurable contact to mitigate the effects of misplaced sensors or changing skin conditions. The second is to leverage signal processing techniques for accurate physiological parameter estimation despite the presence of motion artifacts.

In this research, the sensor contact problem was specifically addressed for dry-contact electroencephalography (EEG). The proposed novel extension to a popular existing EEG electrode design enabled reconfigurable contact to adjust to variations in sensor placement and skin conditions over time. Experimental results on human subjects showed that reconfiguration of contact can reduce the noise in collected EEG signals without the need for manual intervention. To address the motion artifact problem, a particle filter based approach was employed to track the heart rate in cardiac signals affected by the movements of the user. The algorithm was tested on cardiac signals from

human subjects running on a treadmill and showed good performance in accurately tracking heart rate. Moreover, the proposed algorithm enables fusion of multiple modalities and is also computationally more efficient compared to other contemporary approaches.

DEDICATION

To my loving and supporting family, especially my parents Raja and Nithya.

ACKNOWLEDGEMENTS

I would like to thank my committee chair, Dr. Jafari, and my committee members, Dr. Amato, Dr. Coté and Dr. Gutierrez-Osuna, for their guidance and support throughout the course of this research.

Thanks also go to my friends and colleagues and the department faculty and staff for making my time at Texas A&M University a great experience.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supervised by a dissertation committee consisting of Professor Jafari, [advisor] and Professors Amato and Gutierrez-Osuna of the Department of Computer Science and Engineering and Professor Coté of the Department of Biomedical Engineering.

All work for the dissertation was completed independently by the student.

Funding Sources

Graduate study was supported by a research assistantship from Texas A&M University. This work was also made possible in part by the National Science Foundation, under grants CNS-1150079 and EEC-1648451, and by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding organizations.

TABLE OF CONTENTS

Page

LIST OF FIGURES

Page

LIST OF TABLES

Page

1. INTRODUCTION

1.1 Wearable Sensors

In recent times, a significant amount of research has been dedicated to realizing pervasive, around-the-clock personal physiological monitoring systems that are both easy to use and accurate [1-6]. The standard practice for healthcare is for individuals to rely primarily on doctor and hospital visits. This model could be significantly augmented with the help of dedicated wearable sensors. The market for wearable medical devices was valued at approximately $3.3 billion in 2015, and is projected to rise to $7.8 billion by 2020 [7].

Apart from generally being more convenient and cheaper than corresponding devices in a hospital, continuously monitoring wearable devices will be invaluable for identifying health-related events for which sporadic monitoring is insufficient. For example, wearable sensors can help manage patients with chronic obstructive pulmonary disease through early detection of 'exacerbation episodes' such as increased dyspnea or cough [8]. The data collected from wearable sensors can help better keep track of illnesses that require longer-term monitoring; for example, monitoring the vital signs of patients with congestive heart failure [8]. The collected data could also allow physicians to better predict outcomes and prevent more serious complications [9]. For instance, in a previous investigation, parameters such as heart rate and respiration rate collected over a number of days were used to build mortality prediction models [10]. Moreover, monitoring outside of clinics could also drive down healthcare costs [11]. This would be of significant benefit particularly in the United States where in 2015, health care

expenditure was about $3.2 trillion (*i.e.* $9,990-per-person), which accounted for 17.8% of the GDP [12].

There are a number of key requirements to maximize the usefulness and effectiveness of wearable sensor systems. Sensors must be robust to changes in placement by the user as well as changes in skin condition over time. The sensors must unobtrusively monitor the required parameters without unduly restricting the user's choice of movements or activities. Finally, the signal processing by the sensor system must be power-aware to enable prolonged use without the need for a battery recharge.

With the growth of the wearable medical device market, one also has to take into account the heterogeneity of the available devices. Heart rate for example could be measured by bio-potential, optical, bio-impedance and even motion sensors as in the case of the seismocardiogram. There could be multiple sensors simultaneously associated with a user. A successful signal processing framework can take advantage of this by not only being capable of working with different sensing modalities, but also through the fusion of multiple sensors.

1.2 Summary of Noise Sources

When it comes to body-worn physiological sensors, there are two key sources of noise:

a) Improper contact with the sensing site on the body

b) Motion artifacts introduced into the signals due to the user's movements

The research presented in this dissertation tackles these two problems in the context of two different physiological signals. For the problem of improper sensor contact, the focus was on dry-contact bio-potential electrodes worn on the head for electroencephalography (EEG). For the problem of overcoming motion artifacts in the signal, the focus was on two signal modalities used to estimate heart rate: the bio-potential based electrocardiogram (ECG) and the optical photoplethysmogram (PPG). The motivation and context for each of these two problems – sensor contact and motion artifacts – are presented in this section.

For both the problems, the proposed solution targets the scenario of continuous, long-term monitoring with wearable sensors outside of a clinical setting as the user goes about their daily lives. The problems of sensor contact and motion artifacts are particularly exacerbated in this scenario due to the need for convenient dry contact and the natural activities of the user, respectively. This is one of the reasons these types of sensors are currently not deemed reliable enough to be medical grade equipment for healthcare [13].

We first describe the sensor contact problem. For wearable sensors outside of a clinical setting, the user is expected to be able to use these sensors independently and will not necessarily have the requisite expertise or inclination to rigorously ensure proper placement of the sensor each time. Furthermore, since these sensors are expected to be worn continuously as the user goes about their daily activities, the original placement as well as the skin conditions at the sensor location can vary over time as well. For bio-

potential sensors, improper placement can mean a higher impedance contact with the skin. This in turn typically leads to an increased amount of noise in the signal.

EEG, a recording of the brain's electrical activity, is measured by bio-potential sensors. When the brain initiates a task, the stimulus to complete the required action is communicated through a network of neurons, which are electrically excitable cells. One neuron can electrically excite its neighbors, thus building a chain of electrical signaling. When the electrical excitation is high enough, a voltage difference called the "action potential" is achieved and this traverses the neurons in the network. The electrodes placed on the head are designed to read this wave of electrical activity going across the scalp. Different tasks produce different patterns in different parts of the scalp and there are certain readily identifiable patterns. For example, upon closing the eyes there is a steady rhythm in the EEG in the frequency range of 8 – 12Hz, known as the alpha rhythm.

EEG is especially sensitive to the quality of sensor contact. Currently the only solution to the sensor contact problem for dry electrodes is to take additional care in the initial placement and hope that the electrodes are robust enough to remain in the same position throughout the necessary period. Continuous monitoring of patients' EEG outside of clinical settings using dry electrodes will be valuable for detecting the onset of medical conditions such as epilepsy, and for facilitating communication with a brain-computer interface (BCI). BCI provides an excellent non-muscular avenue for a user to interface with an external device. A good BCI infers the user's intentions by interpreting the measured brain signals, perhaps elicited as a result of a specifically designed

paradigm, and performs the required action or actuation. This method of communication and control would be a very valuable asset in improving the quality of life of otherwise physically disabled individuals such as patients living with ALS or other paralyzing conditions [14]. This requires the development of a wearable dry-contact EEG system that takes into account not only the signal quality but also the robustness of the system for everyday use.

Wet or gel-based electrodes still represent the standard for medical grade EEG equipment [15]. Several noise sources associated with EEG acquisition have been linked to high impedance contact [16, 17] and so the gel solution in wet electrodes helps to create a suitably low impedance interface with the scalp. The signal-to-noise ratio (SNR) is consequently excellent for wet electrodes. However, most wet electrode based systems require long preparation times; and regular reapplication of the gel is needed to prevent the contact interface from drying out after long time periods. Wet electrode systems also sometimes require scalp abrasion to improve contact, causing scalp irritation and posing an infection risk. For all these reasons, wet electrodes are not typically considered a convenient, viable option for a wearable system capable of long term EEG monitoring [18].

Dry contact electrodes on the other hand, are generally easy for non-experts to put on the user, or for the users to put on themselves independently when capable. They require no special preparation and are suitable for continuous long term use without significant degradation of the quality of the material making contact with the skin.

However the primary drawback compared to wet electrodes is the relatively significant increase in noise pickup.

A typical design for dry electrodes is a number of conductive fingers, arranged in a circle, that are meant to penetrate through the hair and make a resistive contact with the scalp [19-21]. Examples of such electrodes are shown in Figure 1.1. The plurality of fingers is primarily for stability and redundancy of contact. Ultimately all the fingers are electrically shorted together to give one channel of EEG per electrode.

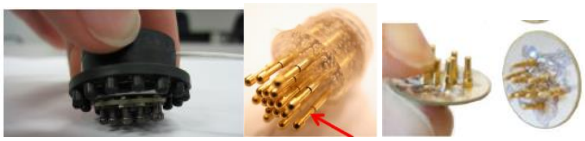

Figure 1.1 - Finger based dry electrodes. Individual images taken from (left to right) [19-21]. Reprinted from Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, Viswam Nathan and Roozbeh Jafari, IEEE Transactions on Biomedical Circuits and Systems, ©2015, IEEE

The other source of noise for wearable sensors is motion artifacts. In the context of bio-potential sensors, motion can result in additional action potentials from the underlying tissue appearing in the target signal. As an example, Figure 1.2 shows two versions of the same electrocardiogram (ECG) segment, one clean and the other corrupted by motion artifacts with 0 dB SNR. The ECG and the motion artifact noise are taken from the MIT-BIH Arrhythmia Database and the MIT-BIH Noise Stress Test

Database respectively [22, 23]. In both figures, the true ECG R-peaks are marked in blue. The figure on the right annotates all the peaks detected by a wavelet-based peak-detection algorithm, including false positives that are marked in black. Thus, we can see how a heart rate estimation algorithm based on peak detection would suffer due to these artifacts. For the optical PPG sensor, motion of the sensor with respect to the skin can result in changes in the optical path of the incident light, or changes in the reception of ambient light, both of which can lead to unpredictable variations in the signal.

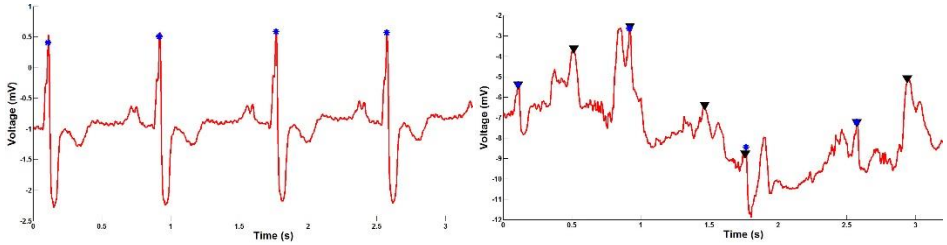

Figure 1.2 - Clean ECG with R-peaks marked (left) and Motion Artifact affected ECG (right). Reprinted from A Particle Filter Framework for the Estimation of Heart Rate from ECG Signals Corrupted by Motion Artifacts, Viswam Nathan, Ilge Akkaya and Roozbeh Jafari, IEEE Engineering in Medicine and Biology Conference, ©2015, IEEE

Among the various physiological measures, there has been significant interest in wearable continuous heart rate (HR) monitoring, and it would be useful for both health and fitness applications. For example, heart rate variability is a condition that could be an indicator of myocardial ischemia [24], and continuous monitoring of the heart rate would increase the chances of a successful diagnosis. Wearable heart rate monitors could

also enable general fitness applications. For example, users would be able to tailor their exercise routines according to how their body is responding to the workload. Thus, this represents an application where overcoming motion artifacts is key to accurately recording an important physiological parameter.

1.3 Research Overview

1.3.1 Research Approach for Sensor Contact Problem

To our knowledge, there has been no previous systematic study looking into the trade-offs involved in choosing the number, and arrangement, of fingers in a given design for dry contact EEG electrodes. The objective for this part of the research was to compare different designs of electrodes with each other to see if there were relative advantages associated with certain designs. One criterion for relative comparison was the impedance of contact under various placement conditions, with the assumption that a lower impedance contact results in a lower noise signal. Another criterion for comparison was SNR for a specific EEG task, which would indicate effectiveness in a BCI. Thus, the first part of the study provides experimental results that shed light on which designs perform relatively better in these tests for robustness and effectiveness.

There has also been no previous study analyzing the signals from each individual finger and the dynamics of the mixed signal obtained when all the fingers are shorted together on a dry EEG electrode. It is reasonable to expect that the different fingers on one electrode have different contact qualities for reasons such as skewed placement of the electrode or the interference of hair or sweat causing local effects on only some

fingers. One question that follows is whether these local noise effects on a subset of fingers can adversely affect the overall EEG signal acquired from the combination of all fingers. If this is indeed the case, a natural next step would be to investigate whether excluding these 'bad contact' fingers from the circuit improves the overall EEG signal. In other words, the objective for this part of the research was to show that the SNR of the EEG signal from a single electrode improves when using only a subset of 'good contact' fingers instead of all the fingers.

This research aims to grant a deeper insight into the electrode skin interface noise for dry electrodes, provide guidance for the improvement of existing finger-based electrode designs and propose a novel extension to the design that enables dynamic reconfiguration of the fingers to continually strive for better SNR under varying scalp conditions and sensor placements.

1.3.2 Research Approach for Motion Artifact Problem

Motion artifacts likely cannot be avoided altogether; it is difficult to imagine a sensor design that is eminently wearable and convenient, and at the same time resistant to motion artifacts. Therefore, we aim to design a signal processing approach that can estimate the target phenomenon despite the presence of motion artifacts.

The primary objective of this second part of the research is to develop a signal processing algorithm to estimate heart rate accurately from a noisy, motion artifact-affected signal, with results validated on ECG and PPG signals.

However, a secondary objective is to develop this algorithm such that it is not wholly tied down to signal features specific to ECG and PPG, and it is also capable of fusing information from multiple signal sources. The reason for this is that wearable sensors to be worn by the user around-the-clock outside of the clinical setting can be in many form factors and utilize other signal modalities (such as bio-impedance, motion sensors or vision sensors).

Finally, an important consideration when designing wearable sensors is the computational power. Due to form factor and battery lifetime constraints, the processors in these systems are relatively modest, and hence it is important to ensure that the computational load due to signal processing algorithms is minimized as much as possible [25]. Optimizing computations to reduce power consumption was not the focus for this research, so the estimation algorithm is not built from the ground-up in a power-aware fashion. However, once the highest level of accuracy for estimation was established, we were interested in tuning parameters of the algorithm to reduce computation time and observing how this affected the accuracy. Moreover, we wanted to ensure the computation time was comparable to contemporary algorithms for the same application.

This research presents the use of a particle filter to estimate heart rate from a variety of noisy signal modalities, as well as the fusion of these. The particle filter is a sequential Monte Carlo routine that probabilistically estimates the true state of a given system by updating the weights and redistributing a set of ‘particles’. More details on the

working of the particle filter will be given in a subsequent section of this document and can also be found in other works [26, 27].

Based on sensor observations, the implementation of the particle filter in this research tracks multiple possibilities for the target parameter and rewards those that are consistent over time. We know that the human heart rate is relatively steady over short time intervals, and so this design is suitable for tracking heart rate accurately. This assumption of consistency and quasi-periodicity applies to other phenomena as well, such as respiration rate or continuous arterial blood pressure (ABP). As such, not only is the proposed implementation independent of specific signal modalities or features, it is also potentially adaptable for other applications in the same domain.

1.3.3 Research Contributions

The contributions of the research are as follows:

1. Novel experiments comparing robustness (consistency of low impedance contact under different conditions) and effectiveness (SNR for EEG BCI task) of different arrangements of fingers on a finger-based dry EEG electrode.
2. Findings on the local noise effects and differences among fingers of a single dry EEG electrode using a custom design that could separately examine signals from each individual finger.
3. Introduced the concept of finger selection and reconfiguration of contact on finger-based dry EEG electrodes to improve SNR. An innovative reconfigurable

electrode was designed and tested on real human subjects to prove the effectiveness of the concept.

4. A novel particle filter formulation for estimation of heart rate from noisy sensor streams without dependence on specific signal features.
5. Demonstration of the accuracy of the proposed heart rate estimation technique using actual motion-artifact affected ECG and PPG data.
6. Proof of the particle filter implementation's potential for fusion of multiple modalities by using simultaneous ECG and PPG streams as well as motion information from an accelerometer to improve heart rate estimates.
7. Demonstration of improved computational efficiency of the particle filter based heart rate estimation solution compared to contemporary related works.

1.4 Organization of this Dissertation

The rest of this dissertation document is organized as follows: Section 2 details the necessary background information for the theoretical concepts and hardware systems used during the course of this research. Section 3 lays out the various published related works to give a sense of the current status of research contributions to the problems being studied and how it contrasts with the contributions of this dissertation. Section 4 describes the research looking into the comparison of different arrangements of fingers on a finger-based dry EEG electrode, whereas Section 5 describes the research investigating the noise effects on individual fingers of a finger-based dry EEG electrode and the possibility of improving the overall signal through finger selection on an

individual electrode. Section 6 describes the research pertaining to the use of a particle filter to estimate heart rate in noisy cardiac signals corrupted by motion artifacts. Finally, Section 7 summarizes and highlights the primary conclusions of this work.

2. PRELIMINARIES

This section expands on some background topics to better contextualize the research described in subsequent sections. We first describe the known theory regarding the noise sources at the skin-electrode interface that cause the issues associated with improper sensor contact. This is then followed by a description of the EEG acquisition system used in this study to collect EEG data in order to validate the various ideas for improving sensor contact for finger-based dry EEG electrodes. Then, the following sub-section presents the known theory surrounding motion artifacts and how they affect signals measured by wearable sensors. This is followed by a brief summary of the human heart rate model and associated assumptions, which encapsulates domain-specific knowledge that will be leveraged in the particle filter algorithm for heart rate estimation. The following sub-section defines the various hardware systems used to collect experimental data for the heart rate estimation study. Finally, the last sub-section provides a high-level overview of particle filters in general.

2.1 Skin Electrode Noise

The circuits and electronics for physiological monitoring have developed to the extent that they are no longer the dominant source of noise in bio-potential measurements using dry electrodes; instead it is the various noise effects at the electrode skin interface [28], especially in the frequencies of interest for physiological applications from 1 to 100Hz [29]. This noise is electrochemical in origin, due to the non-stationary conditions of the ionic skin surface [30]. The amplitude of this noise can be as high as

$60\mu V_{p\text{-}p}$, which is much higher than several target evoked potentials for EEG BCI applications [16]. Moreover, this noise has a frequency characteristic that is anywhere between 1/f and $1/f^2$, meaning it has a larger effect on lower frequencies which are usually the frequencies of interest for physiological signals such as ECG and EEG. Studies have also shown that these noise sources, while not yet fully understood, have a definite correlation with impedance of contact. The half-cell effect, for example, happens when the metal of the electrode comes into contact with an electrolyte such as salt and sweat on the skin, causing an exchange of ions and the development of a potential difference directly proportional to the impedance of contact [16, 17]. This suggests that studying and monitoring the impedance of contact is crucial to the proposed investigations of electrode design and reconfigurable contact. There are other external sources of interference such as capacitively coupled 50/60Hz power line noise [31], but in this study, the electrodes are designed to be buffered and shielded to mitigate external interference. Moreover these are at higher frequencies outside the range of interest and it is assumed they can be digitally filtered out for most applications. In any event, even for external interference, lower impedances of electrodes usually means better impedance matching and hence better rejection of common mode interferences such as 50/60Hz noise.

In summary, the underlying mechanisms of skin electrode noise are not yet fully understood, yet they have a tangible, unpredictable and unavoidable effect on wearable bio-potential sensors. It has been established however, that the impedance of contact is correlated with the noise from the skin-electrode interface.

2.2 EEG Acquisition System

The specifics of the custom electrode designs for the investigations into the effects of the arrangement of fingers will be presented in Section 4, and the designs for analysis of signals from individual fingers and selection of subsets of fingers will be presented in Section 5. In this section we describe the general overall EEG acquisition system that was used to collect data for this study.

The platform used to acquire the EEG signals is a custom PCB designed by our lab that is shown in Figure 2.1. The platform incorporates two daisy-chained TI ADS1299 analog front ends for 16-channel EEG signal acquisition, condition and digitization, the TI MSP430 microcontroller for controlling the communication of the digitized results, and a BlueRadios dual mode Bluetooth radio for wireless communication of the data to a PC. A gain of 1 is used on the ADS1299 differential amplifiers so as not to saturate them. Additional gain is not necessary as the 24-bit ADC has a resolution of roughly 0.4 μV. The board is battery powered and can recharge the battery using a micro-USB interface. The board also incorporates an active driven right leg (DRL) circuit for better common mode rejection based on the design in a previous work [15].

We assembled our own custom dry electrodes which consist of two pieces that snap together as shown in Figure 2.2. The first piece shown to the left in the figure has an LMP7721 low-noise buffer that receives the measured signals through the snap and then transmits them through the wires. Consequently these buffered signals are not affected by motion artifacts or other external noise affecting the cables carrying the

signals to the board. This interfaces with the piece that actually makes contact with a scalp: a circular PCB with multiple spring loaded fingers; the fingers are the same ones used and verified in a previous work [21]. Each of these electrodes is connected to a single input channel of the EEG acquisition board. This overall setup allows dry contact based EEG signal acquisition, whereas custom changes to the electrode design facilitates our research contributions to be described in Sections 4 and 5.

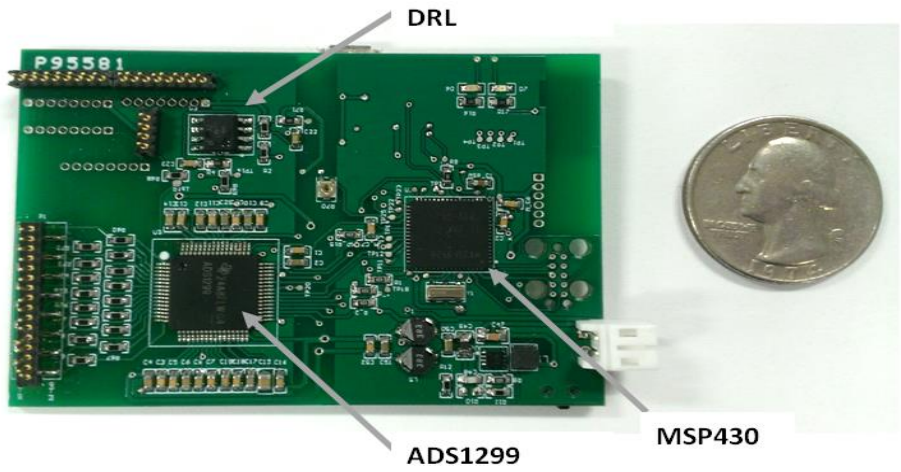

Figure 2.1 - 16-channel EEG platform designed by our lab

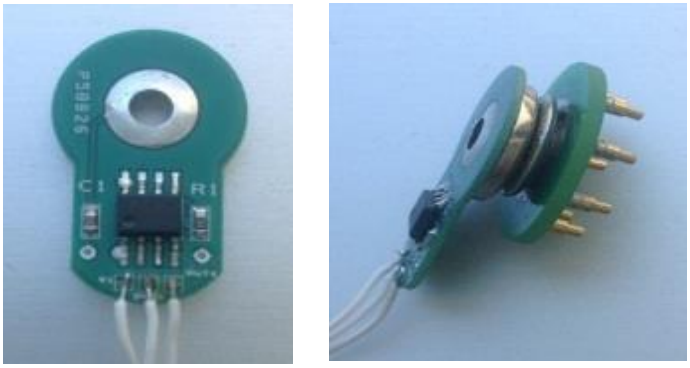

Figure 2.2 – Electrode snap with buffer circuit (left) attached to finger based dry electrode (right)

2.3 Motion Artifacts

Motion artifacts are a well-known and challenging problem for most wearable sensor applications. First, there are the obvious ill-effects of the sensor itself moving without having a stable contact, which is especially the case when there is a need to use dry contact instead of adhesive sensors which are more stable. This is especially problematic for optical sensors, wherein even small movements can result in large and non-deterministic changes in the sensed PPG waveform due to changes in the optical path.

However, even if the sensor is housed in such a way that it is relatively stationary, bio-potential sensors are still susceptible to motion artifacts. This is because motion inherently causes additional unwanted potentials which are picked up by the sensors. Much like the skin electrode artifacts described earlier, there is no scientific consensus in the literature for the true underlying causes of these potentials, but there are some theories. Sweat on the surface of the skin can act as an electrolyte, causing junction

potentials between the electrode and the skin that are sensitive to motion [31]. There is also an inherent stable potential difference between the outer and inner layers of the skin, and any kind of stretching or deformation can change this potential which in turn is picked up by differential voltage sensors [31]. According to one study, this change could be the result of an influence on active metabolic processes, or simply a change in impedance due to the deformation of the skin [32]. This study also showed that even if the skin returns to its stable state and impedance, the changes in potential can continue to last beyond this. Other theories include piezoelectric effects of skin tissue and the varying alignment of Donnan potentials in layers of charged proteins on the skin [33].

Another important point to note is that the frequency spectrum of these motion artifacts overlaps with those of the typical signals of interest, so simple filtering is not a feasible solution. In summary, motion artifacts in wearable systems likely cannot be avoided through hardware design solutions, and sophisticated signal processing is required to monitor physiological parameters during periods of user motion.

2.4 Assumptions on Human Heart Rate

The exact model governing the change in human heart rate over time is complex, and there is still no complete definition or consensus in the literature. There is some evidence that the distribution of beat to beat intervals has multifractal and non-Gaussian properties [34-36]; however, perhaps understandably due to the many factors affecting this physiological process, there is no closed form equation defining the fluctuations in heart rate over time. However, our hypothesis is that the particle filter can provide a

good approximation of this underlying distribution through the observations from the sensors. Moreover, even though we do not know the dynamics of the heart rate exactly, we can assume that the cardiac signals such as the ECG and PPG are quasi-periodic in nature, as was done in a couple of previous works [37, 38].

2.5 Cardiac Signal Acquisition Systems

For the research aiming to estimate heart rate accurately from motion artifact affected cardiac signals, a significant portion of the data is taken from a freely available database. However, we also collected some ECG and PPG signals ourselves, and the systems used to acquire them are described in this section.

2.5.1 ECG Acquisition System

Since the ECG is also a biopotential signal like the EEG, we leveraged the same EEG acquisition platform described in section 2.2. The only change is that we did not use the finger-based contact, which is intended to penetrate through hair on the scalp. Instead, for a dry contact we used the first piece of the electrode alone with the metal of the snap making contact with the skin. When a wet contact was needed, to provide more reliable ECG for reference, we used Arbo pre-gelled electrodes. These electrodes had an adhesive gel underside that made contact with the skin, and incorporated a metal button on the topside that snapped into our electrode just as the finger-based contact for EEG.

2.5.2 PPG Acquisition System

To acquire PPG signals from the wrist while the subject performed activities such as walking or running, we used a system previously designed by our lab called BioWatch. The BioWatch incorporated the TI AFE4400 analog front-end for PPG signal acquisition. The underside of the watch houses the PPG sensor, which includes green LEDs and a photodiode operating in reflective mode. This sensor remains flush against the skin when the watch is worn, and is capable of measuring the PPG accurately when the person is at rest, as shown in some of our previous works using the BioWatch [39-41]. Images of the BioWatch are shown below in Figure 2.3. In addition, the BioWatch also incorporates the MPU-9150, which is a 9-axis inertial sensor. We use the accelerometer data from this sensor as an additional source of information in the proposed particle filtering framework to estimate heart rate.

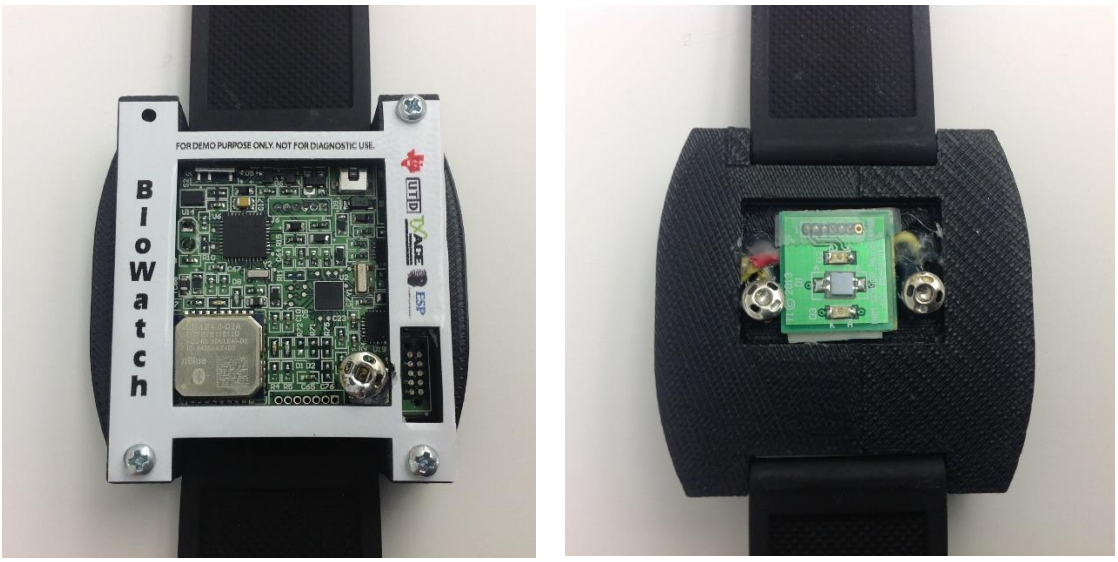

Figure 2.3 - BioWatch top (left) and bottom (right). Reprinted from BioWatch — A wrist watch based signal acquisition system for physiological signals including blood pressure, Simi Susan Thomas *et al*, IEEE Engineering in Medicine and Biology Conference, ©2014, IEEE

2.6 Particle Filtering

Particle Filtering is a tool used to solve the state estimation problem, wherein a certain state changes over time, and the sensor used to measure the state is noisy and imperfect. State estimation problems are widely found in the real world, where the state being estimated can be anything from the position of an object in a room to the price of a stock.

These problems are typically formulated with two models: a system model or state transition model that describes how the true state changes over time, as well as an observation model that describes how the observations from the noisy sensor relate to the state being estimated. If one assumes that these two models are linear functions and

the posterior probability density of the true state is Gaussian, then the Kalman filter provides an optimal solution to the state estimation problem. However, in real-world scenarios often these assumptions do not hold true, due to non-linearities. The Extended Kalman Filter attempts to circumvent this by approximating the non-linearities through local linearization, but this still may not be a suitable and sufficient approximation for certain highly non-linear models.

The particle filter on the other hand, requires no assumptions about the shape of the posterior distribution and can accept arbitrary functions for the state transition and observation models. The algorithm works by instantiating a number of 'particles' in the state space and weighting them such that they altogether represent the posterior probability distribution of the state to be estimated. The weights are updated iteratively according to the observations from the noisy sensor in each time step. It is in a sense emulating a Bayesian filter through Monte Carlo simulations. It has been proven that as the number of particles in the system approaches infinity, the distribution of particles converges to the true posterior probability distribution of the state, or in other words the optimal Bayesian estimate. The particle filter has already been used in a number of different application domains that require non-linear modeling, such as robot localization.

One key oft-cited drawback for particle filters is the computational cost. If the state being estimated is multi-dimensional then the particle filter suffers from the so-called 'curse of dimensionality'. There is always a trade-off between the accuracy and computational cost when choosing the number of particles. This is an especially

important consideration for wearable health monitors wherein processed results may be required in near real-time and the processing itself may need to be done on resource-constrained devices. Therefore this research aimed to leverage the particle filter to handle the non-linearity of the heart rate estimation problem, while remaining mindful of the computational constraints of this application space.

3. RELATED WORK

3.1 EEG Electrode Design

The finger based dry electrode design itself has been presented in a few previous studies [19, 20, 42-45]. In fact, the part used for the fingers and the active electrode design for this work were borrowed from [21]. However, each of these works had a pre-defined, fixed number of fingers in the electrode design, and this number and arrangement varied among the different papers in the literature. There was no real reasoning given for each design decision, whether it be effectiveness in penetration through hair or quality of elicited EEG signals. By contrast, in this work we compare different versions of finger-based dry electrodes to each other and attempt to draw some conclusions to aid design decisions. One previous work that attempted to compare different 3D-printed designs with varying numbers of fingers did so only on a head phantom that does not replicate the effects of hair in a practical scenario, nor allow validation on real EEG signals [46]. For instance, they concluded that the impedance of contact decreased with surface area of contact, but this does not account for the possibility that highly dense designs may not be able to effectively penetrate through the hair in the first place. Another work also developed a titanium nitride (TiN) coated electrodes with differing numbers of fingers [47]. However this work was tested only on a single subject, and the only tangible conclusion was again that the impedance decreases always with an increase in the number of fingers, again not taking into consideration the effectiveness of penetration through hair in a practical scenario.

When it comes to reconfiguration of EEG and BCI systems, there have been previous works that allowed reconfiguration of front-end circuitry such as programmability of the amplifier gain [48] and analog filters [49]. Other works also work on channel selection, which in practice means selecting a subset of the total number of electrodes in order to exclude irrelevant or noisy channels from the targeted analysis, and also potentially allow a reduction in the total number of electrodes necessary in order to improve convenience and comfort. This channel selection can be done through analog spatial filters [50] or regression and classification techniques on the collected signals from all electrodes [51]. One previous work also incorporated impedance measurements similar to what is proposed in this dissertation, but this was again done at a higher level for electrode selection among a number of electrodes [52]. The analysis of individual finger contacts and the idea of reconfigurable contact through selection of a subset of fingers have not been explored before by others. In other words, rather than channel selection, we attempt finger selection within a single electrode, and to our knowledge this has not been previously investigated. However the impedance measurement technique [53] and the assumptions about the associated noise at the electrode-skin interface [15, 31] are gleaned from previous studies.

3.2 Heart Rate Estimation

There have been many proposed approaches in the literature to obtain an accurate heart rate estimate from a noisy cardiac signal. Several works have been based on the use of an adaptive filter, but such techniques always rely on the presence of an external

reference signal, such as accelerometer data [54-56] or electrode tissue impedance [57], which may not always be available. Moreover, different reference signals may be better correlated with different types of motion artifacts and thus a system based on only one reference signal may not represent a generalized solution to handle artifacts from a variety of user actions.

Methods based on a Kalman filter do not rely on an external reference [58] but these techniques assume that the signal and observation models are linear functions and that the noise is Gaussian, which is not always the case for biomedical applications [59]. The extended Kalman filter was introduced to circumvent the disadvantage of the linearity assumption [60], but just like the regular Kalman filter it still suffers from the fact that only unimodal Gaussian distributions can be tracked [59]. In other words, only one possibility for the true state can be tracked at a time and if the estimate diverges from the true state, it may continue to diverge beyond recovery.

Apart from these, there have been a few works that successfully combine several signal processing techniques along with heuristic knowledge of signal characteristics to build a heart rate detection algorithm. There have been three recent related works of note that tackle the problem of heart rate estimation in the presence of extreme motion artifacts when running. The first, dubbed TROIKA [61], involves primarily singular value decomposition, an optimization approach to find a sparse signal representation of the PPG frequency spectrum and finally spectral peak tracking approaches to estimate the heart rate. The second technique developed by the same author, called JOSS [62], has a similar approach except it jointly estimates the spectra of the accelerometer as well

as the PPG and does away with certain steps to save on computation time. The third and final work for PPG signals with motion artifacts [63], which will be referred to as 'Robust EEMD', is based on ensemble empirical mode decomposition (EEMD), followed by a recursive least squares (RLS) adaptive filter using the accelerometer signal as reference. These two techniques are followed by several spectral peak tracking approaches as well as heuristic conditional steps to track the heart rate frequency. As part of the research, we performed a direct comparison of our proposed particle-filter based approach with these three works, both in terms of estimation accuracy as well as computational efficiency.

The particle filter is a probabilistic method that does not depend on any external reference signal nor assume a specific distribution for either the signal or the noise as is the case for the Kalman filter. It is robust and has the potential to recover from incorrect estimates since it can keep track of multiple possibilities. It is generalizable and can be adapted to handle a variety of signal and noise models. It is also straightforward to adjust the number of particles in use, to trade-off between computation time and accuracy depending on the application.

The particle filter has been previously employed in other similar applications, such as identifying the various segments of an ECG in stationary conditions. However, apart from not dealing with motion artifacts, these works usually incorporate a complex dynamical model for the ECG that involves several state dimensions, which in turn increases the computational cost [64-66]. Another work based on an ECG model has a

much reduced dimensionality for the state space; however it is only tested for ECG contaminated by white or pink noise [67].

A particle filter has also been employed for muscle artifact affected ECG de-noising, however this also relies on a sophisticated model that is specific to the progression of ECG with multi-dimensional states and does not seem to be validated on ECG signals with a significant amount of noise [68]. Moreover, in all of the above, the approach that relies on the use of a single rigid and specific mathematical model may not be generalizable to be used for a wider variety of signals from different subjects [69]. The key difference in our proposed framework is that the heart rate itself is directly used as the state to be estimated in the particle filter model equations, and we design the observation densities such that the particle filter simply rewards those observations that are consistent with the expected behavior of a human heart rate. Moreover, the use of only a single state dimension in the formulation greatly eases the computational load compared to previous particle filter implementations in this domain.

4. EEG ELECTRODE DESIGN INVESTIGATION*

In this section is described the experiments to compare different configurations of finger-based dry electrode designs to each other, to see if a certain number or arrangement of fingers translates to better penetration through the hair on the scalp for a lower impedance contact and a better SNR for the sensed EEG signal. After descriptions of the physical characteristics of the different electrodes tested, this section elaborates on the methods used to measure the impedance of contact for each electrode, descriptions of the different *in vivo* experiments, and finally presentation and discussion of the experimental results.

4.1 EEG Electrode Designs

Figure 4.1(a-e) shows the five different electrode configurations evaluated in this work. For all the electrode designs, the distance from the center to the outermost ring of fingers remains constant at 7.21mm. We used spring loaded fingers of height 4.5mm for all the designs. We also compared our designs to the state-of-the-art g.SAHARA dry electrode by g.tec shown in Figure 4.1(f) to ensure that the results from our custom designed electrodes are comparable to that of a commercially available one. The 'Default 8' configuration as shown in the figure matches the g.SAHARA electrode, and represents one of the sparsest arrangements of fingers we have seen among finger-based designs in the literature. Conversely, the 'Circle 20' mimics the densest configurations

seen in works such as those of Matthews *et al* [19] and Liao *et al* [20]. The other configurations were heuristically chosen to cover the spectrum between these two extremes. As mentioned in Section 3, none of the previous investigations with finger-based electrodes explicitly justified the relative sparsity or density of their arrangements. The objective of this part of the research was to directly compare these arrangements with each other, with all other factors such as material, thickness and height of fingers being equal. The criteria for comparison in this research were low impedance contact under varying placement conditions, as well as SNR of EEG in a BCI task.

Our working assumption throughout this study was that the impedance of contact of the dry electrode finger on the scalp is directly proportional to the amount of noise induced on the acquired signal. Most noise sources at the scalp electrode interface, such as the half-cell effect, have been shown to be directly proportional to the impedance of contact by previous studies [31]. Moreover, if the contact with the scalp is disrupted by hair, this would result in a high impedance contact and poor pickup of EEG. Hence, in sections 4.2 and 4.3 are described two different methods to estimate the impedance of contact. These were used in the study evaluating the relative merits of using different numbers of fingers in the electrode design.

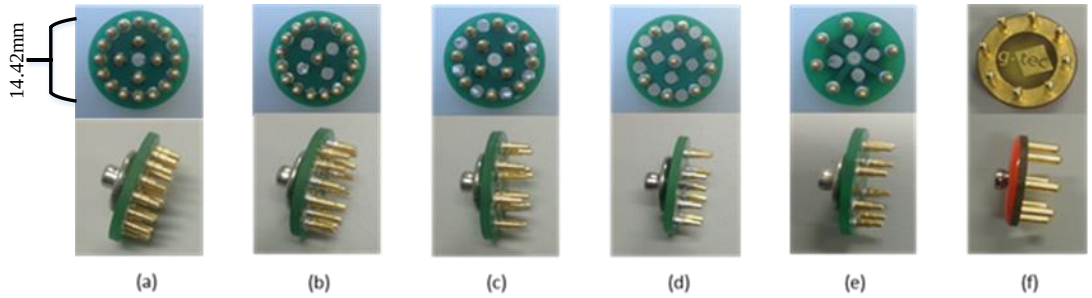

Figure 4.1 – Electrode Configurations. (a) Circle 20: Outer ring of 16 fingers and an inner ring of 4 fingers (b) Circle 17: Outer ring of 16 fingers and one finger at center (c) Spread12: Outer ring of 8 fingers and an inner ring of 4 fingers (d) Center 9: Outer ring of 8 fingers and one finger at center (e) Default 8: Outer ring of 8 fingers (f) g.Sahara: Dry electrode by g.tec[9]. Reprinted from Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, Viswam Nathan and Roozbeh Jafari, IEEE Transactions on Biomedical Circuits and Systems, ©2015, IEEE

4.2 Impedance Measurement

Scalp electrode impedance can be used to compare the effectiveness of different electrode types. A lower impedance contact results in a better signal-to-noise ratio (SNR) and hence more accurate measurements of EEG. Figure 4.2 shows the details of the impedances involved in the contact with the scalp [15, 20]. Z_{ES} denotes the electrode scalp impedance due to the contact of the fingers. Z_S denotes the impedance of the epidermis layer of the skin, whereas Z_D denotes the impedance of the inner dermis layer. The overall impedance faced by the electrode is given by the series combination:

$$Z_{electrode} = Z_{ES} + Z_S + Z_D \quad (4.1)$$

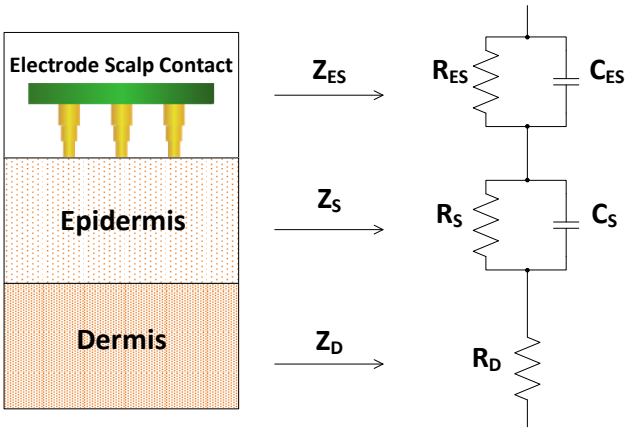

Figure 4.2 - Electrode skin circuit model. Reprinted from Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, Viswam Nathan and Roozbeh Jafari, IEEE Transactions on Biomedical Circuits and Systems, ©2015, IEEE

In this work, we investigate the characteristics of the electrode finger design that minimizes Z_{ES}, thus facilitating a better EEG recording. Since the type of finger used and the overall size of the electrodes remain constant in our designs, we hypothesize that any major differences in the performances across electrode types will be due to the effect of high resistance contact with hair for one or more individual fingers and electrode designs that avoid hair consistently will perform better.

Figure 4.3 shows the overall circuit model when two differential electrodes are used to measure EEG on the scalp. The positive or 'signal' electrode measures V_{sig} with respect to ground, whereas the negative or 'reference' electrode measures V_{ref} with respect to ground. If $Z_{overall}$ denotes the overall impedance between the two electrodes, we have:

$$Z_{overall} = Z_{elec,sig} + Z_L + Z_{elec,ref} \quad (4.2)$$

The first term $Z_{elec,sig}$ gives the impedance faced by the signal electrode, which is the same as $Z_{electrode}$ in Equation (4.1). Similarly, $Z_{elec,ref}$ refers to the impedance faced by the reference electrode. In this work, the type of dry electrode used as the signal electrode is varied, but the reference electrode is always a pre-gelled adhesive patch electrode.

Consequently, in the model for the reference electrode, instead of the finger impedance Z_{ES} we have the impedance of the gel, Z_G. So we can re-write Equation (4.2) as:

$$Z_{overall} = (Z_{ES} + Z_S + Z_D) + Z_L + (Z_G + Z_S + Z_D) \quad (4.3)$$

The additional term Z_L represents the impedance of the length of scalp between the two electrodes, thus completing the circuit. Using an adhesive electrode for the reference is evidently not feasible for a wearable EEG system, however for the purposes of comparing different types of dry electrodes we did not want to introduce the uncertainty of using a dry electrode for the reference as well. We assume all impedances except Z_{ES} described in Equation (4.3) remain constant during experiments for the different dry signal electrodes.

To estimate scalp electrode impedance, we inject current at the signal electrode, shown as I_a in Figure 4.3. When this applied current is a sinusoid of known frequency f_o then the frequency response of V_{out} at f_o is dominated by the voltage drop across the impedance of the circuit due to the injected I_a. The power spectral peak of the signal V_{out} at f_o is hereby termed the 'impedance excitation response', and is directly proportional to the impedance faced by the current I_a.

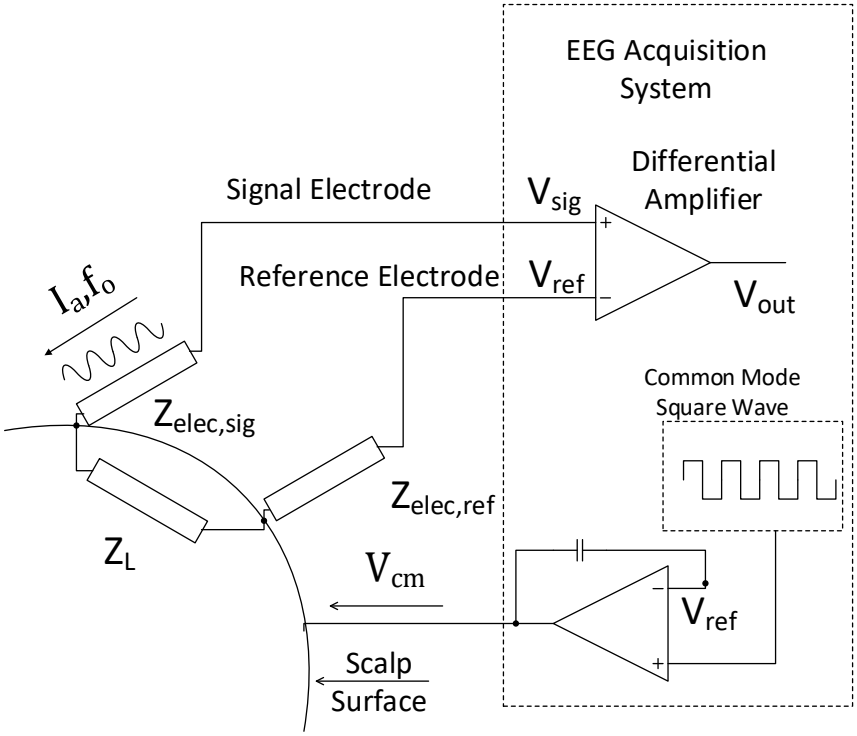

Figure 4.3 - Scalp impedance circuit with signal and reference electrodes. Reprinted from Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, Viswam Nathan and Roozbeh Jafari, IEEE Transactions on Biomedical Circuits and Systems, ©2015, IEEE

4.3 Common Mode Rejection-based Contact Measurement

Another indirect measure of the scalp electrode contact impedance is the common mode rejection ratio (CMRR) of the circuit described in Figure 4.3. CMRR is a measure of how well the differential amplifier can reject signals that are common to both V_{sig} and V_{ref}. The output of the differential amplifier can be written as:

$$V_{out} = G_{diff}\left(V_{sig} - V_{ref}\right) + G_{cm}\frac{\left(V_{sig}+V_{ref}\right)}{2} \tag{4.4}$$

Where G_{diff} and G_{cm} represent the differential gain and common mode gain of the amplifier, respectively. EEG signals of interest are in the *difference* between the two measured signals given by the first term $G_{diff}(V_{sig} - V_{ref})$, and the common mode between the two electrodes is given by the second term and this is typically comprised of uniformly received sources of noise or other undesirable signals.

In ideal situations this common mode is rejected by the amplifier since G_{cm} is very small. In the ADS1299 analog front end used in this work for example, the ratio of G_{diff} to G_{cm} is at least 110dB [70]. However, CMRR decreases when there is an impedance mismatch between the two electrodes. For example, we consider the measurement of steady state visually evoked potentials (SSVEP), the brain's response to being presented with a visual stimulus at a regular frequency. These potentials are strongest at the occipital region at the back of the head [71] and so the signal electrodes are placed here while the reference electrode is placed in a region with suitably lower magnitude of SSVEP such as the right mastoid. EEG signals appearing elsewhere on the scalp, in the frontal region for example, as well as muscle artifacts are assumed to be picked up equally by both the signal and reference electrodes and become a part of the common mode. If the signal electrode faces a different impedance than the reference electrode, then it measures the 'common' signals differently too. This in turn means that the frontal signals and the artifacts are no longer part of the common mode and become part of the differential mode instead. Thus, in the differential signals measured, the frontal signals and artifacts will be mixed in with the desired occipital signals.

In our experiments, one of the two electrodes is always a wet electrode with 'ideal' skin contact. This means that the relatively poor contact of the dry signal electrode results in an impedance mismatch with the reference electrode which in turn causes an increase of the common mode at the final output. In other words, as the contact impedance of the dry electrode gets better, and matches that of the ideal wet contact reference, the CMRR of the circuit improves. In practice, it is not easy to measure CMRR as the common noise entering the circuit at any given time is not easily characterized. However, a signal with known characteristics intentionally added to the common mode can serve as a common mode measurement trace. The CMRR can then be readily estimated by measuring the amount of rejection of this known signal at the output.

For bio-potential measurements, apart from the signal and reference electrodes, a third bias electrode is also attached to the body as part of an active driven right leg (DRL). This bias is applied equally to both the electrodes and hence is also a part of the common mode. In our design, shown in Figure 4.3, we add an AC square wave at a known frequency on top of this DC bias which will serve as the CMRR measurement trace. Any mismatch in contact impedance between electrodes would result in an increased presence of the square wave signal in the output. The power of the known common mode square wave frequency in V_{out} is thus directly proportional to the impedance faced by the signal electrode. This method of contact impedance estimation was previously shown in [53].

4.4 Experimental Setup

In this section we describe the objectives of the various experiments as well as the protocols. Subsection 4.4.1 aims to compare the different electrode types using impedance measurements on the scalp. Subsection 4.4.2 compares those same electrode types with the performance in a BCI task as the criteria. Subsection 4.4.3 repeats the impedance measurements of the electrodes on the hairless forearm to contrast with the effects of hair on scalp.

4.4.1 Impedance Measurement on Scalp

We defined three different use cases:

- *No Adjust*: The EEG cap is put on and no efforts are made to adjust the contact of the electrode
- *Adjust*: After putting on the cap, the electrode is twisted and pushed downwards in an effort to penetrate through the hair and make good contact
- *Headband*: After making efforts to penetrate the electrode through hair, we add a tight headband on top of the electrode to provide an additional downward push for better contact.

The reasoning behind defining these three use cases was to determine whether certain electrode designs would exhibit more robustness in the face of varying capping scenarios. The EEG cap used was a commercial one which was a part of the g.SAHARA system by g.tec. An example of the setup for measurement at scalp location FT8 with and without the headband is shown in Figure 4.4.

Figure 4.4 – Impedance measurement at FT8 without headband (left) and with headband (right). Reprinted from Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, Viswam Nathan and Roozbeh Jafari, IEEE Transactions on Biomedical Circuits and Systems, ©2015, IEEE

One trial of the experiment consisted of about 10 seconds of current injection, with the impedance excitation response collected for each of the 10-second trials under each of the three use cases defined above. Three such trials were conducted for each electrode type in both the temporal FT8 position and the frontal AFZ position according to the 10-20 electrode placement system [72]. The cap was taken off and the hair was readjusted between trials to randomize the contact each time. The impedance excitation response was collected on six human subjects, with varying amounts of hair across subjects. A 24nA sinusoidal AC current at a frequency of 30.5Hz was injected and the peak of the power spectral density (PSD) estimate at that frequency was taken as the

impedance excitation response in units of mV^2/Hz. This in-band frequency is a constraint of the hardware being used; the current injection circuits are internal to the ADS1299 chip and the frequency of current injection is fixed to be in-band.

4.4.2 SSVEP SNR and CMRR Measurement

SSVEPs, introduced in Section 4.3, are a well-known EEG response used in BCI tasks. We wanted to see if certain electrode designs elicited better, *i.e.* higher SNR, SSVEPs on average due to better contact. One session consisted of 4 separate trials of 10 seconds each with the subject fixating on the target flashing LED (18Hz frequency). Three such sessions were collected for each electrode type, with the cap taken off and put back on between sessions to randomize the contact. The data was collected for seven subjects in the 'Headband' case for the best possible contact scenario. The signal electrode was placed at the occipital location OZ, referenced to a wet electrode at the right mastoid. Successfully captured SSVEPs would result in a peak in the PSD of the EEG data at the target frequency. For SSVEP, SNR was defined as the ratio of the target frequency PSD peak to the peaks in the nearby non-target frequencies. The common mode square wave of 6mV amplitude was added at 61Hz throughout the tests and the PSD peak at this frequency was noted as the common mode signal power measured in units of $\mu V^2/Hz$. Since the common mode injection is done by our own custom circuit, the frequency of the signal is controllable and is set to be out-of-band for measurements simultaneously with EEG.

4.4.3 Impedance Measurement on Forearm

One of the major reasons we predicted for some electrode designs performing better than others was the ability to avoid hair. In order to confirm this effect, we designed a control experiment by measuring the impedance excitation response for the various electrode types with the signal electrode on the forearm in an area with little to no hair for seven subjects.

4.5 Experimental Results

4.5.1 Scalp Impedance Measurement

Tables 4.1 and 4.2 show the average impedance excitation responses (defined in Section 4.2) for the different electrode types at the FT8 and AFZ scalp locations respectively. The data is ordered according to the three use cases: 'No Adjust', 'Adjust' and 'Headband' defined in Section 4.4.1. For both the scalp locations, the electrodes with lower density of fingers show better impedances in all use cases.

When comparing the high density electrodes' (Circle 17 and Circle 20) impedance samples with those of the remaining low density electrodes (Default 8, Center 9, Spread 12 and g.SAHARA), the one-sided t-test showed a p-value < 0.004, thus rejecting the null hypothesis that the lower density electrodes show equal or higher impedance. This can be considered a statistically significant result since there are more than 100 samples of impedance for each electrode type when the data from all subjects and capping conditions is aggregated. It can also be observed that the higher density configurations show exceptionally high impedance excitation responses for the 'No

Adjust' case and these responses are drastically reduced for the 'Headband' case. Repeating the above one-sided t-test for only the 'Headband' impedance data shows p-values as high as 0.22. This shows that these high density electrode types depend on effective preparation of scalp electrode contact by adjusting the electrode to penetrate through the hair, after which they could perform similarly to the low-density ones. Conversely, the low density configurations do not show as much variance between the 'No Adjust' and 'Headband' cases which indicates that they are more robust in the face of varying capping conditions.

Table 4.1 - Average impedance excitation response at FT8

	Average Impedance Excitation Response (mV^2/Hz)		
Electrode Type	**No Adjust**	**Adjust**	**Headband**
Default 8	5,637.5	1,591.3	311.3
Center 9	8,663.5	1,923.1	287.0
Spread 12	8,895.9	2,102.5	174.8
Circle 17	21,617.3	7,679.6	470.3
Circle 20	46,327.5	19,853.5	1,029.4
g.SAHARA	9,037.6	3,059.8	511.4

Table 4.2 - Average impedance excitation response at AFZ

	Average Impedance Excitation Response (mV^2/Hz)		
Electrode Type	**No Adjust**	**Adjust**	**Headband**
Default 8	6,812.7	1,753.61	268.2
Center 9	9,822.0	831.4	67.9
Spread 12	6,046.5	738.9	38.3
Circle 17	28,294.4	4,420.1	692.5
Circle 20	68,994.5	8,026.9	376.0
g.SAHARA	10,157.2	3,864.9	150.8

4.5.2 SSVEP SNR and CMRR

Table 4.3 shows the average SSVEP SNR values (defined in Section 4.4.2) across all trials for seven subjects using the different electrode types as well as their respective common mode signal powers (defined in Section 4.3).

Table 4.3 - Average SSVEP SNR and Common Mode Power

Electrode Type	Average SNR	Average Common Mode Signal Power ($\mu V^2/Hz$)
Default 8	4.097	1107.42
Center 9	4.035	818.28
Spread 12	3.927	2906.09
Circle 17	4.240	1637.51
Circle 20	3.781	4298.59
g.SAHARA	4.635	601.05

In most cases the SSVEP SNR of each electrode type is strongly correlated with its corresponding common mode signal power. As noted before, the SSVEP experiments were all conducted under ideal 'Headband' conditions, so the disparity in contact impedance between high and low density electrodes is not too large. The Circle 17 common mode power is suitably low due to this preparation, but the Circle 20 continues to have relatively poorer contact and this in turn adversely affects its SSVEP SNR. The data in general validates the assumption that EEG task performance is inextricably linked to the contact quality, and an electrode is unlikely to show a high performance with poor impedance contact. There was a negative correlation of -0.797 between average SNR of SSVEP and the corresponding common signal power. We also noted a significant performance difference between the Default 8 and g.SAHARA configurations despite their similar structure. The probable reason for this is the increased height and thickness of the fingers on the commercial g.SAHARA compared to our designs.

4.5.3 Forearm Impedance Measurement

Table 4.4 shows the average impedance excitation response collected for each of the different electrode types as they were placed on the forearm of the seven subjects in an area with no hair. The results support the hypothesis that high density configurations suffer in performance primarily due to the effect of hair. Circle 17 and Circle 20 show a marked improvement in impedance compared to the other types.

Table 4.4 - Average Impedance Excitation Response on forearm

Electrode Type	Average Impedance Excitation Response (mV^2/Hz)
Default 8	8285.49
Center 9	8502.23
Spread 12	7485.19
Circle 17	2777.65
Circle 20	3181.43

4.6 Conclusions

Our findings indicate that there are some trade-offs to be considered when designing a dry finger based electrode. The higher density finger configurations showed poor scalp electrode impedances for the different use cases as shown by a one-sided paired t-test with lower density electrode impedances resulting in a p-value < 0.004. This was due to ineffective penetration through hair, as confirmed by the fact that the denser configurations showed better impedance when measured in a region of skin with no hair. However with effective preparation of the contact, SNR performance for the SSVEP EEG task improved. We can conclude that the sparser configurations are more robust to varying conditions, but the advantage is not as significant with proper electrode capping or if the subject has a lower amount of hair. An open research question is the effect of varying the height and thickness of the fingers, as this seems to provide an advantage to the commercial g.SAHARA electrode.

5. RECONFIGURABLE CONTACT ELECTRODES[*]

This section describes the research efforts to build novel EEG electrodes that are capable of active reconfiguration of contact to establish the best possible physical interface with the scalp. This is done through selection of a subset of fingers within each electrode, such that only the fingers with a good contact contribute to the overall EEG signal sensed by that electrode. First, the method used to establish quality of contact is described, followed by physical and functionality descriptions of the novel electrode designs developed to study the signals from each individual finger and actively switch between subsets of fingers respectively. The section then described the experiments conducted to confirm our hypotheses, followed by analysis and discussions of the results, including those of *in vivo* experiments. The objective for this part of the research was to develop a novel hardware-centered approach to solve the problem of improper sensor contact for wearable sensors, and validate the ideas on real human EEG signals. Specifically, the target would be to show that using only a subset of the fingers with a good contact on a dry EEG electrode, as opposed to using all the fingers, results in better correlation with an ideal 'gold-standard' EEG signal from a wet electrode.

[*]© 2015 IEEE. Reprinted, with permission, from Viswam Nathan, Roozbeh Jafari, Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, IEEE Transactions on Biomedical Circuits and Systems, October 2015

5.1 Common Mode-Based Contact Measurement

To estimate contact impedance for this part of the study, we used the same common mode-based contact measurement method described in Section 4.3. Since this method was based on injection of an out-of-band signal, it allowed contact measurement simultaneous with EEG measurement, which was crucial for this part of the study. Direct measurement of the impedance through current injection, as described in Section 4.2, was not an option because the available hardware only allowed current injection at an in-band frequency which would affect the collection of EEG signals.

5.2 Individual Finger Channel (IFC) Electrode

In order to analyze the local noise effects for each finger contact of the EEG electrode, we built a custom electrode that isolated the signals from each finger into separate channels. The electrode built for this purpose (Figure 5.1) has an overall PCB height of 3.76cm, and consists of 8 fingers arranged and spaced in the exact same 'Default 8' configuration described earlier in Section 4.1. However, the difference with this design is that the signal from each finger is immediately buffered before being sent through the cables into separate channels of the ADC on the EEG board. The two ICs that can be seen on the left image of Figure 5.1 are two TI LMP2234 chips, each of which include 4 buffers, thus allowing individual buffering and separation of the 8 fingers on board. Figures 5.2 and 5.3 illustrate the difference between a traditional finger-based electrode and this custom IFC electrode respectively.

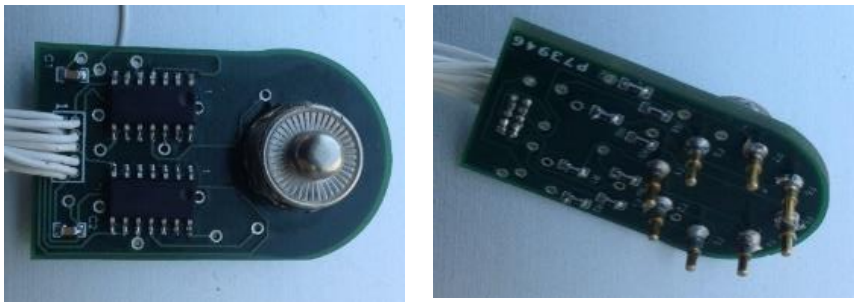

Figure 5.1 – IFC electrode front and back. Reprinted from Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, Viswam Nathan and Roozbeh Jafari, IEEE Transactions on Biomedical Circuits and Systems, ©2015, IEEE

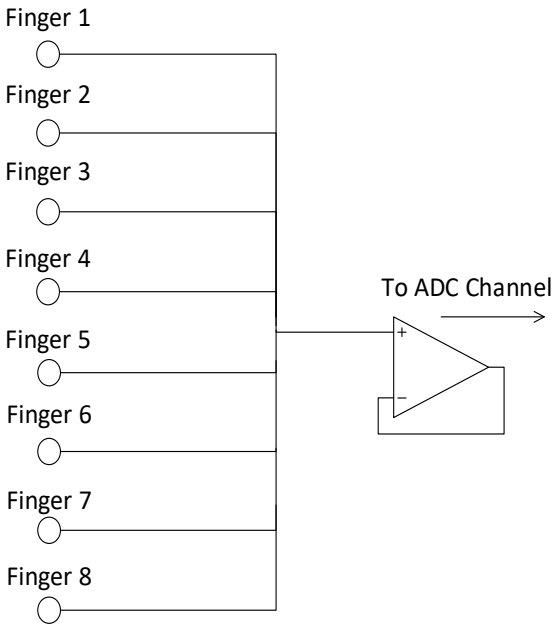

Figure 5.2 – Traditional finger electrode circuit. Reprinted from Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, Viswam Nathan and Roozbeh Jafari, IEEE Transactions on Biomedical Circuits and Systems, ©2015, IEEE

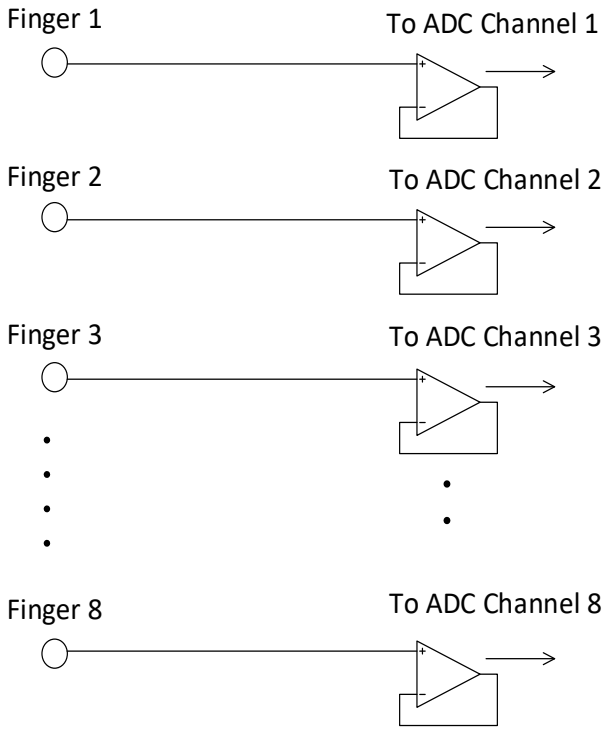

Figure 5.3 – IFC electrode circuit. Reprinted from Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, Viswam Nathan and Roozbeh Jafari, IEEE Transactions on Biomedical Circuits and Systems, ©2015, IEEE

Our hypothesis is that the overall signal from the traditional electrode in Figure 5.4 could be improved by rejecting the individual signals from one or more fingers in the circuit if they are picking up too much noise. Using an estimate of the impedance on each finger, we can generate the mixed signal that would be obtained from any combination of fingers on the IFC electrode. This can be used to test the hypothesis, and the methods to do so are described in detail in Section 5.2.1. It must be stressed that this electrode is not meant to be a practical replacement for existing EEG electrodes, and is

merely intended for exploratory analysis looking into the difference between individual fingers' signals on one electrode.

5.2.1 Exhaustive IFC Electrode Combination Analysis

Since the IFC electrode acquired signals from each individual finger independently and simultaneously, it afforded us the opportunity to uniformly compare the performances of all possible combinations of fingers for the same epoch of EEG. Exhaustively looking at all possible combinations is acceptable in this case since this electrode is meant only for experimental analysis offline and we wanted to ensure we could study the characteristics of the true optimum combination of fingers.

In our experiments, a wet electrode was placed right next to the IFC electrode to provide a ground truth signal. The noise on the wet electrode was assumed to be negligible and it is considered an ideal electrode. This is not strictly true since any electrode will be subject to some amount of noise, but we can safely assume that the amount of noise on the wet electrode is much lower than that on dry contacts and it is the best baseline available for EEG. Therefore, for each set of signals we have:

$$NoiseRMS_i = FingerRMS_i - WetRMS \tag{5.1}$$

Where

i indicates a given finger on the electrode,

$FingerRMS_i$ is the root mean square (RMS) of the signals from Finger i,

$WetRMS$ is the RMS of corresponding signals from the wet electrode and

$NoiseRMS_i$ is the RMS noise magnitude for Finger i.

In the traditional finger electrode, we can model the finger contacts as impedances connected in parallel. Using the contact quality measure, we can obtain the impedance Z_i, and consequently the admittance Y_i, for each finger. Then for any given combination of fingers, we can calculate the contribution of each finger to the resulting parallel circuit. For example, for Fingers 1, 2 and 3 in parallel, the contribution fraction of Finger 1 is given by:

$$C_1 = \frac{Z_2 \| Z_3}{Z_1 + (Z_2 \| Z_3)} = \frac{Y_1}{Y_1 + Y_2 + Y_3} \tag{5.2}$$

In truth, the impedance is composed of both a resistance and capacitance so it would be frequency dependent; but we are only interested in the low frequency region where the electrode-skin interface noise, which is 1/f in nature, would dominate. Therefore, we can ignore the effect of the capacitance. Once we have the contribution fractions of each finger in a given combination, we can compute the overall RMS noise power for that finger combination as:

$$\sqrt{(C_1 \times NoiseRMS_1)^2 + \cdots + (C_n \times NoiseRMS_n)^2} \tag{5.3}$$

where

C_n is the contribution of finger *n* in the current combination and

$NoiseRMS_n$ is the RMS noise magnitude of individual finger *n* defined earlier.

Similarly, we can also estimate the overall EEG signal, as opposed to just the noise, by applying the same contribution fraction on the signals from any combination of fingers:

$$Signal_x = C_1 \times Finger_1 + \cdots + C_n \times Finger_n \tag{5.4}$$

where

$Signal_x$ is the mixed EEG signal for finger combination *x*,

C_n is the contribution of finger *n* in the current combination and

$Finger_n$ is the EEG signal from Finger *n*.

5.3 Multiplexer (MUX) Electrode

In order to validate the analysis from the IFC electrode, we also designed a MUX electrode that can actually switch between using any combination of 8 fingers using the Analog Devices ADG738 IC that includes an 8:1 analog MUX. Figure 5.4 and Figure 5.5 show images of the electrode and the circuit schematic respectively. The overall PCB height is 3.3cm, and again the same Default 8 configuration is used. As opposed to the IFC electrode which was built purely for experimental study, this MUX electrode can be considered a prototype to show the functionality of finger selection added to a typical electrode design.

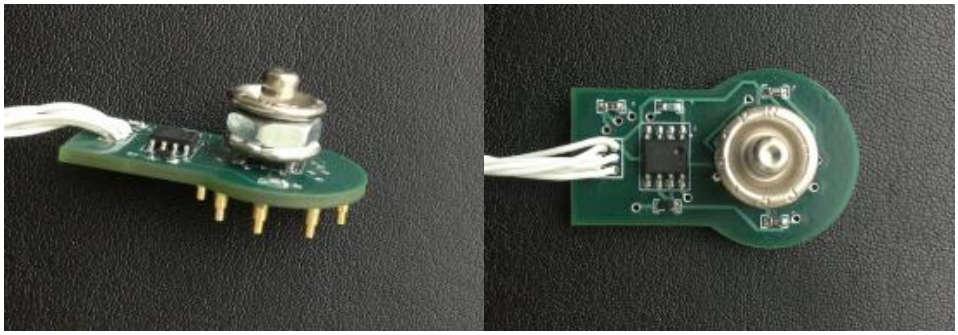

Figure 5.4 – Images of MUX electrode. Reprinted from Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, Viswam Nathan and Roozbeh Jafari, IEEE Transactions on Biomedical Circuits and Systems, ©2015, IEEE

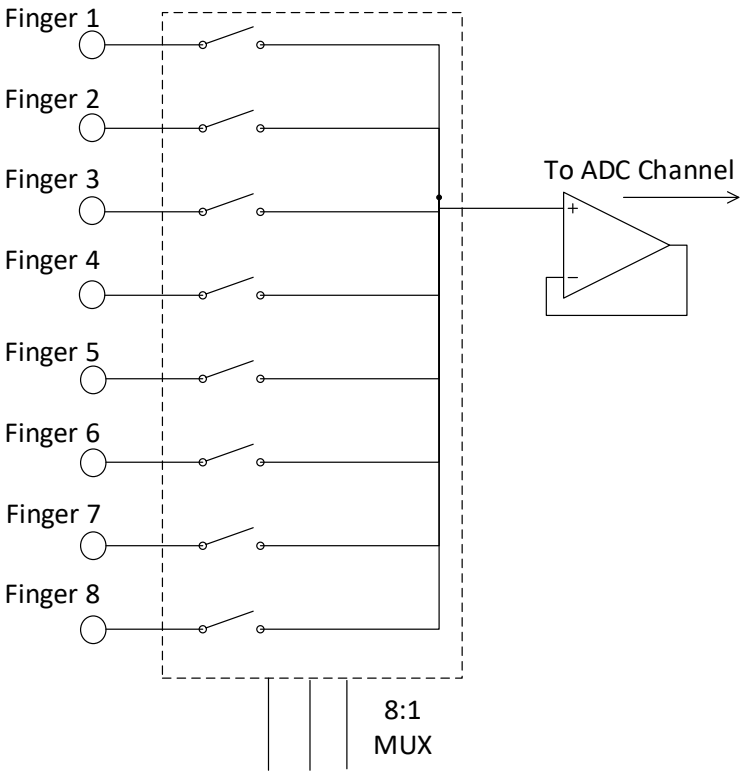

Figure 5.5 – MUX electrode circuit. Reprinted from Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, Viswam Nathan and Roozbeh Jafari, IEEE Transactions on Biomedical Circuits and Systems, ©2015, IEEE

5.3.1 MUX Electrode Finger Selection

A key difference between the IFC electrode and the MUX electrode is that different combinations of the MUX electrode fingers cannot be uniformly compared to each other due to the time delay involved in switching between the different combinations. This is an issue because the EEG signals they are measuring as well as the noise levels are non-stationary. So an exhaustive comparison of all possible finger combinations is not feasible. However, we found through IFC electrode experimental

results, to be shown in Sections 5.5.1 and 5.5.2, that individual fingers that exhibit poor impedance contact are likely to be noisier as well. Therefore, at the start of every MUX electrode experiment we heuristically identified a few good and bad fingers, based on the contact quality measure, and defined combinations that avoided the bad fingers. This process is further detailed in the experimental protocol of Section 5.4.2.

5.4 Experimental Setup

We used the commercial Arbo pre-gelled ECG electrode as the reference electrode in all experiments in this work.

5.4.1 IFC Electrode Experiments

Alpha rhythms are strong EEG waveforms with known characteristics observed when a subject closes his/her eyes [18]. In contrast to SSVEP, this EEG response can be recorded on almost all parts of the scalp. This allowed us to use a wet electrode as an ideal electrode placed on the forehead in a region with no hair to ensure a clean gold standard signal. The IFC electrode is placed right next to the wet electrode in a region with hair for a more realistic, noisy and uneven contact among the different individual fingers. We can safely assume that the EEG alpha signal will be almost identical between the two electrodes at this distance [19], but at the same time the fingers on the IFC electrode will pick up varying amounts of noise due to high impedance contact. Data was collected from 5 subjects with two sessions per subject. Each session involved 20 trials of about 10 seconds of eyes-closed alpha, with the electrode randomly

readjusted between sessions to generate a larger variety of contacts to support the data set.

5.4.2 MUX Electrode Experiments

The MUX electrode experiments were designed in a similar manner to the IFC electrode experiment: 5 subjects performed the alpha task with the MUX electrode placed on the forehead close to a wet electrode for comparison. At the start of every session, on the MUX electrode, one finger was switched on while the rest of them were disconnected in order to measure that individual finger's contact impedance. The process was repeated in turn for each finger until we had the contact impedance of every individual finger. We then heuristically identified 4 or 5 'good' fingers, *i.e.* those fingers that had a relatively low impedance of contact, among the 8 available and then identified three combinations for each subject that involved only some subset of these good fingers. The objective was to see if any of these 'good combinations' performed better than the default case of using all eight fingers. So the experiment involved alternating between one trial of alpha for the good combination and one trial with the 'all fingers' combination. There were 30 trials for each combination leading to a total of 90 trials for each subject. Over this relatively large number of trials, we assume the better combination of fingers would on average show better correlation with the nearby wet electrode.

5.5 Experimental Results

5.5.1 Individual Finger Analysis

In order to test the hypothesis that selecting a subset of fingers on a given electrode could improve the overall signal, we first attempted to show that the signals obtained from each individual finger could indeed be significantly different. We plotted the power spectrum of the signals from each of the fingers on the IFC electrode placed on the scalp as well as the signals from a wet electrode placed nearby for comparison. An illustrative case for one session of alpha on Human Subject 1 is shown in Figure 5.6. Only 7 fingers of the IFC electrode were used since the 8th channel of the EEG acquisition board was reserved for the clean signals from the wet electrode.

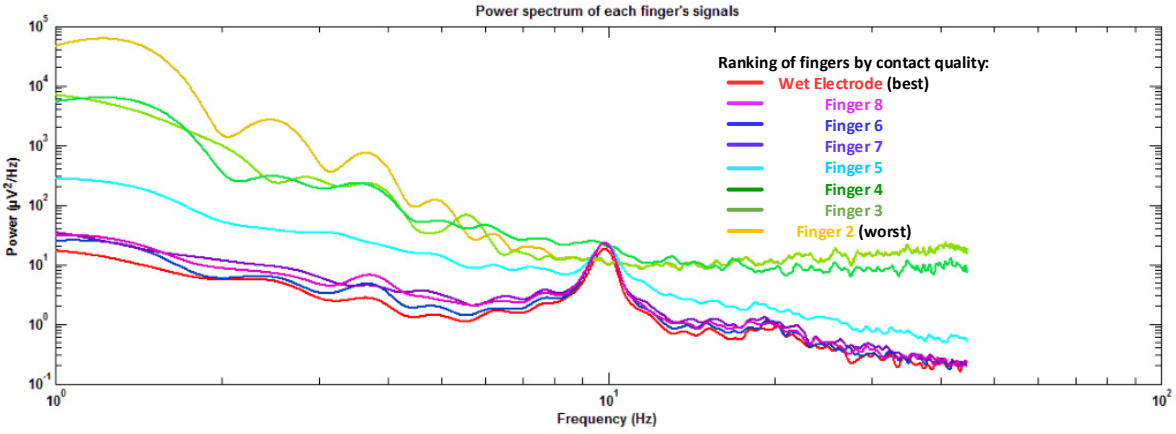

Figure 5.6 – Log scale plot of power spectrum from each of the fingers of IFC electrode and wet electrode. Reprinted from Design Principles and Dynamic Front End Reconfiguration for Low Noise EEG Acquisition with Finger Based Dry Electrodes, Viswam Nathan and Roozbeh Jafari, IEEE Transactions on Biomedical Circuits and Systems, ©2015, IEEE

Evidently there are significant differences in the frequency spectrum, with at least four fingers – Fingers 2, 3, 4 and 5 – showing significantly higher noise levels compared to the others. The wet electrode has the least amount of noise and hence the lowest power as well. The peak at approximately 10Hz corresponds to the EEG response in the state of alpha. Among the three 'good' fingers - Fingers 6, 7 and 8 - the frequency response is mostly overlapped with that of the wet electrode in the higher frequencies, but there is a noticeable separation in the lower frequencies. This confirms that the noise experienced by the dry electrode fingers is 1/f in nature, which agrees with previous findings on skin interface noise [31]. Another observation is that the fingers with higher noise also showed worse contact quality. The example in Figure 5.6 shows in the upper right corner the corresponding rankings for the contact quality of the fingers according to the common mode based impedance measure described in Section 5.1. This confirms that the noise is related to the impedance of the contact and allowed us to use the common mode based contact quality measure as the basis for finger selection.

5.5.2 IFC Electrode Experiments

On the IFC electrode, the mixed signals from all possible combinations of fingers were exhaustively generated using the techniques described in Section 5.2.1. The Noise RMS for each of these combinations of fingers was generated using Equation (5.3). For all subjects and sessions, there were always several combinations better than the 'all fingers' combination in terms of noise. The top four combinations in terms of noise magnitude as well as the combination corresponding to all fingers, averaged across all

sessions and subjects are shown in Table 5.1. Note that 'Best Combination #1' refers to the averaged noise magnitude from the very best combination from each session, but the identity and number of fingers that correspond to the best combination varies for different subjects and sessions due to the varying contact of the electrode.

Table 5.1 - Average RMS Noise magnitude for best finger combinations compared with 'all fingers' combination

Finger Combinations	Average RMS Noise magnitude (μV)
Best Combination #1	2.160
Best Combination #2	2.180
Best Combination #3	2.225
Best Combination #4	2.246
All Fingers	3.635

When comparing the averaged best combination on each subject to the corresponding RMS noise on the 'all fingers' case, the improvement is about 1.5μV. To put this in context, we can assume the wet electrode signal to be ideal and hence the average RMS of this, 13.91μV, estimates the magnitude of the true EEG signal. This means that in our experiments, the noise from the 'all finger' configuration constitutes 26% of the EEG signal, and using the best combination of fingers reduces this noise level by about 40% on average. This is also a statistically significant trend, as proved by the fact that a one-sided paired t-test between the noise from the best combination of fingers and the noise from the 'all fingers' case showed a p-value < 0.05 for 20 trials,

thus invalidating the null hypothesis that choosing fewer fingers does not reduce the noise.

We also generated the overall signal from each of the combinations as described by Equation (5.4). Signals from finger combinations with lower noise are expected to correlate better with the clean wet signal in the time domain. Again, for every session of data there were a few combinations of fingers that showed better correlation than the 'all fingers' case. The averaged correlation coefficients for the best four combinations of fingers, as well as the combination with all fingers, across all subjects and sessions are shown in Table 5.2. On average, the improvement in correlation with the wet electrode for each session was about 12.6% when using the best subset of fingers.

Table 5.2 - Average correlation coefficient with wet electrode for best finger combinations compared with 'all fingers' combination

Finger Combinations	Average Cross Correlation Coefficient
Best Combination #1	0.815
Best Combination #2	0.812
Best Combination #3	0.808
Best Combination #4	0.808
All Fingers	0.724

5.5.3 MUX Electrode Experiments

With the MUX electrode, we can directly obtain signals from different combinations of fingers without the need for mixing based on electrode impedances. For each of the five subjects, at least one of the attempted combinations with fewer fingers showed better performance in terms of average correlation with the wet electrode, as shown in Table 5.3. These were also statistically significant performance improvements over the 15 trials of each session, with p-values < 0.05 in one-sided paired t-tests with the correlation coefficients of the nearest 'all fingers' trials for 4 of the 5 subjects studied. The corresponding p-value for Subject 4 was 0.1 and a possible reason for this was that the correlations with the wet electrode for this subject were relatively low for all trials and combinations.

However it must be noted that these combinations were found somewhat heuristically. In order to more extensively validate performance improvement on the MUX electrode we need to first overcome two challenges. Firstly, there is no a priori guarantee that the combination being selected for comparison is actually better than the 'all fingers' scenario. In other words, we are only looking at the contact qualities of the individual fingers relative to each other. There is no evidence yet to suggest a hard limit on the impedance measure, above which a finger must be considered to have a 'bad contact' and be excluded from any good combination. In fact, even during the MUX electrode experiments of this work, it might well be the case that none of the 'good' combinations we tried were the actual true optimum combination. Exhaustively trying all possible combinations in each new session is not an option; however a machine

learning approach that trains over an extensive data set of different combinations of fingers and contact qualities may be a more feasible approach to predict a good combination. The second challenge is that the strength of the alpha response on a given subject may fluctuate during the course of one experiment. So the noise level of two different configurations cannot be compared uniformly if the desired signal level itself is changing dynamically.

Table 5.3 - Performance comparison between selected MUX combinations vs the 'all fingers' combination for all five subjects

Subject	Average Correlation Coefficient (custom finger combination)	Average Correlation Coefficient (all 8 fingers selected)	p-value from one-sided paired t-test
Subject 1 Combo 1	0.903 (4 out of 8 fingers)	0.877 (all 8 fingers)	0.0297
Subject 2 Combo 1	0.951 (3 out of 8 fingers)	0.935 (all 8 fingers)	0.0086
Subject 2 Combo 2	0.962 (4 out of 8 fingers)	0.939 (all 8 fingers)	0.0151
Subject 3 Combo 1	0.912 (2 out of 8 fingers)	0.882 (all 8 fingers)	0.0425
Subject 4 Combo 2	0.699 (3 out of 8 fingers)	0.662 (all 8 fingers)	0.1258
Subject 5 Combo 1	0.746 (2 out of 8 fingers)	0.687 (all 8 fingers)	0.0458

5.6 Conclusions

In this part of the research, we investigated the possibility of reducing the noise level of EEG signals by reconfiguration of finger-based dry electrodes. A custom electrode was built and served as a test platform to isolate and analyze the signals from each individual finger on an electrode. We observed significant differences in the noise levels across different fingers on the same electrode, and confirmed previous findings on the nature of the skin-electrode interface noise. Mixed signals from different combinations of fingers were generated using this electrode, and we observed that using the signals from only a subset of the available fingers could reduce the amount of noise, by 1.6μV on average, due to the exclusion of noisy fingers. A validation of this conclusion was provided through experimental results that showed that the overall signal from a given electrode can be improved by selecting only a subset of fingers with a good contact on the scalp.

6. PARTICLE FILTERING FOR HEART RATE ESTIMATION*

In this section is described the research looking into using a particle filter based state estimation technique to accurately estimate human heart rate while remaining resistant to motion artifacts in the associated cardiac signals during periods of exercise. The key criterion when evaluating the algorithm was mean absolute error in the estimated heart rate over time. This error was to be minimized while ensuring that the formulation did not rely on signal-specific features and retained the ability to work with multiple signal modalities and their fusion. Moreover, while computation time was not specifically optimized for, we took steps to ensure it remained at least comparable to contemporary algorithms and was not completely compromised in the pursuit of accuracy.

The section begins with the formulation of the estimation problem in the context of the particle filter, and descriptions of the various signal modalities tested, followed by a detailed description of the implemented algorithm used to solve this problem. Subsequently is described a method to fuse the information from multiple signal sources to further improve the estimates, followed by a section describing various additional improvements to the algorithm as a result of specific domain knowledge of this problem. Finally, the section describes the experiments conducted on real human subjects during periods of exercise, followed by analysis of the experimental results, including comparison with contemporary related works that have targeted the same problem. The

objective of this part of the research was to formulate a statistical signal processing framework to solve the problem of estimation in the presence of motion artifacts, which is fundamental to most wearable sensors for daily use.

6.1 Problem Formulation

In order to formulate the state estimation problem for heart rate detection, we first define the state space representation:

$$\mathcal{X}_t \sim \pi_x(\mathcal{X}_t) \qquad \text{(initial distribution)}$$

$$Z_t|\mathcal{X}_t \sim g(\mathcal{X}_t) \qquad \text{(observation density)}$$

$$\mathcal{X}_{t+1}|\mathcal{X}_t \sim f(\mathcal{X}_t) \qquad \text{(transition density)}$$

where $\mathcal{X}_t$ denotes the true system state, *i.e.*, the true heart rate at time t, $\pi_x(\mathcal{X}_t)$ denotes the initial distribution of the system states based on some prior knowledge, Z_t denotes a set of discrete *observations*, $g(\cdot)$ is a function representing the observations conditioned on the true heart rate, and $f(\cdot)$ is the state dynamics or transition model that characterizes the heart rate dynamics as a function of time. In essence, the function $g(\mathcal{X}_t)$ denotes the likelihood of observations given the true state, and the function $f(\mathcal{X}_t)$ describes the progression of the true state due to its own dynamics over time. Instances of these functions for the specific purpose of heart rate estimation will be shown in Section 6.2.

The state estimation problem can be delegated to a particle filter, which is a sequential Monte Carlo method that solves the problem by maintaining a set of weighted

particles, each being a candidate state estimate, its weight being proportional to the likelihood of that particle representing the true state. At each step of the particle filtering problem, the goal is to estimate the posterior state distribution ($p(\mathcal{X}_t|Z_t)$), *i.e.*, the probability distribution of the current true state given a set of observations. This is estimated by the weighted sum:

$$p(\mathcal{X}_t|Z_t) = \sum_{p=1}^{N_p} W_{X_t^p} \delta(\mathcal{X}_t - X_t^p) \tag{6.1}$$

where,

X_t^p is the p^{th} particle at window t,

$W_{X_t^p}$ denotes the weight of particle X_t^p,

N_p is the total number of particles and

$\delta(\cdot)$ is the Dirac delta function, used to place a mass at the particle's location in the posterior probability density function.

Once this posterior probability distribution is updated in each time instance, a suitable method can be used to best estimate the target state at each time. In this work, we use the *maximum a posteriori* (MAP) estimate.

6.1.1 Problem Characteristics

There are a number of reasons why the particle filter is a good fit for the particular problem of heart rate estimation in noisy signals, when compared to other similar techniques. For example, if we consider the problem of heart rate estimation using peak detection on the ECG, a common source of noise is motion that causes spike-

like artifacts. These could lead to false positives for a peak detection algorithm. If we consider a specific instance with a true heart rate of 60bpm, if there is a false positive peak between two true peaks, then the average estimated heart rate for that period becomes, say 120bpm. Thus, it is clear that the noise cannot be modeled as a Gaussian distribution around the true value. When using a peak detection based algorithm on an ECG signal with several peaks that may or may not be true R-peaks, the probability distribution of the heart rate is in fact multi-modal with several distinct possible heart rates in the probability space. This is precisely why, as mentioned earlier, it may be unsuitable to use the Kalman filter which assumes linear Gaussian models, and the Extended Kalman filter that can track only one of these multiple possible modes. Moreover, given this multi-modal probability distribution space where the different modes can be very far apart in terms of heart rate, we decided to minimize the average error by taking the MAP estimate.

The key insight is that the human heart rate is typically a steady, consistent signal over short time windows; in the following sections we will describe the observation mechanisms that allow the particle filter to essentially become a structure that amasses particles in state space regions that show more consistency. In the dimension of time, since the current distribution of the particle filter depends on previous distributions, there is an inherent sense of 'memory' to facilitate rewarding of consistency. In the dimension of state space, N different particles can track N different possible heart rates, thus allowing a parallel search for consistency which reduces the chances of permanently going off track.

Another salient point to note is that the particle filter, as implemented in this work, is decoupled from the signal characteristics. In other words, the particle filter simply receives noisy observations of heart rate, but is agnostic to how these observations were obtained and to what signal modalities and features were used. Thus, the particle filter is not married to the particular observation mechanisms described in this work, and any changes to these mechanisms – for example to add sophistication, or make it more suitable for the given sensor, environment and application scenario – can be easily integrated into the same particle filter framework. More importantly, other arbitrary signal modalities for heart rate detection, such as the ballistocardiogram (BCG), seismocardiogram (SCG), bio-impedance or a camera sensor, could also fit into the same framework and fused with estimates from existing sensors.

6.1.2 Signal Characteristics

Electrocardiogram Signal:

The ECG is a representation of the electrical activity of the heart. In this work we are interested only in the heart rate, and one of the most common ways to estimate the heart rate from ECG is using the R-peaks. The R-peak denotes the point of electrical depolarization of the ventricles of the heart at the start of each beat. The time between successive R-peaks can be used to calculate the beat-to-beat heart rate (Figure 6.1). However, as mentioned before, ‘spike’-like effects caused by motion artifacts could be falsely identified as R-peaks. This naturally leads to overestimation of the heart rate when using peak detection based approaches.

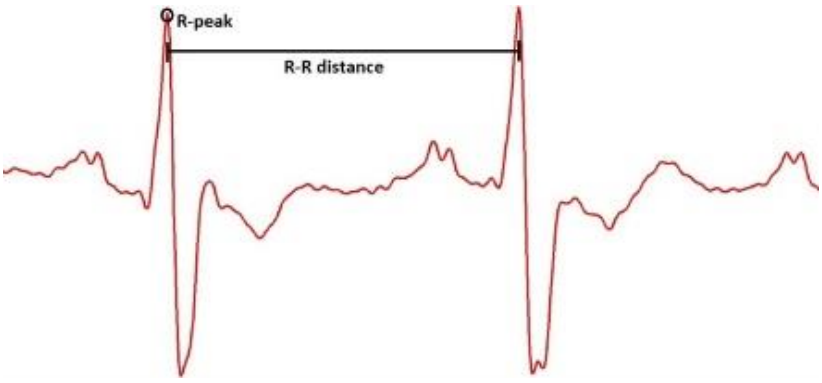

Figure 6.1 - R-peak and R-R interval in an ECG waveform. Reprinted from Particle Filtering and Sensor Fusion for Robust Heart Rate Monitoring using Wearable Sensors, Viswam Nathan and Roozbeh Jafari, IEEE Journal of Biomedical and Health Informatics, ©2017, IEEE

Photoplethysmogram Signal:

The PPG is obtained by transmitting light of suitable wavelength into the skin and measuring the reflected response using a photodiode. Blood flowing in the spot being probed absorbs some of the light and hence affects the optical response of the photodiode. In effect, this means that the periodic flow of blood corresponding to the heart rate will be mirrored in the measured optical response as a pulsatile waveform (Figure 6.2). Unlike ECG, since there is no clear time domain feature, we instead use frequency domain observations for the PPG as will be described in Section 6.2.1.

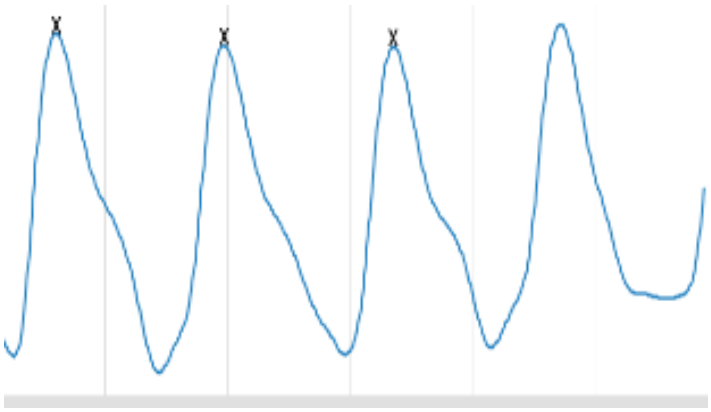

Figure 6.2 – PPG waveforms showing periodic heartbeat. Reprinted from Particle Filtering and Sensor Fusion for Robust Heart Rate Monitoring using Wearable Sensors, Viswam Nathan and Roozbeh Jafari, IEEE Journal of Biomedical and Health Informatics, ©2017, IEEE

Accelerometer Signal

Frequency analysis of the readings from an accelerometer placed close to the heart rate monitoring sensor could prove very useful. It is quite probable that frequencies due to the cadence of walking or running motions would be prominently present in both ECG and PPG signals. Since these frequencies would also be present in the accelerometer's data, we can leverage this to better inform the estimation process. This is particularly important in instances where the motion results in a 'periodic noise' in either the ECG or PPG signal. The particle filter is designed to exploit the assumed quasi-periodicity of the heartbeat and randomness of motion artifacts; thus in the specific instance of noise due to periodic motion, the accelerometer observations can prove critical to distinguish this from periodic heart rate.

6.2 Algorithm Description

6.2.1 Observation Mechanism

In each iteration of the particle filter, the sensors provide multiple observations of the heart rate, and the particles update their weights according to each of these.

Photoplethysmogram Observations

The PPG signal is first bandpass filtered between 0.5 and 15Hz to remove baseline wander and unrelated high frequency noise. Subsequently, we use a spectrogram based approach to get a set of observations of the heart rate. This involves taking moving, overlapping windows of the PPG stream and applying the short-time Fourier transform. The window size was set to be 8 seconds, with an overlap of 2 seconds between successive windows. The frequency spectrum from a window of PPG constitutes one set of observations.

Electrocardiogram Observations

When processing the time domain ECG signal, we use two back-to-back non-overlapping windows dubbed W_{start} and W_{end}. For the purposes of calculating heart rate, we only consider peak-to-peak intervals that begin with a peak in W_{start} and end with a peak in W_{end}. All such intervals taken together constitute a set of observations for a given time window. An example of this windowing approach on an ECG signal is shown in Figure 6.3.

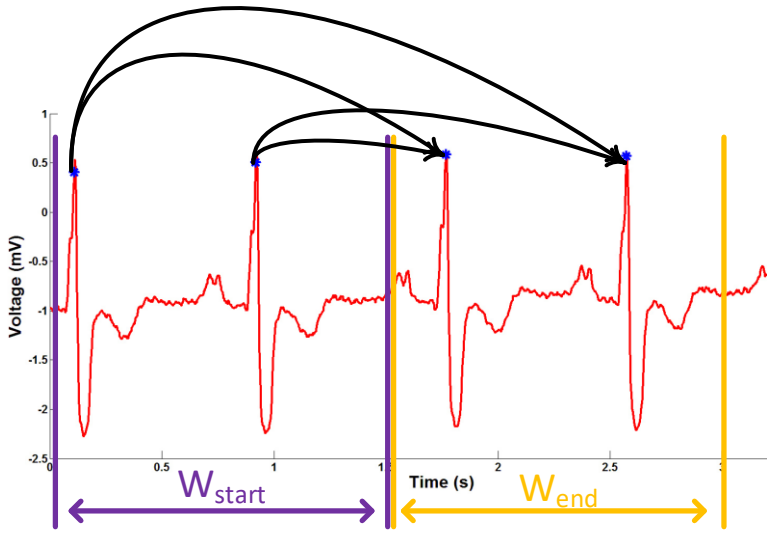

Figure 6.3 - Windowing illustration on ECG signal. Reprinted from Particle Filtering and Sensor Fusion for Robust Heart Rate Monitoring using Wearable Sensors, Viswam Nathan and Roozbeh Jafari, IEEE Journal of Biomedical and Health Informatics, ©2017, IEEE

Note that the peaks that constitute these pairs may or may not be artifacts caused by motion or other noise sources. In Figure 6.3, four different heart rate observations will be considered based on the peak-to-peak pairs shown.

We use windows of size 2 seconds in this work, with a step size of approximately 0.27 seconds. This step size was chosen to accommodate the fact that we expect heart rates as high as 220 beats per minute, and a step size bigger than this could potentially mean skipping true peak-peak observations in those scenarios. The peak detection is then done as follows:

$$findpeaks(W_i, A_{min}, T_{min}) \rightarrow \{P_1, P_2, \dots, P_k\} = P_{W_i} \quad (6.2)$$

$findpeaks(W_i, A_{min}, T_{min})$ finds the time of occurrence of all peaks in the signal W_i that have amplitude at least A_{min} and such that no two peaks are within T_{min} time of each other,

P_{W_i} is the set of peak locations in time, $\{P_1, P_2, \dots, P_k\}$, returned by the 'findpeaks' function. This function is the default implementation found in MATLAB.

This peak detection on its own however is somewhat naïve, so we add an additional step in the procedure for ECG to reduce the number of false positives. We used the continuous wavelet transform (CWT) on the ECG signal with the Mexican Hat wavelet, a center frequency of 0.25Hz and a scale of 5.29 as suggested by a previous work [73]. This helped to accentuate peaks that more closely resemble an R-peak and diminish other trivial peaks. The step described in Equation (2) is then performed on the wavelet transformed signal to obtain the peak locations. It must be noted that this merely reduces the number of false positives but does not eliminate them. Peak-detection based heart rate estimation based solely on the CWT estimate would still overestimate due to false positives, as shown in our previous work [74].

The heart rate observations are then obtained as follows:

$$P_{W_{start}} = \{P_1, P_2, \dots, P_k\} \tag{6.3}$$

$$P_{W_{end}} = \{P_1, P_2, \dots, P_m\} \tag{6.4}$$

$$PP_t \triangleq Set\ of\ all\ (P_b - P_a), \forall P_a \in P_{W_{start}}, P_b \in P_{W_{end}}$$

$$Z_t^n = \left(\frac{fs}{PP_t(n)}\right) \times 60 \tag{6.5}$$

$P_{W_{start}}$ and $P_{W_{end}}$ refer to the sets of peak locations in a starting and ending window respectively,

fs is the sampling rate,

PP_t is the set of peak-to-peak intervals with the first peak in a starting window and second peak in an ending window,

Z_t^n is the n^{th} heart rate observation of window t expressed in beats per minute (bpm).

Taking all Z_t^n in a given window corresponds to the set of observations Z_t referred to in Section 6.1 when describing the particle filter's observation model. It must be noted that we take steps to avoid duplicate observations, *i.e.*, preventing the same two peaks taken as a pair in multiple time windows. We also take steps to ensure any observation included in the set is consistent with other observations of similar heart rate from the same time window already in the set; for example, when we have multiple false peaks we could very well have 4 observations of 50bpm within a 2 second window, but it is clearly impossible in reality for all these observations to be true. So in this example, only those pairs of peaks corresponding to 50bpm that are consistent with each other are taken into the set. This step is necessary to avoid an undue preponderance of lower heart rate observations just because of the nature of our relatively naïve observation mechanism.

Accelerometer Observations

The accelerometer data (denoted as ACC) is processed using the same spectrogram approach used for the PPG signal, with identical windowing procedures. Since there are 3 axes on the accelerometer, we strived to combine them into a single

spectrogram to provide a unified source of observation for the motion noise over time. Using only one axis on its own was not an option because there was no certainty about which axis captured the most activity across the different subjects in the database. This is presumably due to variations in sensor placement and running styles among the different subjects.

We computed the spectrogram for each of the 3 axes, and then stitched together a combined spectrogram that always included only the maximum of the three available powers for each of the frequencies in each time window. This greedy approach allows to always capture the motion frequencies without unduly diminishing their relative power.

6.2.2 Particle Filter Implementation

The initial distribution for the particle filter, $\pi_x(\mathcal{X}_t)$, is defined as follows:

$$\mathcal{X}_t \sim \pi_x(\mathcal{X}_t) = U(HR_{min}, HR_{max}) \tag{6.6}$$

Where $U(\cdot)$ denotes a uniform distribution between HR_{min} and HR_{max}, the assumed lower and upper limits of the heart rate defined by reasonable human physiological bounds. We made the initial distribution uniform since we have no prior knowledge on the initial heart rate, other than extreme limits.

For the PPG, the probability of an observation with respect to a given state of heart rate is computed as follows:

$$\varphi_t^i = S_t^i / \sum_{n=1}^{F} S_t^n, \forall i \in (1, F) \tag{6.7}$$

$$p(Z_t | \mathcal{X}_t) = g(\mathcal{X}_t) = \varphi_t^d \tag{6.8}$$

S_t^i is the i^{th}element of the vector of observed power spectrum amplitudes (measured as described in Section 6.2.1) in time window t for the PPG signal,

F is the total number of frequencies under consideration,

φ_t^i is the i^{th} element of φ_t, the probability density function that results from normalizing the values of the observed power spectrum to be between 0 and 1 in time window t,

d, refers to the frequency in the power spectrum that is closest to the heart rate $\mathcal{X}_t$.

φ_t^d is the probability of the event that the corresponding frequency represents the true heart rate.

This formulation is based on the assumption that a higher power at a given frequency means the more likely it is that that frequency represents the heart rate. However, we know that with motion artifacts there could be high power at certain frequencies as a result of the cadence of motion. This is where the observations from the accelerometer sensor come in; we formulate the accelerometer observation function such that we reduce the likelihood of a given frequency representing the true heart rate if it is present in high power in the accelerometer power spectrum. The formulation is as follows:

$$\tilde{\varphi}_t^i = \sum_{i=i-1}^{i+1} \tilde{S}_t^i / \sum_{n=1}^{F} \tilde{S}_t^n \,,\ \forall i\ \in (1, F) \tag{6.9}$$

$$p(Z_t|\mathcal{X}_t) = g(\mathcal{X}_t) = \left(1 - \tilde{\varphi}_t^d\right) \tag{6.10}$$

$\tilde{S}_t^i$ is the i^{th} element of the vector of the observed power spectrum amplitudes in time window t of the accelerometer spectrogram,

F is the total number of frequencies under consideration,

$\tilde{\varphi}_t^i$ is the i^{th} element of $\tilde{\varphi}_t$, the probability density function for motion noise that results from normalizing the values of the observed accelerometer power spectrum to be between 0 and 1 in time window t,

d, is the index of the power spectrum corresponding to the frequency that most closely matches the heart rate $\mathcal{X}_t$.

$\tilde{\varphi}_t^d$ is the probability of the event that the corresponding frequency is not the heart rate, which for our purposes means it is noise.

For the ECG, in order to create a continuous probability distribution out of the discrete observations, we fit Gaussian distributions around each of the observations resulting in a Gaussian mixture. The probability of a set of observations is then computed as follows:

$$p(Z_t|\mathcal{X}_t) = g(\mathcal{X}_t) = \sum_{n=1}^{O_t} p(Z_t^n|\mathcal{X}_t) = \sum_{n=1}^{O_t} N(Z_t^n, \mathcal{X}_t, \sigma_z) \qquad (6.11)$$

Z_t^n refers to the n^{th} heart rate observation in window t,

O_t is the total number of observations in window t.

$N(Z_t^n, \mathcal{X}_t, \sigma_z)$ denotes a Gaussian distribution with mean equal to the heart rate $\mathcal{X}_t$ in window t, and standard deviation σ_z reflecting the maximum tolerable deviation between the true heart rate and the observation, evaluated at Z_t^n. σ_z is heuristically set to be 3bpm in this work, to ensure that a given particle is reasonably close to an observation to gain weight. Making this parameter too high would mean even unrelated particles gain weight from a given observation, whereas making it too low would too strictly require particles to exactly match the observation to gain weight.

These probabilities guide the update of weights of particles based on closeness to the observations. The particle filter is initialized as follows:

$$X_0^p = U(HR_{min}, HR_{max}) \tag{6.12}$$

$$W_{X_0^p} = \frac{1}{N_p} \tag{6.13}$$

$$\forall p \in (1, N_p)$$

X_0^p is the p^{th} particle sampled from the uniform distribution between HR_{min} and HR_{max}, defined to be 40 and 220 bpm respectively for this work, at time $t = 0$,

$W_{X_0^p}$ is the initial weight of particle p at time $t = 0$.

N_p is the total number of particles, set to be 300 in this work. Choosing the number of particles affects a trade-off between estimation accuracy and computation time, which we will elaborate further on in Section 6.7.2.

After this initialization, with each succeeding time window, the particle weights are updated as shown in (14). Note that we use the so-called 'bootstrap filter' wherein the state transition density is used as the importance distribution, making the weights of the particles directly proportional to the observation density [26]. We chose to do this to simplify the computational load considering the application domain.

$$W_{X_t^p} = p(Z_t|X_t^p) = \begin{cases} \sum_{n=1}^{O_t} N(Z_t^n, X_t^p, \sigma_z), for\ ECG \\ \varphi_t^d, for\ PPG \\ (1 - \tilde{\varphi}_t^d), for\ ACC \end{cases} \tag{6.14}$$

$$\forall p \in (1, N_p)$$

X_t^p is the p^{th} particle of window t,

$W_{X_t^p}$ is the weight of particle X_t^p,

$N(Z_t^n, X_t^p, \sigma_z)$ is the value of a Gaussian distribution with mean X_t^p and standard deviation σ_z evaluated at Z_t^n,

φ_t^d is the probability of the event that the frequency corresponding to X_t^p represents the true heart rate.

$\tilde{\varphi}_t^d$ is the probability of the event that the frequency corresponding to X_t^p is not the heart rate, which for our purposes means it is noise.

The weights are all then normalized to be between 0 and 1:

$$\widehat{W}_{X_t^p} = W_{X_t^p} / \sum_{r=1}^{N_p} W_{X_i^r} \tag{6.15}$$

$$\forall p \in (1, N_p)$$

$\widehat{W}_{X_t^p}$ is the weight of the p^{th} particle of window t normalized so the weights form a probability mass function.

Once the particle weights are calculated the well-known sampling importance resampling (SIR) procedure is employed to prevent particle degeneracy [27]:

$$M_t^p = \sum_{r=1}^{p} \widehat{W}_{X_t^r}, \forall p \in (1, N_p) \tag{6.16}$$

$$u = \underset{a}{argmin} | R_U \sim U(0,1) \leq M_t^a \tag{6.17}$$

$$X'^p_t = X_t^u, \forall p \in (1, N_p) \tag{6.18}$$

M_t^p is the p^{th} element of a cumulative sum vector of the normalized particle weights

X'^p_t is the updated state of the p^{th} particle of window t after resampling, and

R_U is a randomly selected number from the uniform distribution between 0 and 1.

After this step, the distribution of particles approximates the posterior probability distribution of the true heart rate state. To get an estimate for the heart rate in the current time window, as mentioned before, we use the MAP estimate. Since the particle weights are now equalized, we instead look to the distribution of particles to capture the most likely estimate.

We cluster the particles belonging to a similar heart rate together, and can say that the largest cluster represents the most likely state as it is analogous to taking the highest weight particle without the SIR procedure. The clusters and the heart rate estimate are thus calculated as follows:

$$C_n \triangleq Set\ of\ all\ X'^m_t \left| \left| X'^m_t - X'^n_t \right| < CS,\ \forall m \in (1, N_p), \forall n \in (1, N_p) \right.$$

$$E_t = \sum_i C^i_{max} / |C_{max}| \tag{6.19}$$

C_n is the n^{th} cluster of particles

CS is the maximum spread of a cluster (set to be 3 bpm)

C^i_{max} refers to the i^{th} member of the largest cluster C_{max}, and

E_t is the estimate for time window t. For this specific application, the estimate is heart rate in bpm.

The final step in a given iteration of the particle filter is the model-based update that reflects the state transition model defined earlier in Section 6.1. Essentially, as time progress the true heart rate is expected to be dynamic to an extent, and not remain constant. Therefore, the particles are updated accordingly at the end of each time window to approximate this behavior. We assume that the model governing the human heart rate changes over time is a normal distribution:

$$X_{t+1}^{p} \sim f(X_t^p) \sim N(X'^{p}_{i}, \sigma_x) = X'^{p}_{i} + (\sigma_x \times R_N \sim N(0,1)) \quad (6.20)$$

$$\forall p \in (1, N_p)$$

R_N is a randomly generated number from the standard normal distribution, and σ_x is the standard deviation capturing the expected change in heart rate from one window to the next. With the window step size being 2 seconds, σ_x is heuristically set to be 6 bpm.

The window then shifts to a new section of the signal and the particle filter continues to track the heart rate in this manner iteratively over successive windows.

6.2.3 Particle Weighting Assumptions

It can be seen from the formulation for ECG that for a given set of observations in one time window, we consider all the observations as equally likely. We deemed it more generalizable to not rely on any specific features among a set of observations to differentiate them. Instead, we assume that the true heart rate for the subject would make relatively smooth, continuous and gradual changes over time. Leading on from this, we also assume that the observed heart rates as a result of false positive peaks are more random and inconsistent. With these assumptions, our expectation is that even though all observed heart rates are considered equally likely, the particles will build over the correct heart rate as that is observed more consistently over successive time windows.

6.3 Fusion with Multiple Sensors

Since we have formulated the particle filter with only the heart rate as the state to be estimated, multiple signal modalities and their observation mechanisms can be fused in the same framework. The particle weighting for an arbitrary number of observation sources, *i.e.*, sensors, is given by:

$$W_{X_t^p}^{fusion} = \prod_{s=1}^{S} p(Z_t^s | X_t^p) \tag{6.21}$$

$W_{X_t^p}^{fusion}$ is the weight assigned to particle X_t^p when fusing the information from multiple sources of observation

S is the total number of observation sources

Z_t^s is the set of observations in time window t from source s

In essence we assume that since the different sources are observing the same target phenomenon, particles corresponding to states that are observed with higher weight across multiple sources should be rewarded. Conversely, it is unlikely that the same false state would be observed with high probability across multiple sources. In other words, it would be rare for a source of noise to affect sensors with different modalities placed in different locations in the same way.

In this work, for the fusion of ECG, PPG and ACC sensors, the particle weighting process is modified as follows:

$$W_{X_t^p}^{fusion} = p(Z_t^{ECG} | X_t^p) \times p(Z_t^{PPG} | X_t^p) \times p(Z_t^{ACC} | X_t^p) \tag{6.22}$$

$$p(Z_t^{ECG} | X_t^p) = \sum_{n=1}^{O_t} N(Z_t^n, X_t^p, \sigma_z) \tag{6.23}$$

$$p(Z_t^{PPG}|X_t^p) = \varphi_t^d \tag{6.24}$$

$$p(Z_t^{ACC}|X_t^p) = \left(1 - \tilde{\varphi}_t^d\right) \tag{6.25}$$

Similarly, in the database to be described in more detail in Section 6.5.1, there are two separate PPG sensors in addition to the accelerometer in a watch-like device; so the formulation above is modified by simply replacing the observations from ECG with the observations from the second PPG sensor.

With this formulation, we can also get an idea of the contribution of each sensor or signal modality to the overall particle filter estimate in each time window, as shown below:

$$\beta_t^s = \sum_{X_t^p \in C_{max}} p(Z_t^s|X_t^p), \forall s \in S \tag{6.26}$$

β_t^s is the contribution of sensor *s* to the particle filter estimate in time window *t*

X_t^p is a particle in the maximum clique C_{max} for window *t*

Z_t^s is the set of observations in time window t from sensor s

S is the set of all sensors or signal modalities

This contribution can then be normalized with respect to all the sensors in the system and expressed as a percentage:

$$\beta'^s_t = (\beta_t^s / \sum_{m \in S} \beta_t^m) \times 100 \tag{6.27}$$

This can help determine if one or more sensors in the system are noisy. This follows from the fact that sensors producing random, noisy observations will likely have a low contribution to the overall particle filter estimate, thus potentially informing

dynamic adjustments to the contributions of individual sensors based on perceived signal quality in real time. Moreover, prior knowledge of the increased reliability of one sensor could allow increased weightage of observations originating from that sensor. In this initial work however, we keep it simple and do not assume that any one signal sensor is inherently better than the other. The advantage of this overall method of fusion is that it is simple and generalizable and can easily be reused for different applications as well as an arbitrary number of sensors.

6.4 Additional Improvements

While the particle filter framework is complete with the implementation described so far, we found during the course of our experiments with the data that we could make additional improvements to the algorithm to further reduce error for this specific scenario of estimating HR for a running subject. These improvements are detailed in this section.

We assume that the power of a frequency in the accelerometer spectrum is directly proportional to the probability of that frequency representing motion noise. However, in a few subjects' data there was a harmonic of the movement frequency that was somewhat low in strength but still high enough to mislead the particle filter. Therefore, we modified the ACC probability function to remove from consideration an observation if the power of the corresponding frequency was greater than 10% of the maximum power observed in the accelerometer for that time window.

There were a few instances wherein the dominant ACC frequency happened to overlap with the true heart rate frequency. This would be especially problematic with the hard thresholding introduced above. Therefore, we implemented a rough frequency margin around the expected heart rate, and if the ACC frequency under consideration was within this zone, we did not perform the thresholding. This ensured that we did not effectively remove from consideration particles corresponding to the heart rate simply because the ACC frequency was close. The bounds for the margin were set by taking the average of the previous 3 heart rate estimates in Hz and making a conservative bound of +/- 0.1 Hz. This corresponds to an assumption that the heart rate would not change by more than 6 bpm in either direction from one time window to the next.

In all of the data, the subject starts at rest at least for a few seconds before beginning any activities. It makes little sense to include the accelerometer observations in these states. Therefore, we first find the magnitude of acceleration in each time window as follows:

$$\tau = \sqrt{(a_x(t))^2 + (a_y(t))^2 + (a_z(t))^2} \tag{6.28}$$

$a_x(t)$ is the x-axis acceleration for the given time window

$a_y(t)$ is the y-axis acceleration for the given time window

$a_z(t)$ is the z-axis acceleration for the given time window

The observations of the accelerometer are taken into consideration for the final heart rate estimate only if this magnitude was above a certain threshold. The threshold was heuristically determined to be 1.04 g by examining the data. This parameter has to

be heuristically set this way in the absence of more sophisticated activity detection algorithms.

When the subject transitions from a resting state to walking or running, there is naturally a sudden increase in heart rate in response to the increased workload. In most cases the particle filter tracked these increases as well, but there were a few instances where it was slow to catch up simply because there happened to be continued observations corresponding to the slower resting state heart rate which turned out to be false observations. In these situations, there is an error in heart rate for a few time windows because the particle filter already had a preponderance of particles around the resting heart rate and continued to see observations consistent with that heart rate. So in a sense there may be some 'latency' for the particle filter estimates to catch up to the true heart rate when there is an abrupt change in the dynamics.

Therefore we wanted to introduce the notion of context-awareness and have multiple operating modes for the particle filter. When we have the accelerometer, we have an independent source of information that provides additional context for the user's current state. When the subject is at rest or running steadily, we do not expect rapid changes in heart rate and so the particle filter model state update (described in equation (20)) will be conservative. Conversely, when the subject's activity level increases rapidly, we can accordingly adjust the model for state update to temporarily allow for greater changes. A similar idea for this adaptive changing of model equations based on the current context has been previously implemented in other application areas [75].

For our problem, when the subject was previously at rest (as determined by the magnitude threshold) and the ACC magnitude from (28) changes by a significant margin from one time window to the next, we can assume that increased activity has begun. Then for the next few time windows, instead of using the state update equation described in Equation (6.20), we use the following:

$$X_{t+1}^{p} \sim f(X_t^p) \sim N(X'^{p}_{i}, \sigma_x) = X'^{p}_{i} + (R_N \sim N(\alpha, \sigma_x)) \quad (6.29)$$

$$\forall p \in (1, N_p)$$

R_N is a randomly generated number from the normal distribution with mean α and standard deviation σ_x

α is a positive bias meant to indicate that on average, the heart rate is expected to increase. It is set to 6bpm in this work.

σ_x is the standard deviation capturing the change in heart rate from one window to the next. It is set to be 10bpm, a larger number to reflect the possible rapid changes in heart rate.

The threshold for required change in acceleration magnitude is set to be 0.04 g, and when such a change occurs the alternate update equation (6.29) is used for a period of 5 time windows. Note that we do not assume the heart rate definitely must increase whenever higher ACC activity is detected, as that is placing too much trust in a rudimentary activity detection approach. Rather, we simply allow the particles to be more spread out than usual for a few time windows when we detect a possible sign of volatility in the heart rate. In other words, there will be more particles than usual in the

higher heart rate regions in anticipation of a sudden increase, but there will continue to be particles corresponding to the previous steady state, lower heart rates, and everything in between. This is one of the core advantages of the particle filter, wherein particles can track multiple possible states in parallel. With this context-aware mode-switching approach, the instances of particle filter estimate latency due to sharp heart rate changes was reduced. We did not observe this latency effect for longer than 3 time windows or 6 seconds across all subjects tested in this work, and there was no latency at all for many subjects.

6.5 Experimental Setup

6.5.1 Database Description

Motion artifact affected PPG data was taken from the database used as part of the 2015 IEEE Signal Processing Cup (SP Cup) [61]. This data was recorded at a sampling rate of 125Hz using a wrist-worn dual PPG sensor (*i.e.*, two simultaneous channels of PPG) from 12 subjects. The sensor also included a 3-axis accelerometer. Each trial for a subject consisted of 30 seconds of resting, followed by four stages of activity each for 1 minute, and finally 30 seconds of rest again. The four periods of activity consisted of alternating between relatively slower (6km/h or 8km/h) and faster (12km/h or 15km/h) treadmill speeds. ECG was also simultaneously recorded from the chest using wet sensors, and this is used to obtain the ground truth heart rate. Since the three related works mentioned earlier – TROIKA, JOSS and Robust EEMD – also all worked on the same dataset, we can directly compare the average errors in heart rate estimation.

At this juncture, we note that the JOSS work resampled the data to 25Hz (presumably to ease the computational load) and also truncated the data for 6 of the 12 subjects; this is because that algorithm is entirely dependent on a clean start for the tracking, and half the subjects had signals with some noise to varying degrees even at the initial resting stage. Therefore, to compare with JOSS we also perform the resampling as well as the truncations described in that paper [62]. However, one of the key advantages of the particle filter is the increased ability to recover from going off-track due to the presence of multiple different particles in the state space. So we will also present the results from the un-truncated datasets and show that the particle filter effectively recovers from these 'false starts'.

6.5.2 ECG Database with simulated noise

As the dataset described in the previous section had only clean ECG, we wanted to find another solution to obtain motion artifact affected ECG to test the particle filter on the estimation of heart rate from the fusion of simultaneous noisy ECG and noisy PPG. To the best of our knowledge there was no existing database that provided simultaneously recorded ECG and PPG data that were affected by real motion artifacts and also had the ground truth heart rate available.

Therefore, we turned to the MIT-BIH Noise Stress Test Database to get real motion artifact noise and add it to the existing clean ECG signals from the aforementioned Signal Processing Cup database [22, 23]. The MIT-BIH database, including the techniques to synthetically introduce realistic motion artifact noise, is well

respected and has been used in several previous works. The owners of the database themselves provided a technique to add calibrated amounts of motion noise data to any given ECG record from their own database such that the desired SNR level is obtained. We have simply adapted this approach to inject the motion artifact noise into the ECG data from the SP Cup database.

In order to test the fusion approach, we used the ECG data injected with motion artifacts in conjunction with the PPG data that is already present in the same database with real artifacts due to the running activity. The particle filter estimates from these fused observations are compared to the heart rate from the unaffected clean ECG data in the database. We ensure that the added motion artifact noise for ECG is proportionally increased in intensity as and when the running speed increases in a given data record. We chose SNR levels of 3dB and -3dB respectively for the slower and faster speeds.

Of course, it would be unfair to compare the results from this fusion with the other related works, since they did not have the noisy ECG; however, we believe the results on their own will be illustrative of the ability of the particle filter to effectively fuse information from multiple modalities. We will compare the performance of using just the individual modalities vs. using the fusion of them to illustrate this.

6.5.3 Experimental Data Collection

Even though we believe the methodology of adding noise to the ECG described in the previous section is sound, we readily concede that the ideal scenario would be to have simultaneously collected ECG and PPG data that were both affected by real motion

artifacts during the course of the data collection. Since such a database is lacking in the literature to our knowledge, we conducted a limited data collection of our own to bolster our experimental conclusions. We used a previously developed system called the BioWatch [41] that collected one channel of PPG signals from the wrist and also included an accelerometer. For the ECG, we used a custom platform based on the TI ADS1299, an analog front end for bio-potential signals. One channel of ECG was recorded from the chest using adhesive gel electrodes in a Lead II configuration to be used as the ground truth. A second channel was recorded using a dry electrode that was secured to the forearm just above the BioWatch using medical tape. This was meant to provide an ECG signal that was more susceptible to motion artifacts. Both devices sampled data at a rate of 125Hz and transmitted data to a PC wirelessly using Bluetooth. Data was collected from 5 subjects running on a treadmill after informed consent and protocol approval by the IRB at Texas A&M University (IRB2016-0193D). The experimental protocol was designed to be similar to that of the database described earlier: 30 seconds of rest, followed by four 1-minute periods alternating between walking and running, and 30 seconds of rest at the end. Examples of the signals from our system after pre-processing (0.5 to 15Hz bandpass for PPG and 0.5 to 30Hz bandpass for ECG), for both standing and running scenarios are shown in Figures 6.4 and 6.5 respectively.

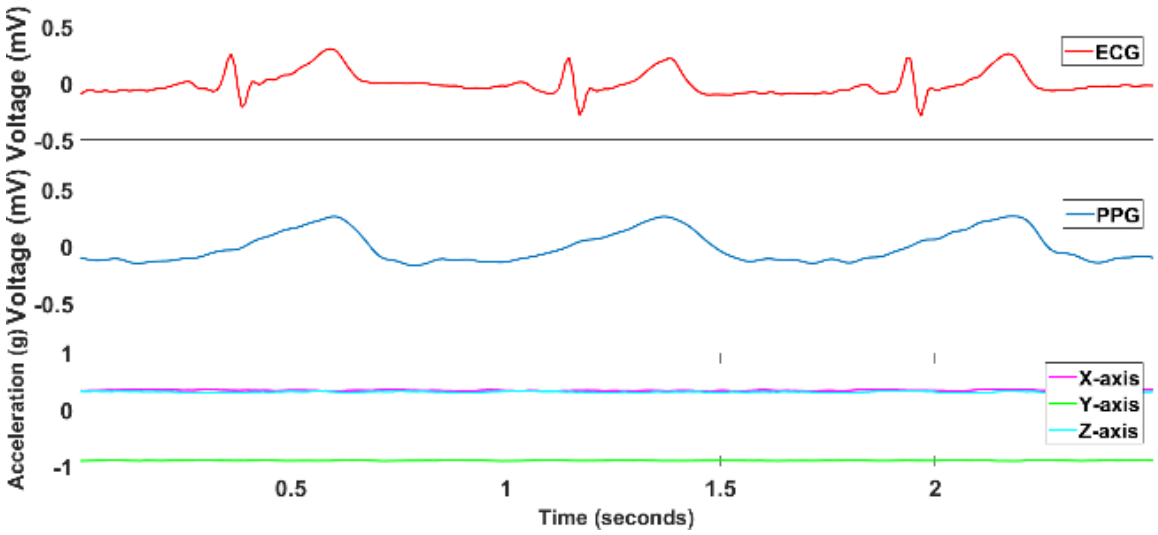

Figure 6.4 – ECG, PPG and Accelerometer signals with subject at rest. Reprinted from Particle Filtering and Sensor Fusion for Robust Heart Rate Monitoring using Wearable Sensors, Viswam Nathan and Roozbeh Jafari, IEEE Journal of Biomedical and Health Informatics, ©2017, IEEE

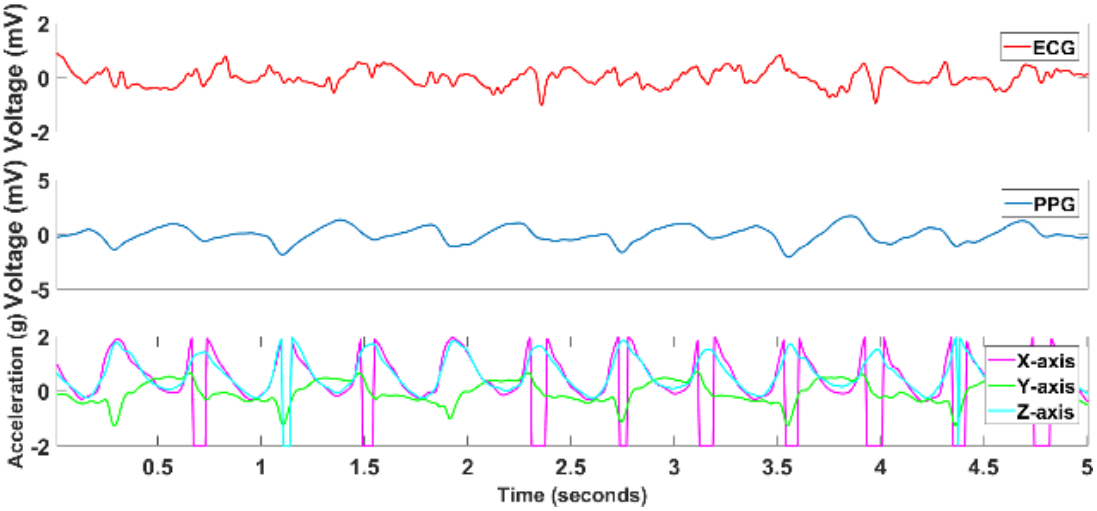

Figure 6.5 – ECG, PPG and Accelerometer signals with subject running. © 2017 IEEE. Reprinted from Particle Filtering and Sensor Fusion for Robust Heart Rate Monitoring using Wearable Sensors, Viswam Nathan and Roozbeh Jafari, IEEE Journal of Biomedical and Health Informatics, ©2017, IEEE

6.6 Experimental Results

6.6.1 Heart Rate Estimation Accuracy – PPG Database

Table 6.1 shows the average heart rate estimation error in bpm for each of the 12 subjects in the SP Cup database as well as the overall mean and standard deviation of error. We can see that the average error is < 2bpm for most subjects. Also shown for comparison are the corresponding results from the JOSS, TROIKA and Robust EEMD works. Note that Table I shows the results for the truncated data, and results are presented for our proposed work as well as the Robust EEMD at both 25Hz and 125Hz sampling rate. The average errors are more or less similar for the different methods, with the 'Robust EEMD' marginally better, whereas the proposed method at 125Hz shows the lowest standard deviation of error. The results for the un-truncated data are in Table 6.2, and we can see that the error from the particle filter estimates are hardly affected despite the noisy initial periods that prohibited the use of the JOSS algorithm.

In Figure 6.6 is shown the Bland-Altman plot for the particle filter estimates' agreement with the ground truth at the full 125Hz sampling rate. The limits of agreement (LOA) were defined following standard practice as $[\mu - 1.96\sigma, \mu + 1.96\sigma]$, where μ is the average difference and σ is the standard deviation, 2.35 bpm in this case. The LOA were [-4.75, 4.45] bpm, and 95% of the difference values were within this confidence interval.

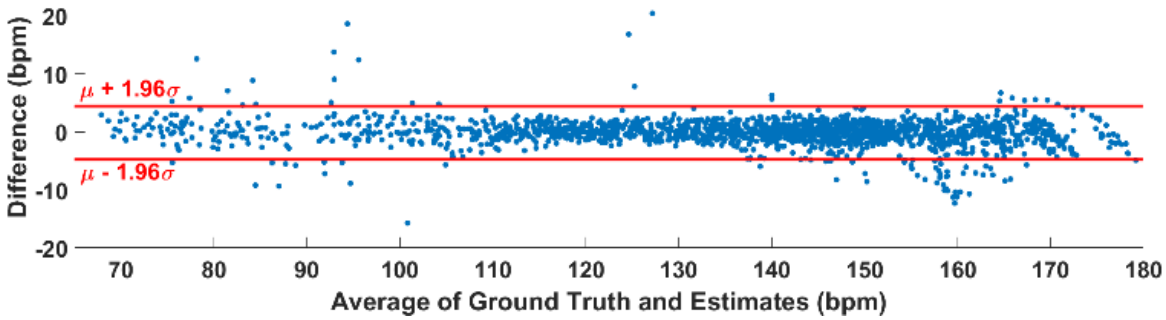

Figure 6.6 – Bland-Altman plot for particle filter agreement with ground truth. © 2017 IEEE. Reprinted from Particle Filtering and Sensor Fusion for Robust Heart Rate Monitoring using Wearable Sensors, Viswam Nathan and Roozbeh Jafari, IEEE Journal of Biomedical and Health Informatics, ©2017, IEEE

6.6.2 Heart Rate Estimation Accuracy – ECG Database and Fusion

Table 6.3 shows the estimation error when using the particle filter to estimate heart rate from the noisy ECG simulated as described in section 6.5.2. For comparison, we show the average estimation error for heart rates as computed by our implementation of the well-respected Pan-Tompkins algorithm, which was designed specifically to estimate heart rate from ECG signals [76]. Of course, the Pan-Tompkins algorithm was not designed for this intensity of motion artifacts, but we included it to show the extent of noisiness in the ECG which causes significant issues for an established algorithm. We can see how the particle filter also works well with this different modality with low error rates. In addition, also shown in the table are the results of fusion of this noisy ECG with the two noisy PPG channels and the accelerometer. We can see how the fusion almost always improves the accuracy, showing how the particle filter was able to effectively

reward the consistent true observations across the different sources and make the best of the sensors available.

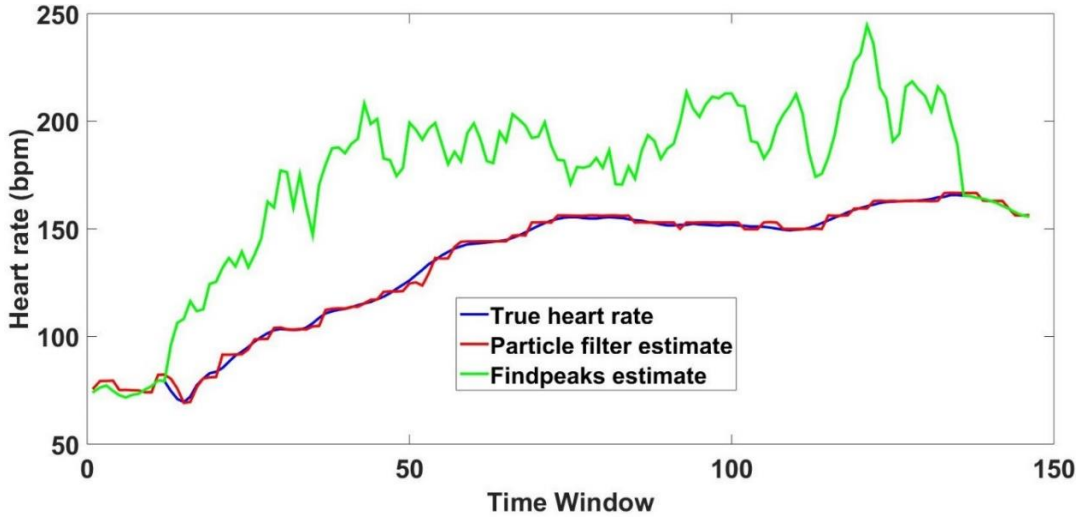

Figure 6.7 – Heart rate estimation performance on a single subject. Reprinted from Particle Filtering and Sensor Fusion for Robust Heart Rate Monitoring using Wearable Sensors, Viswam Nathan and Roozbeh Jafari, IEEE Journal of Biomedical and Health Informatics, ©2017, IEEE

The particle filter tracking over time for Subject 1 is also shown in Figure 6.7 for illustrative purposes. In this figure, 'Findpeaks estimate' refers to the heart rate estimate based solely on the CWT-based peak observation method on ECG, and it can be seen how it tends to overestimate as soon as the motion starts, whereas the particle filter continues to keep track even as the subject's heart rate changes substantially during periods of motion activity.

Table 6.1 - Mean Absolute Heart Rate Estimation Error (in bpm) for the Various Algorithms on the Truncated Datasets

Subject #	1	2	3	4	5	6	7	8	9	10	11	12	Mean ± SD
JOSS [62] (25Hz)	1.33	1.75	1.47	1.48	0.69	1.32	0.71	0.56	0.49	3.81	0.78	1.04	1.28 ± 2.61
TROIKA [61] (25Hz)	3.05	3.31	1.49	2.03	1.46	2.35	1.76	1.43	1.28	5.08	1.8	3.02	2.34 ± 2.86
Robust EEMD [63] (25Hz)	1.7	0.84	0.56	1.15	0.77	1.06	0.63	0.53	0.52	2.56	1.05	0.91	1.02 ± 1.79
Particle Filter (25Hz)(Our method)	2.21	1.71	1.11	1.71	1.1	1.72	1.11	1.29	1.12	3.5	1.68	1.57	1.65 ± 2.07
Robust EEMD [63] (125Hz)	1.83	0.85	0.63	1.21	0.65	1.03	0.7	0.5	0.47	2.83	1.14	0.9	1.06 ± 2.02
Particle Filter (125Hz) (Our method)	1.91	1.3	1.08	1.63	1.06	1.64	1.09	1.25	1.1	3.41	1.65	1.59	1.56 ± 1.73

Table 6.2 - Mean Absolute Heart Rate Estimation Error (in bpm) for the Various Algorithms on the Un-truncated Datasets

Subject #	1	2	3	4	5	6	7	8	9	10	11	12	Mean ± SD
TROIKA [61] (25Hz)	3.05	3.49	1.49	2.03	1.46	2.35	1.76	1.42	1.28	5.73	1.79	3.02	2.41 ± 3.45
Robust EEMD [63] (25Hz)	1.64	0.81	0.57	1.44	0.77	1.06	0.63	0.47	0.52	2.94	1.05	0.91	1.07 ± 2.17
Particle Filter (25Hz) (Our method)	2.21	1.55	1.41	1.65	1.1	1.72	1.11	1.24	1.12	3.63	1.65	1.57	1.66 ± 2.17
TROIKA [61] (125Hz)	2.29	2.19	2	2.15	2.01	2.76	1.67	1.93	1.86	4.7	1.72	2.84	2.34 ± 0.82
Particle Filter (125Hz) (Our method)	1.91	1.46	1.39	1.61	1.06	1.64	1.09	1.25	1.1	3.58	1.73	1.59	1.62 ± 2.01

Table 6.3 - Mean Absolute Heart Rate Estimation Error (in bpm) for the ECG and Fusion (ECG+PPG) Particle Filters, and Pan-Tompkins

Subject #	1	2	3	4	5	6	7	8	9	10	11	12	Mean ± SD
ECG Particle Filter (Our method)	1.49	1.83	2.31	1.2	1.05	2.42	1.91	1.53	1.44	1.13	1.04	1.34	1.56 ± 2.02
ECG+PPG Particle Filter (Our method)	1.26	1.17	0.85	1.11	0.84	1.03	0.87	0.93	0.87	2.24	1.08	1.16	1.12 ± 1.32
Pan-Tompkins[76]	26.1	17.5	19.9	23.5	23.3	24.6	22.7	18.6	18.2	33.9	25.2	24.4	23.2 ± 20.02

6.6.3 Heart Rate Estimation – Experimental Data Collection

In Table 6.4 we also present the results of heart rate error from the fusion particle filter on the dataset collected ourselves, which guarantees real simultaneous ECG and PPG affected by motion artifacts. This shows that the particle filter performance continues to be effective even in this scenario. Again, for comparison is shown the error rates when using the Pan-Tompkins algorithm on the noisy ECG. Note that for Subject 4 the Pan-Tompkins algorithm's adaptive parameters completely went off track early on in the data record due to excessive noise, and did not recover estimates thereafter.

Table 6.4 - Mean Absolute Estimation Error for Fusion Particle Filter and Pan-Tompkins on Our Experimental Dataset

Subject #	Error for Particle Filter (bpm)	Error for Pan-Tompkins [76] (bpm)
1	1.55	13.73
2	1.63	19.57
3	1.25	11.46
4	1.12	N/A
5	1.47	10.4
Mean ± SD	1.4 ± 1.55	13.79 ± 17.35

6.6.4 Fusion Contribution Analysis

In order to further illustrate how the fusion of modalities works, we take a closer look at the performance on Subject 10 from the database. As can be seen in Tables 6.1

and 6.2, estimation performance on this subject is noticeably worse, for our algorithm as well as those of other previous works. This suggests that the PPG signals themselves were relatively more unreliable for this subject. However, we see that in Table 6.3 when using the noisy ECG the performance is much better; so we can assume in this instance that the ECG is a more reliable signal at least for certain segments of the data.

Figure 6.8 shows the relative contribution of each modality – ECG and the two PPG sensors – over time for Subject 10, computed as described in equations (6.26) and (6.27). In this figure, we plot only a subset of the time windows, spanning about 1 minute. Moreover, overlaid in red is the particle filter heart rate estimation error for each of those windows. The error rises to almost 20 beats per minute around window 10, but soon after this the contribution of the ECG to the overall estimate increases. It is clear that the particle filter fusion rewards the more consistent observations from the ECG, and correspondingly the overall error drops sharply. We see a similar trend on a smaller scale around time window 40, where the error is relatively high until the ECG contributions become higher and the overall estimation performance becomes better. In future work, we aim to implement techniques that can recognize these trends of quality of observations and explicitly re-weight individual modalities in the fusion formulation.

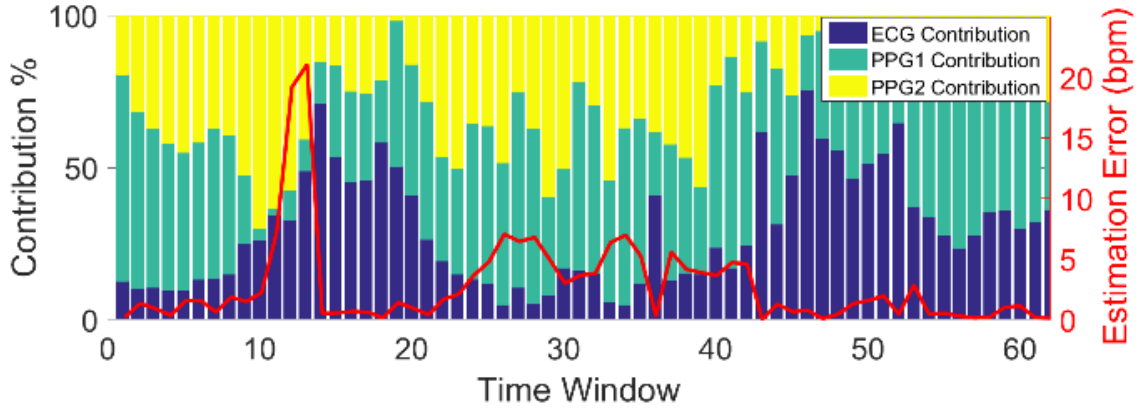

Figure 6.8 – Relative contribution of ECG and PPG modalities to overall fusion particle filter estimate over time for Subject #10. Reprinted from Particle Filtering and Sensor Fusion for Robust Heart Rate Monitoring using Wearable Sensors, Viswam Nathan and Roozbeh Jafari, IEEE Journal of Biomedical and Health Informatics, ©2017, IEEE

6.7 Comparison with Related Works

6.7.1 Discussion of Estimation Performance

The estimation errors are low, but in order to provide further context, we have compared the results to those of recent state-of-the-art works on heart rate estimation in the presence of motion artifacts. The estimation error levels are comparable to the most recent related works in the area. We note that the other related works were specifically developed and optimized for the objective of heart rate monitoring using PPG signals with several heuristics; for instance, TROIKA and JOSS use heuristics such as a rigid artificial bound on the variability of reported heart rate estimates from one window to the next, thresholds for what constitutes a big enough peak in the PPG frequency spectrum, and polynomial curve fitting based on previous heart rate estimates to predict the next estimate when the tracking does not return a satisfactory result. Similarly, the

‘Robust EEMD’ work, in addition to using EEMD and an adaptive filter, has an arbitrary ‘absolute criterion’ to designate a ‘reliable peak’ in the PPG spectrum for heart rate and thresholds for what constitutes a strong enough peak in the PPG spectrum. The algorithm also deletes or removes segments of the signal from consideration if the corresponding accelerometer magnitude is too high. Moreover, with the EEMD approach the user is required to manually detect in a training phase which of the several intrinsic mode functions has the pertinent heart rate frequency information, and this also changes with sampling rate. It is therefore notable that the relatively more generalized particle filter framework introduced here with minimal heuristics or rule-based steps, no requirement for clean start, no deletion of data, which can work with other signal modalities as shown with ECG, and can also be applied to other physiological signal estimation problems, exhibits comparable performance to contemporary works that were purpose-built for the heart rate estimation problem on PPG signals. Moreover, as will be noted in the next section, this comparable estimation performance is achieved with an algorithm that is far more computationally efficient compared to these works.

6.7.2 Analysis of Computation Time

In this work, due to the formulation with the heart rate state, we mitigated the computational load by tracking only one state dimension with just 300 particles. Indeed, the contemporary works we can compare this to are significantly more computationally intensive. The authors of the ‘Robust EEMD’ work [63] note that the TROIKA algorithm takes about 17 minutes and 30 seconds on average to complete heart rate

estimation on a single subject at a sampling rate of 125Hz; whereas the Robust EEMD algorithm itself takes about 55 seconds per subject. Similarly, at a sampling rate of 25Hz, the JOSS algorithm takes about 25s on average per subject and the corresponding Robust EEMD algorithm takes about 16s. When we measured the execution time of our particle filter implementation on MATLAB, the average time per subject was only about 1.04 seconds for the 25Hz sampling rate, and 1.18 seconds for the 125Hz sampling rate. It must be noted that the execution times reported above for the related works were gathered from a work that used MATLAB 2013a, whereas we use MATLAB 2017a. However, this alone cannot account for the highly significant difference in computation time. Furthermore, the machine used to extract these results has similar specifications to the one used to report the results for the related works [63]. In particular, we used a Windows 10 64-bit PC with an Intel i7-6700 processor at 2.60 GHz and 16GB of RAM.

We also analyzed the trade-off between the accuracy and computational cost as a function of the number of particles. Figure 6.9 shows a comparison of the error rates and computation time per minute of data for our particle filter as the number of particles is varied for a single subject. As a reminder, we used N = 300 particles in our work. While the estimation performance does improve as we increase the number of particles, as expected, it is likely that the higher values of N would make it impractical to compute these estimates in real-time, especially on wearable sensors. On such systems, one can easily adjust the number of particles subject to the availability of computational resources.

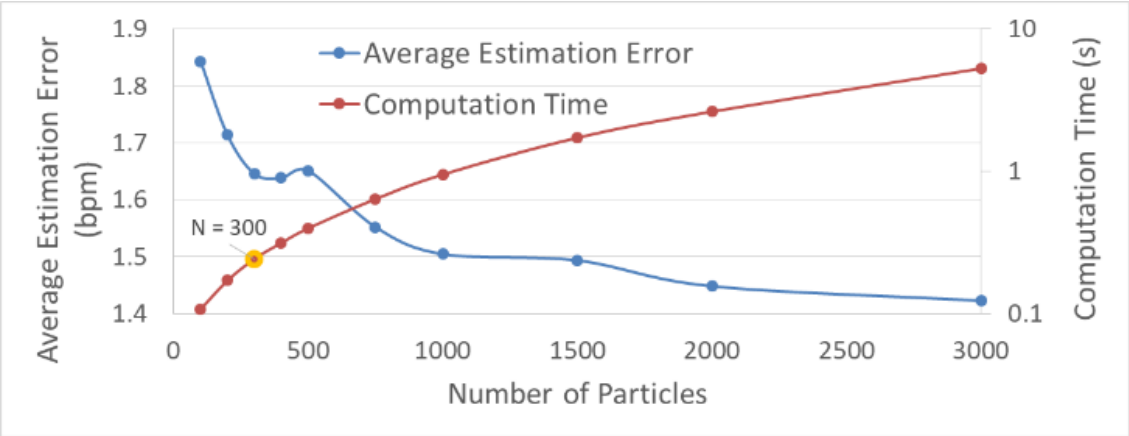

Figure 6.9 – Changes in average estimation error and computation time per minute of data on a single subject as the number of particles is changed. Reprinted from Particle Filtering and Sensor Fusion for Robust Heart Rate Monitoring using Wearable Sensors, Viswam Nathan and Roozbeh Jafari, IEEE Journal of Biomedical and Health Informatics, ©2017, IEEE

6.8 Limitations

We did not test on patients with heart rate variability or other cardiac conditions; this will likely require some tuning of the parameters, but this would be applicable to other contemporary signal processing techniques as well. Testing on subjects with abnormal cardiac activity will be left for future work. We also note that precise computational benchmarking is not the primary goal of this work; the previous section was only meant to provide a rough guide indicating a definite computational advantage over contemporary related works in the area. Deployment of the algorithm on a system is out of the scope of this work; however we submit that the design of such a system is eminently feasible, especially if we leverage cloud computing resources or other techniques to circumvent the computational constraints on typical wearable systems.

6.9 Conclusions

This part of the research introduced a generalized particle filter framework that can be used to track heart rate and proved the feasibility of the technique on real world PPG and ECG signals affected by motion artifacts. Furthermore, it was shown how the particle filter can be used to successfully improve estimation accuracy by combining information from multiple modalities simultaneously measuring the same target phenomenon or even the noise associated with the target. This will prove useful in the context of the upcoming IoT ecosystem where there are multiple wearable and environmental sensors continuously monitoring the physiological status of the user.

7. CONCLUSION

This dissertation document described the research that tackled two primary sources of noise pertaining to wearable sensors for continuous non-invasive health monitoring: improper sensor contact and motion artifacts.

The problem of improper sensor contact was explored in the context of EEG signal monitoring, since this physiological signal modality is particularly susceptible to any deterioration in the skin-electrode interface, and existing solutions do not provide an automated, user-friendly approach that maintains the convenience of wearable dry electrode solution. The research first explored the relative merits of varying the number and density of fingers in a finger-based dry electrode design, since this aspect has not been previously studied despite this being a popular design architecture. Experimental results on human subjects led to the conclusion that denser configurations with a high number of fingers required extra efforts and preparation to properly penetrate through the hair and make contact with the scalp, whereas sparser configurations penetrated the hair more readily. However, with proper manual adjustment and placing, higher density configurations showed lower impedance of contact which led to higher quality signals.

The next part of the EEG electrode research explored a possible solution to adaptively reconfigure the electrodes automatically to continually strive for the best possible contact without manual intervention. The noise effects of individual fingers were studied on a custom-built finger-based electrode, leading to the conclusion that indeed some individual fingers had worse contact and hence more noise compared to other fingers. These findings were then leveraged in another custom-built electrode that

could actively select any subset of fingers on an electrode, and this was used to prove that selecting only the subset of fingers with a good contact could improve the overall EEG signal when compared to using all of the fingers as is standard practice with existing dry finger-based EEG electrodes.

The problem of motion artifacts was tackled in the context of continuous heart rate monitoring, as this is a relevant metric especially during periods of intense motion such as exercise. The developed solution was based on the state estimation technique known as particle filtering, with a novel formulation in the context of heart rate estimation. The developed algorithm was tested on real motion-artifact affected ECG and PPG signals both from an existing database as well as our own experimental data collections. The algorithm showed very effective tracking performance, with average heart rate estimation error < 2bpm over all subjects tested despite interleaving periods of running on a treadmill and associated motion artifacts and heart rate changes. Furthermore, the proposed solution was developed to take advantage of multiple sensors in the IoT, and was shown to be capable of fusing the information from ECG, PPG and Accelerometer sensors to improve the estimates. Finally, the proposed algorithm showed comparable accuracy of estimation with respect to contemporary related works, while remaining more generalized than those works with substantially faster computational performance.

The findings of this work can help alleviate the most important roadblocks to continuous, pervasive physiological signal monitoring using wearable sensors.

REFERENCES

[1] S. C. Mukhopadhyay, "Wearable sensors for human activity monitoring: A review," *IEEE Sensors Journal*, vol. 15, pp. 1321–1330, March 2015.

[2] G. Appelboom, E. Camacho, M. E. Abraham, S. S. Bruce, E. L. Dumont, B. E. Zacharia, R. D'Amico, J. Slomian, J. Y. Reginster, O. Bruyère, and E. S. Connolly, "Smart wearable body sensors for patient self-assessment and monitoring," *Archives of Public Health*, vol. 72, p. 28, Aug 2014.

[3] M. M. Rodgers, V. M. Pai, and R. S. Conroy, "Recent advances in wearable sensors for health monitoring," *IEEE Sensors Journal*, vol. 15, pp. 3119–3126, June 2015.

[4] S. R. Steinhubl, E. D. Muse, and E. J. Topol, "The emerging field of mobile health," *Science Translational Medicine*, vol. 7, no. 283, pp. 283rv3–283rv3, 2015.

[5] Y. L. Zheng, X. R. Ding, C. C. Y. Poon, B. P. L. Lo, H. Zhang, X. L. Zhou, G. Z. Yang, N. Zhao, and Y. T. Zhang, "Unobtrusive sensing and wearable devices for health informatics," *IEEE Transactions on Biomedical Engineering*, vol. 61, pp. 1538–1554, May 2014.

[6] P. J. Soh, G. A. E. Vandenbosch, M. Mercuri, and D. M. M. P. Schreurs, "Wearable wireless health monitoring: Current developments, challenges, and future trends," *IEEE Microwave Magazine*, vol. 16, pp. 55–70, May 2015.

[7] Medgadget, "Wearable medical device market to reach USD 7.8 billion with 17.7% CAGR forecast to 2020," 2017.

[8] S. Patel, H. Park, P. Bonato, L. Chan, and M. Rodgers, "A review of wearable sensors and systems with application in rehabilitation," *Journal of NeuroEngineering and Rehabilitation*, vol. 9, p. 21, Apr 2012.

[9] H. Banaee, M. U. Ahmed, and A. Loutfi, "Data mining for wearable sensors in health monitoring systems: A review of recent trends and challenges," *Sensors*, vol. 13, no. 12, pp. 17472–17500, 2013.

[10] B. M. Marlin, D. C. Kale, R. G. Khemani, and R. C. Wetzel, "Unsupervised pattern discovery in electronic health care data using probabilistic clustering models," in *Proceedings of the 2Nd ACM SIGHIT International Health Informatics Symposium*, IHI '12, (New York, NY, USA), pp. 389–398, ACM, 2012.

[11] A. Pantelopoulos and N. G. Bourbakis, "A survey on wearable sensor-based systems for health monitoring and prognosis," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 40, pp. 1–12, Jan 2010.

[12] Centers for Medicare and Medicaid Services, "NHE fact sheet," 2017.

[13] S. Ha, C. Kim, Y. M. Chi, A. Akinin, C. Maier, A. Ueno, and G. Cauwenberghs, "Integrated circuits and electrode interfaces for noninvasive physiological monitoring," *IEEE Transactions on Biomedical Engineering*, vol. 61, pp. 1522–1537, May 2014.

[14] E. W. Sellers, T. M. Vaughan, and J. R. Wolpaw, "A brain-computer interface for long-term independent home use," *Amyotrophic Lateral Sclerosis*, vol. 11, no. 5, pp. 449–455, 2010. PMID: 20583947.

[15] Y. Chi, T.-P. Jung, and G. Cauwenberghs, "Dry-contact and noncontact biopotential electrodes: Methodological review," *Biomedical Engineering, IEEE Reviews in*, vol. 3, pp. 106–119, 2010.

[16] E. Huigen, A. Peper, and C. A. Grimbergen, "Investigation into the origin of the noise of surface electrodes," *Medical and Biological Engineering and Computing*, vol. 40, pp. 332–338, May 2002.

[17] E. S. Kappenman and S. J. Luck, "The effects of electrode impedance on data quality and statistical significance in erp recordings," *Psychophysiology*, vol. 47, no. 5, pp. 888–904, 2010.

[18] T. Sullivan, S. Deiss, T.-P. Jung, and G. Cauwenberghs, "A brain-machine interface using dry-contact, low-noise eeg sensors," in *Circuits and Systems, 2008. ISCAS 2008. IEEE International Symposium on*, pp. 1986–1989, May 2008.

[19] R. Matthews, N. J. McDonald, H. Anumula, J. Woodward, P. J. Turner, M. A. Steindorf, K. Chang, and J. M. Pendleton, "Novel hybrid bioelectrodes for ambulatory zero-prep eeg measurements using multi-channel wireless eeg system," in *Proceedings of the 3rd International Conference on Foundations of Augmented Cognition*, FAC'07, (Berlin, Heidelberg), pp. 137–146, Springer-Verlag, 2007.

[20] L.-D. Liao, I.-J. Wang, S.-F. Chen, J.-Y. Chang, and C.-T. Lin, "Design, fabrication and experimental validation of a novel dry-contact sensor for measuring electroencephalography signals without skin preparation," *Sensors*, vol. 11, no. 6, pp. 5819–5834, 2011.

[21] Y. Chi, Y.-T. Wang, Y. Wang, C. Maier, T.-P. Jung, and G. Cauwenberghs, "Dry and noncontact eeg sensors for mobile brain-computer interfaces," *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, vol. 20, pp. 228–235, March 2012.

[22] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000 (June 13). Circulation Electronic Pages: http://circ.ahajournals.org/cgi/content/full/101/23/e215 PMID:1085218; doi: 10.1161/01.CIR.101.23.e215.

[23] G. B. Moody, W. Muldrow, and R. G. Mark, "A noise stress test for arrhythmia detectors," *Computers in Cardiology*, vol. 11, pp. 381–384, 1984.

[24] G. E. Prinsloo, H. L. Rauch, and W. E. Derman, "A brief review and clinical application of heart rate variability biofeedback in sports, exercise, and rehabilitation medicine," *The Physician and Sportsmedicine*, vol. 42, no. 2, pp. 88–99, 2014.

[25] T. R. Bennett, J. Wu, N. Kehtarnavaz, and R. Jafari, "Inertial measurement unit-based wearable computers for assisted living applications: A signal processing perspective," *IEEE Signal Processing Magazine*, vol. 33, pp. 28–35, March 2016.

[26] O. Cappe, S. Godsill, and E. Moulines, "An overview of existing methods and recent advances in sequential Monte Carlo," *Proceedings of the IEEE*, vol. 95, pp. 899–924, May 2007.

[27] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *Signal Processing, IEEE Transactions on*, vol. 50, pp. 174–188, Feb 2002.

[28] N. Gandhi, C. Khe, D. Chung, Y. M. Chi, and G. Cauwenberghs, "Properties of dry and non-contact electrodes for wearable physiological sensors," in *2011 International Conference on Body Sensor Networks*, pp. 107–112, May 2011.

[29] H. J. Scheer, T. Sander, and L. Trahms, "The influence of amplifier, interface and biological noise on signal quality in high-resolution EEG recordings," *Physiological Measurement*, vol. 27, no. 2, p. 109, 2006.

[30] C. Gondran, E. Siebert, S. Yacoub, and E. Novakov, "Noise of surface bio-potential electrodes based on nasicon ceramic and Ag-AgCl," *Medical and Biological Engineering and Computing*, vol. 34, pp. 460–466, Nov 1996.

[31] E. Huigen, "Noise in biopotential recording using surface electrodes," Master's thesis, University of Amsterdam, 2000.

[32] H. de Talhouet and J. G. Webster, "The origin of skin-stretch-caused motion artifacts under electrodes," *Physiological Measurement*, vol. 17, no. 2, p. 81, 1996.

[33] W. Vos, P. Bergveld, and E. Marani, “Low frequency changes in skin surface potentials by skin compression: Experimental results and theories,” *Archives of Physiology and Biochemistry*, vol. 111, no. 4, pp. 369–376, 2003. PMID: 15764077.

[34] L. A. Nunes Amaral, P. C. Ivanov, N. Aoyagi, I. Hidaka, S. Tomono, A. L. Goldberger, H. E. Stanley, and Y. Yamamoto, “Behavioral-independent features of complex heartbeat dynamics,” *Phys. Rev. Lett.*, vol. 86, pp. 6026–6029, Jun 2001.

[35] P. C. Ivanov, L. A. N. Amaral, A. L. Goldberger, S. Havlin, M. G. Rosenblum, Z. R. Struzik, and H. E. Stanley, “Multifractality in human heartbeat dynamics,” *Nature*, vol. 399, pp. 461 EP –, Jun 1999.

[36] K. Kiyono, Z. R. Struzik, N. Aoyagi, and Y. Yamamoto, “Multiscale probability density function analysis: non-Gaussian and scale-invariant fluctuations of healthy human heart rate,” *IEEE Transactions on Biomedical Engineering*, vol. 53, pp. 95–102, Jan 2006.

[37] D. Cyrill, J. McNames, and M. Aboy, “Adaptive comb filters for quasiperiodic physiologic signals,” in *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE Cat. No.03CH37439)*, vol. 3, pp. 2439–2442 Vol.3, Sept 2003.

[38] F. Scholkmann, J. Boss, and M. Wolf, “An efficient algorithm for automatic peak detection in noisy periodic and quasi-periodic signals,” *Algorithms*, vol. 5, no. 4, pp. 588–603, 2012.

[39] S. S. Thomas, V. Nathan, C. Zong, P. Aroul, L. Philipose, K. Soundarapandian, X. Shi, and R. Jafari, “Demonstration abstract: Biowatch – a wrist

watch based physiological signal acquisition system," in *IPSN-14 Proceedings of the 13th International Symposium on Information Processing in Sensor Networks*, pp. 349–350, April 2014.

[40] S. S. Thomas, V. Nathan, C. Zong, E. Akinbola, A. L. P. Aroul, L. Philipose, K. Soundarapandian, X. Shi, and R. Jafari, "Biowatch - a wrist watch based signal acquisition system for physiological signals including blood pressure," in *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 2286–2289, Aug 2014.

[41] S. S. Thomas, V. Nathan, C. Zong, K. Soundarapandian, X. Shi, and R. Jafari, "Biowatch: A noninvasive wrist-based blood pressure monitor that incorporates training techniques for posture and subject variability," *IEEE Journal of Biomedical and Health Informatics*, vol. 20, pp. 1291–1300, Sept 2016.

[42] Y. J. Huang, C. Y. Wu, A. M. K. Wong, and B. S. Lin, "Novel active comb-shaped dry electrode for EEG measurement in hairy site," *IEEE Transactions on Biomedical Engineering*, vol. 62, pp. 256–263, Jan 2015.

[43] G. Edlinger, G. Krausz, and C. Guger, "A dry electrode concept for SMR, p300 and SSVEP based BCIs," in *Complex Medical Engineering (CME), 2012 ICME International Conference on*, pp. 186–190, July 2012.

[44] S. Lee, Y. Shin, S. Woo, K. Kim, and H. N. Lee, "Dry electrode design and performance evaluation for EEG based BCI systems," in *2013 International Winter Workshop on Brain-Computer Interface (BCI)*, pp. 52–53, Feb 2013.

[45] P. Fiedler, P. Pedrosa, S. Griebel, C. Fonseca, F. Vaz, F. Zanow, and J. Haueisen, “Novel flexible dry pu/tin-multipin electrodes: First application in EEG measurements,” in *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 55–58, Aug 2011.

[46] S. Krachunov and A. J. Casson, “3D printed dry EEG electrodes,” *Sensors*, vol. 16, no. 10, 2016.

[47] P. Fiedler, S. Brodkorb, C. Fonseca, F. Vaz, F. Zanow, and J. Haueisen, “Novel TiN-based dry EEG electrodes: Influence of electrode shape and number on contact impedance and signal quality”, pp. 418–421. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010.

[48] A. A. Alhammadi, T. B. Nazzal, and S. A. Mahmoud, “A CMOS EEG detection system with a configurable analog front-end architecture,” *Analog Integrated Circuits and Signal Processing*, vol. 89, pp. 151–176, Oct 2016.

[49] T. Zhao, X. Liu, G. Zhang, and Y. Su, “Design of a programmable and low-frequency filter for biomedical signal sensing applications,” in *2016 9th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, pp. 1746–1750, Oct 2016.

[50] C.-C. Lo, T.-Y. Chien, Y.-C. Chen, S.-H. Tsai, W.-C. Fang, and B.-S. Lin, “A wearable channel selection-based brain-computer interface for motor imagery detection,” *Sensors*, vol. 16, no. 2, 2016.

[51] K. Colwell, D. Ryan, C. Throckmorton, E. Sellers, and L. Collins, "Channel selection methods for the p300 speller," *Journal of Neuroscience Methods*, vol. 232, no. Supplement C, pp. 6 – 15, 2014.

[52] Y. Zou, V. Nathan, and R. Jafari, "Automatic identification of artifact-related independent components for artifact removal in EEG recordings," *IEEE Journal of Biomedical and Health Informatics*, vol. 20, pp. 73–81, Jan 2016.

[53] T. Degen and H. Jackel, "Continuous monitoring of electrode–skin impedance mismatch during bioelectric recordings," *Biomedical Engineering, IEEE Transactions on*, vol. 55, pp. 1711–1715, June 2008.

[54] T. Tamura, Y. Maeda, M. Sekine, and M. Yoshida, "Wearable photoplethysmographic sensors—past and present," *Electronics*, vol. 3, no. 2, p. 282, 2014.

[55] C. Yang and N. Tavassolian, "Motion noise cancellation in seismocardiogram of ambulant subjects with dual sensors," in *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 5881–5884, Aug 2016.

[56] D. Jarchi and A. J. Casson, "Estimation of heart rate from foot worn photoplethysmography sensors during fast bike exercise," in *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 3155–2158, Aug 2016.

[57] N. V. Helleputte, M. Konijnenburg, J. Pettine, D. W. Jee, H. Kim, A. Morgado, R. V. Wegberg, T. Torfs, R. Mohan, A. Breeschoten, H. de Groot, C. V.

Hoof, and R. F. Yazicioglu, “A 345 μW multi-sensor biomedical SoC with bio-impedance, 3-channel ECG, motion artifact reduction, and integrated DSP,” *IEEE Journal of Solid-State Circuits*, vol. 50, pp. 230–244, Jan 2015.

[58] A. Galli, G. Frigo, C. Narduzzi, and G. Giorgi, “Robust estimation and tracking of heart rate by PPG signal analysis,” in *2017 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, pp. 1–6, May 2017.

[59] K. Sweeney, T. Ward, and S. McLoone, “Artifact removal in physiological signals - practices and possibilities,” *Information Technology in Biomedicine, IEEE Transactions on*, vol. 16, pp. 488–500, May 2012.

[60] S. S. Bisht and M. P. Singh, “An adaptive unscented Kalman filter for tracking sudden stiffness changes,” *Mechanical Systems and Signal Processing*, vol. 49, no. 1–2, pp. 181 – 195, 2014.

[61] Z. Zhang, Z. Pi, and B. Liu, “TROIKA: A general framework for heart rate monitoring using wrist-type photoplethysmographic signals during intensive physical exercise,” *Biomedical Engineering, IEEE Transactions on*, vol. 62, pp. 522–531, Feb 2015.

[62] Z. Zhang, “Photoplethysmography-based heart rate monitoring in physical activities via joint sparse spectrum reconstruction,” *IEEE Transactions on Biomedical Engineering*, vol. 62, pp. 1902–1910, Aug 2015.

[63] E. Khan, F. A. Hossain, S. Z. Uddin, S. K. Alam, and M. K. Hasan, “A robust heart rate monitoring scheme using photoplethysmographic signals corrupted

by intense motion artifacts," *IEEE Transactions on Biomedical Engineering*, vol. 63, pp. 550–562, March 2016.

[64] S. Kim, L. A. Holmstrom, and J. McNames, "Tracking of rhythmical biomedical signals using the maximum a posteriori adaptive marginalized particle filter," *British Journal of Health Informatics and Monitoring*, vol. 2, no. 1, 2015.

[65] C. Lin, M. Bugallo, C. Mailhes, and J.-Y. Tourneret, "ECG denoising using a dynamical model and a marginalized particle filter," in *Signals, Systems and Computers (ASILOMAR), 2011 Conference Record of the Forty Fifth Asilomar Conference on*, pp. 1679–1683, Nov 2011.

[66] C. Lin, G. Kail, A. Giremus, C. Mailhes, J.-Y. Tourneret, and F. Hlawatsch, "Sequential beat-to-beat p and t wave delineation and waveform estimation in ECG signals: Block Gibbs sampler and marginalized particle filter," *Signal Processing*, vol. 104, pp. 174 – 187, 2014.

[67] G.-J. Li, X. na Zhou, S. ting Zhang, and N.-Q. Liu, "ECG characteristic points detection using general regression neural network-based particle filters," in *Bioelectronics and Bioinformatics (ISBB), 2011 International Symposium on*, pp. 155–158, Nov 2011.

[68] G. Li, X. Zeng, J. Lin, and X. Zhou, "Genetic particle filtering for denoising of ECG corrupted by muscle artifacts," in *Natural Computation (ICNC), 2012 Eighth International Conference on*, pp. 562–565, May 2012.

[69] S. Edla, N. Kovvali, and A. Papandreou-Suppappola, "Sequential Markov chain Monte Carlo filter with simultaneous model selection for

electrocardiogram signal modeling," in *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pp. 4291–4294, Aug 2012.

[70] Texas Instruments, *Low-Noise, 8-Channel, 24-Bit Analog Front-End for Biopotential Measurements*, 2012.

[71] G. Burkitt, R. Silberstein, P. Cadusch, and A. Wood, "Steady-state visual evoked potentials and travelling waves," *Clin. Neurophysiol.*, vol. 111, no. 2, pp. 246–258, 2000.

[72] H. H. Jasper, "Report of the committee on methods of clinical examination in electroencephalography: 1957," *Clin. Neurophysiol.*, vol. 10, pp. 370–375, 1958.

[73] I. R. Legarreta, P. S. Addison, M. J. Reed, N. Grubb, G. R. Clegg, C. E. Robertson, and J. N. Watson, "Continuous wavelet transform modulus maxima analysis of the electrocardiogram: beat characterisation and beat-to-beat measurement," *International Journal of Wavelets, Multiresolution and Information Processing*, vol. 03, no. 01, pp. 19–42, 2005.

[74] V. Nathan, I. Akkaya, and R. Jafari, "A particle filter framework for the estimation of heart rate from ECG signals corrupted by motion artifacts," in *Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE*, pp. 6560–6565, Aug 2015.

[75] F. Caron, M. Davy, E. Duflos, and P. Vanheeghe, "Particle filtering for multisensor data fusion with switching observation models: Application to land

vehicle positioning," *IEEE Transactions on Signal Processing*, vol. 55, pp. 2703–2719, June 2007.

[76] J. Pan and W. J. Tompkins, "A real-time QRS detection algorithm," *IEEE Transactions on Biomedical Engineering*, vol. BME-32, pp. 230–236, March 1985.