Low-Power Strategies for Signal Compression in Ambulatory Healthcare

by

Darren Craven, B.E.

in fulfilment of the requirements for the degree of

Doctor of Philosophy

in the subject of

Electrical and Electronic Engineering

Discipline of Electrical and Electronic Engineering
College of Engineering and Informatics
National University of Ireland, Galway
Ireland

Supervisor

Dr. Edward Jones

Co-supervisors

Dr. Martin Glavin

Mr. Liam Kilmartin

August 2016

Table of Contents

Abstract

Advances in signal processing are enabling promising solutions for ambulatory monitoring in healthcare applications. A critical objective of such systems is low-power operation to enable long-term battery life in wearable/mobile devices. This project addresses the low-power objective through the use of lossy compression techniques, particularly focusing on Compressed Sensing (CS).

CS is a recently-introduced paradigm for compression that has significant potential for deployment in Body Area Networks (BAN) compared to existing compression methods. A particular advantage of CS in this context is the possibility of very simple encoding implementations in the wearable sensors, with greater computational complexity in signal reconstruction. However, the performance of CS remains limited in terms of signal reconstruction quality when compared to existing approaches based on Nyquist sampling. In this thesis, novel encoding strategies for CS are developed and novel Dictionary Learning (DL) techniques are utilised for reconstruction to challenge the current state-of-the-art in this area.

The proposed algorithms are analysed in terms of signal fidelity for a given Compression Ratio (*CR*) and results demonstrate their ability to outperform existing CS approaches. Furthermore, there is increasing interest in Computer Aided Diagnosis (CAD) systems that tolerate some loss in signal fidelity at higher *CR*s but where diagnostic integrity of the signal is maintained. The proposed algorithms are also evaluated in terms of their ability to maintain performance in CAD systems at higher *CR*s. Finally, this thesis analyses the power consumption of CS and compares it with a state-of-the-art lossy compression algorithm in terms of overall system power consumption (including power consumption of signal acquisition, digital processing, and wireless transmission) in a BAN. The results confirm the substantial benefits of employing CS as a low-energy encoding implementation.

Acknowledgements

Firstly, I would like to thank my project supervisors Dr. Edward Jones, Dr. Martin Glavin and Mr. Liam Kilmartin who have provided invaluable assistance and advice to me over the course of my research. I would also like to acknowledge the help of all staff in the Electrical and Electronic Engineering department at NUI Galway, with a particular mention to Mr. Martin Burke, Mr. Myles Meehan and Mrs. Catherine Feeney for their assistance throughout.

I would like to thank all members past and present of the Signal Processing and Communications Research Laboratory for sharing their ideas and feedback with me over the years. A special word of thanks must go to Dr. Brian McGinley who has provided constant support for the duration of the project.

I am extremely grateful of the love and support offered by my parents Gerard and Annette, brother Brian and sisters Karen, Ruth and Hazel who have always been there for me over the years.

Finally, I would like to thank Prof. Peter McHugh for his work and dedication in the formation, and ensuring the success of, the structured PhD programme in Biomedical Engineering and Regenerative Medicine (BMERM) which provided my funding. This programme was funded by the Programme for Research in Third Level Institutions (PRTLI) Cycle 5.

Declaration of Authorship

I, **Darren Craven**, declare that this thesis titled, **"Low-Power Strategies for Signal Compression in Ambulatory Healthcare"** and the work presented in it is my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this university.
- Where any part of this thesis has previously been submitted for degree or any other qualification at this university or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signature: 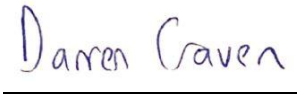

Darren Craven

List of Publications

Journal Publications

D. Craven, B. McGinley, L. Kilmartin, M. Glavin, and E Jones, "Compressed Sensing for Bioelectric Signals: A Review," *IEEE Journal of Biomedical and Health Informatics*, vol. 19, pp. 529-540, 2015.

D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Impact of compressed sensing on clinically relevant metrics for ambulatory ECG monitoring," *Electronics Letters*, vol. 51, pp. 323-325, 2015.

D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Adaptive Dictionary Reconstruction for Compressed Sensing of ECG Signals," accepted for publication in *IEEE Journal of Biomedical and Health Informatics*, February 2016.

D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Energy-Efficient Compressed Sensing for Ambulatory ECG Monitoring," *Computers in Biology and Medicine*, vol. 71, pp. 1-16, 2016.

D. Craven, M. O'Halloran, B. McGinley, R. C. Conceicao, L. Kilmartin, E. Jones and M. Glavin, "Compressive Sampling for Time Critical Microwave Imaging Applications," *Healthcare Technology Letters*, vol. 1, pp. 6-12, 2014.

Conference Publications

D. Craven, L. Kilmartin, M. Glavin and E. Jones, "Potential for Extended Battery Life in Mobile Healthcare with Bluetooth Low Energy and Signal Compression," in *2012 IET Irish Signals and Systems Conference*, 2012, pp. 1-6.

D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, " Effects of Non-Uniform Quantization on ECG acquired using Compressed Sensing," in *2014 EAI 4th International Conference on Wireless Mobile Communication and Healthcare (Mobihealth)*, 2014, pp. 79-82.

List of Figures

List of Tables

List of Acronyms

AD	Adaptive Dictionary
ADC	Analog-to-Digital Converter
AEC	Average Execution Cycle
ASF	Activity Scaling Factor
AUC	Area Under Curve
AZTEC	Amplitude Zone Time Epoch Coding
BAN	Body Area Network
BLE	Bluetooth Low Energy
BP	Basis Pursuit
BSBL	Block-Sparse Bayesian Learning
CAD	Computer Aided Diagnosis
CORTES	Coordinate Reduction Time Encoding Scheme
CoSaMP	Compressive Sampling Matching Pursuit
CR	Compression Ratio
CS	Compressed Sensing
DCT	Discrete Cosine Transform
DL	Dictionary Learning
DWT	Discrete Wavelet Transform
ECG	Electrocardiogram
EZW	Embedded Zerotree Wavelet
FN	False Negative
FOM	Figure-Of-Merit
FP	False Positive

GDP	Gross Domestic Product
HF	High Frequency
HRV	Heart Rate Variability
i.i.d.	independently identically distributed
JPEG	Joint Photographic Experts Group
LASSO	Least Absolute Shrinkage and Selection Operator
LF	Low Frequency
LZ	Lempel-Ziv
meanNN	Mean Normal-to-Normal
MMB	Modified Model-Based
MMB-CoSaMP	Modified Model-Based Compressive Sampling Matching Pursuit
MMB-IHT	Modified Model-Based Iterative Hard Thresholding
MOD	Method of Optimal Directions
NIHT	Normalised Iterative Hard Thresholding
NN	Normal-to-Normal
OMP	Orthogonal Matching Pursuit
PDF	Probability Density Function
PRD	Percent Root-mean-squared Difference
PSim	Percentage Similarity
PVC	Premature Ventricular Contraction
RD	Random Demodulator
RIP	Restricted Isometry Property
RMPI	Random Demodulator Pre-Integrator
ROC	Receiver Operator Characteristic

ROLS	Regularised Orthogonal Least Squares
RRM	Redundancy Removal with Mean
RRM-DL	Redundancy Removal with Mean - Dictionary Learning
SAR	Successive Approximation Register
SD	Standard Dictionary
SDNN	Standard Deviation Normal-to-Normal
SE	Sensitivity
SNR	Signal-to-Noise Ratio
SP	Specificity
SPIHT	Set Partitioning In Hierarchical Trees
SRMPI	Spread Spectrum Random Modulation Pre-Integrator
SRR	Standard Redundancy Removal
SVD	Singular Value Decomposition
TH-DWT	Thresholding - Discrete Wavelet Transform
TP	True Positive
WTS	Wireless Transmission Savings

List of Mathematical Notation

$x \in \mathbb{R}^N$	x is an N-dimensional vector in the space $\mathbb{R}$
$X \in \mathbb{R}^{M \times N}$	X is an M x N matrix in the space $\mathbb{R}$
$[X]_{M,N}$	X is an M x N matrix
$\|.\|_1$	l_1 norm
$\|.\|_2$	l_2 norm
$\langle x, y \rangle$	Inner product of vectors x and y
$\|x\|$	Absolute value of x
$\bar{x}$	Mean of x
$\sum_{i=1}^{N} x_i$	Sum of x_i as i goes from 1 to N
$\|.\|_F$	Frobenius norm
$\|.\|_0$	l_0 norm
[1:N]	Sample 1 to N
$x \approx N$	x is approximately equal to N

CHAPTER 1

Introduction

1.1. Introduction

Recent advances in mobile and communications technology are enabling increased functionality and ensuring a rapid transformation of the healthcare industry. These technology-driven innovations are having a global impact on the methods by which healthcare is provided. The exponential increase of healthcare expenditure as a percentage of each nation's Gross Domestic Product (GDP) [1] and the constant ageing of the world's population, ensure the importance of the development and employment of intelligent strategies that provide affordable healthcare solutions and improve the quality of life for a patient. The growing costs associated with the demand for long-term continuous care, particularly for the elderly, has led healthcare authorities to focus more on ubiquitous and ambulatory monitoring instead of hospital-centred services [2]. Ambulatory monitoring allows for efficient management of a patient's health information, by focusing on continuous monitoring of a patient's physiological signals, providing the potential for real-time analysis allowing possible health problems to be identified quickly and thus reducing the need for clinic-based care. In fact, off-site patient monitoring has been identified as an effective way for screening of several chronic diseases such as cardiovascular disease and diabetes [3]. Such off-site patient monitoring uses low-cost diagnostic

testing that can help alleviate reliance on highly specialised expensive monitoring equipment found in hospitals and frees up resources in terms of number of hospital admissions, beds used and clinical staff time.

An example of a paradigm that offers such ambulatory healthcare is a wireless Body Area Network (BAN). BANs consist of wirelessly-enabled sensors that monitor important physiological signals in or around the human body with the capability for further data processing and wireless transmission of the relevant data. The aim of a BAN is to extend conventional healthcare provided by bedside monitoring in order to provide care to patients both inside and outside the hospital environment [4]. BANs are still relatively uncommon in terms of actual deployment but there has been significant focus on them from a research and development perspective in recent times [5-7]. The notion that a patient can be constantly monitored (both in and out of a clinical environment) is an attractive prospect for healthcare providers and has significant potential to reduce healthcare costs. A typical BAN architecture consists of sensors that continuously monitor the patient and communicate the measured patient data back to an external care site, or base station such as a smartphone or server, for viewing by a clinician or the patient themselves. A clinical expert can visually inspect the patient data, or Computer Aided Diagnosis (CAD) can be used to analyse the data and issue immediate feedback or treatment recommendations for the patient. Alternatively, the data can be stored for future clinical observation.

BANs ensure the patient is not compelled to endure lengthy hospital stays and the wireless nature of their implementation increases patient mobility, unobtrusiveness and provides a more cost-efficient application [8]. The likelihood of user acceptance to wireless BANs has increased in recent years in conjunction with the miniaturisation of sensors, electronic devices and battery components [9]. Constant decreases in the costs of smart bio-sensors and the increasing availability of smart devices and smartphones are pushing towards the realisation of “pervasive healthcare”, that is, the capability to provide healthcare to anyone, anytime, anywhere [10]. This is aided by the ever-increasing number of mobile phone subscribers worldwide and the growing percentage of the world’s population being

covered by a commercial wireless signal, thus facilitating easy transmission of data between patient and clinician. Additionally, the availability of low-power wireless transmission protocols such as Zigbee [11] and Bluetooth Low Energy (BLE) [12] is increasing the viability of providing efficient communication of patient data.

Despite presenting various advantages, the implementation of BANs still face many key research and development challenges that need to be addressed to ensure their widespread deployment in modern healthcare. Security in terms of patient privacy, data integrity and authentication, wireless routing, biocompatibility of sensors (particularly for implantable devices) and maintaining the unobtrusiveness of the architecture were identified as some of the key challenges to BANs in [9]. However, one of the most significant requirements is low-power implementation. As BANs are generally battery powered it is essential that the architecture is designed with the limitations in terms of computational power, storage and battery life in mind. Therefore, in such ambulatory monitoring environments there is a strong need to efficiently manage the potentially large quantities of real-time signal data being generated.

There are three main operations contributing to the overall power consumption in a BAN: signal sampling, digital processing and wireless transmission. Wireless transmission is generally considered to be the operation that expends the most power in such a system [13, 14]. In many biomedical applications continuous monitoring of patient data is an important requirement and this places a constant burden on the sensor and also increases the amount of data that have to be wirelessly transmitted. For complete signal fidelity, lossless compression may be used [15], however, lossless compression is generally limited in terms of the compression rate, and hence energy savings, that can be achieved. To address this, lossy compression has been proposed to manage large quantities of data and this is considered in this thesis. The fundamental objective of lossy compression algorithms is to efficiently maximise the Compression Ratio (*CR*) whilst minimising the associated distortion of the reconstructed signal. Lossy compression can significantly reduce the amount of data required for storage and wireless transmission with the trade-off of a decrease in signal fidelity post reconstruction. A well-known lossy

compression algorithm that has been widely used in the field is Set Partitioning In Hierarchical Trees (SPIHT) [16, 17], which consists of a wavelet analysis of the signal, followed by organisation of the wavelet coefficients into ordered lists based on their significance.

In the context of the challenge to achieve low-energy operation in ambulatory healthcare, the primary focus of this thesis is the development of intelligent strategies for signal acquisition in a BAN, combined with the use of lossy compression to reduce the data storage and wireless transmission requirements, hence improving the overall power efficiency of such architectures. The particular case study that is the focus of this research is the use of the electrocardiogram (ECG) signal, which is of wide importance for a range of clinical conditions.

Of particular interest is a recently-proposed low-power compression technique: Compressed Sensing (CS). CS is a lossy compression paradigm that enables sub-Nyquist rate sampling of signals on the condition that the signal has a sparse representation in a particular domain or dictionary [18-21]. Convex optimisation procedures are then generally used to reconstruct such signals using very few measurements [22]. CS is considered advantageous in low-power ambulatory monitoring environments due to the low complexity nature of the compression or encoding process. Several additional strategies such as removal of redundancy existing in ECG signals, low bit level quantisation and Huffman coding are investigated in this thesis to further enhance the *CR* with CS. Dictionary Learning (DL) techniques and an Adaptive Dictionary (AD) reconstruction framework are also proposed to further improve CS-based ECG compression performance. The performance objective with any lossy compression algorithm is to maximise the *CR* without significantly impacting on the reconstruction quality of the signal, and signal fidelity is used in this thesis as a fundamental performance measure, as it is widely used in the field. Additionally, the performance of the algorithm is analysed in the context of CAD systems, and specifically the degree to which diagnostic information is retained in the signal after compression. This is evaluated in terms of the QRS detection accuracy metric, a metric that is considered clinically-relevant in CAD systems [23, 24]. Furthermore, in keeping with the low-power operation goals of

modern BAN architectures, the power consumption of the proposed approach is analysed in detail. This analysis includes the power required to sample the signal, to perform digital signal compression, and for wireless transmission considering the use of CS techniques. Where appropriate, CS-based approaches are compared with the SPIHT algorithm as a baseline.

1.2. Thesis Contributions

The principal contributions from this thesis can be summarised as follows:

1. Optimisation of CS encoding strategies through a combination of redundancy removal, quantisation and Huffman encoding (Chapter 3).

2. Development of a CS architecture for ECG-based compression that outperforms the existing CS algorithms in the literature (Chapter 3).

3. Proposal of a novel AD reconstruction scheme to improve performance of DL with ECG compression (Chapter 4).

4. Classifying the impact of compression with CS on the performance of the diagnostically-relevant task of QRS detection and examining the maximum *CR* and acceptable signal distortion achievable in order to maintain diagnostic integrity (Chapter 5).

5. Characterising the proposed CS architecture in terms of the overall computational and power requirements for implementation in a BAN, and comparing with existing approaches (Chapter 6).

1.3. Thesis Structure

The remainder of the thesis is structured as follows:

Chapter 2 - Background and Literature Review

This chapter provides an introduction to BANs, and highlights the benefits and challenges of ambulatory monitoring. A comprehensive review of CS for use in biomedical signal compression is also detailed, with a particular focus on ECG

signals. The experimental approach and performance objectives for the proposed techniques are then considered.

Chapter 3 - Compressed Sensing Methodology

The proposed CS architecture is introduced in this chapter. A number of different optimisations are developed and their effect on performance examined. The DL reconstruction methodology is also presented, including an analysis of optimal parameters for dictionary creation.

Chapter 4 - Adaptive Dictionary Reconstruction

This chapter details the proposed novel AD reconstruction framework. The performance is then compared to standard dictionary reconstruction with CS, existing CS approaches and state-of-the-art lossy approaches to ECG compression.

Chapter 5 - Computer Aided Diagnosis

This chapter presents the assessment of CS in terms of QRS detection accuracy and several Heart Rate Variability (HRV) performance metrics. The proposed CS algorithm is compared with the SPIHT lossy compression algorithm as a benchmark. Finally, visual evaluation of QRS detection during certain abnormal ECG events is provided.

Chapter 6 – Sensor Power Consumption Comparison for CS and SPIHT

In this chapter power consumption profiles are created for both CS and SPIHT compression algorithms in terms of signal acquisition, digital signal compression and wireless transmission. The performance is then assessed in terms of signal reconstruction quality and QRS detection accuracy to quantify the energy efficiency of each algorithm for implementation in an ECG BAN monitoring architecture.

Chapter 7 – Conclusions and Future Work

The main contributions of the thesis are summarised and ideas for future work are discussed.

CHAPTER 2

Background and Literature Review

2.1. Introduction

This chapter provides a detailed background to the main topics relevant to the work described in this thesis. A brief description of the BAN architecture for ECG monitoring used as a scenario in this thesis is first provided, and the work of this thesis is placed in the context of this architecture. A background to signal compression on biomedical signals is then given, followed by a discussion on the merits and importance of signal compression in such a monitoring scenario. The main body of the chapter consists of a detailed description of Compressed Sensing (CS). The theory behind CS is outlined, and the different approaches to CS performed in the literature are then comprehensively reviewed. Finally, the database and performance metrics used throughout are presented and the broad performance goals for the proposed CS algorithms are discussed.

2.2. Ambulatory ECG Monitoring

An electrocardiogram (ECG) signal is a representation of the electrical activity of the heart throughout the cardiac cycle. The ECG is used mainly as a diagnostic tool for general clinical cardiology and can provide a clinician with useful information about the function and structure of the patient's cardiovascular system. It is commonly used by clinicians for the detection of cardiac arrhythmias [25],

myocardial ischemia [26] or atrial fibrillation [27] and therefore, continuous monitoring of ECG signals can be useful for screening of such common cardiovascular disorders.

The ECG is obtained non-invasively by placing electrodes on a patient's body that detect changes in the electrical activity of the heart caused by muscle depolarization. An example of a standard ECG signal is illustrated in Figure 2.1 with the P, Q, R, S and T sub-waveforms annotated. The human cardiac cycle is a process where blood is pumped throughout the body and each wave is representative of a stage in the cycle of a normal heartbeat. The process begins with the P wave which occurs as a result of atrial depolarization and is representative of atrial contractions, as blood is pumped into the ventricles. The QRS complex occurs during ventricular depolarization as ventricular contractions pump the blood to the lungs and the rest of the body. Finally, ventricular repolarization results in the formation of the T wave as the ventricles are relaxed and the process of the cardiac cycle resumes again.

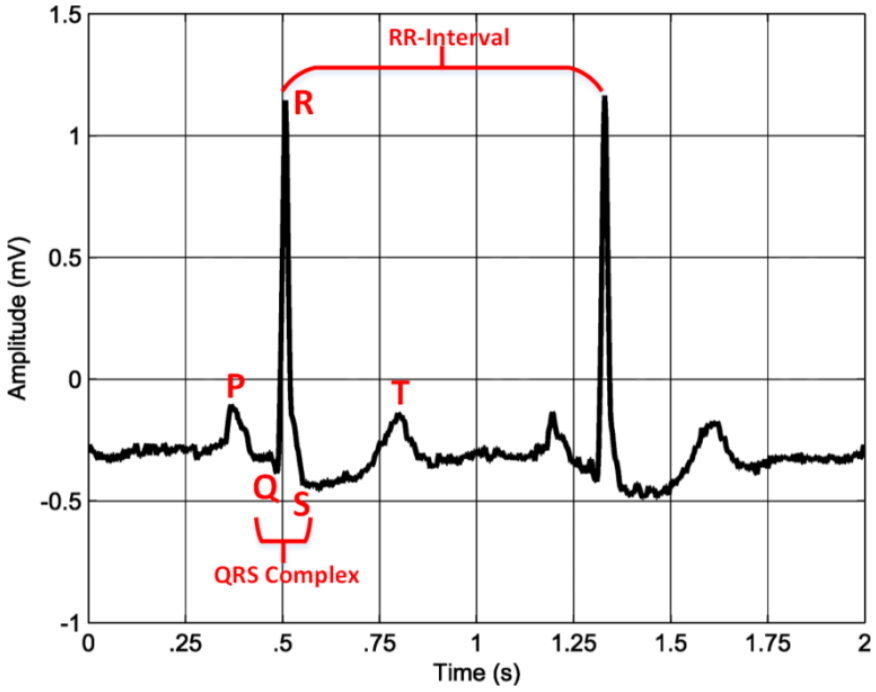

Figure 2.1: Sample of a single-channel ECG signal with the QRS complex and RR-interval annotated.

For clinical interpretation, each heartbeat is represented by the spike in the signal known as the QRS complex. The prominence of this complex in normal healthy ECGs makes it a relatively easy feature to extract for signal processing applications, and it is useful in calculating metrics such as the RR-interval which enables monitoring of heart rate or Heart Rate Variability (HRV) [28, 29]. Another advantage of the unique shape of the ECG signal is that the most relevant information in the signal (the QRS complex) is contained in a narrow window of the time domain since the R-peaks have much larger amplitudes than the P and T waves as illustrated by Figure 2.1.

The ubiquity and benefit of ECG monitoring in the context of cardiac care mean that ECG monitoring is a common application of BANs. A diagram of the ECG BAN architecture assumed throughout this research is shown in Figure 2.2. The ECG signal is acquired on a battery-driven wearable device. The device has processing capabilities to perform signal compression operations and storage of sampled data prior to wireless transmission. Wireless transmission of the compressed data is performed to transfer data for resynthesis in a centrally-located data centre that is considered unconstrained in terms of computational resources, and hence power consumption. Once the signal is reconstructed, it is ready for clinical observation or CAD. At this point a clinician may make a decision about the patient, store the ECG signals for archiving, or decide to run the signals through a diagnostic decision support system.

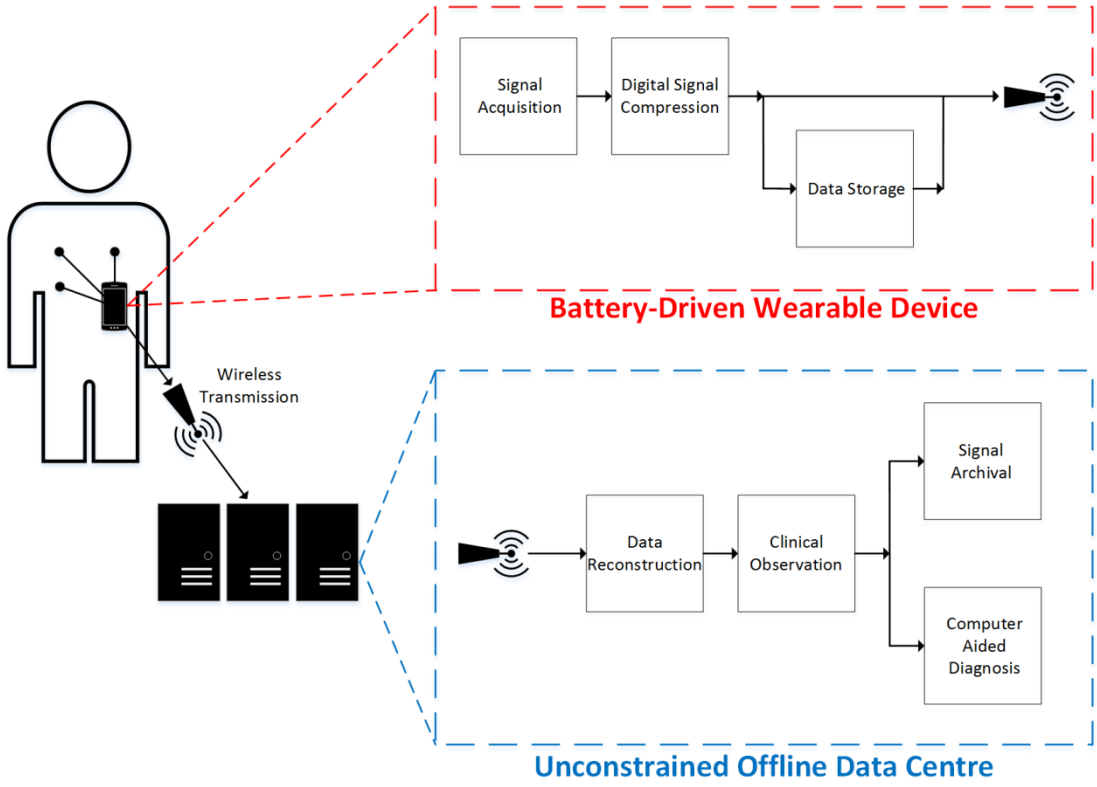

Figure 2.2: Diagram illustrating a typical wireless ECG BAN monitoring system.

The main aspects of novelty presented in this thesis lie in the signal acquisition and CS reconstruction strategies. Based on Figure 2.2 the two key performance requirements of a compression algorithm implemented in such a BAN are that the reconstructed signals are suitable for clinical usage; and that implementation of the compression algorithm does not significantly impact on the power consumption of the battery-driven wearable device. Therefore, both will be considered when evaluating the proposed techniques. The performance of the reconstructed signals are quantified with energy-based distortion metrics such as Percentage Root-mean-squared Difference (*PRD*) and the ability to keep below acceptable *PRD* levels for an increasing *CR*. On the other hand, in relation to diagnostically-relevant metrics that would be considered in CAD systems, an analysis is carried out on the effects of compression in terms of QRS detection accuracy. Here, Sensitivity (*SE*) and Specificity (*SP*) of QRS detection accuracy are investigated. Finally, in keeping with

the ultimate goal in the design of such a BAN system, that the proposed algorithm is an energy efficient implementation, the proposed methods are validated by analysing overall power consumption.

2.3. Signal Compression

Digital signal compression is the process of reducing the amount of data required to represent a signal, and is used in a wide range of applications. Typically, it is a two-stage process where an acquired signal is compressed, reducing the bit-length, before the inverse operation is performed to reconstruct the signal. Two fundamental types of compression exist: lossless and lossy compression.

Lossless compression is a completely reversible operation in the sense that it allows the signal to be compressed and reconstructed with no errors in the recovery of the signal. Generally, it aims to exploit any statistic redundancies that exist in a signal to reduce the total bit length. A common format of lossless compression is a codebook-based approach, where the most frequently occurring values are assigned a short binary code word and the less frequent values are assigned a longer binary description [30]. The sequence of values can then be encoded by linking the binary code word for each value in the signal. The original signal is recovered through the use of a look-up table or codebook to access the original information. Examples of such techniques include Arithmetic or Huffman coding and Lempel-Ziv (LZ) coding [31]. While these dictionary-based implementations have been applied with ECG signals in the literature, alternative methods have been investigated recently such as prediction-based lossless approaches [32-34] or block-sorting techniques [35].

While presenting a scheme where no information is lost, the obtainable levels of compression with lossless implementations are generally low. Lossy compression can achieve higher compression rates and is typically employed in applications where a certain level of signal distortion can be tolerated. As BANs are battery-driven, the use of lossy compression can provide essential extensions to battery life by reducing the costs associated with both data storage and wireless transmission.

Therefore, due to the potential to provide far superior compression gains in this thesis the main interest is in lossy compression techniques.

Lossy signal compression techniques can be further divided into the following categories: direct, transform and parameter extraction methods. Direct methods operate in the time domain and attempt to exploit redundancies within the signal. The advantage of direct methods is their low complexity operation and the small error associated with their implementation. Some of the common techniques in the literature that have been employed with ECG signals include the Coordinate Reduction Time Encoding Scheme (CORTES) [36], the Amplitude Zone Time Epoch Coding (AZTEC) [37], Delta Coding [38] and the Fan Algorithm [39]. The main limitation of such time domain compression approaches occurs when the signal is not sparse in the time domain, which is true for most biomedical signals, including ECG, particularly when collected in a BAN environment where noise may be present in the signal. This results in a degradation of signal reconstruction performance.

Transform techniques typically focus on the energy distribution of the signal in a domain other than the time domain, and can enforce sparsity through a linear orthogonal transformation of the ECG samples from one domain to another. Once in the transformed domain the compression algorithm is performed in that domain on the newly sampled signal coefficients. Transforms such as the Discrete Cosine Transform (DCT), Discrete Wavelet Transform (DWT) and the Fourier Transform are commonly used with ECG signals. Wavelet transform-based methods provide the most promising of transform methods for ECG signals due to their good localisation properties in both the time and frequency domain [17]. That is, having a fine frequency domain resolution and a coarse time domain resolution for low frequency signals and the opposite for high frequency signals. Examples of the associated compression techniques which exploit the redundancy in ECG signals offered by the wavelet transform are Set Partitioning In Hierarchical Trees (SPIHT) [16], CS, Joint Photographic Experts Group (JPEG) [40, 41] and thresholding methods which zero a fixed percentage of wavelet coefficients [42].

Parameter extraction methods do not allow for complete reconstruction of the signal and instead focus on the extraction of suitable characteristics of the signal. These characteristics are then used for classification purposes against prior knowledge of such signal features. Examples of such approaches include peak picking [43] and a long-term prediction approach [44]. A disadvantage of the parameter extraction methods is that the complete signal information is not available at reconstruction which is not suitable in many clinical scenarios. In general, transform-based methods are the primary focus of this thesis as they can maintain signal reconstruction performance at significantly higher *CR*s compared to direct techniques which is critical to the low-power functionality of a BAN. As noted previously, the SPIHT algorithm is used as a benchmark from the literature; the algorithm is described in more detail in Section 2.7.3.

2.4. Overview of Compressed Sensing

2.4.1. Overview

Compressed Sensing (CS) is a lossy compression technique that involves sampling of analog signals at sub-Nyquist rates while avoiding the issue of aliasing and allowing for effective compression of signals during the sampling process. To achieve this compression gain, CS benefits from two fundamental concepts: sparsity of the signal and incoherence. It is critical to CS that the signal being acquired is sparse in at least one domain i.e. the majority of entries in the signal are zero in that domain (such as the frequency domain). Many real world signals meet this criterion [18-21]. If a signal possesses this sparse property, an accurate reconstruction can be obtained by taking a relatively small number of measurements when sampling. The incoherence principle states that the measurement matrix, which is used to acquire the signal, must be incoherent with the dictionary that represents the signal sparsely. In general the more sparse the signal being measured, the more CS exploits the signal characteristics and the better the reconstruction quality.

The CS paradigm allows for the reconstruction of N samples of a signal x by the acquisition/transmission of only M samples of compressed data representing x, where $M << N$. Therefore the amount of data for storage or transmission is reduced, and the acquisition process itself is very simple. A trade-off does exist though, which is an increase in the computational complexity of the reconstruction. To recover the N samples of data from the original signal with only M samples of data available, a convex optimisation problem must be solved by finding the minimum l_1 norm solution [18]. However, this computationally expensive stage can be offloaded to another system which may have a much larger power budget or computational capability. This matches well with the BAN architecture considered in this thesis, where low-computation acquisition may be carried out on a resource-constrained wearable device, while reconstruction of the signal may be carried out on a central server to which the signal samples are wirelessly transmitted, and which is not computationally-constrained. This allows for the acquisition stage of the system to be low in computational complexity and power, which is a desirable feature of ambulatory monitoring.

2.4.2. Signal Acquisition and Reconstruction

The basic CS sampling model measures compressed data $y \in \mathbb{R}^M$, through linear combinations of an original signal $x \in \mathbb{R}^N$ with a random sensing matrix $\Phi \in \mathbb{R}^{M \times N}$, where $M << N$, demonstrated by (2.1) and Figure 2.3.

$$[y]_{M,1} = [\Phi]_{M,N} [x]_{N,1} \tag{2.1}$$

This random matrix Φ is known as a *sensing* or *measurement* matrix. Note that the sensing matrix remains constant and a version must be available at the decoding stage to enable reconstruction. The measurements y of length M are a representation of the signal x and are transmitted to the receiver in place of x.

Compression is achieved by having the number of elements *M* in *y* be much less than the original number of samples *N*.

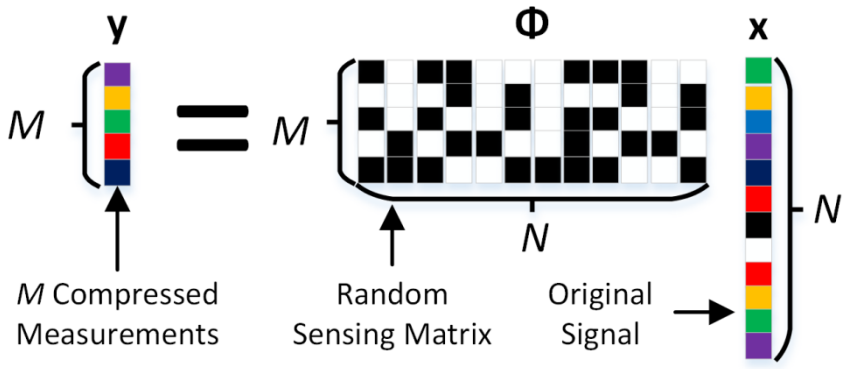

Figure 2.3: Illustration of CS acquisition process that uses a random sensing matrix to compute compressed measurements.

In order to reconstruct the original signal, a fundamental requirement of CS is that a sparse vector $\alpha \in \mathbb{R}^P$ exists as in (2.2), where α is a sparse representation of the reconstructed signal *x′* in a *sparse dictionary* $\Psi \in \mathbb{R}^{N \times P}$ as shown by Figure 2.4. The sparse dictionary is not necessarily a square matrix and the number of columns *P* can be increased above *N* to create overcomplete dictionaries which can enhance signal sparsity [45].

$$[x]_{N,1} = [\Psi]_{N,P}[\alpha]_{P,1} \tag{2.2}$$

Thus, the aim of signal recovery is to produce a sparse version of the vector α (of length *P*) from the vector *y* (of length *M*). The equation for *y* can then be expanded, as in (2.3):

$$[y]_{M,1} = [\Phi]_{M,N}[\Psi]_{N,P}[\alpha]_{P,1} \tag{2.3}$$

The equation to solve for α is in the form of the classic linear algebra problem $Ax = B$, where x is unknown. Since the set of equations consists of N unknowns in M equations, it can be a computationally expensive system to solve or approximate for sparse real world signals. The underlying theory behind CS states that, if the signal is sparse in the dictionary Ψ, the probability of solving the underdetermined set of linear equations in (2.3) is high [18, 19, 21].

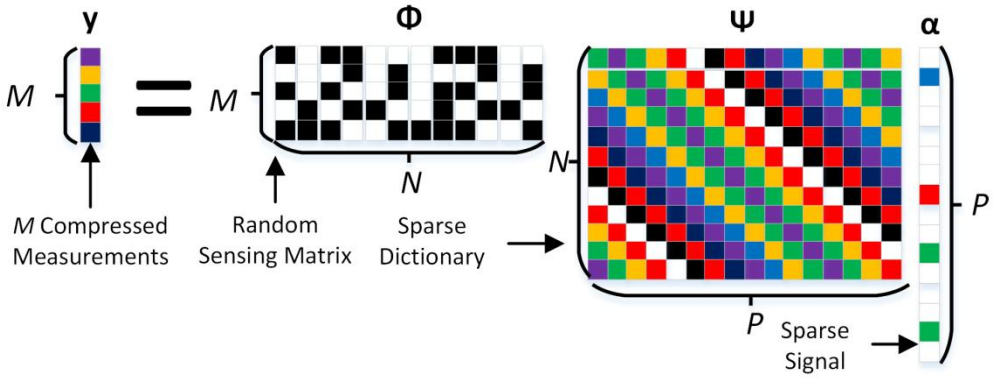

Figure 2.4: Illustration of CS reconstruction process that uses a sparse dictionary to sparsely represent the sampled signal.

The most common method for CS recovery is l_1 norm minimisation, where the vector with the minimum l_1 norm will correspond to a close approximation of the signal x, provided that enough measurements M have been taken. The sparse optimisation is described in (2.4) and the original x can then be recovered by using (2.2).

$$Min \ \|\alpha\|_1 \ \text{subject to} \ y = \Phi\Psi\alpha \tag{2.4}$$

There are also two important conditions that the sensing matrix Φ must satisfy to ensure accuracy and robustness in signal recovery: the *Restricted Isometry Property* (RIP) [20] and *mutual coherence* with the sparse dictionary [46]. Once

these properties are satisfied, CS ensures sparse recovery of the signal from the linear measurements y by solving the optimisation problem described by (2.4). The RIP condition for the sensing matrix is defined in (2.5):

$$(1 - \delta_k)\|\alpha\|_2 \leq \|\Phi\Psi\alpha\|_2 \leq (1 + \delta_k)\|\alpha\|_2 \tag{2.5}$$

where δ_k, the isometry constant of Φ, must be smaller than 1, and the smaller the value of δ_k the higher the probability of an exact reconstruction.

In practice, the RIP is difficult to verify and instead the coherence between the sensing matrix and the sparse dictionary can be measured. The coherence μ between the sensing matrix and the dictionary measures the largest correlation between any two elements of Ψ and Φ, defined in (2.6) [46]. Ideally the coherence will be small i.e. the two are incoherent, as the value for μ is effectively proportional to the number of measurements required when sampling.

$$\mu(\Phi, \Psi) = \sqrt{N}.\max_{1\leq k,j\leq N}\left|\langle \Phi_k, \Psi_j \rangle\right| \tag{2.6}$$

Commonly, the matrix Φ contains random components which are known to be independently identically distributed (i.i.d.). These i.i.d. entries are often chosen from a Gaussian or Bernoulli distribution and as such matrices have very small isometry constants, they are incoherent with very high probability, with a fixed dictionary Ψ. It has been shown that with using these i.i.d. entries only $M > C_k S\log(N/S)$ measurements are required for accurate reconstruction (where C_k is an empirical constant and S relates to sparsity and refers to the number of non-zero signal entries) [20]. Another common sensing matrix formation is a random discrete Fourier matrix. Fourier matrices also have a low isometry constant and require $M > C_k S log(N)^6$ measurements for accurate and stable reconstruction [47].

The l_1 norm is chosen ahead of l_0 norm and l_2 norm minimisation for a number of reasons [20]. Firstly, l_0 norm minimisation is guaranteed to exactly recover the original signal. However, this procedure is a non-convex optimisation problem and has no numerical solution as it is NP-hard [47]. The l_2 norm minimisation can be easily solved, but the reconstructed signal in this case is not sparse and carries a large error. Furthermore, l_2 minimisation returns non-sparse solutions as its energy is distributed across all the signal elements rather than being condensed into a small number of non-zero signal entries as is the case with l_1 norm minimisation.

While the convex optimisation Basis Pursuit (BP) [48] l_1 method effectively solves (2.4) and has been shown to give an accurate reconstruction, there are other greedy methods which have been implemented for signal reconstruction that iteratively solve (2.4) by making locally optimal decisions, with the aim of locating the non-sparse coefficients to enable quicker reconstruction [49, 50]. Algorithms such as Orthogonal Matching Pursuit (OMP) [51, 52] and Compressive Sampling Matching Pursuit (CoSaMP) [53] generate a solution that may not be globally optimal but is computed more quickly than the standard l_1 norm BP method.

2.5. Hardware Implementations

This section will now consider the practical implementation of CS. At this point, it is important to note that current approaches to CS are not as effective as other Nyquist rate state-of-the-art lossy compression techniques such as SPIHT, when considering only *CR* vs. signal reconstruction quality. Therefore, the choice of CS depends on its ability to deliver a low-power implementation of lossy compression. This section focuses on the design, development and evaluation of hardware capable of efficiently performing the sub-Nyquist acquisition which is currently the focus of significant research interest. Firstly, designs to realise the signal acquisition process of CS, whilst providing the power advantages CS promises over Nyquist sampling are reviewed. Secondly, several innovative strategies, applied

mainly with ECG signals, to improve the efficiency of CS implementations are discussed.

2.5.1. Sub-Nyquist Architectures

One of the early publications on the implementation of hardware in CS was Duarte *et al.*'s 2008 paper [54] proposing single pixel imaging using CS. Even at this early stage, Duarte *et al.* indicated that CS has the potential to "substantially increase the performance and capabilities of data acquisiton and processing" while still being aware of the "clear trade-offs and challenges" in its implementation. Since then, a number of significant advances have been made. Considering an application where the number of samples in an acquired frame N is equal to the sampling frequency, conventional sampling using an Analog-to-Digital Converter (ADC) will sample an input signal every $1/N$ seconds. The theory of CS suggests the ability to reduce the effective sampling rate of the acquisition process to a rate of $1/M$. In order to achieve this sub-Nyquist sampling in an efficient manner, the design of the signal acquisition system requires careful attention. The CS hardware architectures in the literature can be divided into two main approaches: Digital CS and Analog CS. Digital CS in essence consists of Nyquist sampling prior to linear multiplication with the sensing matrix to obtain the M measurements; hence, the sub-Nyquist sampling property of CS is not exploited. Analog CS is where M measurements are obtained in the analog domain prior to being digitised. Analog CS is an area of much current research interest since it has the greatest potential for delivering the low-power acquisition possibilities that CS offers.

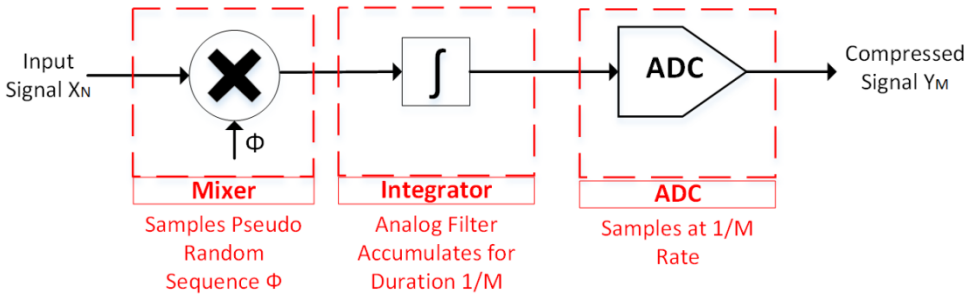

Figure 2.5: Block diagram of a typical RD scheme for CS consisting of a mixer, integrator and an ADC.

Analog CS hardware acquisition systems in the literature are generally variations of a Random Demodulator (RD) architecture [55-57] shown in Figure 2.5. The RD architecture is composed of three main stages: a mixer, an integrator and the ADC. The mixer or demodulator multiplies the input signal by a pseudorandom sequence. This step corresponds to the multiplication of an input signal x by a sensing matrix Φ. This multiplication must be continuous at a rate of at least the Nyquist rate. The integrator or low pass filter accumulates the output voltage of the demodulator. Functionality must exist to reset the value of the integrator after each sample is taken. These samples correspond to the compressed measurements. The final component is the ADC which now samples at a rate of $1/M$ instead of $1/N$ (which is the case with digital CS).

The RD is the fundamental architecture for Analog CS acquisition, however it has certain drawbacks when operating in isolation. The coherence in the matrix Φ is not at a minimum due to the fact that Φ is populated with many 0's and the rows of the matrix all use the same random sequence. As the compression increases, the coherence will also increase. A common practical advancement for CS comprises parallel blocks of the RD, known as the Random Demodulator Pre-Integrator (RMPI) shown in Figure 2.6. This architecture allows for the signal to be multiplied by different random sequences thereby increasing the incoherence as each parallel block produces its own compressed measurements. RMPI provides a further reduction in the ADC requirements with the trade-off being that more accumulators and mixers

operating at the Nyquist rate are required. The RMPI architecture is the most commonly employed version of Analog CS and many implementations exist in the literature [58-61].

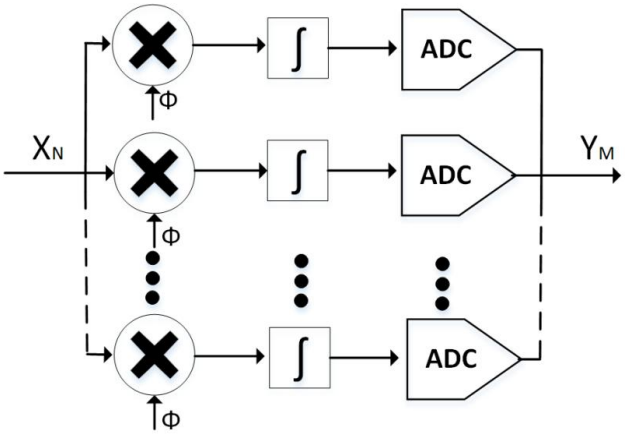

Figure 2.6: Block diagram of RMPI architecture showing mixers, integrators and ADCs operating in parallel channels.

Chen *et al.* presented a hardware implementation of CS with a view to implementing the algorithm on wireless sensor nodes [60, 61], providing a detailed comparison of both Analog CS and Digital CS implementations at the encoder stage of CS [60]. The analog encoder consisted of an RMPI architecture. The digital encoder model consisted of an amplifier and ADC followed by an accumulator and XOR (the XOR is used with the carry-in of the accumulator for the multiplication of the signal by the ±1 entries in the measurement matrix) to generate the compressed measurements. The digital encoder first acquires the signal using a Nyquist rate ADC before obtaining the smaller number of measurements through multiplication by the measurement matrix. Their analysis includes a comparison of the power consumption between the digital and analog implementations. Each implementation was tested with the aim of obtaining a target *CR* of 10. The power consumptions of each implementation shows the digital implementation to be the preferred approach for wireless sensor applications. The main problem with the analog design that Chen *et*

al. presented is that the parallel channels are designed so that the number of channels equals the number of desired measurements. Mamaghanian *et al.* noted that the implemenentation uses very high frequency mixers which results in the analog implementation consuming more power than is neccessary, and proposed their own low-power analog design for CS in [62]. This paper compares a traditional RMPI architecture with digital CS (Nyquist sampling) and their own Spread Spectrum Random Modulation Pre-Integrator (SRMPI). The SRMPI is an innovative design which introduces a mixer before the traditional parallel RD blocks as illustrated in Figure 2.7. This allows only a single mixer to operate at the Nyquist rate but then allows the other mixers in the design to reduce their sampling rate, therefore providing a more efficient implementation compared to the RMPI. The results in [62] demonstrated that the SRMPI approach consumes less power than the Nyquist sampling implementation. The RMPI can reduce the power consumption by 63% whereas the SRMPI can reduce it by 75% when compared to Nyquist sampling in an 8 channel architecture. Therefore, the SRMPI architecture is considered the benchmark for Analog CS acquisition in this thesis and will be used later in Section 6.2 when the overall power consumption of the proposed algorithms are assessed.

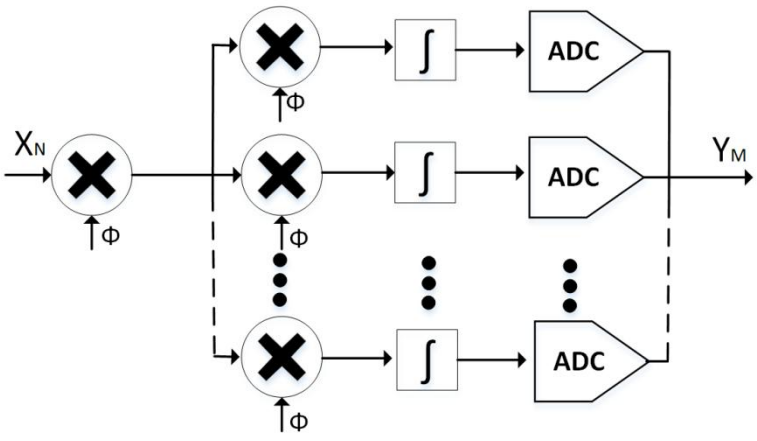

Figure 2.7: Block diagram of the SRMPI architecture proposed by Mamaghanian *et al.*.

2.5.2. Low-Power Sampling Strategies

While the previous sub-section deals with the analog acquisition of compressed measurements in CS, strategies exist in the literature to enable an additional extension in *CR* prior to wireless transmission. For example, Mamaghanian *et al.* [63] proposed a redundancy removal technique to eliminate the significant redundancy inherent in adjacent compressed ECG measurements. The redundancy removal operation (which will be referred to henceforth as Standard Redundancy Removal (SRR), to distinguish it from a novel redundancy removal technique that will be presented in this thesis) focuses on the minimal variance of the compressed measurements of consecutive ECG frames due to a fixed sensing matrix being used and the pseudo-periodic nature of ECG signals. The approach computes the difference between consecutive measurements and performs Huffman coding on the difference, providing an increase in the level of compression. The nature of their approach allows for a lossless extension in *CR* in terms of signal reconstruction performance and has been further implemented in many CS-based implementations [64, 65].

A second common modification to CS aquisition is the design of different sensing matrices for improved signal reconstruction quality or to reduce the memory footprint or power consumption of sampling. As stated earlier, CS algorithms typically create their measurement matrix using random methods such as Bernoulli or Gaussian distributions. Ansari-Ram *et al.* proposed a new method for creating the sensing matrix which improves the *PRD* vs. *CR* trade-off for ECG signals [66]. The proposed non-uniform binary sensing measurement matrix is designed specifically for ECG signals and takes into account the region of interest in a healthy ECG signal (i.e. the QRS complex). A drawback of this method is the requirement to alter and transmit the sensing matrix in order to ensure it aligns with the QRS complex for each frame. It could also be expected that abnormalities in the ECG would significantly affect this approach. Dixon *et al.* also considered additional sensing matrices (which are variations of Bernoulli random sampling): Toeplitz, Circular and

Triangular for low complexity acquisition [67]. The use of sparse binary sensing matrices, which only contain a small number of non-zero elements in each column for efficient hardware computation, was proposed in [63] and further implemented in [64, 68]. Sparse binary matrices allow for a simple hardware arrangement and further reduce the complexity of CS signal acquisition. However, Polania *et al.* concluded in [64], whilst acknowledging the potential of the sparse binary matrices, that Bernoulli distributed sensing matrices outperform the sparse binary sensing matrices in terms of signal reconstruction quality. Therefore, in this thesis, the sampling considered is an Analog CS design where the sensing matrix entries are Bernoulli distributed ±1 entries.

2.6. Compressed Sensing for ECG Compression

After the initial surge of interest in the CS paradigm many researchers investigated sparse representations of ECG signals in domains other than the time domain, mainly using wavelet transforms. The use of such a wavelet basis to create sparse representations of ECG signals has been commonly employed with CS [63, 66, 69-72]. There are, however, well known performance limitations with the use of wavelet transforms and many new strategies have been explored in the literature. For example, methods focusing on the block sparsity or common support of wavelet coefficients, modifications of reconstruction algorithms and Dictionary Learning (DL) techniques are approaches that have been proposed or employed; such techniques will be further analysed in this thesis. This section will outline the current state of CS ECG compression in the literature.

2.6.1. Wavelet-Based Approaches

Mishra *et al.* [71, 73, 74] quantified the selection of the best wavelet basis function for compressing ECG using CS. The following wavelet families were tested and evaluated based on their performance and their ability to create sparsity in an ECG signal: Coiflets, Haar, Symlets, Daubechies, Biorthogonal and Reverse Biorthogonal. Signal reconstruction quality was analysed at numerous *CR*s ranging

from 2 to 10, and at each *CR* the best wavelet family was identified. For the majority of *CR*s it is suggested that the Reverse Biorthogonal wavelets rbio3.7 and rbio3.9 perform best with ECG compression based on the *PRD* values.

Chae *et al.* [72] investigated the performance of CS using a wavelet basis, against a thresholding DWT (TH-DWT) compression method on noisy signals (typical of an ambulatory ECG environment) and artefacts consistent with the occurrence of body movement. Initially, the two approaches were compared on non-noisy signals. The TH-DWT method achieved a *PRD* of 9% at a *CR* of 5 whereas the CS approach obtained a *CR* of 1.67 for the same *PRD*. On the noisy signals, the TH-DWT method maintained a constant *PRD* up to a *CR* of 2.5 before decreasing slightly, while CS underwent a sharper fall, with performance decreasing from a *CR* of 1.25. The authors in [72] acknowledged that CS is attractive due to its encoder simplicity but recommend exercising caution when considering CS for ECG compression, where noisy signals would be expected.

In [67] Dixon *et al.* added white Gaussian noise to yield a Signal-to-Noise Ratio (*SNR*) of 80 dB (a level deemed typical of an analog front-end ECG sensor). The results present the *SNR* of the reconstructed signal at *CR*s up to 16 and the ability of the reconstructed signal to maintain this *SNR* at each *CR* is reported. As expected with higher levels of sparsity, higher *CR*s are achieved while maintaining the *SNR*. Another area which was explored by Dixon *et al.* was CS reconstruction using greedy algorithms and their performance with ECG compression [67, 75, 76]. They used ECG compression to compare BP convex optimisation with four greedy reconstruction algorithms and tested them on ECG signals for accuracy, reliability and computational time at different sparsity levels [75]. The algorithms tested were BP convex optimisation [48], OMP [52], CoSaMP [53, 77], Regularised Orthogonal Least Squares (ROLS) [78] and Normalised Iterative Hard Thresholding (NIHT) [79]. The results showed OMP to be the best algorithm for ECG CS reconstruction based on its reconstruction accuracy and computational time. This appears to be the only paper in the CS literature that advocates the use of OMP over BP for optimal reconstruction accuracy.

The final wavelet-based implementation of interest demonstrates the potential of applying CS in BANs. Mamaghanian *et al.* [63] explored the role of CS in wireless BANs for acquisition and compression of ECG signals. The technique was compared against a DWT-based lossy compression technique in terms of energy consumption on a Shimmer device [80], which was being used as a sensor node. The energy consumption analysis focused on each of the three different stages: sampling, compression and wireless transmission. The battery lifetime of the sensor node was compared in three scenarios: no compression, DWT-based compression and CS-based compression. The results demonstrated a 37.1% extension in the battery life for the CS approach over the DWT-based approach. This was mainly due to the low complexity nature of CS's simultaneous sampling and compression operations. However, reconstruction with a Daubechies wavelet basis resulted in their CS implementation being outperformed by the DWT-based compression algorithm from a signal reconstruction quality perspective. The energy consumption of wireless transmission for the DWT-based approach is lower than that of CS as higher *CR*s are achieveable within the acceptable *PRD* values. The ECG signals were initially stored on the memory of the Shimmer device and hence it could be argued that a more thorough analysis should take into account the power consumption of compression occuring at the sensor prior to the ADC. This analog front-end power profiling analysis has since been performed by Mamaghanian *et al.* [62] with the proposed SRMPI CS acquisition architecture that was described in the previous sub-section.

It is interesting to consider the performance of CS while keeping signal distortion within the constraints of what would be considered "acceptable" in clinical applications. Zigel *et al.* [81] classified different *PRD* ranges in terms of the quality of the reconstructed signal as perceived by a clinician. This research demonstrated that *PRD*s of less than 9% for ECG signals are either "good" or "very good" quality. Based on the classification proposed by [81], the literature where wavelets are employed has shown that CS can generally reach *CR*s of around 4 before exceeding this 9% *PRD* measure. Ansari-Ram *et al.* [66] obtained a *PRD* of 8.527% at a *CR* of 4 and [63] demonstrated an average *PRD* of less than 9% at a *CR* of approximately

3.5. In [82], the authors achieved a *CR* of around 8 before going above 9% *PRD*, however the method for calculating the *CR* differs from the traditional methods. Casson *et al.* [69] obtained a *PRD* of 2.6% at a *CR* of 2.5, while Chae *at al.* [72] obtained a *PRD* of 9% at a *CR* of 2.5.

2.6.2. Advanced CS Approaches

The focus will now turn to more recent CS implementations which can significantly improve on the performance offered by reconstructions relying on sparse wavelet domain representations. Pant *et al.* proposed an algorithm to minimise the l_p pseudo norm of the second-order difference of ECG signals [83]. The technique then utilises patient-specific dictionaries to optimise the reconstruction quality of the l_p norm. Zhang *et al*. [68] employed Block-Sparse Bayesian Learning (BSBL) to reconstruct raw foetal ECG signals. The method exploits the intra-block correlation that exists in the sparsity pattern of wavelet coefficients in foetal ECG signals and then uses BSBL in the reconstruction. Polania *et al.* [64] modified two reconstruction algorithms for CS with ECG signals: Iterative Hard Thresholding (IHT) and CoSAMP. Their Modified Model-Based approaches (MMB-IHT and MMB-CoSAMP) exploit prior knowledge of ECG signals [64]. The two key attributes of their model-based approach to CS is that the algorithms exploit both the high fraction of common support between the wavelet coefficients of consecutive ECG signals and the connected subtree structure of the coefficients formed by the largest coefficients across different scales in the wavelet transform output.

Polania *et al.* compared their compression schemes (MMB-IHT and MMB-CoSAMP) to three existing CS approaches and employed SPIHT as a baseline [64]. Their modified reconstruction algorithms (MMB-IHT and MMB-CoSAMP) outperformed the following in terms of reconstruction quality: (i) Mamaghanian *et al.*'s use of a Daubechies wavelet with BP [63], (ii) a BSBL technique for ECG signals as proposed by Zhang *et al.* [68] and (iii) an overcomplete multi-scale wavelet dictionary created with DL [84]. Polania *et al.*'s two approaches achieved better performance than all reported CS-based approaches though they were still

outperformed by SPIHT in terms of reconstruction performance. However, since the performance of the MMB-IHT and MMB-CoSAMP algorithms exceeded alternative CS approaches, they are considered the state-of-the-art for CS and are used as a baseline in this thesis when evaluating the proposed CS implementations.

2.7. Experimental Methods

2.7.1. ECG Database

The ECG database employed in these experiments is the widely used MIT-BIH Arrhythmia database [85]. The database consists of forty-eight patient records, each thirty minutes in length, recorded from both in-patient and out-patient subjects. The recordings are sampled at 360 Hz with 11 bit resolution and were manually annotated by at least two cardiologists. This database was chosen as it is widely used in the literature and it contains a variety of both routine and rare arrhythmia phenomenon which makes it useful for clinical analysis.

In order to analyse the performance against current state-of-the-art CS [64], the first channel from the following patient records used in [64] (a test set originally proposed by [17]) were selected: 100,101,102,103,107,109,111,115,117,118 and 119. This data test set is useful for a direct comparison with [64], but also ensures robust testing of the proposed algorithm since the data contains a mixture of normal ECG, a variety of arrhythmias and abnormal heartbeats. This is the test set used in the majority of tests in the thesis (unless for scenarios where it is explicitly stated otherwise). An MIT-BIH database directory [86] is provided which contains annotations of different physiological events in each record and this will provide a basis for visual assessment of reconstruction performance. Additionally, the annotated QRS complexes allow for evaluation of QRS detection accuracy performance.

As the target application is ambulatory ECG monitoring in a BAN, an additional MIT-BIH Noise Stress Test database will also be used for assessment [87].

The database contains increased levels of noise and morphologies such as baseline wander, electrode motion and muscle artefacts. The ability of the proposed algorithm to maintain performance in such a testing scenario is therefore analysed.

2.7.2. Performance Metrics

The metrics used to quantify the performance of the compression algorithms in this thesis can be divided into different groups. Compression Ratio (*CR*) and Wireless Transmission Savings (*WTS*) are used to evaluate the savings made in terms of data storage and transmission requirements. The quality of the reconstructed signals is also evaluated using signal energy-based distortion metrics known as Percentage Root-mean-squared Difference (*PRD*). Additionally, an analysis focusing on the diagnostic value of reconstructed signals through the accuracy of QRS detection and HRV estimation is employed.

Compression Ratio (*CR*): *CR* is a measure of the reduction in the number of bits needed to represent the original signal *x*.

$$CR = \frac{(N)(B_O)}{B_T} \tag{2.7}$$

In (2.7), N is the frame size of the original signal, B_O is the original bit resolution (11 for this database) and B_T is the number of transmitted bits representing the compressed signal.

Wireless Transmission Savings (*WTS*): *WTS* is a measure of the relative percentage reduction in wireless transmission costs of the proposed algorithm over a benchmark compression algorithm [65]. The savings are a function of the *CR* obtained by each algorithm and are calculated as:

$$WTS = \frac{|CR_{PA} - CR_{BA}|}{CR_{BA}} . 100 \tag{2.8}$$

CR_{PA} is the *CR* obtained by the proposed algorithm and CR_{BA} is the *CR* obtained by the benchmark algorithm.

Percentage Root-mean-squared Difference (*PRD*): *PRD* is a measurement of the distortion between the reconstructed signal *x'* and the original *x.* There are two commonly used definitions for *PRD*:

$$PRD = \frac{\|x - x'\|}{\|x\|} . 100 \tag{2.9}$$

$$PRD^* = \frac{\|x - x'\|}{\|x - \bar{x}\|} . 100 \tag{2.10}$$

Note that in (2.9) and (2.10) || || represents the Euclidean or l_2 norm. With the distortion metric *PRD* in (2.9), the mean of the signal $\bar{x}$ is not removed from the original signal. The non-removal of the mean can result in a DC bias that artificially lowers the distortion results. *PRD** is more commonly used in practice and is the main signal distortion metric used in this thesis. *PRD* is only presented in this thesis for comparison purposes with the state-of-the-art CS results in [64], which also used *PRD*.

In order to further evaluate the lossy compression performance of CS, clinical metrics which ensure that diagnostic integrity is maintained are also used. These metrics provide an understanding on how CS preserves clinically-relevant information in the ECG signal, in contrast to the objective energy-based *PRD* metric which is primarily a measure of time domain signal distortion. In this work, the focus is on the accuracy of QRS detection and HRV-related metrics. These clinical metrics

were chosen due to their wide use in clinical practice with regards to the detection and prevention of many common cardiovascular conditions [88, 89].

The MIT-BIH Arrhythmia database provides clinical annotations for the QRS complexes in the ECG signals, and Sensitivity (*SE*) and Specificity (*SP*) are used to analyse the accuracy of QRS detection. To detect the QRS complexes in the reconstructed signals, Afonso *et al.*'s beat detection algorithm was employed [90]. The algorithm's detected beats were compared with the clinically-annotated QRS data from the MIT-Arrhythmia database. Finally, an overall measure of QRS detection performance is given by the standard area under curve (*AUC*) of the receiver operator characteristic (ROC) plots which combines *SE* and *SP*. The definitions of *SE* and *SP* used here are as follows:

Sensitivity (*SE*): *SE* relates to the ability to correctly detect annotated QRS complexes:

$$SE = \frac{TP}{TP + FN} .100 \tag{2.11}$$

Specificity (*SP*): *SP* measures the ability to correctly identify non-QRS complexes:

$$SP = \frac{TP}{TP + FP} .100 \tag{2.12}$$

where true positives (*TP*) are the number of QRS complexes correctly identified, false negatives (*FN*) are the number of missed QRS complexes and false positives (*FP*) the number of incorrectly identified QRS complexes.

In addition, several Heart Rate Variability (HRV) metrics proposed in [24], were also analysed: *meanNN*, *SDNN*, *LF/HF* ratio and *HF power*. These metrics capture both time and frequency domain characteristics that are representative of the ECG signal. *MeanNN* is the average value and *SDNN* is the standard deviation of the Normal-to-Normal (NN) QRS intervals. The *LF/HF* Ratio is measured in the frequency domain and is the ratio of low frequency (LF) to high frequency (HF) components. *HF power* is the measured power in the HF range of 0.15 to 0.4 Hz. In order to evaluate the impact of compression, these metrics are evaluated on the original and the reconstructed signal and the Percentage Similarity (*PSim*) is calculated:

$$PSim = 100 - \left(\frac{|y - \bar{y}|}{y} .100\right) \tag{2.13}$$

where y is the value of the metric from the original signal and $\bar{y}$ is the value of the same metric on the resulting signal after reconstruction.

2.7.3. Benchmark Algorithms and Broad Performance Goals

Throughout this thesis the performance of the proposed CS methodology will be compared against both existing CS-based approaches and state-of-the-art non-CS methods for lossy ECG compression. The previous section introduced several high performing CS-based implementations that will be used as the benchmark algorithms throughout the thesis, in particular the MMB-IHT, MMB-CoSAMP and the BSBL approaches. Here, a suitable non-CS benchmark is considered.

In terms of state-of-the-art lossy compression implementations the Set Partitioning In Hierarchical Trees (SPIHT) algorithm is chosen for evaluation. SPIHT is a lossy ECG compression technique proposed by Said and Pearlman [17]. The core principles of SPIHT are derived from the Embedded Zerotree Wavelet

(EZW) technique, which exploits the fact that wavelet coefficients from different sub-bands exhibit a temporal relationship between each other. SPIHT coding effectively orders the output of the DWT according to the significance of each coefficient by scanning the coefficients according to a threshold. The output of the DWT is generally quantised before SPIHT coding techniques are applied to construct a compressed binary stream. In order to achieve this encoding, SPIHT maintains the coefficients in lists and recursively performs sorting and refinement passes on these lists to encode the wavelet information. As the bits are encoded in order of significance, simple control of *CR* can be obtained by truncating the bit stream at a desired *CR*.

For this thesis, two implementations of SPIHT are used, each with a different DWT basis. The main SPIHT implementation used was employed in [91, 92] where a CDF9/7 biorthogonal wavelet was used for the DWT using 8 levels of decomposition, and quantisation of the DWT output using 8 bits. This provides the best *CR* vs. *PRD* performance for SPIHT. The second implementation uses a CDF5/3 integer to integer wavelet transform. This implementation also uses an 8 level decomposition but removes the need for signal quantisation. The CDF5/3 wavelet is employed in certain scenarios due its computational efficiency, despite presenting a slight loss in compression performance. Therefore this implementation was used to provide a fairer comparison when comparing overall power consumption with CS-based approaches as the most efficient CS implementations were investigated throughout.

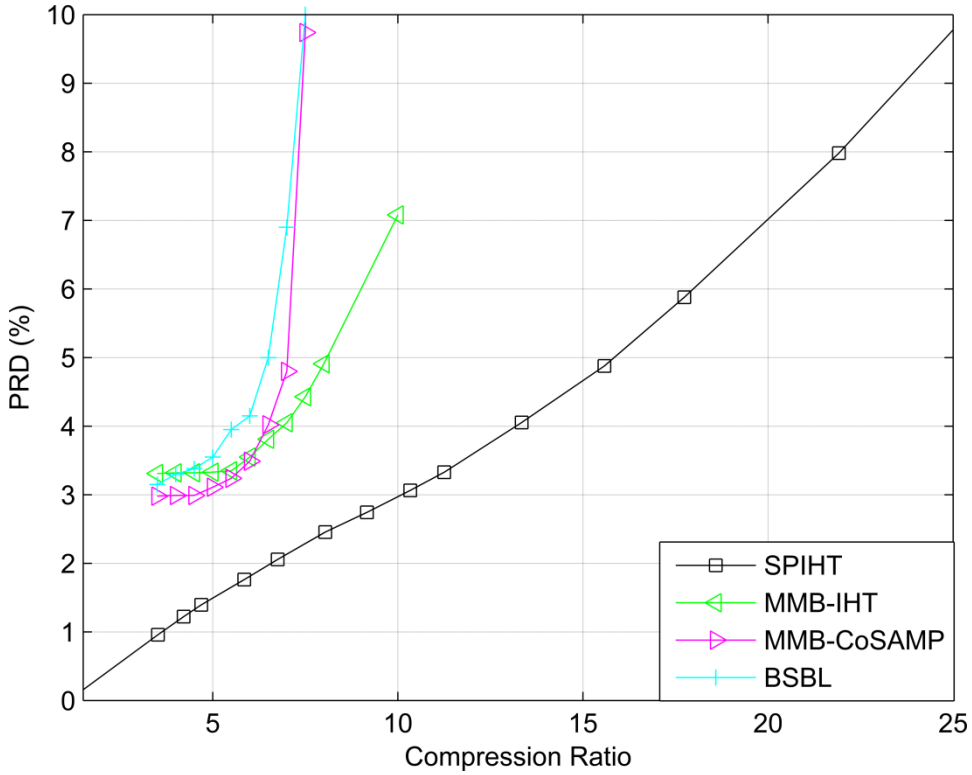

Figure 2.8: Summary of the performance of existing CS implementations in the literature.

Figure 2.8 highlights the current state of the chosen benchmark CS implementations in terms of ECG signal reconstruction performance, when compared with SPIHT. The following conclusions can be drawn. Firstly, despite the advantages CS offers in terms of efficiency of encoder implementation, CS remains outperformed by state-of-the-art lossy compression techniques particularly for applications where the acceptable level of distortion increases. Therefore, the approaches proposed in the thesis have the following broad performance objectives: 1) The proposed algorithms should outperform current state-of-the-art CS techniques, in particular, maintain *PRD* performance at higher *CR*s which is currently limiting CS-based approaches; and 2) in so doing, the proposed algorithms should close the existing performance gap that currently exists between CS and SPIHT, while providing a system-level energy-saving relative to SPIHT.

2.8. Summary

This chapter has provided an overview on BANs and the use of ECG signals in an ambulatory monitoring environment. A discussion of signal compression and the importance of its implementation in terms of reducing data storage and wireless transmission power requirements was given. CS was then comprehensively reviewed. From this, the current benchmark algorithms for CS implementations with ECG signals were identified; these algorithms will be used as a basis for comparison with the novel approaches developed in this thesis. Finally, the ECG database and associated performance metrics used throughout were presented.

CHAPTER 3

Compressed Sensing Methodology

3.1. Introduction

In the previous chapter, the current state-of-the-art in Compressed Sensing (CS) implementations was reviewed. This highlighted the significant potential of CS for employment in low-power budget applications such as BANs where battery resources are limited. The analysis highlighted that there is scope for improvement in the functional performance of CS, which would make it an even more suitable choice for low-power BANs. In particular, there is scope for improvement over current state-of-the-art CS approaches in terms of signal reconstruction quality at higher *CR*s, a limiting factor in current CS implementations. This chapter proposes an approach that incorporates three main optimisations that aim to improve the performance of CS when used to compress ECG signals: (i) scalar quantisation at the encoder to maximise compression; (ii) an improved alternative to the Standard Redundancy Removal technique discussed in Chapter 2 – Redundancy Removal with Mean (RRM); and (iii) reconstruction using a dictionary trained with patient-specific ECG signals. For notational purposes, the proposed approach that employs RRM during signal acquisition and uses Dictionary Learning (DL) in reconstruction will be referred to as RRM-DL for the remainder of this thesis. The RRM-DL system block diagram is presented in Figure 3.1. Each component in the system will be described in detail in the following sections.

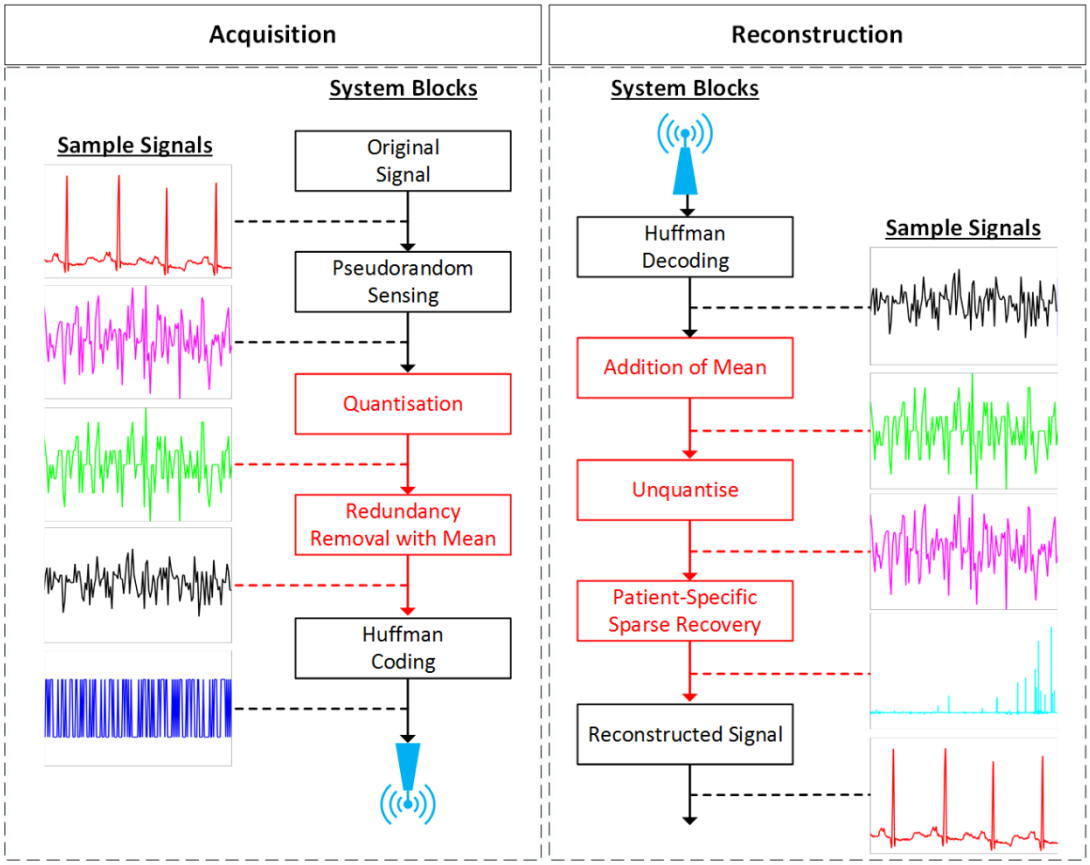

Figure 3.1: Block diagram showing each stage of the proposed RRM-DL implementation. The main optimisations of the RRM-DL method that are covered in detail in this chapter are indicated by red blocks.

There are many performance goals of the proposed RRM-DL algorithm. This chapter will focus on the objective of outperforming existing CS-based technologies in the literature (identified in the previous chapter) in terms of signal reconstruction quality. To achieve this, the proposed technique employs explicit quantisation at previously untested bit levels with a novel redundancy removal module to maximise the *CR* gains. Improving *CR* vs. *PRD* performance particularly at higher *CR*s is enabled by the use of highly overcomplete patient-specific dictionaries.

Moreover, the proposed RRM-DL technique will also allow the reconstructed signals to maintain suitable diagnostic integrity at higher *CR*s, which is of interest in CAD systems [24, 91, 93, 94]. The investigation of the effects of the RRM-DL

approach on the diagnostic integrity of the signals in CAD systems will be discussed in detail in Chapter 5.

3.2. Signal Acquisition

3.2.1. Scalar Quantisation

Quantisation is a relatively simple compression operation that maps a sequence of numbers $y(N)$ to a smaller set of quantised values $y_Q(N)$. Quantisation occurs as signals are digitised and provides an estimate of the signal by mapping each sample to a code word based on its voltage. The decoder then reconstructs the signal $y(N)$ based on this quantised value. This process introduces an error into the signal as the original sample can only be estimated and is therefore not guaranteed to be correctly reconstructed. In general, the accuracy of this decoding mechanism depends on the bit resolution of the quantised signal being used, as higher bit resolutions allow for a smaller step size and a more accurately defined signal representation.

For the work presented here, quantisation of the pseudorandom compressed measurements takes place prior to the redundancy removal module and entropy coding stage of the algorithm as highlighted by Figure 3.1. Mamaghanian *et al.*'s original implementation effectively reduced the bits representing the difference to 9 bits [63] and Polania *et al.* used 8 bit scalar quantisation [64]. However, quantisation below 8 bits with a redundancy removal architecture has not been investigated thus far. Previous work by the authors in [95], focusing solely on the quantisation of compressed measurements, suggested the suitability of quantising with fewer than 8 bits. Quantisation at lower bit levels can boost performance by extending the *CR* without significantly impacting *PRD*. Moreover, where the goal is decision support for a diagnostic function, quantisation is appropriate for applications where some degradation in signal fidelity is acceptable as long as diagnostic integrity is maintained, such as QRS complex extraction applications [95].

The work presented here investigated linear quantisation of CS measurements at 8, 7, 6 and 5-bit resolutions to increase the compression gain, therefore Huffman codebook sizes C of 2^9, 2^8, 2^7 and 2^6 were tested. For example, by quantising the measurements at 8 bits, the difference between subsequent measurements (the output of a redundancy removal module) can then be represented by a maximum of 9 bits. A similar principle exists when employing the proposed redundancy removal with mean method (described in detail in the next sub-section) where the mean vector and CS measurements are also quantised at 8, 7, 6 and 5-bit resolutions.

3.2.2. Redundancy Removal with Mean

CS-based implementations for ECG compression typically employ a redundancy removal module, such as the SRR proposed in [63] which is combined with an entropy coding technique such as Huffman coding during acquisition. SRR is based upon the principle that when a fixed sensing matrix is used, consecutive ECG frames yield similar measurement vectors. Therefore, a vector containing the difference between consecutive measurements is further processed instead of the compressed measurements themselves. The Probability Density Function (PDF) of values present in the difference vectors of ECG signals for the SRR method in Figure 3.2(a) indicates that the distribution of the difference values is non-uniform and hence further compression can be achieved by applying Huffman coding to these difference measurements which can help remove redundancy in the data. The Huffman dictionary is created offline as described below, and remains fixed.

This work proposes to modify the SRR method by subtracting the compressed measurements vector from a *mean measurement vector* instead of the previously sampled measurements as in SRR. The mean can be calculated offline along with the Huffman dictionary; therefore the architecture incurs no additional computational costs during signal acquisition in an ambulatory monitoring environment compared to approaches that utilise the SRR technique, as offline Huffman dictionary creation is also a necessary step with the SRR algorithm.

In order to validate the RRM approach, 11 patient records (divided into frames of N = 1024 non-overlapping samples) were extracted from the MIT-BIH Arrhythmia database [85] and, using a fixed sensing matrix with Bernoulli entries, Huffman codebooks with 512 (2^9) entries were created for varying numbers of measurements M. The training phase, which involved creating the Huffman dictionaries for both approaches and calculating the mean measurements vectors for the RRM approach, used 90% of the signals from each record. The mean measurement vector $\bar{y}$ is calculated by obtaining the mean of each element for the set of training signals as shown in (3.1):

$$\bar{y}_i = \frac{\sum_{j=1}^{T} Y_{t\langle i,j\rangle}}{T} \tag{3.1}$$

where $\bar{y}$ is the mean vector consisting of M entries, Y_t are the training measurement vectors, T is the number of measurement vectors used in training, i is the index of each vector and j indexes through all the training signals from 1 to T.

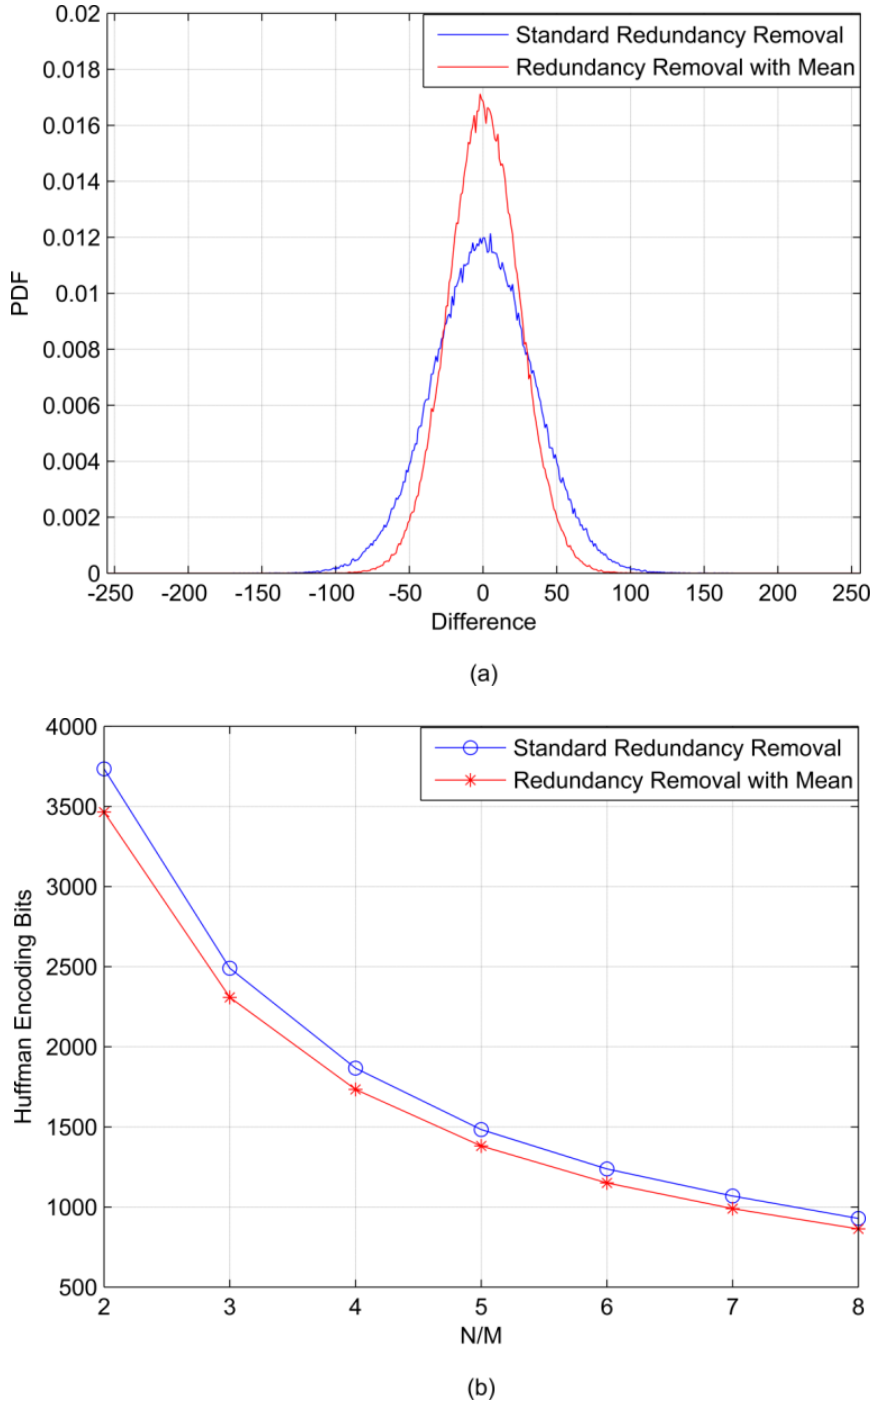

Figure 3.2: Comparison between the SRR and RRM techniques showing (a) PDF curves of the difference values and (b) the average Huffman encoded output bit lengths at different numbers of measurements *M* relative to the fixed frame size *N* of 1024.

Testing utilised the remaining 10% of the signals and Figure 3.2 shows the average results from each record. For each record a tenfold cross validation was

employed to ensure the complete record was tested, whereby 10% of the signals were used for each test. The PDF curves in Figure 3.2(a) show the probability of each difference value in the range [-255,256] (representative of the 9 bit Huffman codebook), based on their probability of occurrence during the training phase. By computing the difference between the measurements vector and the mean of the measurements in RRM, the PDF is narrower than when the difference is calculated by subtracting from adjacent compressed measurements, as performed by SRR (shown in Figure 3.2(a)). This is because the mean vector has been calculated over multiple measurement vectors and on average can provide a closer approximation than the previous measurements alone. This allows for more difference values to be clustered near 0, thereby offering a lower overall variance and a greater potential to remove redundancy when using RRM. Figure 3.2(b) shows the average number of bits after Huffman encoding is performed. As expected, this demonstrates a consistent reduction in the number of encoded bits using the novel RRM approach proposed here. This represents an advance over the current state-of-the-art SRR method for CS ECG acquisition by providing an increase in *CR* without significantly impacting on the signal quality relative to employing SRR. This is because the Huffman encoding itself is a lossless operation and therefore any error introduced is caused by the signal quantisation which would be at the same rate for both algorithms.

3.3. Reconstruction with Dictionary Learning

3.3.1. Background

Dictionary Learning (DL) is a process of constructing a dictionary of signal components known as *atoms* from a set of training signals [96]. The atoms provide good estimates of the training data allowing the DL algorithm to exploit signal characteristics and enable an optimised sparse representation. DL operates as follows. Given a set of J training signals $X_t = [x_1,x_2,x_3...x_J]$, where each signal is of length N (and where x_1,x_2 etc. are the columns in the matrix X_t), the objective is to find an optimal dictionary Ψ (NxP matrix) that can achieve sparse representations of

the class of signal in a set of coefficients β (PxJ matrix). The dictionary Ψ is described as overcomplete when $P >> N$.

In this thesis, CS signal reconstruction uses l_1 norm minimisation, which leverages the sparsity of the original signal x with respect to a dictionary Ψ. The use of highly overcomplete dictionaries, trained with patient-specific signals, allows ECG signals to be represented with very few redundancies [83]. Therefore, these dictionaries can be considered ideal for l_1 norm minimisation and CS signal recovery. In fact, DL has previously been employed with CS [83, 84, 97, 98]. In contrast to previous DL implementations, however, the proposed approach uses overcomplete dictionaries of a much higher order with respect to N for CS signal reconstruction. For example, Pant *et al.* used a square dictionary [83] and the dictionaries employed in [84] were overcomplete by the order of approximately 4N. As dictionary creation and signal reconstruction are performed offline in an unconstrained power environment with no impact on the BAN, in this work overcomplete dictionaries were constructed with up to 100N atoms.

Generally, the aim of DL algorithms is to iteratively improve the dictionary by achieving increasingly sparse representations of the training signals and updating the dictionary based on the current sparse representations. The common method for constructing the dictionary solves the optimisation problem:

$$\min\left\{\left\|X_{t_{N,J}} - \Psi_{N,P}\beta_{P,J}\right\|_F\right\} \text{ subject to } \|\beta_i\|_0 < S\ (i = 1,..,K) \tag{3.2}$$

where X_t are the training signals, Ψ is the dictionary, β are the sparse coefficients, $\|.\|_F$ is the Frobenius norm, S is the sparsity level and K is the maximum number of non-zero entries in each sparse representation β. Many algorithms provide efficient implementations of this optimisation process. The K-SVD algorithm [99] and the Method of Optimal Directions (MOD) algorithm [100] are two techniques widely used in the literature [101]. These iterative methods alternate between sparse coding

and dictionary update stages. The sparse coding stage generally looks for approximate solutions to (3.2). As (3.2) is a NP-hard problem, greedy pursuit algorithms such as OMP [51] and Least Absolute Shrinkage and Selection Operator (LASSO) [102] are generally employed. The dictionary update stage is mainly where the K-SVD and MOD algorithms differ. MOD updates all atoms simultaneously using a matrix inversion, whereas the K-SVD algorithm updates the dictionary on an atom by atom sequence. The K-SVD algorithm was chosen for this application as it has been found previously to be a more efficient algorithm and less demanding than MOD [101, 103].

The K-SVD algorithm proposed by Aharon *et al.* [99] is a data-driven learning process for sparse representations, implemented using a singular value decomposition (SVD) approach. The algorithm includes two main steps which constitute one iteration of the algorithm: sparse coding and the dictionary update stage. The implementation of the algorithm used in this study is detailed in [104, 105]. The sparse coding uses the OMP algorithm. The dictionary update is performed on an atom by atom basis optimising the target function for each atom while keeping the remainder of the dictionary fixed. The dictionary update is the main innovation in the K-SVD algorithm. While updating an atom, the update stage only uses the training signals in X_t which are used for that particular atom. This step is based on the SVD procedure with each column sequentially updated. This update is employed to minimise the approximation error and find an improved sparse representation for the training data.

3.3.2. Optimal Parameter Investigation

The main dictionary parameters affecting the performance of DL with CS are related to the dictionary size, in terms of both frame length N and number of atoms P. For these experiments, dictionaries were created with atoms of $N = 256$ and $N = 1024$. Using a smaller atom size ($N = 256$) helps minimise the variability in the ECG frame. For example, with ECG signals, the 256 frame size corresponds to 0.71s of ECG data and there is generally a maximum of one QRS complex per frame (for typical normal heart rates), whereas a frame size of 1024 corresponds to a time

interval four times longer, and could possibly contain up to four QRS complexes per frame. The 1024 sample frame size was investigated to allow for increased *CR*s, since the smaller frame size does not contain enough measurements for reconstruction at higher *CR*s. Examples of atoms contained in dictionaries created for both $N = 256$ and $N = 1024$ are shown in Figure 3.3 for patient record 100 from the MIT-BIH database.

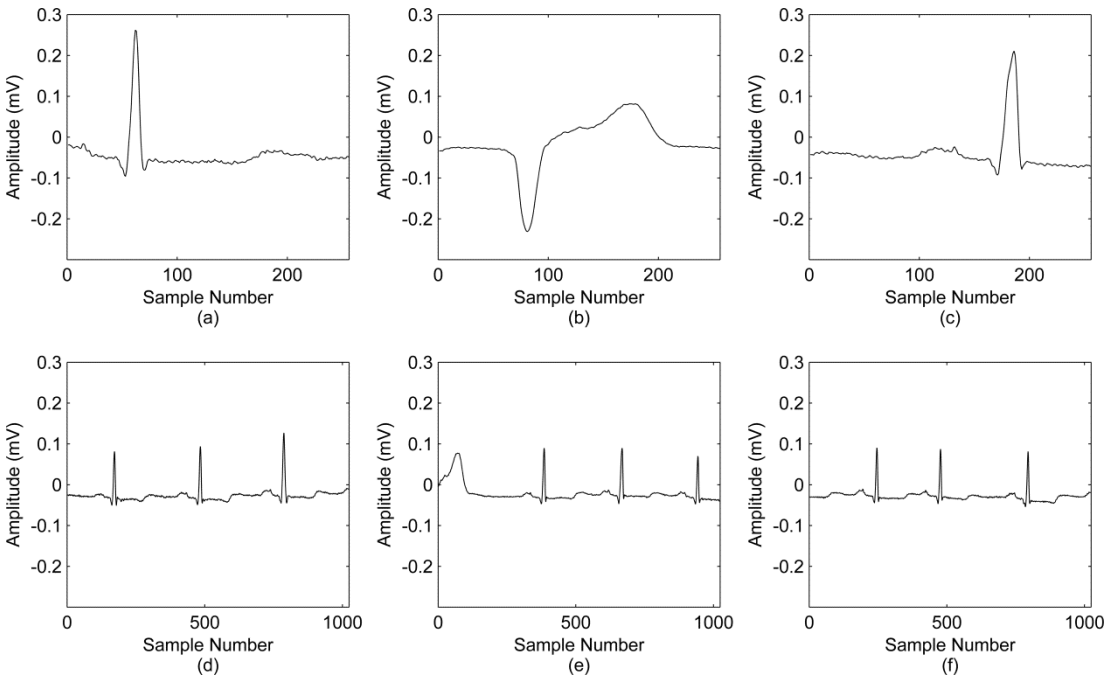

Figure 3.3: Examples of atoms contained in trained dictionaries for $N = 256$ (a, b and c) and $N = 1024$ (d, e and f).

Using $N = 256$ dictionaries as an example, the optimal number of atoms P in each dictionary is illustrated in Figure 3.4(a). Dictionaries with varying number of atoms were created using 90% of the signals for training and the remaining 10% for testing. The results were averaged over the 11 patient records and the signal reconstruction quality (*PRD**) was investigated for different numbers of compressed measurements. Figure 3.4(a) highlights that the difference in performance between the dictionaries tested when the number of atoms $P > 1000$ is marginal. In fact, the performance of the dictionary shows practically no change beyond $P = 20000$ and

therefore this is determined as the optimal number of atoms, as higher values increase the time of both dictionary creation and signal reconstruction and provide only minimal reconstruction accuracy benefit. Figure 3.4(b) shows the effect of different training partition sizes. Promisingly, the training partition size only marginally affects performance and the DL method can still function in scenarios where less training data are available. A training partition of 90% provides optimal performance as shown in Figure 3.4(b), and will be used in this thesis since it does not increase dictionary creation or reconstruction time, unlike the number of atoms. Therefore, for the testing in this section the N = 256 and N = 1024 overcomplete dictionaries contain P = 20000 and P = 100000 (based on similar empirical analysis of optimal dictionary lengths for N = 1024) atoms respectively, and were created using 90% of the signals from each record during the training phase.

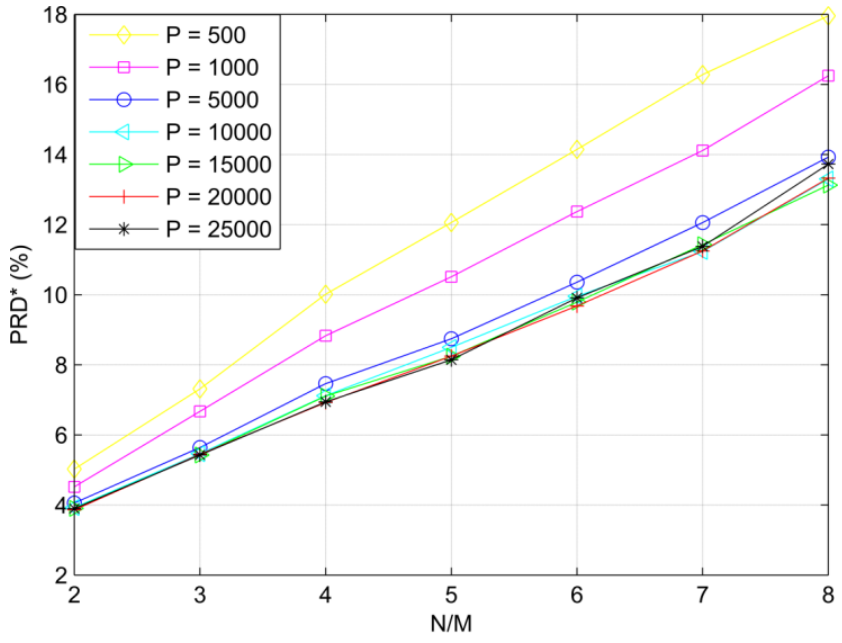

(a)

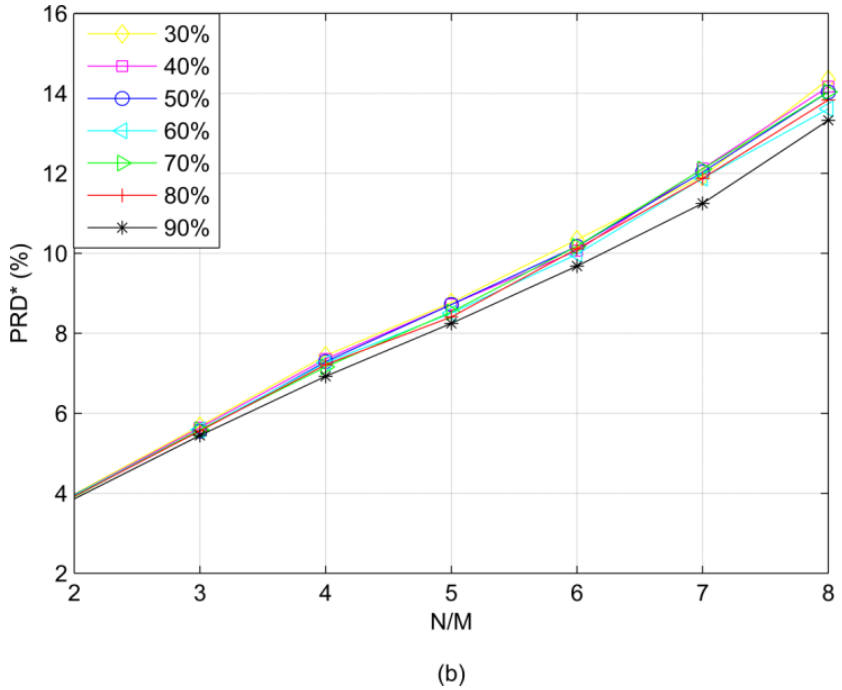

Figure 3.4: Investigation of optimal parameters for dictionary creation for N = 256 size dictionaries. Analysis of (a) optimal number of dictionary atoms P and (b) different training partition sizes.

3.4. Experimental Results

3.4.1. Dictionary Parameters

For the numerical experiments to test the complete RRM-DL architectures, three distinct dictionaries were created from each ECG record using different training data for both the N = 256 and N = 1024 frame sizes. The results for each record were then averaged from the performance of the three dictionaries. The overall distortion metrics and QRS detection rates were then averaged over all 11 records. The dictionaries were created using 90% of the record for training and 10% for testing. Training data were extracted from all parts of the signal for the different dictionaries. In order to support the creation of highly overcomplete dictionaries, additional training frames were created by overlapping successive frames in the training set. The parameters for creating the dictionaries were as follows. The N = 1024 dictionary contained P = 100000 atoms and was created by overlapping frames to ensure J = 150000 signals in the training set. The N = 256 dictionary contained P =

20000 atoms and was trained with J = 30000 training signals. The number of overlap samples between training frames is calculated to be the maximum sample shift that produces the required number of training frames from the training portion of the signal.

3.4.2. Comparison of RRM-DL Approaches

Figure 3.5 illustrates the performance of the proposed architecture, by showing the variation of *PRD** vs. *CR* for the two proposed dictionary atom lengths, and for two different quantisation levels of 7 and 9 bits (reflected in different Huffman codebook sizes). Due to the reduced variability in the number of QRS complexes (generally a maximum of one) in a single 256 sample frame, the performance of the RRM-DL-256 (i.e. RRM-DL with N = 256) method is generally superior at lower *CR*s (approximately < 24), however the performance of the RRM-DL-1024 is superior at higher *CR*s and continues to outperform the RRM-DL-256 approaches at *CR* values greater than 24. This can be attributed to the very small number of compressed measurements associated with higher *CR*s when frame lengths of 256 are used. For example, at a *CR* of 30, only 8 measurements are available using a frame length of 256, whereas for a frame length of 1024, 34 measurements exist. In general, 8 compressed measurements do not contain sufficient information about a frame to facilitate a sufficiently accurate reconstruction of the original signal. This effect on reconstruction quality becomes more pronounced as *CR* increases. In relation to quantisation, the results suggest that the lower quantisation levels improve the overall performance at higher *CR*s. For clarity, the results in Figure 3.5 are presented based on the performance of two quantisation rates. However, this does not change the conclusions of the experiment and the same trends on the effects of quantisation are found when Huffman codebook sizes $C = 2^8$ and $C = 2^6$ are used.

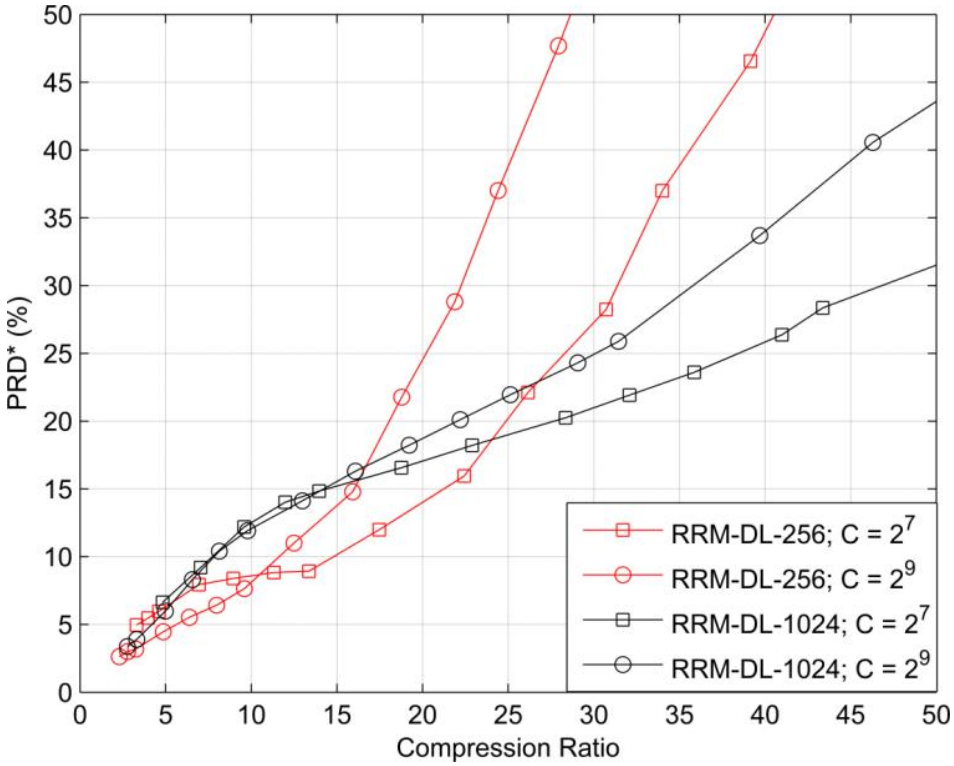

Figure 3.5: *PRD** values for *CR*s up to 50 for the Huffman codebook sizes $C = 2^9$ and 2^7 for frame sizes $N = 1024$ and 256.

A statistical analysis of the performance of the RRM-DL-256 implementation is presented by the Tukey boxplots performed over the 11 records in Figure 3.6. In each box the median is the central line, the box edges signify the 25th and 75th percentiles and the whiskers show the most extreme data points that are not considered outliers. The general trend of Figure 3.6 is that smaller *CR*s exhibit more inter-record variance at lower quantisation levels. However, as *CR* increases, the variation of inter-record performance between the three quantisation levels is low. Some variation between different ECG records is not uncommon; this has also been evident in previous CS studies on the same database [63, 64].

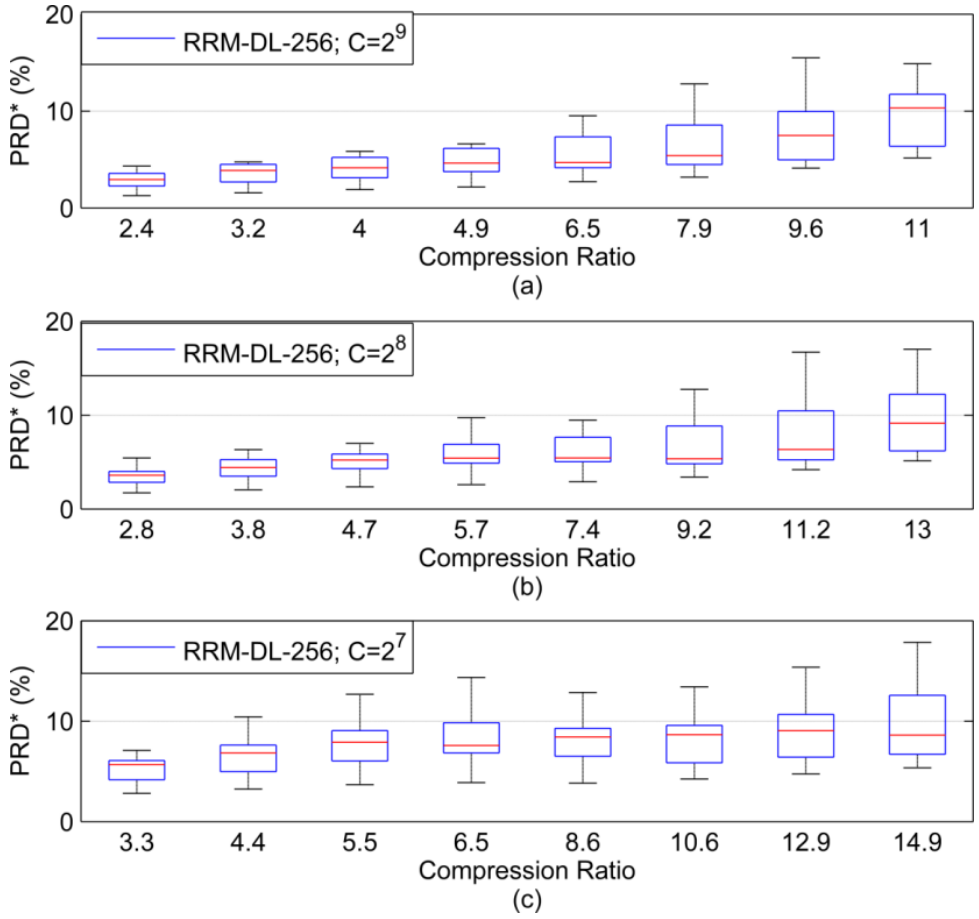

Figure 3.6: Boxplots of *PRD** values at each *CR* averaged over the frames of each patient record for different RRM-DL-256 implementations with (a) $C = 2^9$ (b) $C = 2^8$ and (c) $C = 2^7$ Huffman dictionary sizes.

An additional test was performed to evaluate the significance of each element of the proposed RRM-DL algorithm. For the purposes of this analysis the reconstruction performance of an RRM-DL-256 implementation was compared with the following two scenarios: (1) signal reconstruction with only learned dictionaries (where no RRM or quantisation are employed) and (2) signal reconstruction with learned dictionaries and quantisation (but no RRM). Figure 3.7 illustrates the improvement that quantisation and RRM add to overall performance of the RRM-DL approach in terms of *PRD** for increasing *CR*. The extension in *CR* due to quantisation appears to increase as *CR* increases and for smaller quantisation rates.

The further increase in compression due to the RRM module is more consistent between the 8 and 7 bit quantisation rates as the *CR* increases.

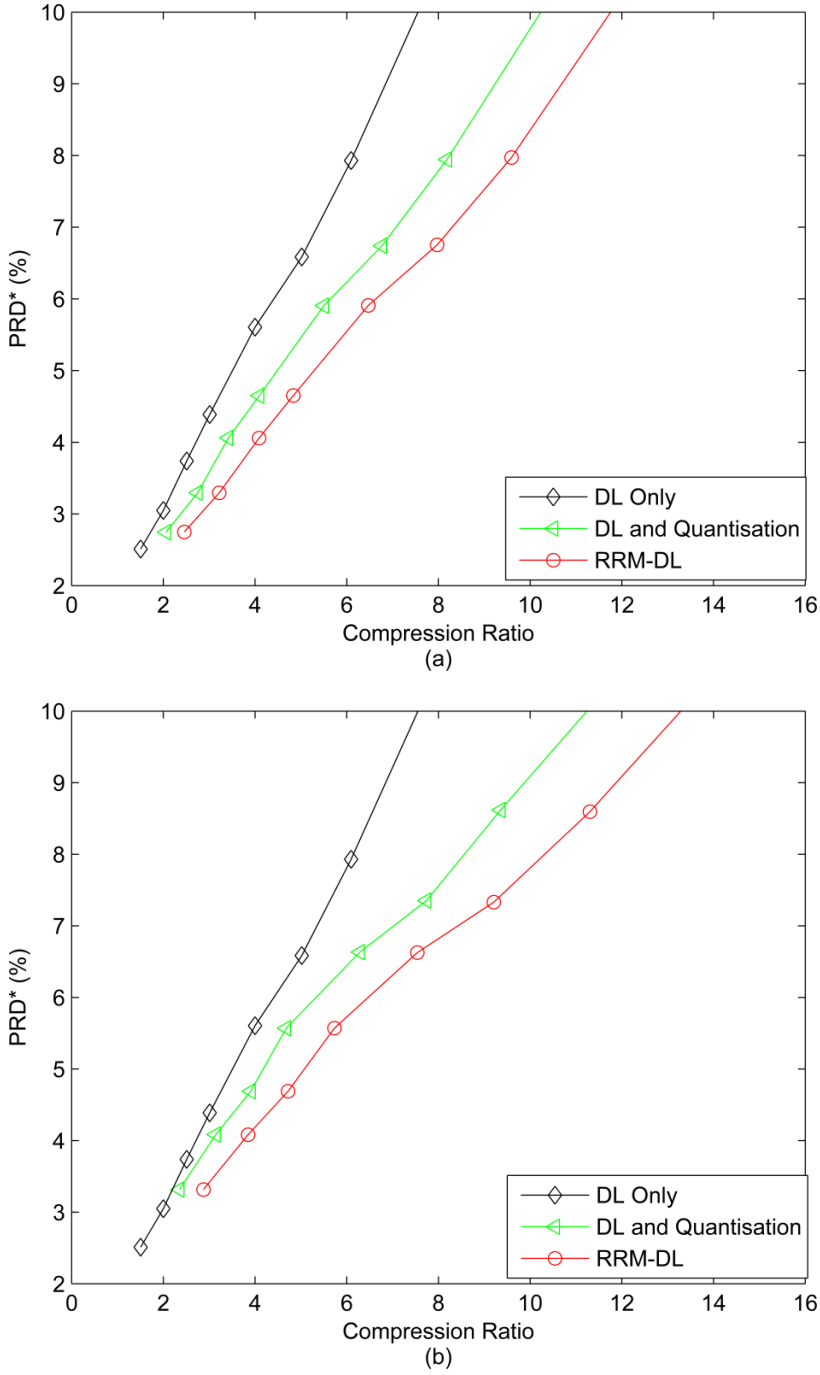

Figure 3.7: Comparison of RRM-DL-256 performance with two different approaches. Firstly, where only DL is employed and secondly, where DL and quantisation are used. Results from two different quantisation levels are shown: (a) 8 bit quantisation and (b) 7 bit quantisation.

Based on Figure 3.5, this research concluded that using smaller frame sizes is beneficial at the smaller *CR*s. Therefore, the RRM-DL-256 algorithm is recommended for applications where the operating $CR < 24$. However, for applications where performance can be maintained at higher *CR*s, RRM-DL-1024 is recommended. This is further explored in the remainder of the thesis, where RRM-DL-256 and smaller frame size based implementations are compared with state-of-the-art CS compression at the lower *CR*s and RRM-DL-1024 is utilised for QRS detection in a CAD system context in Chapter 5. QRS detection is an application where reconstruction errors can be tolerated, while preserving the diagnostic integrity of the signal, and therefore higher *CR*s can be obtained.

3.4.3. Performance Evaluation against Existing CS Approaches

Figure 3.8 shows a comparison between the MMB-CoSAMP and MMB-IHT algorithms taken from Polania *et al.* [64] and the proposed RRM-DL-256 CS approach in terms of *PRD* performance. The range of *CR*s in Figure 3.8 is limited, as no data are provided above a *CR* of 10 in [64]. This analysis focuses on RRM-DL-256, due to its superior performance in this compression range, based on the results presented in Figure 3.5. There is consistency in the ECG data tested as the same patient records used to test the RRM-DL-256 approach were also used in [64].

The RRM-DL-256 method outperforms MMB-CoSAMP and MMB-IHT for all *CR*s shown. By examining the highest *PRD* value shown with MMB-IHT (7%), the RRM-DL-256 method outperforms MMB-IHT by improving the *CR* from 10 to 15.95. At higher *CR*s, the *PRD* for the RRM-DL algorithms tend towards a plateau. This contrasts with the MMB methods that appear to increase rapidly as *CR* rises. Previous work by the authors [95] on the quantisation of compressed measurements indicates that quantisation offers the best balance of bit reduction and signal fidelity at higher *CR*s. Similarly, DL offers best performance at higher *CR*s because of the highly sparse representations in the learned dictionaries and the common support present in ECG signals [97].

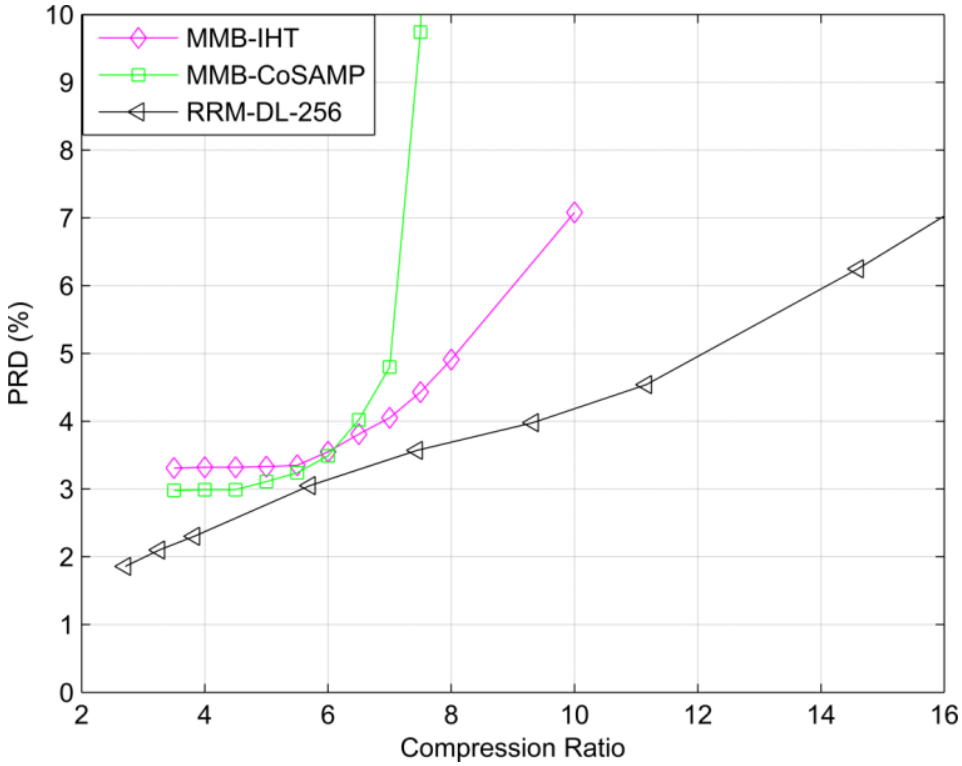

Figure 3.8: *PRD* of MMB-IHT and MMB-CoSAMP algorithms and the proposed RRM-DL-256 approach using 8 bits ($C = 2^8$) to represent the compressed measurement differences.

Further analysis quantifying the potential of RRM-DL-256 in terms of wireless transmission power savings compared to the MMB methods is presented in Table 3.1. *CR* is directly linked to wireless transmission costs and therefore the relative *PRD** advantage of RRM-DL-256 increases as *CR* increases (as illustrated in Figure 3.8). At a *PRD* of 7%, RRM-DL-256 demonstrates a *WTS* of 59.5% over MMB-IHT while at an 8% *PRD*, a *WTS* of 142.1% is afforded over the MMB-CoSAMP implementation. This analysis is significant due to the dominant power consumption of wireless transmission in BANs [4, 14].

Table 3.1: *WTS* of RRM-DL-256 over MMB methods for varying *PRD*

	CR			RRM-DL-256 *WTS* (%) over	
PRD (%)	RRM-DL-256	MMB-IHT	MMB-CoSAMP	MMB-IHT	MMB-CoSAMP
3	5.58	-	4.54	**-**	**22.9**
3.5	7.21	5.88	6.01	**22.6**	**19.9**
4	9.39	6.90	6.48	**36.0**	**44.9**
4.5	11.03	7.57	6.81	**45.7**	**61.9**
5	12.09	8.08	7.02	**49.6**	**72.2**
5.5	13.09	8.54	7.07	**53.2**	**85.1**
6	14.1	9.00	7.12	**56.6**	**98.0**
6.5	15.05	9.47	7.17	**58.9**	**110.0**
7	15.95	10.00	7.22	**59.5**	**121.1**
7.5	16.85	-	7.26	**-**	**132.0**
8	17.75	-	7.33	**-**	**142.1**

3.4.4. Average CR for Clinically-Acceptable Reconstruction

The relationship between *PRD** and desired signal quality perceived by a clinician was considered by Zigel *et al.* in [81]. That work classified compressed ECG signals into different quality groups using a weighted diagnostic distortion signal measure which focuses on how closely the reconstructed signal matches the original and preserves diagnostic information. Different values of *PRD** are then classified based on the quality groups. They considered *PRD** values of 2-9% to represent either "very good" or "good" signal reconstruction quality. For *PRD** values > 9% certain reconstructions are identified as "not good" or "bad". Based on that study, 9% *PRD** will be considered here as the target value for acceptable signal reconstruction quality for clinical examination. Therefore, the performance of RRM-DL-256 in terms of achievable *CR*s at the target 9% *PRD** rate is summarised in Table 3.2, indicating an average *CR* of 15.11 over the test records. To the best of the author's knowledge, a *CR* this high has not previously been reported in the literature for a *PRD** of 9% with a CS implementation for ECG compression. The *CR* obtained if SRR is used instead of RRM is also presented, demonstrating that RRM provides a higher *CR* for each record tested. The optimal quantisation level of each record varies as shown in Table 3.2, however 6 bit quantisation ($C = 2^7$) achieves the best performance most often. The general trend with quantisation is the lower bit levels (5

and 6) are more suitable for records where higher *CR*s can be achieved for the 9% *PRD** rate.

Table 3.2: Optimal *CR* of RRM-DL-256 method for *PRD** = 9%

Patient	RRM *CR*	SRR *CR*	Quantisation Bits
100	**11.32**	10.64	7
101	**15.19**	13.86	6
102	**10.64**	9.29	6
103	**21.15**	19.11	6
107	**23.68**	21.20	5
109	**25.6**	25.47	5
111	**10.61**	9.89	6
115	**14.26**	12.73	6
117	**10.88**	9.75	8
118	**8.06**	8.01	7
119	**14.9**	14.01	6
Average	**15.11**	13.99	-

3.4.5. Patient-Agnostic Testing

The previous sub-section focused on the performance of the proposed RRM-DL architecture on patient-specific dictionaries. In this sub-section, a separate independent analysis is presented showing the performance of RRM-DL when tested on records where no signals from that record are employed in the training phase, known as patient-agnostic dictionaries. The patient-specific approach has merits for long-term monitoring applications based on the high performance offered compared with other CS approaches. However, the advantage of a patient-agnostic implementation is the removal of a patient training phase. For the purposes of this experiment, each record (for each of the 11 patients) was tested using a dictionary created by combining the forty-seven patient-specific dictionaries from all the other records in the database. As this presents a more challenging testing scenario than a patient-specific approach, certain parameters were increased. With $N = 256$, each sub dictionary contained $P = 2000$ atoms, trained with $J = 50000$ training signals.

Therefore, by combining forty-seven sub dictionaries, the resulting dictionary used for testing contained P = 94000 atoms.

The results (Figure 3.9) further highlight the potential of the RRM-DL approach. Previous results in this thesis have shown that for an increasing *CR*, lower quantisation levels provide optimal performance. Therefore, the RRM-DL plots in Figure 3.9 show the best performance over the four tested quantisation levels. As expected, the patient-specific approach outperforms the patient-agnostic approach at all *CR*s. However, the patient-agnostic dictionaries still outperform other CS implementations over the following *CR* ranges: MMB-IHT approach for *CR*s < 5.2 and > 7.6 and the MMB-IHT implementation for *CR*s < 4.65 and > 6.5.

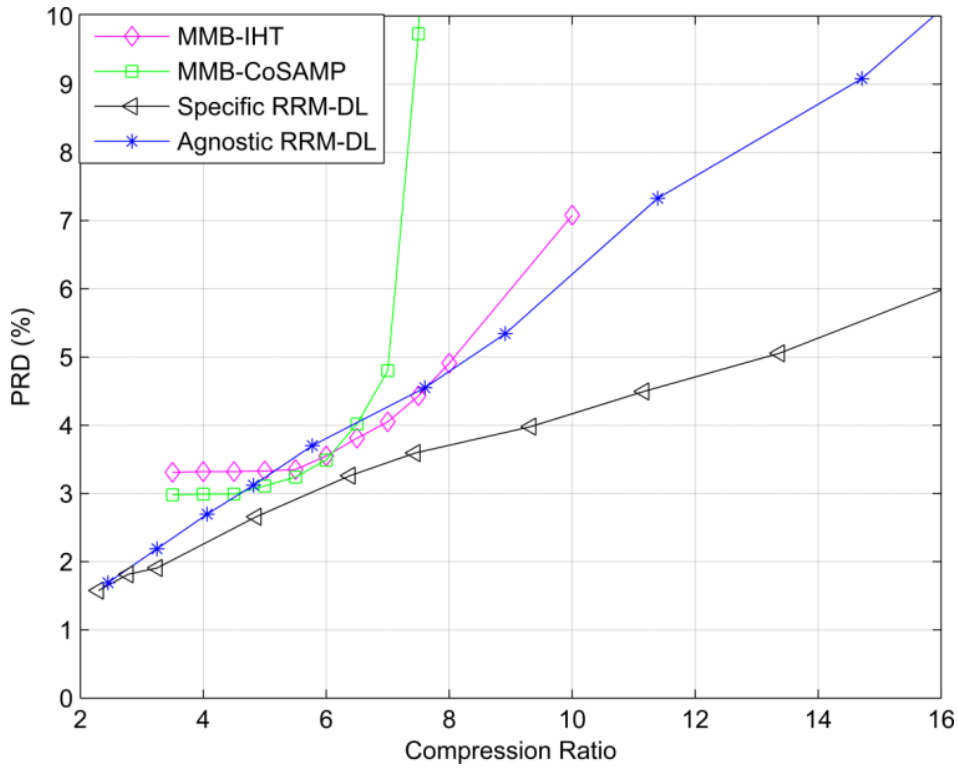

Figure 3.9: *PRD* of MMB-IHT and MMB-CoSAMP algorithms and the proposed patient-specific and patient-agnostic RRM-DL approaches.

3.5. Discussion

The three main optimisations in the RRM-DL algorithm are the addition of Redundancy Removal with Mean (RRM), quantisation, and the use of patient-specific dictionary reconstruction. The RRM approach enables an extension in *CR* and was shown to outperform the existing Standard Redundancy Removal (SRR) method on each record tested. It is important to note that the alternative algorithms in the literature that use SRR also require an offline training phase to create the Huffman codebook. The proposed technique was tested over a wide variety of records containing both routine and rare arrhythmia phenomenon, abnormal heart beats and various signal artefacts. These ECG abnormalities demonstrate that the RRM method is robust to changes in QRS complex structure and shifts in its position across a frame. Quantisation is also beneficial in improving performance. The combination of quantisation with patient-specific DL enables low quantisation bit resolutions to be used without introducing significant signal degradation. At the target *PRD** of 9%, 6 bit quantisation ($C = 2^7$) offers the best reconstruction performance. This combination of approaches enables good performance, particularly at the higher *CR*s in comparison to existing CS-based ECG approaches.

A practical consideration to the deployment of the RRM-DL architecture is the requirement of the one-off training phase per patient. However, the results presented in the previous sub-sections demonstrate that the proposed patient-specific RRM-DL approach outperforms existing CS techniques in the literature in terms of signal reconstruction quality for a given *CR*. Considering that the operation of wireless transmission consumes the most power in a BAN [4, 14] and the costs of transmission are directly linked to *CR*, an ambulatory ECG monitoring architecture will benefit from a low-energy consumption implementation typical of CS-based approaches in the literature once such a training phase is complete. Alternatively, the potential of an overcomplete patient-agnostic dictionary architecture was investigated and can provide competitive performance to existing CS implementations. However, the performance still needs to be improved if it is to be considered for implementation over the patient-specific dictionaries.

The use of a patient-specific training phase is not uncommon in ECG monitoring systems in the literature [84, 97, 106-108] and a number of implementation options exist for the collection of the data required for the training phase. The current implementation utilises 90% of the ECG records for training and therefore implies that approximately 27 minutes of data would suffice for dictionary creation and mean measurement calculation. Firstly, and perhaps the most ideal scenario, the collection of training data could be performed during device set-up in a clinical environment. The advantage is this would ensure no power implications on the wearable device as the dictionary creation would be completed offline in an unconstrained power environment. An alternative online option would be collection and transmission of Nyquist-sampled data for the 27 minute training phase. The dictionary creation and mean measurement values can then be calculated offline at the server. Once this training process is complete the RRM-DL implementation is enabled. The disadvantage of this approach is the wireless transmission costs during training prior to the RRM-DL implementation. However, training partitions smaller than 90% were shown to not have a significant impact on signal reconstruction quality in Figure 3.4(b) and this is useful for applications where power restrictions require a shorter training phase, albeit at a cost of a slight degradation in system performance. A final option is the integration of patient-agnostic dictionaries investigated in this chapter and previously in [97]. Patient-agnostic dictionaries do not require training and could be used for the duration of monitoring. Alternatively the device could operate with patient-agnostic dictionaries prior to a switch to patient-specific dictionaries. In such a situation the implementation would utilise the patient-agnostic dictionaries until enough data are collected for a specific implementation. An advantage of this implementation over Nyquist transmission during training would be the reduced wireless transmission costs as the training data are compressed. However, an obvious drawback of this approach is the patient-agnostic reconstructions would introduce an error into the training signals and would have to be performed at a low *CR*, as the performance of agnostic dictionaries does not match that of specific dictionaries based on Figure 3.8.

3.6. Summary

This chapter has proposed a new CS architecture, RRM-DL, aimed at low-power CS ECG implementations. The algorithm was evaluated in terms of signal reconstruction quality for increasing *CR* and compared with existing CS implementations in the literature. The novel RRM-DL implementation was shown to outperform the best known existing CS approaches in terms of *PRD* for all *CR*s tested. These results are significant as they demonstrate the ability to significantly reduce the wireless transmission costs associated with current CS implementations while making no additional sacrifices to the integrity of the signal. The work described in this chapter was published in [65].

The next chapter (Chapter 4) continues the development of the proposed CS approach by considering adaptive dictionaries as an alternative to the fixed dictionaries considered in Chapter 3.

CHAPTER 4

Adaptive Dictionary Reconstruction

4.1. Introduction

The previous chapter considered the use of DL in CS, and demonstrated the ability of the RRM-DL architecture to outperform existing CS architectures in the literature. In this chapter, the aim is to further improve the performance of a learned dictionary approach for resynthesis through an Adaptive Dictionary (AD) reconstruction scheme. The modified reconstruction process is a two-stage process that leverages information about the signal from an initial reconstruction pass. By identifying whether a QRS complex is present in a particular frame and determining a location estimate of the QRS, a better dictionary is selected and a more refined signal reconstruction can be obtained with a second stage of reconstruction.

The same performance metrics used in Chapter 3 will be used here and the algorithm will be benchmarked against existing CS implementations, SPIHT, and the traditional DL approach from Chapter 3, referred to here as the standard dictionary (SD) approach. Moreover, an additional test on the technique is performed by monitoring its ability to maintain performance on a noisy ECG database, typical of what would be expected in an ambulatory BAN monitoring architecture.

4.2. Adaptive Dictionary Reconstruction Framework

The proposed AD reconstruction framework aims to improve on the performance offered by SD learning based CS implementations. The proposed approach leverages knowledge of the signal based on an initial reconstruction pass with a standard learned dictionary. This signal information allows the reconstruction scheme to select a more suitable dictionary for a second stage refinement signal reconstruction pass. The proposed method operates as follows. Suppose $x_{initial}$ is an initial reconstruction of the signal where a dictionary $\Psi_{initial}$ is used in the convex optimisation operation. The process then exploits the knowledge of $x_{initial}$ to optimise signal reconstruction by allowing a more suitable dictionary to be selected based on two key signal characteristics of $x_{initial}$.

The first signal property is based on whether a QRS complex is detected in the first stage signal estimate $x_{initial}$. If no QRS complex is present in $x_{initial}$, the second stage of the algorithm consists of signal reconstruction with a dictionary Ψ_{NQ}, specifically trained with signals containing no QRS complex and therefore providing a better definition of the signal information in the signal region between QRS complexes. If a QRS complex is detected in the first stage signal reconstruction, then the location of the QRS complex is determined. In this case q sub-dictionaries are used, such that the overall dictionary consists of the collection of sub-dictionaries: $\Psi_Q = [\Psi_{Q1}, \Psi_{Q2}, \Psi_{Q3} \ldots \Psi_{Qq}]$. Each sub-dictionary is created with training signals that contain QRS complexes across individual temporal windows within the frame, each containing N/q samples, shifted across the frame with no overlap between windows. For example, the set of training signals used to create the dictionary Ψ_{Q1} contain QRS complexes in the sample regions $[1{:}N/q]$. In general, this concept allows the signal reconstruction quality to be improved, as the chosen sub-dictionary is trained for sparse representations of signals where the majority of signal energy in the frame is concentrated on the location corresponding to the sub-dictionary. Figure 4.1 demonstrates an example of training signals used to create the dictionaries $\Psi_{initial}$ (a),

Ψ_{NQ} (b) and Ψ_{Q2} (c) (with q = 12 and N = 192). Each individual dictionary is created with the associated training signals using the K-SVD algorithm as in Chapter 3.

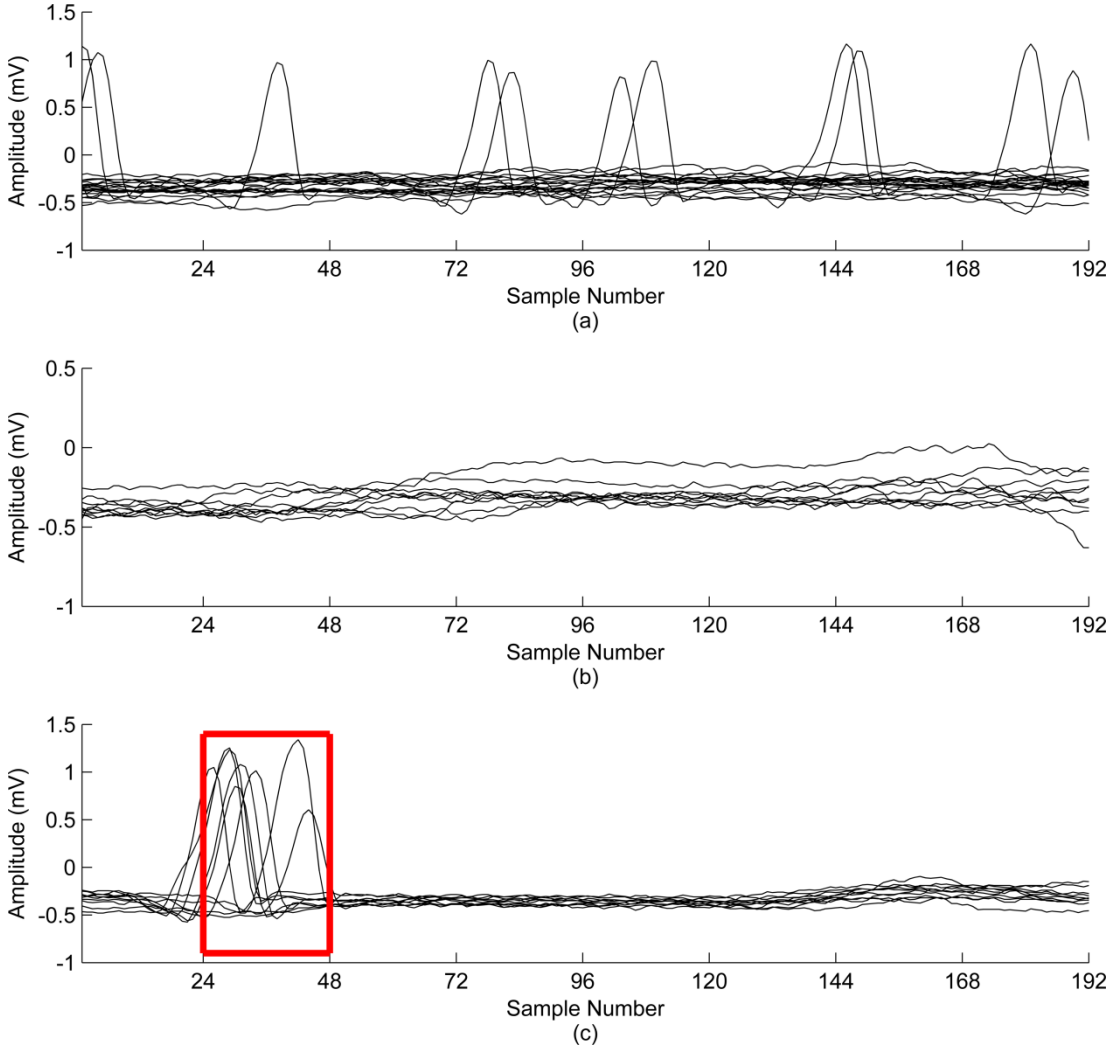

Figure 4.1: Examples of 20 random training signals for various dictionary creations. (a) Standard DL approach where all types of signals are used in training to create $\Psi_{initial}$. (b) Training frames for a sub-dictionary Ψ_{NQ} where no QRS complex is detected in a frame. (c) Training frames for a sub-dictionary Ψ_{Q2} where q = 12 and N = 192. The red box indicates the defined region where the QRS complex is present, for the training signals.

The AD algorithm is illustrated in Table 4.1. There are several practical considerations that need to be considered in the AD approach. The length of the ECG frame N = 192 was chosen to ensure a good balance between the number of frames

containing QRS complexes and non-QRS complex frames. Note that the frame size is dependent on the sampling frequency (360 Hz) and may need to be altered if the ECG signals are sampled at a different frequency. Ideally, the number of sub-dictionaries q is maximised to allow a small temporal window size and good time resolution in terms of each QRS location. However, there must also be sufficient training signals to create high performing overcomplete dictionaries. Preliminary testing of variations of q, found that q = 12 provides optimal performance for N = 192.

Table 4.1: AD Reconstruction Framework

Inputs: Compressed measurements y, sensing matrix Φ, dictionaries $\Psi_{initial}$, Ψ_{NQ} and $\Psi_{Q1}, \Psi_{Q2}, \Psi_{Q3} \ldots \Psi_{Qq}$.
First Stage (SD):
1: $\min_{\alpha_{initial} \in \mathbb{R}^N} \lVert \alpha_{initial} \rVert_1$ s.t. $y = \Phi\Psi_{initial}\alpha_{initial}$ 2: $x_{initial} = \Psi_{initial}\, \alpha_{initial}$
Second Stage (AD):
4: **if** ($d_{max} < Q_{Thresh}$) **then** 5: $\min_{\alpha_{NQ} \in \mathbb{R}^N} \lVert \alpha_{NQ} \rVert_1$ s.t. $y = \Phi\Psi_{NQ}\alpha_{NQ}$ 6: $x' = \Psi_{NQ}\alpha_{NQ}$ 7: **else** 8: **for** k = 1 **to** N **do** 9: **if** $\lvert\, \lvert(x_{initial}(k))\rvert - \overline{x_{initial}} \,\rvert = d_{max}$ **then** 10: $w = round(k/N + 0.5)$ 11: $\min_{\alpha_{QW} \in \mathbb{R}^N} \lVert \alpha_{QW} \rVert_1$ s.t. $y = \Phi\Psi_{QW}\alpha_{QW}$ 12: $x' = \Psi_{QW}\alpha_{QW}$ 13: **return** x'

The QRS detection in each frame is based on the maximum absolute deviation from mean value ($d_{max} = \lvert\, \lvert \max(x_{initial}) \rvert - \overline{x_{initial}} \,\rvert$), whereby if this value exceeds a threshold Q_{Thresh} determined during training of the dictionaries, the frame is deemed to contain a QRS complex. This is calculated on a frame-by-frame basis after the initial signal reconstruction and used to select the most suitable dictionary for second stage reconstruction. It is used in place of a standard QRS detection algorithm for the following reasons. Generally, QRS detection algorithms

require more than $N = 192$ samples to accurately detect QRS locations and function by combining frames in such instances. A potential issue here is the occurrence of a QRS complex at the boundary between two frames. A QRS detection algorithm may correctly identify the sample location of the complex, however the adjacent frame may still contain part of the QRS complex and thus a high energy portion of the signal will remain, but will be treated as a frame with no QRS complex. Using the maximum absolute deviation from the mean ensures that if the amplitude of the portion of the complex contained in the frame is above the threshold as determined during training, it is treated as a QRS complex and this makes the algorithm more robust to such issues.

Despite adding to the computational time of dictionary creation and doubling the reconstruction time compared to SD reconstruction, the proposed AD framework does not result in increased energy consumption in the actual sensors that constitute a BAN. This is because the reconstruction and dictionary creation are completed offline in an unconstrained power environment with no energy impact on the sampling, compression or wireless transmission operations. A further advantage of the AD reconstruction scheme is that it is easily scalable. At present it is designed to detect the presence and location of a QRS complex in the first stage and perform a second reconstruction pass based on this information. However, in theory the algorithm does not need to be limited to two reconstruction passes, and could be extended to more stages for a more sophisticated classification process if necessary. Additionally, alternative classification approaches other than QRS detection could be employed where appropriate.

4.3. Performance Evaluation

4.3.1. Dictionary Training

The signal acquisition architecture assumed in this chapter employs the SRR module and signal quantisation prior to Huffman coding. Patient-specific signals were used to create the Huffman dictionary (for encoding the redundancy removal

output), the reconstruction dictionaries (using the K-SVD algorithm) and the QRS thresholds for the AD operation. Results in Chapter 3 concluded that the impact of the training partition size is not significant and therefore in this chapter only 50% of the records are used as training signals for the one-off training phase to enhance the efficiency of the implementation. The remaining 50% of the record not used during training was used for testing purposes. Dictionary creation was performed using a set of t = 15000 training signals X_T, with each overcomplete dictionary Ψ containing P = 10000 atoms. The results were cross-validated to ensure the complete patient records were used during testing and averaged over each record tested.

4.3.2. Signal Energy-Based Distortion

Figure 4.2(a) shows the average *PRD** of the proposed AD approach compared with SD reconstruction and SPIHT. AD-Q6 (i.e. AD with quantisation of the difference values at 6 bits) outperforms the SD-Q6 reconstruction method across all *CR*s tested. While SPIHT still performs better than CS, the proposed technique reduces the gap in performance with SPIHT. Based on Zigel *et al.*'s assessment of different *PRD** ranges and their suitability for clinical diagnosis as perceived by clinical specialists, a target *PRD** of 9% will be used as a basis for assessment of compression performance as in Chapter 3 [81]. At this *PRD** level the AD approach extends the *CR* over the SD reconstruction from 12.1 to 14.8. SPIHT can achieve a *CR* of 15.9 for this *PRD** level.

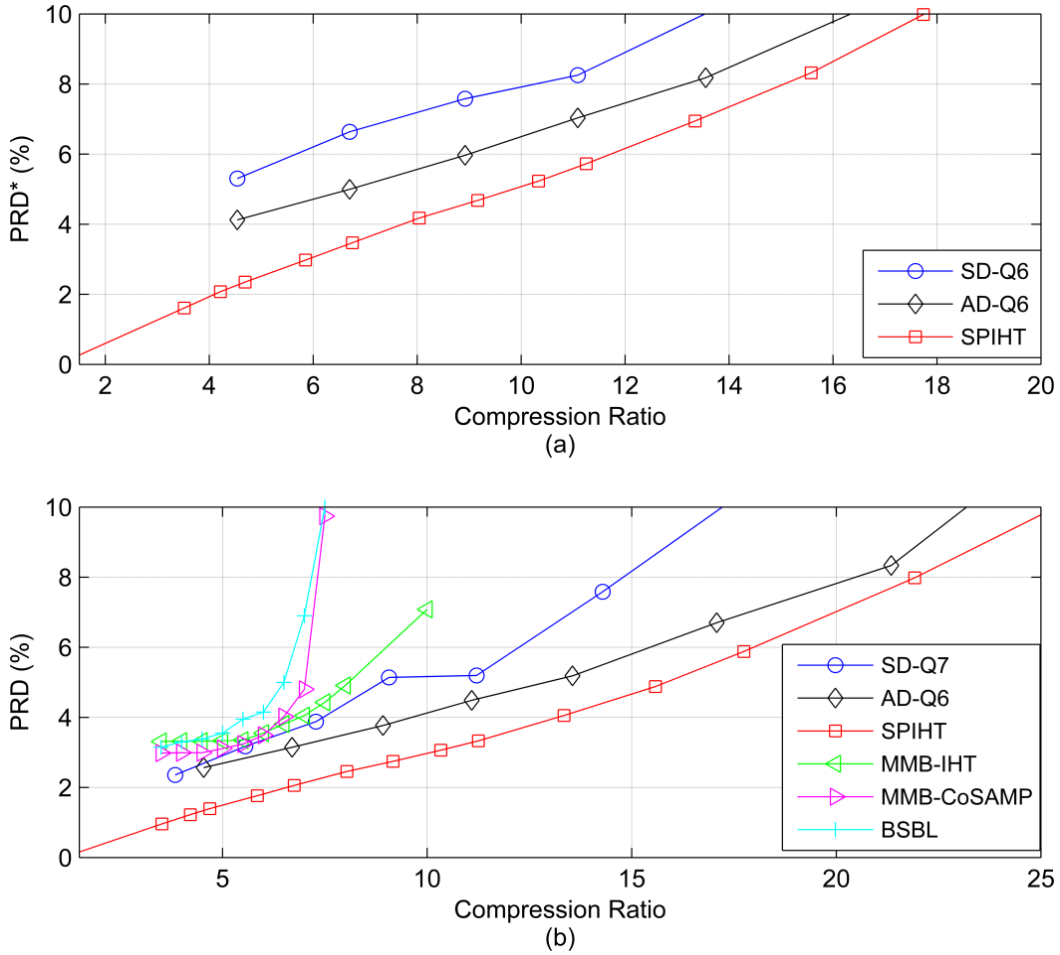

Figure 4.2: Comparison of (a) *PRD** and (b) *PRD* between the SD and AD reconstruction approaches, the MMB-IHT and MMB-CoSAMP, and BSBL approaches in the literature and SPIHT. Q6 and Q7 indicate quantisation of the difference values at 6 and 7 bits respectively.

The proposed AD-Q6, SD-Q7, SPIHT and the results presented in [64] for MMB-IHT, MMB-CoSAMP and BSBL are compared in terms of *PRD* in Figure 4.2(b) (note that *PRD* is used again instead of *PRD** to permit comparison with previously published results based on *PRD* [64]). It is clear AD-Q6 outperforms the MMB methods and BSBL for all *CR*s tested. Despite the use of *PRD* rather than *PRD**, this should not change the general conclusions in terms of comparison with these CS methods. While SPIHT still gives better overall performance, among the CS-based methods, the results obtained with the AD approach proposed here demonstrate a significant improvement in CS implementation for ECG.

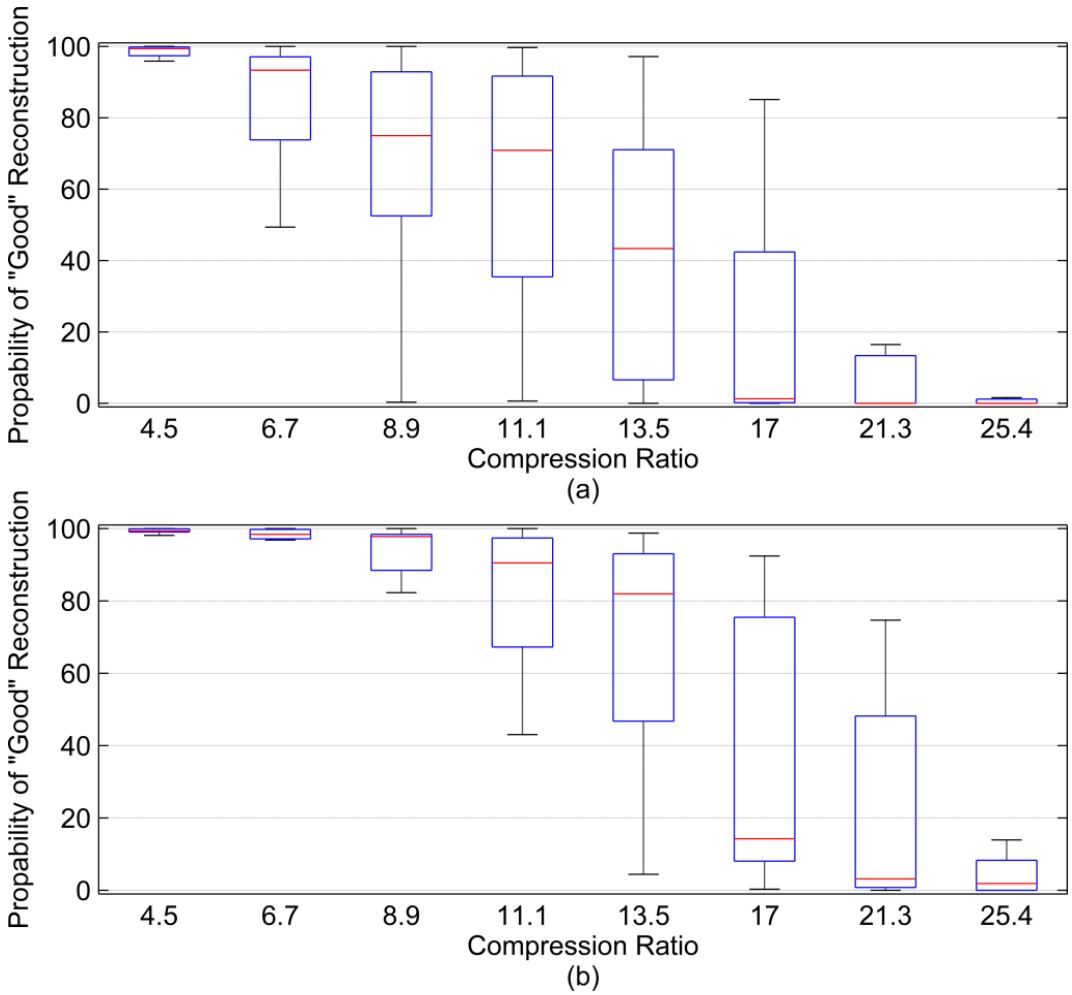

Figure 4.3: Boxplots of (a) SD-Q6 and (b) AD-Q6 showing the probability of a "good" reconstruction.

Boxplots for SD-Q6 and AD-Q6 in Figure 4.3 show the probability of "good" reconstruction (< 9% *PRD**) per 1024 samples averaged over all records. In each box the median is the central line, the box edges signify the 25th and 75th percentiles and the whiskers show the most extreme data points that are not considered outliers. It is clear from Figure 4.3 that the AD improves the probability of "good" reconstruction with a higher mean value at each *CR*. It also reduces the variance in reconstruction performance from record to record at the lower *CR*s.

Figure 4.4 provides a visual example of the ability of the AD method and SPIHT to preserve "good" signal reconstruction quality (< 9% *PRD** in the frames shown) where two premature ventricular contraction (PVC) beats were annotated.

This analysis is useful as it highlights the effectiveness of the AD algorithm to reconstruct potential abnormal events in an ECG signal, providing similar reconstruction quality to SPIHT at a high *CR* of 17.6.

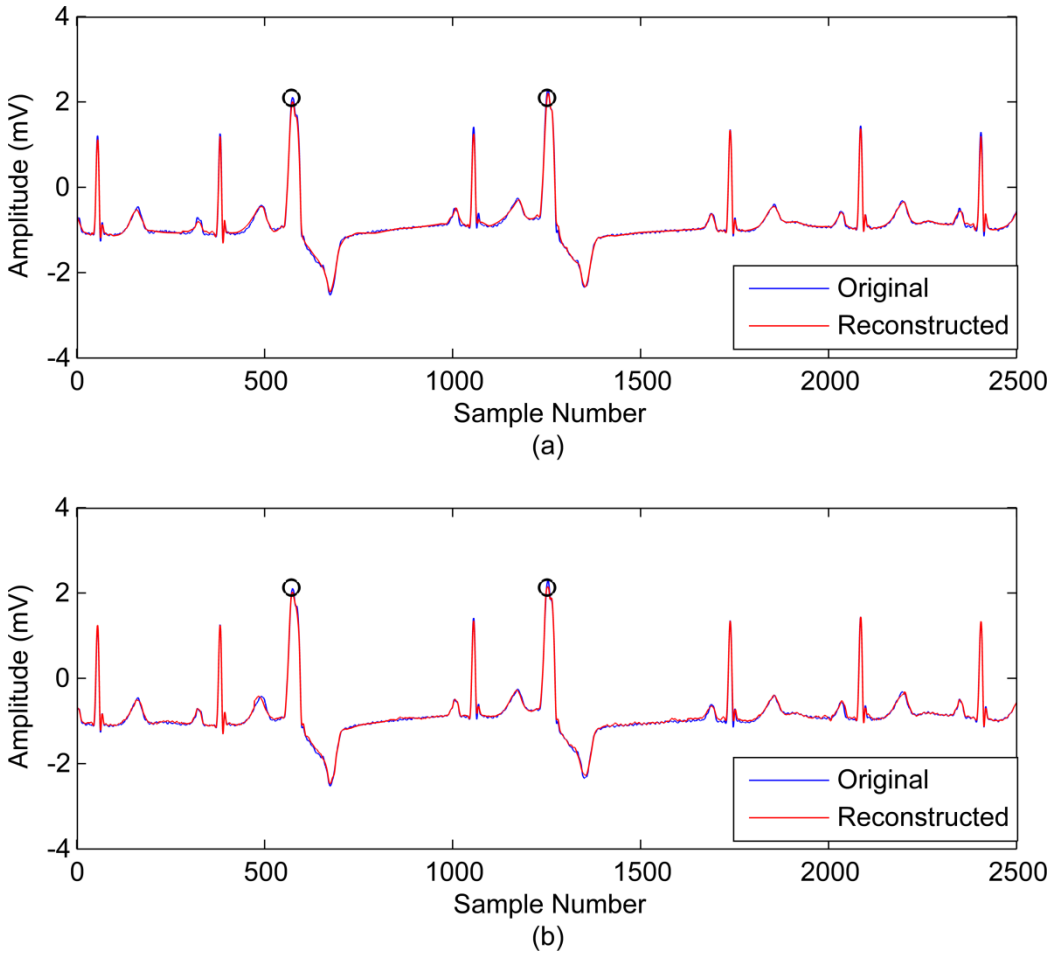

Figure 4.4: Examples of two PVC events in record 119. The original signals are in blue and the reconstructed signals (with a *CR* = 17.6) are in red. (a) SPIHT reconstruction with a *PRD** of 7.05%; (b) AD reconstruction with a *PRD** of 6.66%. The black circle indicates the location of the annotated PVCs.

4.3.3. Noise Stress Test Database

In order to further assess the potential of the proposed AD approach, the algorithm was tested on noisy ECG records typical of those that would be encountered in a BAN environment. Physionet provides a Noise Stress Test ECG

database with signals that contain baseline wander, muscle artefacts and electrode motion artefacts [87].

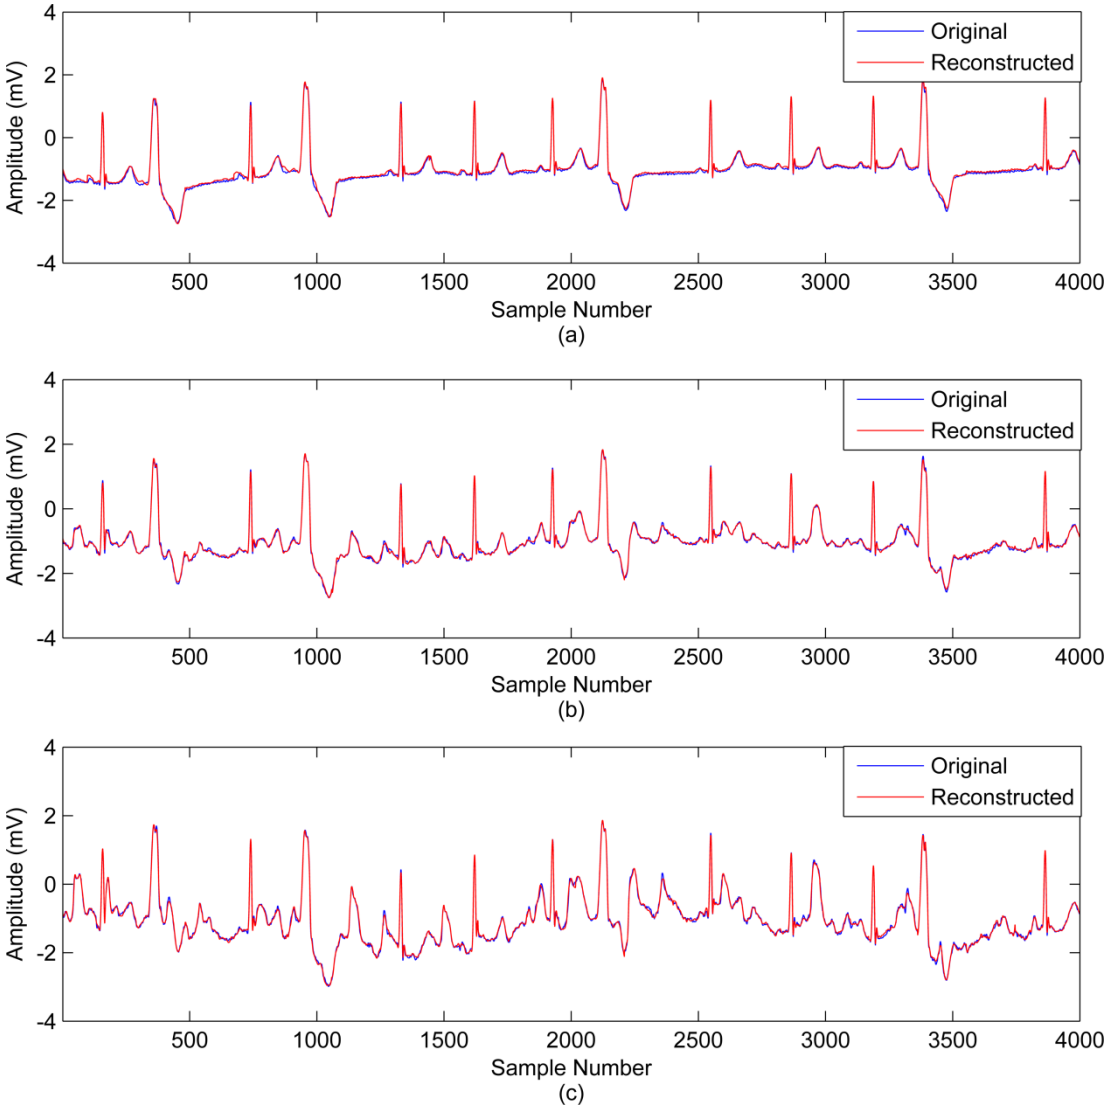

Figure 4.5: Examples of three original and reconstructed signals from record 119 for different levels of *SNR*. Plots shown are: (a) an original clean signal and a reconstructed signal with a *PRD** of 9.3% at a *CR* of 10.97; (b) a signal with a *SNR* of 18 dB reconstructed with a *PRD** of 7.12% at a *CR* of 10.24; and (c) a signal with a *SNR* of 12 dB reconstructed with a *PRD** of 6.81% at a *CR* of 9.54.

These signals represent a more difficult testing scenario for the proposed algorithm than the relatively "clean" signals due to the presence of artefacts as illustrated by Figure 4.5. The signals for the Noise Stress Test ECG database were

created as follows. Firstly, data were collected from physically active volunteers using standard ECG recorders with the electrodes positioned where the subjects' ECG signal was obstructed or not fully visible. Once these data were assembled intervals containing baseline wander, muscle artefacts and electrode motion artefacts were extracted. Then using two "clean" records from the MIT-BIH Arrhythmia database (record 118 and 119), calibrated amounts of the extracted noisy signals were added to the original records to create noisy signals at various *SNR* levels. For the purposes of these experiments, noisy data at *SNR* levels of 24 dB, 18 dB and 12 dB were used for dictionary creation and testing. The database provides 12 minutes of noisy ECG data per record and is useful as it can be directly compared with the performance on the relatively "clean" ECG signal for these two records (118 and 119). Figure 4.5 (b) and (c) show reconstructed signals at 18 dB and 12 dB *SNR* levels, with the noise-free record given in Figure 4.5(a) for comparison.

During these experiments, the signals were reconstructed with the so-called BP Denoise algorithm which is defined by the reconstruction problem in (4.1) [48, 109]. The BP Denoise algorithm is useful in scenarios where the signal is corrupted by noise. The bounding parameter σ balances the optimisation between sparsity and signal reconstruction quality. The parameter σ was set to 10^{-6} and the number of iterations for reconstruction was increased from $10M$ to $20M$ to obtain better performance.

$$\min_{\alpha \in \mathbb{R}^N} \|\alpha\|_1 \ \text{ subject to } \ \|\Phi\Psi\alpha - \mathrm{y}\|_2 \ \leq \ \sigma \tag{4.1}$$

The boxplots in Figure 4.6 show the *PRD** performance for the AD-Q6 and SD-Q6 approaches at varying *CR* and *SNR* levels averaged over each record. There are two important conclusions about the AD algorithm to be drawn from Figure 4.6. Firstly, the introduction of noise (muscle, electrode and baseline wander artefacts) does not significantly impact the performance of the AD-Q6 method. While the performance of the AD-Q6 decreases slightly, even at a low *SNR* of 12 dB, the

algorithm maintains good performance. For example, up to a *CR* of 11.7, the mean *PRD** value is below the desired 9% for all noise levels with AD-Q6. Secondly, the AD algorithm maintains improved performance over the SD approach, and the implementation is robust to the presence of the noise related artefacts and electrode motion in reconstruction to demonstrate the full potential of the method for implementation in low-power BANs.

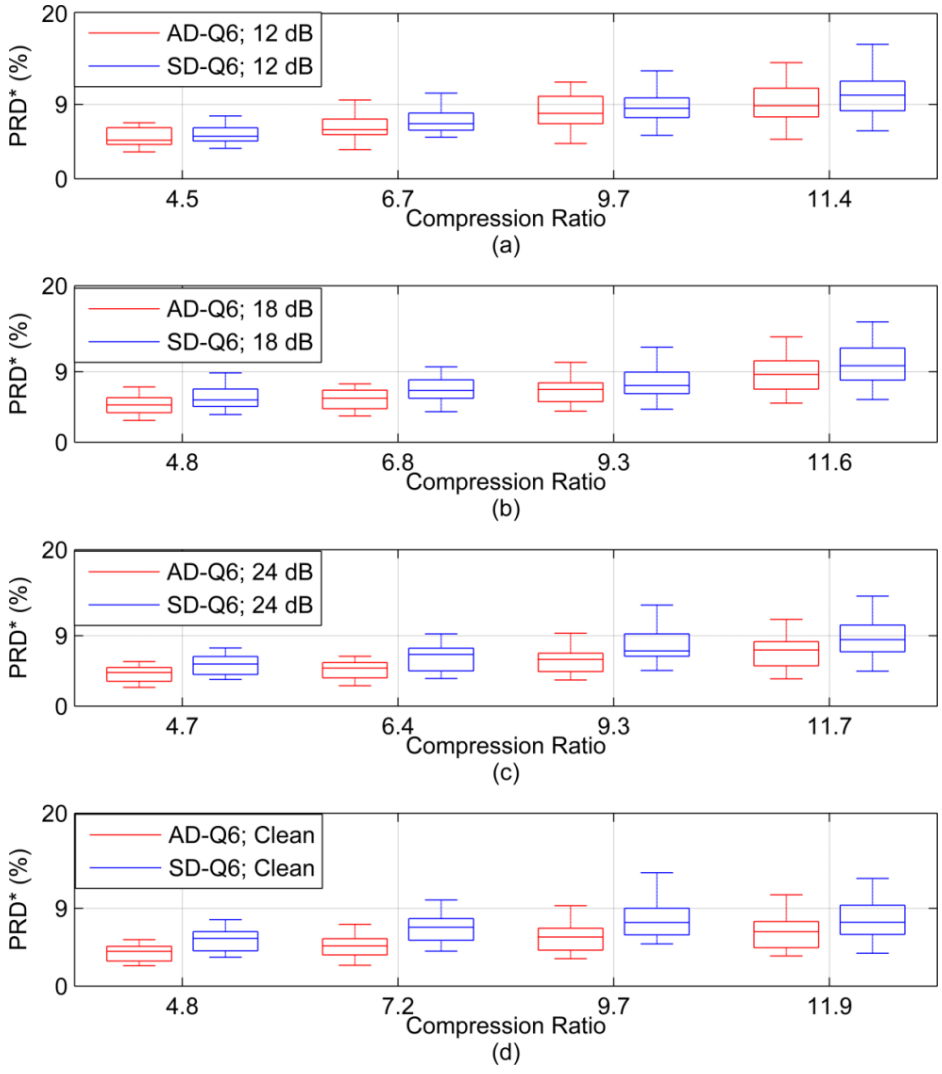

Figure 4.6: Boxplots of the *PRD** of AD-Q6 and SD-Q6 tested at different *SNR* levels for approximately equal *CR*s. The performance at (a) 12 dB (b) 18 dB (c) 24 dB and (d) the clean signals are illustrated.

4.4. Discussion

An important aim of the proposed algorithm was to improve upon existing CS approaches in terms of the trade-off between *CR* and distortion and to reduce the gap in performance between CS and SPIHT. The results show the potential of the proposed framework as it outperforms existing CS techniques in terms of the *PRD* vs. *CR* trade-off. In particular, the comparison with SPIHT shows a promising boost in CS performance compared to current CS implementations in terms of signal fidelity. The nature of the AD approach allows it for flexibility in implementation, and while QRS complexes are used to classify the reconstructed signals in this approach, alternative features for classification could also be employed. The fact that signal reconstruction is performed offline also allows for more customisation of the AD approach whereby the algorithm does not need to be limited to a two-stage process and could be extended to multiple stages based on the consideration of additional signal properties.

The main difference between the proposed DL-based CS implementations (including the RRM-DL architecture proposed in Chapter 3) and the existing CS approaches in the literature is that while existing techniques focus on exploiting the structure of the ECG wavelet coefficients, the proposed approaches leverage signal sparsity during reconstruction. For example, the MMB methods exploit knowledge of signal structures and dependencies that form across the significant coefficients in different sub-bands. On the other hand the BSBL focuses on the inter block structure that exists in the different wavelet decomposition levels of the coefficients. While the most recent high performing CS techniques demonstrate good reconstruction quality, far superior to the traditional wavelet basis approaches, the combination of overcomplete patient-specific dictionaries and the AD framework ensure the proposed CS architecture can outperform these CS techniques. The AD algorithm allows for superior reconstruction qualities at all *CR*s, but in particular as the *CR* increases the proposed technique is more robust than current CS approaches.

4.5. Summary

In this chapter a novel reconstruction framework for CS-based ECG compression was proposed. More specifically, by combining DL with CS, signal reconstruction quality was improved by exploiting knowledge of signal properties based around the prominent QRS complex. The algorithm was evaluated against the benchmark CS algorithms discussed previously, as well as SPIHT. Additionally, tests were performed on a noisy ECG database highlighting the ability of the AD approach to maintain performance in the presence of noise levels that would be expected in a typical BAN. The AD algorithm and the experiments described in this chapter appear in [110].

The next chapter extends the evaluation of the proposed CS approaches by considering their performance when used in CAD systems, where higher signal distortion may be acceptable as long as the diagnostic information in the signal is preserved.

CHAPTER 5

Computer Aided Diagnosis

5.1. Introduction

Computer Aided Diagnosis (CAD) systems interpret biomedical signals to obtain information that can be used to detect or screen for potential important abnormalities with a patient. The use of CAD systems has become more important in recent times, in particular because they can assist a clinician with their analysis of biosignals where large volumes of data may need to be analysed. There are various implementation options for these systems. A common approach is continuous running of a software algorithm on a centrally-located secure server that can detect or screen for potential signal abnormalities. If the algorithm is running in real-time and is receiving continuous biosignals from a worn sensor, this enables continuous monitoring of a patient, and feedback to be sent to the clinician/patient if a potential problem arises. Alternatively, as the data are received and processed, they can be stored for future review when requested by a clinician. In a scenario where the clinician chooses to monitor the signals at irregular intervals, the CAD system can aid the clinician with their signal analysis during the review by highlighting parts of the signal that are worthy of increased clinical attention. Such an architecture has the

advantage of freeing up clinical resources while still providing high value interpretation by an experienced clinician.

In the previous chapters different dictionary reconstruction methods to enhance CS were proposed and their performance was quantified by energy-based signal distortion metrics. This chapter will now focus on the ability of the proposed CS algorithms to maintain diagnostic integrity in the signal for an increasing *CR*. At present there is limited research in the literature specifically related to CS in this area. The metrics used will align more directly with clinical practice, and will give an indication of the clinical quality of the retained signal information and hence the suitability of the proposed CS methods for use with such CAD systems. The diagnostic measures will be mainly based around the detection of the R-peaks in the QRS complexes of reconstructed ECG signals. QRS detection analysis provides an insight into how the compression approach preserves the clinically-relevant information in an ECG signal as opposed to purely energy-based *PRD* metrics that mainly reflect signal fidelity. In fact, QRS complex feature extraction is a critical step for reliable HRV analysis [28, 29], diagnosing cardiac disorders such as cardiac arrhythmia [88], sleep apnea [111] and atrial fibrillation [112], and therefore can be seen as an essential diagnostic tool. In addition, several time domain and frequency domain HRV performance measures will be used.

In this chapter, a comparison between patient-specific and patient-agnostic dictionary reconstruction with CS (presented in Chapter 3) in terms of HRV performance metrics is first presented. Then, using the RRM-DL-1024 technique detailed in Chapter 3 the performance of the proposed algorithm is compared in detail with SPIHT in terms of QRS detection accuracy. The resulting Wireless Transmission Savings (*WTS*) over SPIHT with the proposed algorithm are examined. Finally, visual examples of QRS complexes in reconstructed signals, where the data contains certain abnormal rhythms, are shown to illustrate the behaviour of the algorithm.

5.2. Heart Rate Variability

Heart Rate Variability (HRV) is a physiological method used to describe the variation in the interval between successive normal sinus heartbeats over time. It includes various time or frequency domain measurements that give an indication of the many physiological factors modulating the normal rhythm of the heart and can be used to assess overall cardiac health. The clinical relevance of HRV analysis is that significant changes to heart rate can be predictors for certain cardiac conditions. For example, HRV analysis can provide a means of observing the interplay between the sympathetic and parasympathetic nervous systems which is responsible for regulating cardiac activity, and, for example, reduced HRV may be indicative of myocardial infarction [113]. The use of HRV analysis is not strictly confined to cardiac conditions and it may also be used in the treatment of non cardiac diseases such as epilepsy, Alzheimer's disease and obstructive sleep apnea [89]. While certain diseases may be spotted by visually interpreting the ECG signal on a continuous basis, a clinician may also use HRV analysis as a convenient method for observing the ECG signals long-term. This ensures continuous signal monitoring is not required by the clinician and HRV is becoming a highly used diagnostic tool for in-depth cardiac analysis, particularly in CAD systems [113].

5.2.1. HRV Measures

A number of different parameters may be extracted that capture different aspects of HRV. Traditionally, these parameters are divided into two main categories: time and frequency domain measures. The metrics proposed in [24], were chosen for the analysis presented in this thesis. The time domain measurements are related to the feature extraction of the QRS complex and include *meanNN* and *SDNN*. *MeanNN* is the average value of the Normal-to-Normal (NN) intervals, where the NN interval is a measure of the time between R-peaks of adjacent QRS complexes. Therefore, variations in this parameter can provide a useful indication of an increase or decrease in cardiac rhythm. *SDNN* is the standard deviation of the NN intervals. It captures the total variability of the heart rate and it is the most commonly

used HRV metric. *SDNN* is useful for long-term monitoring and generally monitored over periods of up to 24 hours [114]. In terms of frequency domain measurements *HF power* and *LF/HF* ratio are used for the work presented in this thesis. *HF power* is the measured power in the high frequency (HF) range of 0.15 to 0.4 Hz. The physiological interpretation of the components in the HF band is a representation of respiratory sinus arrhythmia related to the parasympathetic control of the heart rate [115]. Finally, *LF/HF* ratio is the ratio between the power in the low frequency (LF) and HF components as measured in the frequency domain. In this instance the LF band ranges from 0.04 to 0.15 Hz and the LF components reflect the sympathetic control of heart rate [115]. In this work, the percentage similarity (*PSim*) between metrics derived from the reconstructed signals and the original uncompressed signal is investigated as a function of *CR* (*PSim* was introduced in Chapter 2). This gives a measure of the effects CS has on the HRV performance, and in particular, how well the HRV metrics are preserved.

5.2.2. Dictionary Construction

For this analysis two DL approaches are investigated: patient-agnostic and patient-specific. The concept of using patient-agnostic dictionaries was presented in Chapter 3. Patient-agnostic dictionaries are created by training with data from multiple patient records. Testing is then performed on a record not used to train the dictionary. Patient-specific dictionaries are created by training and testing on data from the same patient record. For HRV analysis the complete 48 records of the MIT-BIH Arrhythmia database were used. Forty-eight patient-agnostic and patient-specific dictionaries were created for each record in the database. Each agnostic dictionary was created using 29,798 (47x634) signals from the 47 other patient records. Patient-specific dictionaries were created with 600 adjacent frames from each record. The testing was then performed on the remaining frames. Additional frames were created for training by windowing between successive training frames. To mirror the patient-agnostic testing 29,798 frames were created using this windowing approach. All dictionaries created contained 20,000 atoms as preliminary testing in Chapter 3 indicated that performance gain was minimal beyond this point.

5.2.3. Evaluation of HRV Metrics

The effects of CS DL compression on the HRV diagnostic metrics are shown in Figure 5.1(a) and Figure 5.1(b), which present *PSim* as a function of *CR* for patient-agnostic and patient-specific DL approaches. A biomedical signal processing toolbox by Vidaurre *et al.* which provides the capability to analyse the chosen HRV metrics with reconstructed ECG signals was employed [116]. Although the primary focus of this analysis was the HRV metrics, Sensitivity (*SE*) and Specificity (*SP*) are also included to indicate the accuracy of QRS detection of each algorithm as the *CR* is increased. Note in Figure 5.1 *SE* and *SP* are expressed as normal on the y-axis and *PSim* relates specifically to the HRV metrics. As seen from Figure 5.1, the patient-specific dictionaries provide superior performance to the patient-agnostic dictionaries across all the tested metrics, as might be expected. In general, *SE*, *SP* and *meanNN* are more robust to higher *CR*s and demonstrate that both approaches still maintain a high rate of QRS preservation. At a *CR* of 25, the patient-specific dictionaries report accuracy of at least 89% for these three metrics. The agnostic dictionaries can obtain a *CR* of approximately 14 for similar levels of distortion. However, the *SDNN*, *LF/HF* and *HF power* metrics are less tolerant to errors as *CR* increases, for both patient-agnostic and patient-specific approaches. This indicates that the *SDNN* and HF information between QRS complexes is not accurately preserved at higher *CR*s.

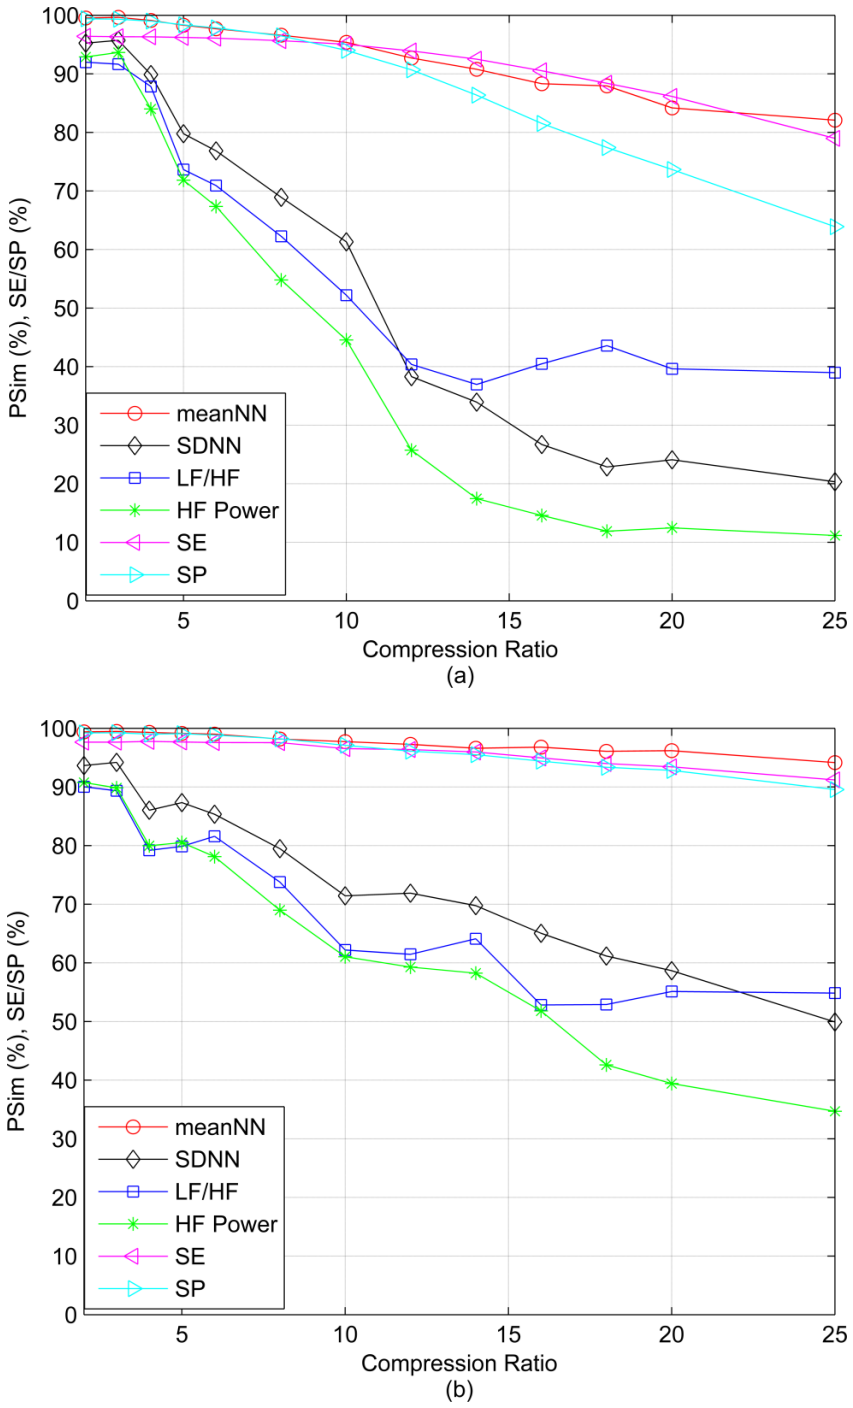

Figure 5.1: Average HRV metrics and *SE*/*SP* performance for (a) patient-agnostic and (b) patient-specific dictionaries. Note that both P_{sim} and *SE*/*SP* values are expressed on a percentage scale.

For the study shown here, no attempt is made to increase the *CR* using redundancy removal techniques or quantisation. Based on the fact that no

optimisations were made on the signal acquisition side for this test, obtaining far superior *CR*s with the employment of such techniques could be expected. The next section investigates this by analysing the *SE* and *SP* performance of the RRM-DL-1024 architecture.

5.3. RRM-DL QRS Detection Performance

In this section the diagnostic performance of the proposed compression algorithm is assessed based on its ability to preserve the QRS complex for an increasing *CR*. An advantage of applications relying on QRS detection is that much higher *CR*s can be obtained, while still maintaining high QRS detection accuracy. The nature of the reconstruction approach considered throughout this thesis is also advantageous compared to implementations that detect R-peaks on the sensor node and solely transmit this information as it supplies the clinician with the full reconstructed signal for review and offers the potential for further real-time CAD.

Afonso *et al.*'s [90] well known algorithm for beat detection was used to detect the QRS complexes in the compressed signals. The MIT-BIH Arrhythmia database [85] (which provides clinically-annotated QRS data) was again used to evaluate whether the algorithm can correctly detect the QRS complex at varying *CR*s. Based on the research of Friesen *et al.* [117] and the approach described in [24], a temporal tolerance of 88 ms was allowed either side of the annotated QRS complex to determine if a particular QRS complex was successfully identified.

Prior to previous work by the authors in [97], assessment of CS largely focused on signal fidelity, with little or no assessment of CS performance in diagnostic applications. SPIHT has previously been assessed by Twomey *et al*. for QRS detection accuracy in [24], which demonstrated that high accuracies of *SE* and *SP* can be reached up to a *CR* of 30. Twomey *et al.* found that for *CR*s above 30 the signal distortion introduced by compression begins to corrupt the detection performance of SPIHT, particularly in *SE* (increased number of false negatives (*FN*)). However, little or no published work exists comparing QRS detection accuracy of CS and SPIHT. Therefore, SPIHT is employed as a baseline state-of-the-

art compression technique to benchmark the performance of RRM-DL CS. To ensure consistency, SPIHT was implemented on the same patient records that the CS method was tested on and the same QRS detection technique was used on the SPIHT reconstructed signals. The SPIHT implementation employed by Higgins *et al.* in [91, 92] was utilised.

The results of the comparison are shown in Figure 5.2 in terms of *SE* (Figure 5.2(a)), *SP* (Figure 5.2(b)) and *AUC* of the ROC (Figure 5.2(c) and Figure 5.2(d)). Due to the superior performance of RRM-DL-1024 over RRM-DL-256 at higher *CR*s, only the RMM-DL-1024 method is used for this comparison. CS with the Huffman dictionaries of size $C = 2^6$ and 2^7 outperform SPIHT in terms of *SE* at all the *CR*s tested. However, SPIHT provides better performance in *SP*, sustaining 99% *SP* before experiencing a fall-off at *CR* above 60. CS provides a *CR* of 44 before encountering a similar fall-off. Both compression algorithms allow for significant savings, in terms of achievable *CR*s, while maintaining high performance in *SE*, *SP* and *AUC*. The 2^6 Huffman dictionary size provides higher *CR*s than SPIHT when both *SE* and *SP* are kept above a certain accuracy threshold for the range of values shown in Figure 5.2.

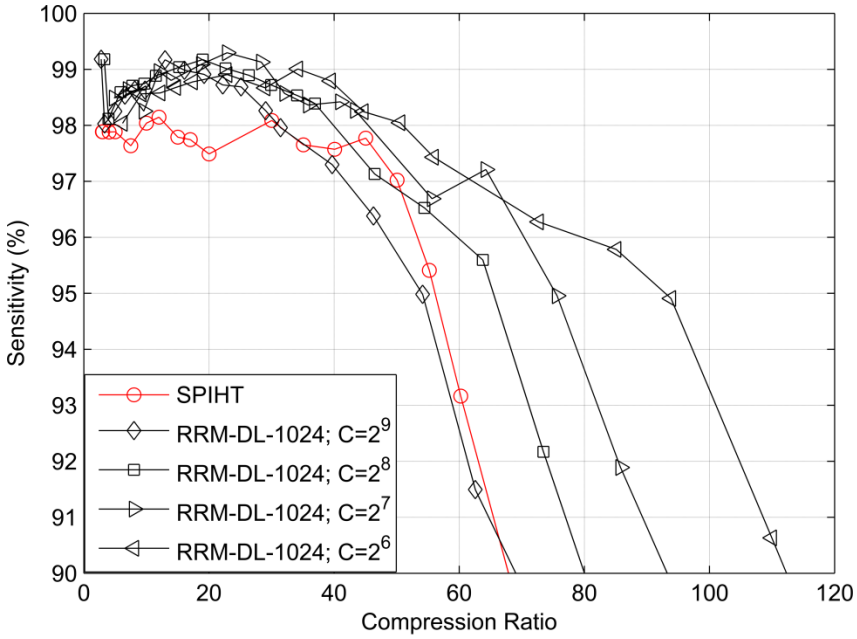
Sensitivity (%)
Compression Ratio
SPIHT
RRM-DL-1024; C=2^9
RRM-DL-1024; C=2^8
RRM-DL-1024; C=2^7
RRM-DL-1024; C=2^6

(a)

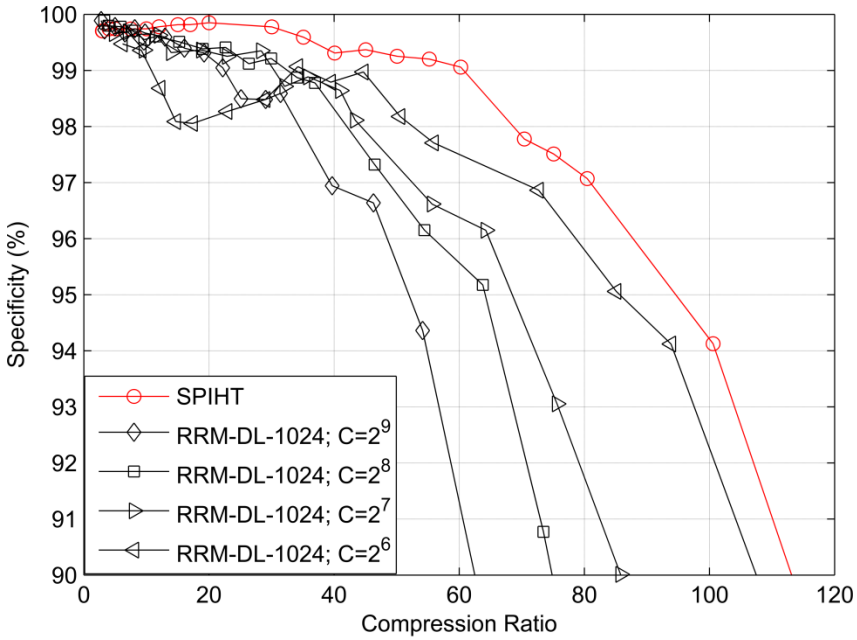
Specificity (%)
Compression Ratio
SPIHT
RRM-DL-1024; C=2^9
RRM-DL-1024; C=2^8
RRM-DL-1024; C=2^7
RRM-DL-1024; C=2^6

(b)

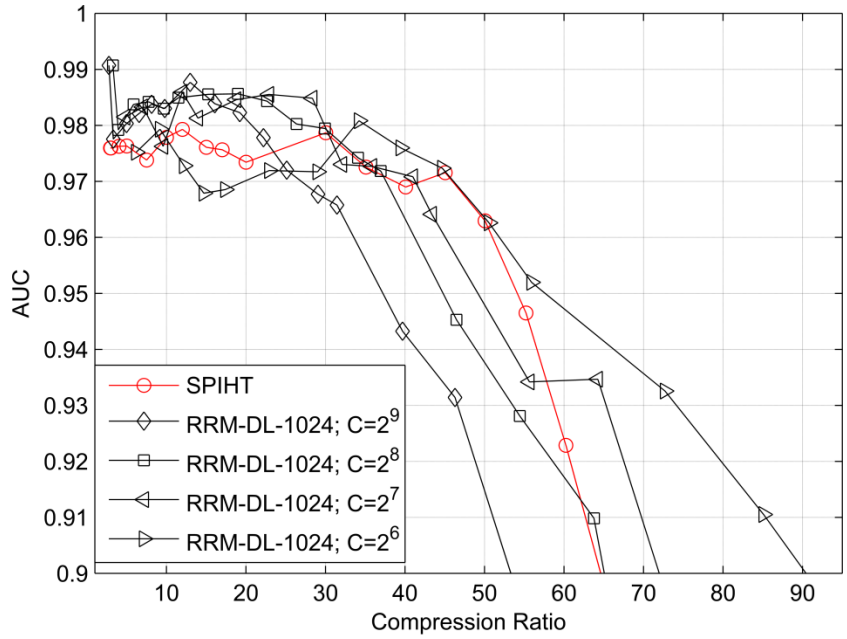

(c)

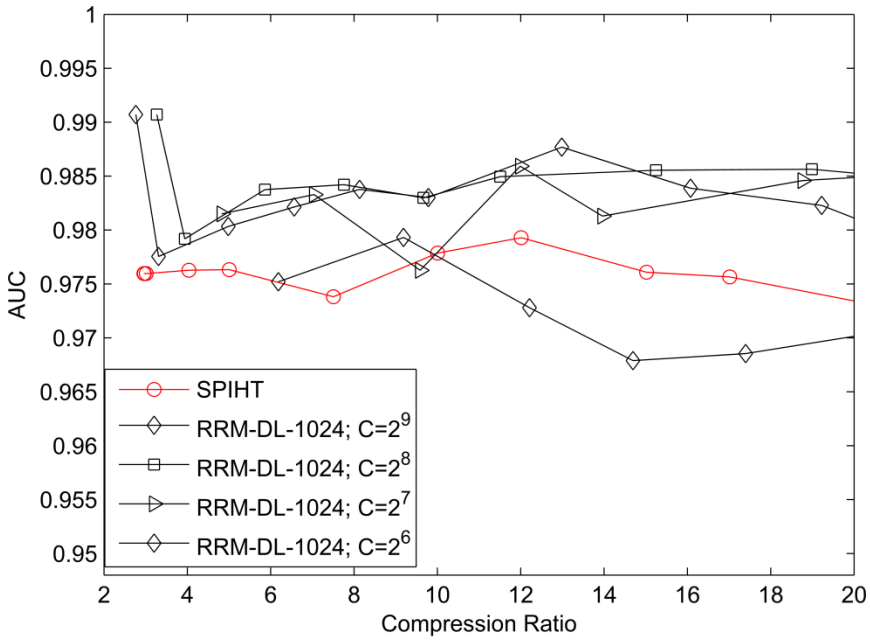

(d)

Figure 5.2: Comparison of RRM-DL-1024 and the SPIHT implementation in terms of *CR* for (a) *SE*, (b) *SP* and (c) *AUC* performance (with (d) showing a zoomed-in version of the high value *AUC* region corresponding to low *CR*).

AUC values in the range from 0.9 to 1.0 are considered to provide "excellent" classification in diagnostic applications, according to [118]. Therefore, the assessment of the proposed RRM-DL implementation will focus on this region. The *AUC* results show the proposed RRM-DL approach can outperform SPIHT in the high accuracy *AUC* region shown in Figure 5.2(d). Considering 0.98 *AUC* as an acceptable operating point (the highest *AUC* value SPIHT reaches in Figure 5.4), the optimal RRM-DL implementation ($C = 2^6$) reaches a *CR* of 35.15 compared to a *CR* of 30.05 with SPIHT. The 2^6 Huffman dictionary size achieves the best performance for all *AUC* values < 0.98 tested and at least one of the four implementations of RRM-DL outperforms SPIHT at all *AUC* values shown. For an application requiring an *AUC* of 0.9 (considered to be at the lower limit of acceptable performance), the CS approach can compress the signals at a *CR* of 90.45 compared to 64.6 with SPIHT. Despite outperforming the 2^6 Huffman codebook size in *SP*, the *AUC* of SPIHT is significantly impacted by its relatively poorer *SE* performance. Analysis of reconstructed signals at higher *CR*s illustrates that the reconstructed SPIHT signals contain more *FN*s (lower *SE*) and fewer false positives (*FP*) (higher *SP*) than the CS implementations. Visual evaluation of reconstructed signals at higher *CR*s (Figure 5.3, $CR \approx 100$) highlight SPIHT's *SE* performance degradation is due to a significant reduction in the amplitude of the QRS peaks resulting in a failure in identifying the distorted QRS complex. The inferior *SP* performance of CS compared to SPIHT results from the presence of noisy high frequency components between QRS complexes which leads to the incorrect detection of QRS complexes with CS, and more *FP*s.

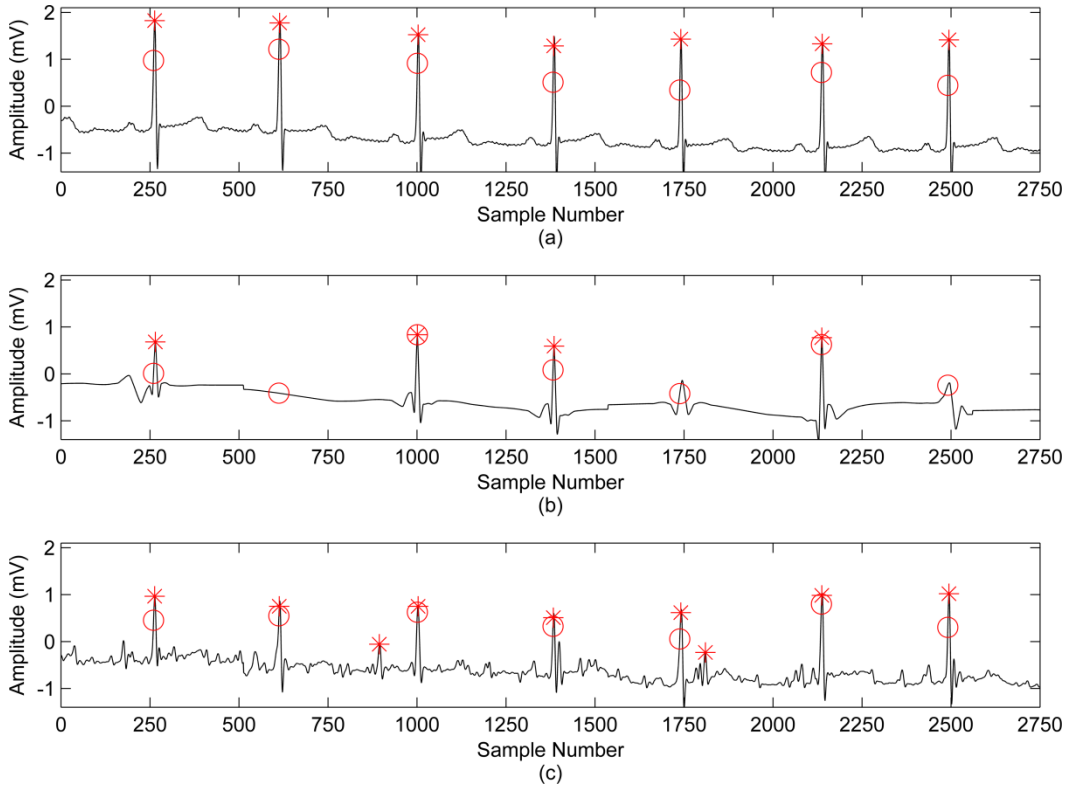

Figure 5.3: Reconstructed ECG segments from record 115. The annotated QRS complexes are denoted with a circle and the detected QRS complexes are denoted with an asterisk. (a) Original signal with no *FP*s or *FN*s. Reconstructed signals at a $CR \approx 100$ for (b) SPIHT showing 3 *FN*s and (c) CS showing 2 *FP*s.

As shown in Table 5.1, RRM-DL-1024 enables a reduction in wireless transmission power expenditure (reflected in *WTS*) over SPIHT, for high values of *AUC* reflecting strong preservation of clinical information. Considering the advantages of CS over DWT-based compression algorithms in terms of acquisition complexity [62] and computational complexity [63], this result (Table 5.1) supports the potential of the RRM-DL-1024 approach over SPIHT for QRS detection applications. For example, at an *AUC* performance threshold of 0.9, the proposed CS algorithm offers approximately 40% *WTS* over SPIHT. This comparative analysis does not account for the additional benefits CS offers in terms of power savings in signal acquisition and compression (which will be analysed later in this thesis).

Table 5.1: *WTS* of RRM-DL-1024 over SPIHT for a varying *AUC*

	CR		RRM-DL-1024 *WTS* (%) over
AUC	RRM-DL-1024	SPIHT	SPIHT
0.97	46.1	45.95	**0.32**
0.96	51.87	50.99	**1.72**
0.95	57.62	54.12	**6.46**
0.94	66.27	56.6	**17.08**
0.93	74.14	58.72	**26.2**
0.92	79.76	60.78	**31.22**
0.91	85.36	62.7	**36.14**
0.9	90.45	64.6	**40.01**

5.4. Compression of ECG with Abnormal Rhythms

The MIT-BIH Arrhythmia database directory [85] provides annotations of the various beat types in each record. Despite mainly containing normal sinus rhythms many of the records contain different rhythms such as atrial fibrillation, bradycardia and tachycardia. The database contains various abnormal beats such as bundle branch block beats, atrial premature beats, fusion beats, paced beats and PVCs. It is of interest to examine the performance of a compression algorithm in reconstructing different beats which differ from normal sinus heartbeats. For clinical diagnostics it is important that the compression algorithm sufficiently preserves the information in these beats. This also signifies the useful nature of the MIT-BIH Arrhythmia database as it ensures the compression algorithms are tested over a wide range of ECG signal characteristics.

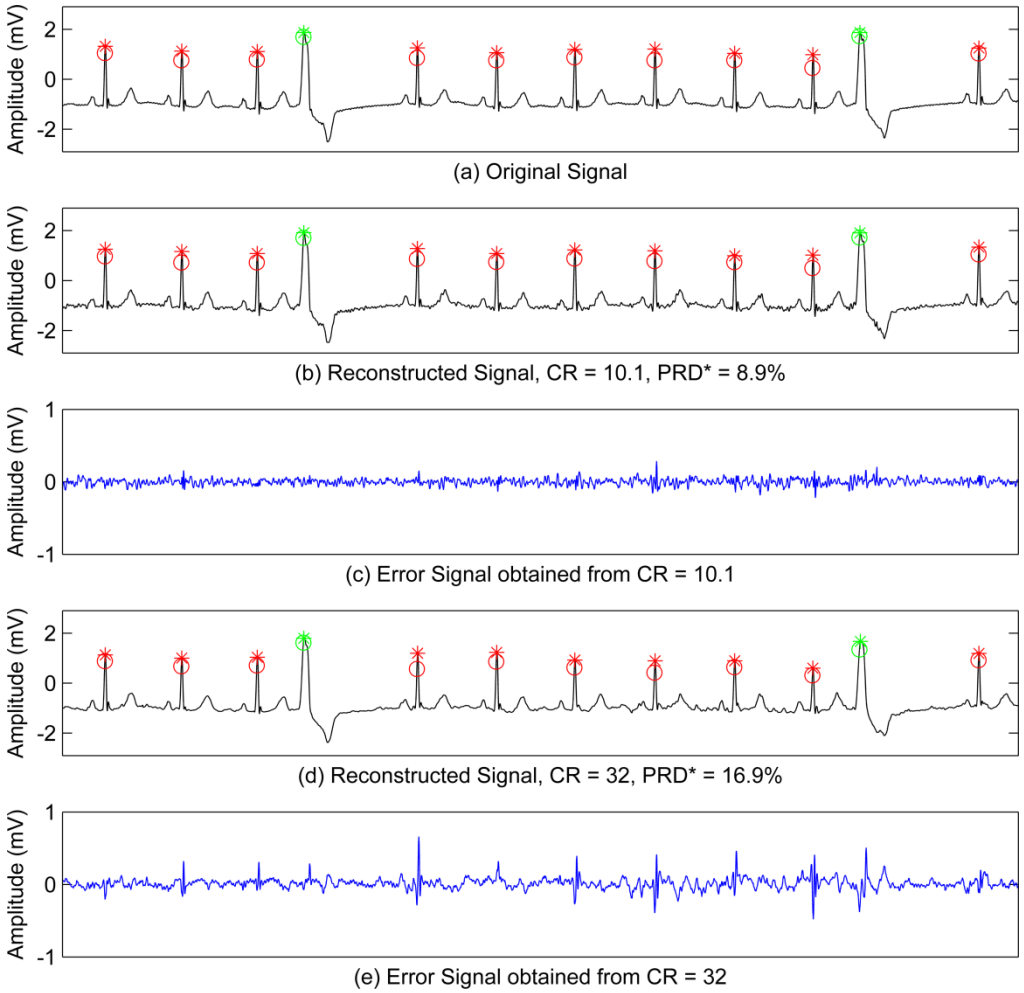

Figure 5.4: Examples of an original signal and two corresponding reconstructed signals over a 10 second time window taken from record 119. (a) Original signal; Reconstructed signals with a *PRD** of (b) 8.9% and (d) 16.9% are shown for comparison. The annotated QRS complexes are denoted with a circle and the detected QRS complexes are denoted with an asterisk. Normal beats are shown in red and the two annotated PVC beats are shown in green. The error signals (c) and (e) corresponding to each reconstructed signal are also shown. For all the plots, the x-axis represents Sample Number.

Figure 5.4 illustrates the performance of the QRS detection algorithm on an original ECG signal and two signals reconstructed with CS at varying *CR* values. The signals shown are over a window of 10 seconds and in each instance all of the QRS complexes are detected. The 10 second duration signal in Figure 5.4(a) contains ten normal beats and two PVC beats. Figure 5.4(b) shows an example of a reconstructed signal for a *CR* of 10.1, which results in a *PRD** close to 9%, with the error shown in Figure 5.3(c). However, Figure 5.4(d) indicates that high detection performance is maintained even at a much higher *CR* of 32 (with *PRD** = 16.9%),

with all QRS complexes detected. Despite the higher *CR* adding more distortion than in the low *CR* case, the signal in Figure 5.4(d) still preserves important clinical attributes of the ECG signal and the occurrence of the two abnormal premature ventricular contraction beats do not significantly impact the reconstruction quality. It can also be noted that the S and T waves are well preserved and this contributes to important diagnostic information about the signal.

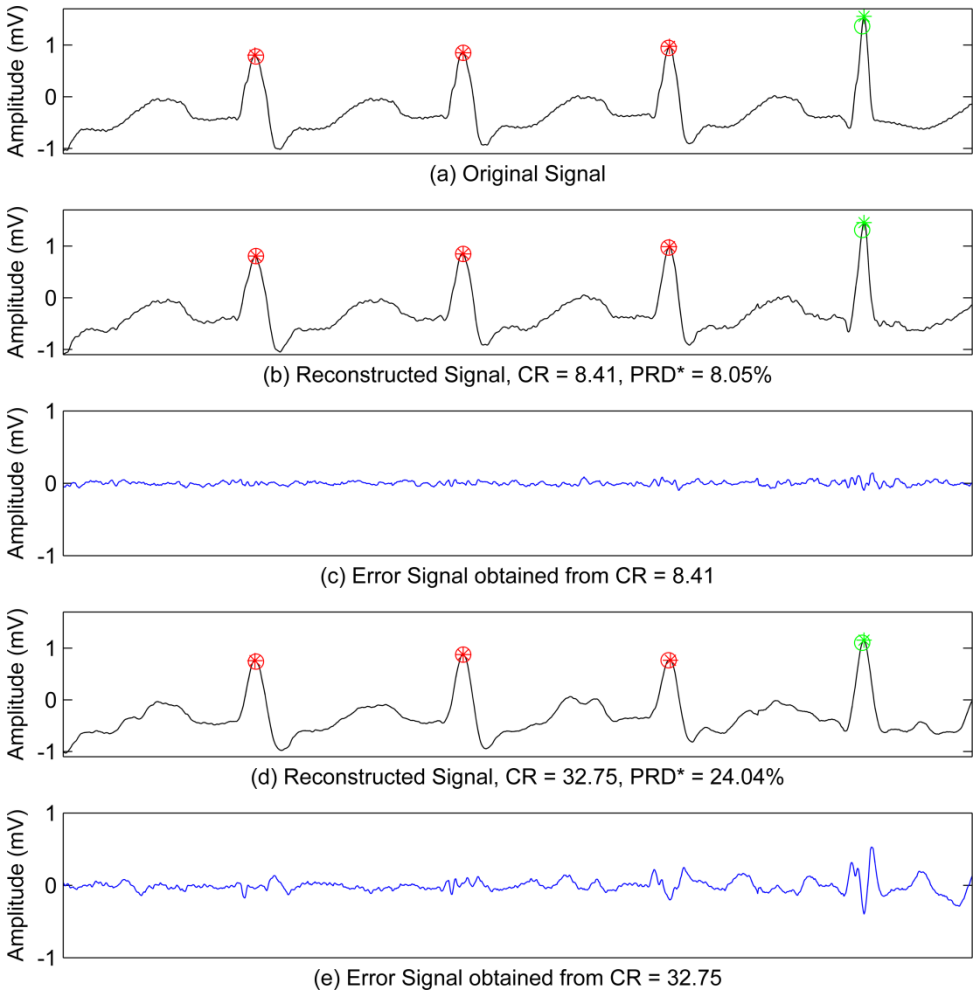

Figure 5.5: Examples of an original signal and two reconstructed signals over a 1 second time window taken from record 109. (a) Original signal; Reconstructed signals with a *PRD** of (b) 8.05% and (d) 24.04% are shown for comparison. The annotated QRS complexes are denoted with a circle and the detected QRS complexes are denoted with an asterisk. Normal beats are shown in red and the correctly-detected ventricular fusion beat is shown in green. The error signals (c) and (e) corresponding to each reconstructed signal are also shown. For all the plots, the x-axis represents Sample Number.

An additional example is shown by Figure 5.5 which shows reconstructed signals over a 1 second window where again all QRS complexes are detected. In this example the existing artefact is a ventricular fusion beat which causes a considerable variation in the height of the QRS complex. The significance of this ventricular fusion beat is that the set of training signals for this particular dictionary contained no annotated fusion beats and therefore demonstrates the ability of the learned dictionaries to adapt to new signal abnormalities that may occur. This is evident in the reconstructed signal from Figure 5.5(b) where the reconstructed error signal (Figure 5.5(c)) shows no significant change at the time of the fusion beat relative to the rest of the signal shown. In the case of the signal in Figure 5.5(d) there is a slight increase in the error of the fusion beat relative to the rest of the signal. However, this is at a much higher *CR* of 32.75 and the QRS complex is still detected at this rate.

5.5. Summary

This chapter examined the proposed CS algorithm based on its ability to maintain diagnostic performance for increasing *CR*. Diagnostic performance was expressed in terms of typical HRV metrics and QRS detection accuracy. The comparison between patient-specific and patient-agnostic dictionaries showed the superior performance of patient-specific dictionaries. For example, the patient-specific approach maintained a minimum of 89% *PSim* in terms of *meanNN*, and similar percentage accuracy for *SE* and *SP*, at *CR*s up to 25. It also provided superior performance to patient-agnostic dictionaries in terms of *SDNN*, *LF/HF* and *HF power*. The HRV results presented in Section 5.2 were obtained where no additional strategies were employed to extend the *CR*. Therefore, better performance could be expected with the use of redundancy removal, quantisation and Huffman coding strategies; this was investigated in Section 5.3. The results showed significant compression gains can be achieved when adding these operations, in particular the *CR* can be approximately doubled for a similar level of performance.

The RRM-DL-1024 implementation was chosen for the analysis of QRS detection accuracy because of the performance it offers at high *CR*s relative to RRM-

DL-256. The performance was compared with SPIHT at all quantisation rates. When analysing QRS detection rate, 5 bit quantisation ($C = 2^6$) provides best performance, superior to that of SPIHT. With regard to QRS complex extraction, RRM-DL compares well with the DWT-based SPIHT lossy compression technique. RRM-DL can achieve a *CR* of 90.45 whilst maintaining an *AUC* of 0.9 in QRS detection accuracy whereas SPIHT provides a *CR* of 64.6 for the same AUC, amounting to a 40% decrease in wireless transmission costs with RRM-DL. The work described in this chapter appears in [65, 97].

The next chapter will analyse in detail the proposed CS methods in terms of overall power consumption. The benchmark algorithms with which the proposed algorithms were compared will also be assessed and used for comparison purposes.

CHAPTER 6

Sensor Power Consumption Comparison for CS and SPIHT

6.1. Introduction

The importance of signal compression in helping to minimise energy costs in BANs has already been identified. Minimising energy consumption to extend battery life is one of the most critical design aspects for any BAN architecture. Therefore, it is essential that compression algorithms can provide high quality reconstruction performance at high *CR*s while ensuring that the algorithm itself does not impact significantly on the overall energy consumption of the wearable device. CS holds many advantages over alternative lossy compression algorithms due to its low complexity signal acquisition operation and the lack of a need for digital signal processing operations (which would typically be implemented in embedded software).

Up to this point, this thesis has focused on the *functional* performance of CS, and in particular has proposed and demonstrated the performance of a number of novel CS approaches that improve on the state-of-the-art in this area. This chapter aims to further confirm the potential of the proposed techniques by analysing the

algorithms in terms of *total system power consumption* (analog and digital) in a wearable sensor used in a BAN. A detailed power comparison is presented over a wide range of *CR*s based on the ability to maintain suitable *PRD* performance between the Adaptive Dictionary (AD) reconstruction technique proposed in this thesis, existing high performance CS implementations, and SPIHT. The power consumption profiles were created by considering (i) recent CS encoding techniques from the literature; (ii) the power consumption of digital signal compression using an Analog Devices Blackfin DSP processor (BF537) as an example platform [119]; (iii) wireless transmission via Bluetooth Low Energy (BLE) using a Texas Instruments CC2540 BLE transceiver [120]. Note that since the focus of interest is increasing battery life in the wearable device, only the power consumption of the wearable sensor is considered here; the power required for reconstruction is not considered, as it is assumed (for both CS and SPIHT) that reconstruction takes place in a server node that is not energy-constrained.

6.2. Power Consumption Profiles

The CS and SPIHT architectures assumed in this thesis for acquisition, compression and transmission are detailed in Figure 6.1. The power consumption of each algorithm is determined by calculating the power required in (i) sampling the signal, (ii) performing digital signal compression and (iii) wireless transmission of the ECG data via a BLE wireless transceiver [120]. The proposed AD algorithm will be compared against the Modified Model-Based (MMB) and Block Sparse Bayesian Learning (BSBL) methods detailed in [63], and the SPIHT compression algorithm [16].

The review of existing CS analog acquisition architectures in Chapter 2 identified the SRMPI design proposed by Mamaghanian *et al.* [62] as a high performing CS acquisition system. Therefore, the random sensing of CS measurements based on SRMPI is assumed here. This architecture was also employed in [63] when the MMB and BSBL methods were analysed, which is useful for comparison purposes. The SRMPI design consists of mixers, integrators and a

sub-Nyquist rate ADC. Additionally, the functional advantages of the use of a redundancy removal module, quantisation and Huffman coding have been established in this thesis and their cost will be included as part of the digital signal compression operations for each CS implementation prior to BLE wireless transmission.

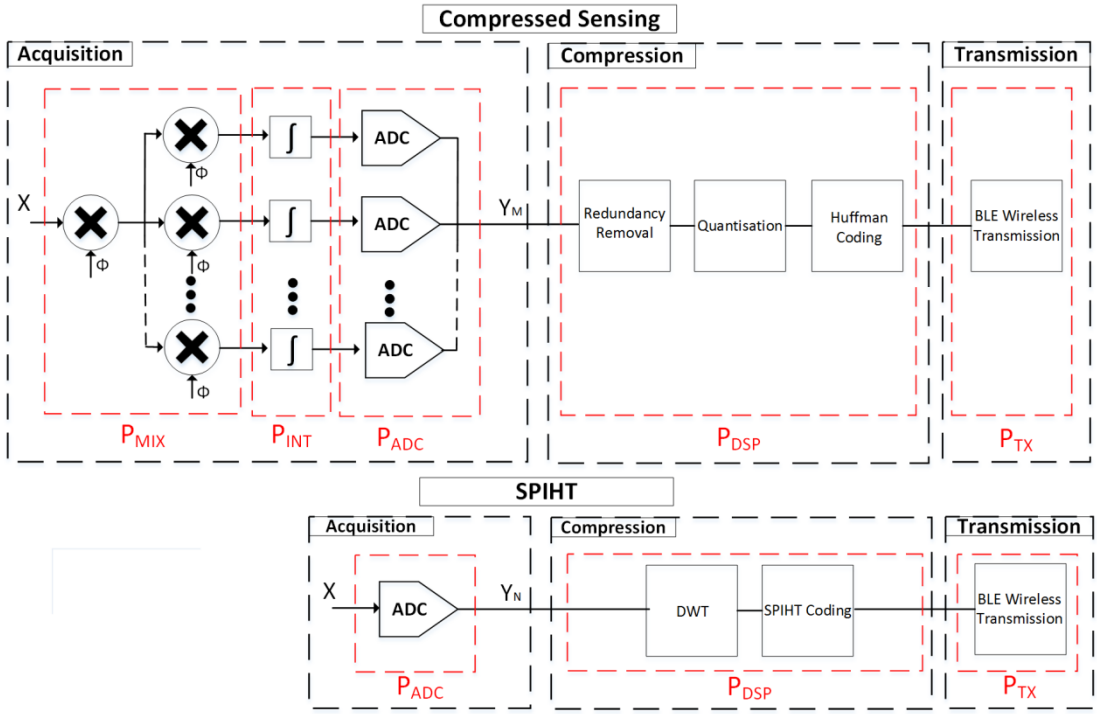

Figure 6.1: Block diagrams for the acquisition, compression and transmission operations for CS and SPIHT.

In relation to SPIHT, the sampling of the signals occurs at the Nyquist rate. As a CS model where many modifications have been made to enhance the energy efficiency of its implementation has been presented, for this comparison a CDF 5/3 biorthogonal lifting-based integer wavelet transform was used for the DWT operation. The wavelet is computationally efficient to implement due to the simplicity of the multiplication steps and the limited number of lifting steps required compared with other wavelets [121, 122]. It also removes the need for quantisation of the signal as the DWT output is integer-based. The DWT was performed using 8

levels of decomposition and SPIHT coding techniques were applied to the output of the DWT to construct a compressed binary stream. This method slightly affects the signal reconstruction quality performance of SPIHT compared to the CDF 9/7 wavelet employed previously. However, the more efficient compression operations ensures that it improves the overall energy efficiency of SPIHT and provides a fairer benchmark for comparison with the CS-based implementations.

The following sub-sections explain the methodology used to estimate power consumption in each major component of the signal compression system: signal acquisition; digital signal compression; wireless transmission. The variables used in the power analysis are defined in Table 6.1, and their values given where appropriate. The values for M, AEC and B_T are not given as they are variable, depending on the level of compression.

Table 6.1: Variables used in Power Calculations, with values where appropriate

Variable	Value	Notes
f_s	360 Hz	ECG sampling frequency [85]
B	11	ECG bit resolution [85]
N	-	ECG frame size in samples (CS-192, SPIHT - 1024)
V_{DD}	0.6 V	CS sampling integrator operating voltage [60]
M_X	0.2938 nJ	Mixer energy per CS conversion (J/sample) [62]
C	16	Number of SRMPI channels
C_P	10 fF	Integrator capacitance [60]
M	-	Number of compressed measurements.
FOM	10 fJ	Figure-Of-Merit for ADCs per conversion [123]. See text for details.
I_{DD}	15 mA	Blackfin baseline dynamic current [124]
ASF	1	Blackfin Activity Scaling Factor [124]. See text for details.
V_{BF}	0.8 V	Blackfin operating voltage [124]
AEC	-	Average Execution Cycle in VisualDSP++
CLK	100 MHz	Blackfin core clock frequency [124]
B_T	-	Number of bits transmitted per frame

6.2.1. Power Consumption of Signal Acquisition

Mixer: Mamaghanian *et al.*'s [62] SRMPI design implemented mixers using the single pole, double throw CMOS switch (ADG636) from Analog Devices. The

sub-Nyquist rate mixers were operating at half the Nyquist sampling frequency and a similar approach is utilised for the analysis here. The mixer power (P_{MIX}) is defined in (6.1), based on the assumptions used in [62] (the quantities on the right hand side of the equation have been defined in Table 6.1):

$$P_{MIX} = (f_s)(M_X) + (C)\left(\frac{f_s}{2}\right)(M_X) \tag{6.1}$$

Integrator: The integrator power estimation (P_{INT}) is taken from Chen *et al.*'s CS implementation [60].

$$P_{INT} = \frac{(C)\left(V_{DD}{}^2\right)(10pi)(f_s)(C_P)}{16} \tag{6.2}$$

ADC: The SPIHT power model assumes Nyquist rate (f_s) sampling. In contrast, the CS model effectively operates at a sub-Nyquist rate of (f_s)(M/N). The Figure-Of-Merit (*FOM*) is a design specification commonly used in the literature for comparison of ADCs and for calculating the power efficiency of a converter [60, 125]. Based on a recent survey of ADCs [123] that concluded the most energy efficient ADCs are Successive Approximation Register (SAR) ADCs which can provide a *FOM* of < 10fJ per conversion. The ADC power for each algorithm (P_{ADC}) is then expressed as in (6.3). Note that with SPIHT, M is fixed and is equal to the frame size N. For CS, M is alternating and is reduced as the *CR* increases.

$$P_{ADC} = (M)(FOM)(2^B)(f_s) \tag{6.3}$$

Overall Acquisition Power: The total power (P_{ACS} and P_{ASPIHT}) equations for each acquisition implementation are given in (6.4) and (6.5).

$$P_{ACS} = P_{MIX} + P_{INT} + P_{ADC} \tag{6.4}$$

$$P_{ASPIHT} = P_{ADC} \tag{6.5}$$

6.2.2. Power Consumption of Digital Signal Compression

The computational complexity of each compression algorithm is assessed by determining the average number of clock cycles to compress 1 second of ECG data for each *CR* tested. The number of clock cycles was calculated for the Analog Devices Blackfin BF537 DSP using the Analog Devices Visual DSP++ code execution profiling tool. Different operating characteristics associated with the Blackfin BF537 DSP were derived from [124]. It is reported in [124] that the "typical activity" of a Blackfin processor with an operating voltage (V_{BF}) of 0.8 V and a core clock frequency (*CLK*) of 100 MHz will have a dynamic current (I_{DD}) of 15 mA and an Activity Scaling Factor (*ASF*) of 1.

The SPIHT implementation can be divided into two main processes: the DWT of the input signal and the SPIHT encoding operation. The CS implementation entails a redundancy removal operation, quantisation and Huffman coding. The Average Execution Cycle (*AEC*) count of each operation at different *CR*s was obtained and the power required (P_{DSP}) for each algorithm was calculated as in (6.6), based on [124].

$$P_{DSP} = (I_{DD})(ASF)(V_{BF})\left(\frac{AEC}{CLK}\right)\left(\frac{f_s}{N}\right) \tag{6.6}$$

6.2.3. Power Consumption of Wireless Transmission

Siekkinen *et al.* [126] measured the power consumption of a Texas Instruments CC2540 BLE transceiver [120] as 84mW during wireless transmission. The recorded duration to transmit one byte of data was 8μs. It can then be estimated the energy required to transmit one bit of data is 84nJ using BLE. Note this analysis

assumes the transceiver is switched off between transmissions. Therefore, the power consumed during wireless transmission (P_{TX}) for each algorithm is dependent on the level of compression and is calculated as follows:

$$P_{TX} = (84nJ)(B_T)\left(\frac{f_s}{N}\right) \tag{6.7}$$

6.3. Performance Assessment

This section is divided into three parts. Firstly, the power consumption of CS and SPIHT will be given in terms of signal sampling, digital signal compression and wireless transmission. Secondly, the performance of the chosen CS and SPIHT algorithms in terms of signal distortion metrics are introduced to analyse the performance of each algorithm in terms of *CR*. Finally, the effects of patient-specific training on the overall energy efficiency of a CS implementation are investigated.

6.3.1. Power Consumption

Figure 6.2 illustrates the power consumption of the three algorithm operations (sampling, compression and transmission) for CS and SPIHT at varying *CR*s (note these CS values apply to both SD and AD reconstruction approaches). The impact of wireless transmission is significant for both algorithms particularly at the lower *CR*s. With regard to SPIHT, at *CR*s beyond 15, P_{DSP} consumes the most power. In fact, the most significant difference between the two algorithms is P_{DSP}. Despite the decreasing *AEC* for SPIHT with increasing *CR*, the DWT element remains constant irrespective of *CR*, ensuring that the drop-off in P_{DSP} is gradual. In contrast, with CS, each of the quantisation, Huffman coding and redundancy removal processes decrease in complexity as the number of measurements is reduced. The average clock cycles and power dissipation of each digital operation at varying *CR*s are shown in Table 6.2, assuming a Blackfin BF537 platform implementation.

Table 6.2: Average Clock Cycles per 1 second for each operation

	CR	DWT	SPIHT Coding	P_{DSP} (µW)
SPIHT	2	49034	244563	**35.2**
	5	49034	232510	**34.1**
	10	49034	151287	**24.0**
	20	49034	81680	**15.7**
	N/M	RR and Huffman Coding	Quantisation	P_{DSP} (µW)
CS	2	20717	96493	**14**
	5	8212	48980	**6.8**
	10	3990	33680	**4.5**
	20	1879	24699	**3.1**

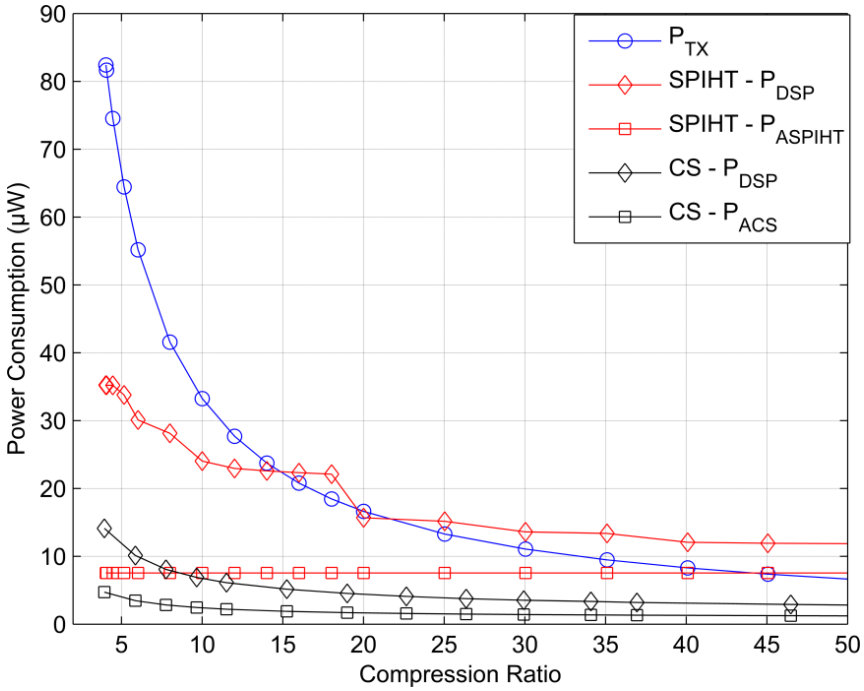

Figure 6.2: Power consumption profiles of individual components for each compression algorithm. The power consumption of each process is shown for SPIHT and CS. Note that the wireless transmission power (P_{TX}) is the same for both algorithms, for a given *CR*.

The acquisition power profiles for each algorithm are also shown in more detail in Table 6.3. SPIHT consumes 7.55 µW during signal acquisition (constant, as Nyquist sampling is used). At a *CR* of 2, CS consumes 4.73 µW and P_{ACS} decreases

as compression increases. This is due to a reduction in the number of measurements (*M*) required which reduces P_{ADC} for CS.

Table 6.3: Acquisition Power Consumption for SPIHT and CS

	N/M	P_{MIX} (µW)	P_{INT}	P_{ADC} (µW)	$\boldsymbol{P_{ACS}/P_{ASPIHT}}$
SPIHT	-	-	-	7.55	**7.55**
CS	2	0.95	0.04	3.77	**4.73**
	5	0.95	0.04	1.51	**2.46**
	10	0.95	0.04	0.75	**1.70**
	20	0.95	0.04	0.37	**1.33**

6.3.2. Signal Distortion

In the previous sub-section, the power consumption of each operation of the SPIHT and CS algorithms for various *CR*s was analysed. This section now compares the AD approach with the MMB approaches, BSBL and SPIHT in terms of overall system power consumption by using signal reconstruction performance as a guide for suitable obtainable *CR*s. The ECG signals in the MMB and BSBL implementations in [63] are acquired and compressed using the redundancy removal, quantisation and Huffman encoding techniques described in this thesis and therefore there is consistency in the comparison. The AD-Q6 implementation was considered for this analysis. The results are shown in Figure 6.3. As AD-Q6 was shown to outperform the existing CS approaches in terms of *PRD* in Chapter 4, ultimately leading to a reduction in wireless transmission costs, as expected the results show that it outperforms other CS approaches in terms of overall power consumption (as the acquisition and compression techniques are considered the same). SPIHT has the advantage of the minimum cost of wireless transmission, however the costs associated with acquisition and in particular digital signal compression compared to the CS approaches significantly impact on its relative power consumption.

For all *PRD* values considered, AD outperforms the other existing CS techniques. Also for *PRD* values > 3%, AD-Q6 is the best performing algorithm offering superior energy efficiency to SPIHT. However, SPIHT performs best in

terms of power consumption for *PRD*s less than 3%. The decline in reconstruction quality for increasing *CR*s is sharper for the MMB and BSBL techniques and SPIHT outperforms MMB-CoSAMP and BSBL at all *PRD*s shown despite the advantages in lower power sampling and compression operations with these CS variants. For *PRD* values > 4.5% MMB-IHT outperforms SPIHT for the *PRD* range shown. However, it is clear the performance of the AD is superior to the CS techniques tested and this advantage increases as the level of *PRD* increases.

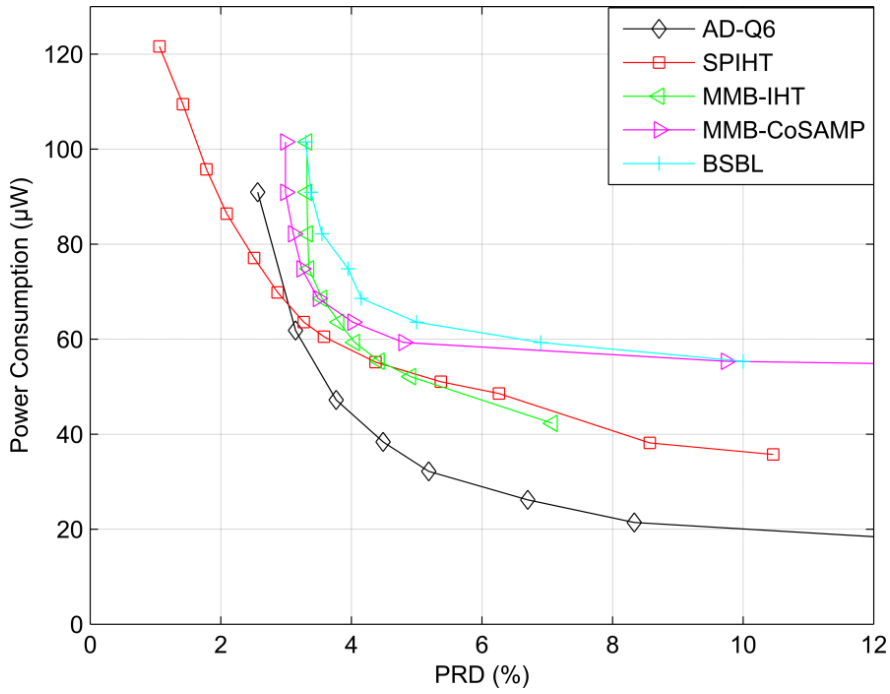

Figure 6.3: Total system power consumption of the AD-Q6 implementation, SPIHT, MMB-IHT, MMB-CoSAMP and BSBL for an increasing *PRD* value.

6.3.3. Signal Training - Use Case Analysis

A practical consideration of the AD technique is the requirement for the collection of data for a training phase. Here two use cases for training with the AD approach are presented and the effects on long-term monitoring are analysed by estimating the average power consumption over various time periods. AD-Case 1 refers to the case where training is performed in a clinical environment prior to

implementation, ensuring there are no power costs associated with training to the BAN. AD-Case 2 is the scenario where training is performed on the device itself. The architecture considered in this thesis requires a training phase involving the collection of 15 minutes of ECG data. Therefore, AD-Case 2 assumes Nyquist rate sampling for 15 minutes prior to the proposed AD implementation being employed. For comparison purposes, the two use cases will be compared with the MMB-IHT algorithm (as a benchmark CS algorithm) and SPIHT. It should be noted the MMB approaches utilise Huffman dictionaries in their implementation which require offline creation. However, this is not included in this analysis as training periods are not specified in [63, 64] and straightforward implementation of the MMB-IHT technique is assumed. Similarly, there is no requirement for a training period for SPIHT so the average power consumption stays constant in that case. At a *CR* of 14.8 the AD-Q6 method achieves an average *PRD** of 9% (deemed acceptable for clinical inspection in [81]) and *PRD* of 5.5%. As only average *PRD* (rather than *PRD**) values for were reported MMB-IHT [64] (around the target *PRD**), a *PRD* of 5.5% will be considered the operating point for this use case analysis. As demonstrated by Figure 6.3, at the target *PRD* of 5.5%, AD-Q6 consumes 30.9 µW, SPIHT consumes 50.65 µW and MMB-IHT consumes 49.4 µW.

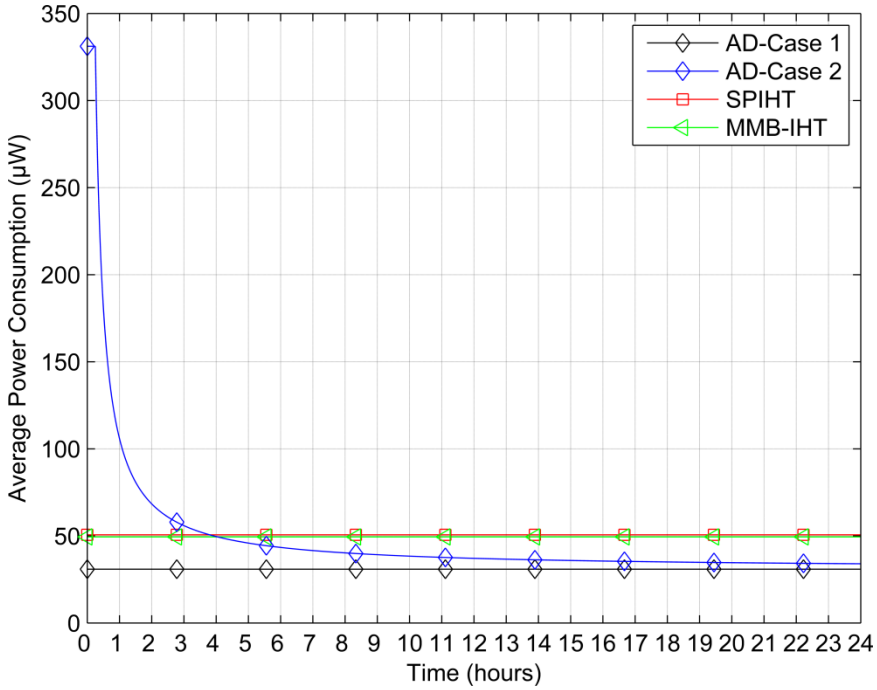

Figure 6.4: Comparison of two use cases presented with MMB-IHT and SPIHT over various time durations at an average *PRD* of 5.5%.

Figure 6.4 highlights the average power consumption, showing the effects of training where applicable (AD-Case 1), of the two use cases, MMB-IHT and SPIHT over a period of one day (24 hours). Significantly the results in Figure 6.4 demonstrate the effects of the additional power consumption used in training are minimal in applications where long-term monitoring is required. The average power consumption of AD-Case 2 converges on the ideal AD-Case 1 average power as time elapses. From a period of approximately 3 hours and 45 minutes, the analysis suggests AD-Case 2 outperforms SPIHT in terms of average power consumption. It takes approximately 4 hours for the power consumption of AD-Case 2 to drop below that of the MMB-IHT implementation. However, once this once-off training phase is complete and sufficient monitoring time has passed the efficiency of AD-Case 2 compared to MMB-IHT and SPIHT increases due to the achievable *CR* and low complexity sampling and compression as shown by Figure 6.2. This analysis supports the advantages of the proposed AD approach to reduce power consumption, compared to both SPIHT and existing CS algorithms.

6.4. Summary

This chapter has examined several CS compression implementations and SPIHT in terms of overall power consumption. The power consumption required to compress ECG data was analysed at various *CR* and *PRD* performance levels. The results verified the value of CS as a low-power compression algorithm suitable for BAN implementation. In terms of signal acquisition and digital signal compression, the results presented verified that the CS methods tested consumed less power than high-performance non-CS lossy algorithms such as SPIHT. The evaluation of power consumption in terms of signal distortion showed that the proposed AD algorithm provides a more energy efficient implementation than all existing CS implementations. Similarly, for *PRD* values > 3%, AD outperforms the SPIHT implementation presented here. It was estimated that at a *CR* of 14.8 the AD-Q6 method achieves an average *PRD** of 9% (deemed acceptable for clinical inspection in [81]) and *PRD* of 5.5%. At the target *PRD* of 5.5%, AD-Q6 consumes 30.9 μW, SPIHT consumes 50.65 μW and MMB-IHT consumes 49.4 μW. Based on these savings, an analysis of the power consumption of different signal training approaches was presented. The results provide a positive outcome showing that, for long-term monitoring applications, the impact of power consumption due to training decreases with increasing monitoring time, therefore the CS approach would be suitable in long-term low-power ambulatory monitoring scenarios. The power consumption comparison presented in this chapter appears in [110].

CHAPTER 7

Conclusions and Future Work

7.1. Introduction

This final chapter of the thesis summarises the work presented and highlights the primary contributions made to the field. Each of the core contributions is summarised, and several possibilities for future work are identified.

7.2. Summary of Thesis

This thesis has explored the use of lossy compression in a Body Area Network (BAN) to provide a low-power operation which can enable an increase in the battery life of mobile or wearable devices. More specifically, in relation to the monitoring of electrocardiogram (ECG) signals, different implementations of Compressed Sensing (CS) were proposed. A critical performance goal of this research was the ability of the proposed compression techniques to maintain diagnostic performance in the reconstructed signals despite the reduction in the amount of data required to represent the signal.

The opening chapters of the thesis provided an overview of BANs and highlighted the importance of the use of signal compression in supporting an extension in battery life of a wearable device/sensor. Different lossy compression

methods were discussed and compared based on their performance with ECG signals in terms of preserving signal reconstruction quality. A detailed review of CS was also presented identifying the potential it possesses for use in low-power ambulatory monitoring architectures.

In Chapter 3 the proposed RRM-DL CS architecture, which combines a novel Redundancy Removal with Mean (RRM) technique with quantisation and Dictionary Learning (DL) for signal reconstruction, was presented. The RRM-DL architecture was compared with the best known existing CS approaches in terms of signal reconstruction quality for a range of Compression Ratios (*CR*s). Further to this approach a novel Adaptive Dictionary (AD) reconstruction framework for CS-based ECG compression was proposed in Chapter 4. The AD method also uses trained dictionaries for reconstruction but aims to exploit knowledge of the signal from an initial reconstruction pass, to determine a more suitable dictionary for a refined reconstruction. By assessing whether a QRS complex is present in an individual frame and if so, determining the location of that QRS complex, an appropriate dictionary for a second pass providing more refined reconstruction is selected. Similar to the RRM-DL technique, the AD technique was compared with the existing CS methods in the literature.

The next body of research focused on the ability of the proposed CS algorithms to maintain diagnostic performance for an increasing *CR*, in the context of Computer Aided Diagnosis (CAD) systems. The results were presented by monitoring Heart Rate Variability (HRV) metrics and QRS detection accuracy. For this analysis different dictionary architectures, patient-agnostic and patient-specific, were compared. The RRM-DL architecture was then compared with SPIHT in terms of QRS detection accuracy. Finally, the proposed AD technique was compared with SPIHT and existing CS approaches in terms of overall power consumption. The power consumption of each algorithm was calculated in terms of signal acquisition, digital signal compression and wireless transmission.

7.3. Main Contributions

The principal contributions from this thesis are as follows. A list of the publications arising from the work described in this thesis was given at the beginning of the document.

1. **Optimisation of CS encoding strategies through a combination of redundancy removal, quantisation and Huffman encoding:** The concept of removing the redundancy that exists between consecutive ECG frames has been shown to provide useful lossless extensions in *CR* for CS. Here the technique is modified by proposing a novel redundancy removal architecture that includes the mean value. This RRM architecture was shown to outperform standard existing redundancy removal modules based on its ability to extend the *CR*. To further enhance the energy efficiency of CS signal acquisition the use of quantisation at previously untested low-bit levels with this redundancy removal architecture was explored. The results demonstrated that ECG dictionary-based CS implementations have a high tolerance to quantisation and that the impact on signal fidelity of quantisation is minimal. Furthermore, a resulting advantage of the proposed encoding strategies is the increase in *CR* prior to wireless transmission, allowing for significant savings in the amount of data wirelessly transmitted. This consequently results in improving the overall energy efficiency of a CS ECG implementation. This work was published in [65].

2. **Development of a CS architecture for ECG-based compression that outperforms existing CS algorithms in the literature:** The proposed RRM-DL CS architecture combines the encoding techniques discussed above with the use of highly overcomplete dictionaries for CS signal recovery. The strategy of using such patient-specific reconstruction dictionaries presents an improvement over existing CS implementations. The performance of the RRM-DL architecture was quantified by monitoring signal reconstruction quality performance against the best known existing CS implementations over a range of *CR*s. The results presented show

significant performance improvements over the benchmark CS algorithms for ECG compression as it outperforms each algorithm at all *CR*s. At higher *CR*s, the *PRD* for the proposed RRM-DL algorithms tend towards a plateau which contrasts with the MMB methods that appear to increase more rapidly as *CR* rises. For example, RRM-DL demonstrated *WTS* of 22% over the MMB-based approaches at a *PRD* of 3% a figure which rises to 142% for a *PRD* of 8%. These results were published in [65].

3. **Proposal of a novel AD reconstruction scheme to improve performance of DL with ECG compression:** A further advancement to DL for CS signal reconstruction was proposed in Chapter 4 through the employment of a novel AD reconstruction scheme. Here a two-stage reconstruction process that exploits knowledge of the reconstructed signal based on the presence of a QRS complex and its location within a frame was proposed. Leveraging this signal information allows the selection of a better dictionary for the second reconstruction pass. Significantly, the results of the AD reconstruction scheme demonstrated improved performance over standard dictionary (SD) reconstruction and therefore existing CS techniques for all *CR*s tested. For example, at a *PRD** level of 9% the AD approach extends the *CR* over the SD reconstruction from 12.1 to 14.8. The proposed AD scheme also presented a reduction in the performance gap that existed between CS and the SPIHT compression algorithm. The robustness of the AD reconstruction scheme was tested by compressing signals that contain various noise and motion artefacts. Again the AD approach maintained good performance for increasing levels of noise and the presence of noise did not have any more impact on the AD method compared to the SD reconstruction. This AD reconstruction scheme has been accepted for publication [110].

4. **Classifying the impact of compression with CS on the performance of the diagnostically-relevant task of QRS detection and examining the maximum *CR* and acceptable signal distortion achievable in order to maintain diagnostic integrity:** There had been limited focus in the literature on the effects of CS on the

diagnostic performance of reconstructed signals. Therefore, in Chapter 5 the focus switched from energy-based distortion metrics to an analysis of the diagnostic integrity of the reconstructed signals. The ability to preserve the QRS complex information which is a fundamental element of many cardiac CAD systems was evaluated for increasing levels of compression. Additionally, the effects of *CR* on several HRV metrics were assessed. The proposed RRM-DL architecture demonstrated the ability to maintain QRS detection accuracy performance at high *CR*s. The results also showed good performance when compared with the state-of-the-art lossy compression algorithm SPIHT, with significant power savings. For example for an *AUC* of 0.9 the RRM-DL architecture demonstrated 40% *WTS* over SPIHT. The research described in this thesis focusing on the effects of compression on diagnostic integrity was published in [65, 97].

5. **Characterising the proposed CS architecture in terms of the overall computational and power requirements for implementation in a BAN, and comparing with existing approaches:** The importance of low-power BAN implementations is widely recognised in ambulatory healthcare. Therefore, this thesis presented an analysis of the overall power consumption of the proposed CS algorithm compared with both the current state-of-the-art CS approaches and SPIHT. Power consumption profiles for each algorithm were created by considering the power required to acquire, perform digital signal compression operations, and wirelessly transmit the compressed ECG data. The results highlighted the ability of the algorithm to be more energy efficient than all the existing benchmark CS approaches. Additionally, the AD-Q6 architecture demonstrated improved energy efficiency over SPIHT mainly due to the costs of DSP operations and the benefits in decreased wireless transmission costs with the proposed algorithms. For *PRD* values > 3% the proposed algorithms consumed less power than the SPIHT implementation and at all *PRD* values the best CS implementations from the literature were outperformed. Using a target *PRD* of 5.5% (*PRD** of approximately 9%), the results showed that the AD-Q6 method consumes 30.9 μW, while SPIHT consumes 50.65

μW and the MMB-IHT CS algorithm consumes 49.4 μW to sample, compress and wirelessly transmit the ECG data. The work described in Chapter 6 has been accepted for publication [110].

7.4. Future Work

The rapidly evolving nature of CS ensures that there is constant scope for research in the area. The potential of CS for implementation in low-power ambulatory monitoring networks is well established. There are many areas however where CS can be improved on to further strengthen the overall energy efficiency. As a result, there is a demand to improve signal reconstruction quality performance by proposing novel CS recovery strategies. Similarly, advances in the design of low-power analog hardware to implement the CS acquisition paradigm requires careful consideration but has the potential to offer significant benefits to BANs employing CS. While acknowledging the many potential areas of research with CS, the following suggestions will specifically focus on the prospective aspects of work emanating from the results presented in this thesis.

1. The use of highly overcomplete patient-specific dictionaries has been shown to be an appropriate method for CS reconstruction based on the results presented in this thesis. The impact on overall power consumption of training during a patient-specific dictionary implementation was shown to be minimal in Chapter 6. Nevertheless there are certain considerations that must be made and a patient-agnostic implementation has benefits in situations where a training phase is not possible. Therefore, further expansion of the implementation of patient-agnostic dictionaries may prove beneficial. The ultimate goal of employing patient-agnostic dictionaries would be to achieve similar performance as patient-specific implementations. As the training would generally be implemented offline where it is not limited by power consumption, theoretically the amount of signal data used for creating an agnostic dictionary does not need to be limited. Additionally, data could be taken from multiple databases to expose the dictionaries to various other abnormalities and heart rhythms that may occur in testing.

2. The concept of the AD reconstruction framework presents a promising reconstruction mechanism but there are alternatives approaches that can be taken to its implementation:

(i). The design of the AD technique allows it to be easily customisable. An expansion of the AD method whereby alternative metrics for classification are utilised could provide performance improvements. At present the scheme determines whether a QRS complex exists in a frame and based on its location determines a suitable dictionary for second stage reconstruction. However, exploiting other characteristics of the ECG signal may lead to improved performance, such as additional morphological features or frequency information of the signal.

(ii). The technique could also potentially expanded beyond the current two-stage process in terms of the number of refinement passes as this does not affect the power consumption of the wearable device.

3. Finally, the acquisition aspects of the RRM-DL architecture could be combined with the AD reconstruction scheme. The results presented in this thesis for the AD technique used the SRR module mainly to allow for a more exact comparison with existing CS methods, but the RRM-DL technique could be added to AD to improve the compression performance.

References

[1] A. P. Dhawan, W. J. Heetderks, M. Pavel, S. Acharya, M. Akay, A. Mairal, B. Wheeler, C. C. Dacso, T. Sunder, N. Lovell, M. Gerber, M. Shah, S. G. Senthilvel, M. D. Wang, and B. Bhargava, "Current and Future Challenges in Point-of-Care Technologies: A Paradigm-Shift in Affordable Global Healthcare With Personalized and Preventive Medicine," *IEEE Journal of Translational Engineering in Health and Medicine*, vol. 3, pp. 1-10, 2015.

[2] A. Benharref and M. A. Serhani, "Novel Cloud and SOA-Based Framework for E-Health Monitoring Using Wireless Biosensors," *IEEE Journal of Biomedical and Health Informatics*, vol. 18, pp. 46-55, 2014.

[3] W. Xiaoliang, G. Qiong, L. Bingwei, J. Zhanpeng and C. Yu, "Enabling Smart Personalized Healthcare: A Hybrid Mobile-Cloud Approach for ECG Telemonitoring," *IEEE Journal of Biomedical and Health Informatics*, vol. 18, pp. 739-745, 2014.

[4] C. Huasong, V. Leung, C. Chow and H. Chan, "Enabling technologies for wireless body area networks: A survey and outlook," *IEEE Communications Magazine*, vol. 47, pp. 84-93, 2009.

[5] M. Chen, S. Gonzalez, A. Vasilakos, H. Cao and V. C. Leung, "Body Area Networks: A Survey," *Mobile Network Applications*, vol. 16, pp. 171-193, 2011.

[6] G. V. Crosby, T. Ghosh, R. Murimi and C. A. Chin, "Wireless Body Area Networks for Healthcare: A Survey," *International Journal of Ad Hoc, Sensor & Ubiquitous Computing (IJASUC)*, vol. 3, pp. 1-8, 2012.

[7] S. Ullah, H. Higgins, B. Braem, B. Latre, C. Blondia, I. Moerman, S. Saleem, Z. Rahman, and K. S. Kwak, "A Comprehensive Survey of Wireless Body Area Networks," *Journal of Medical Systems*, vol. 36, pp. 1065-1094, 2012.

[8] B. Latré, B. Braem, I. Moerman, C. Blondia and P. Demeester, "A survey on wireless body area networks," *Wireless Networks*, vol. 17, pp. 1-18, 2011.

[9] C. A. Chin, G. V. Crosby, T. Ghosh and R. Murimi, "Advances and challenges of wireless body area networks for healthcare applications," in *International Conference on Computing, Networking and Communications (ICNC)*, 2012, pp. 99-103.

[10] U. Varshney, "Pervasive healthcare," *Computer*, vol. 36, pp. 138-140, 2003.

[11] ZigBee Alliance, [Online]. Available: *http://www.zigbee.org*. Accessed: January 2016.

[12] Bluetooth Low Energy, [Online]. Available: *https://www.bluetooth.com/what-is-bluetooth-technology/bluetooth-technology-basics/low-energy*. Accessed: January 2016.

[13] M. Tao, P. L. Shrestha, M. Hempel, P. Dongming, H. Sharif, and C. Hsiao-Hwa, "Assurance of Energy Efficiency and Data Security for ECG Transmission in BASNs," *IEEE Transactions on Biomedical Engineering*, vol. 59, pp. 1041-1048, 2012.

[14] S. Feng-Tso, C. Kuo and M. Griss, "PEAR: Power efficiency through activity recognition (for ECG-based sensing)," in *5th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, 2011, pp. 115-122.

[15] K. Sayood, *Introduction to Data Compression (Fourth Edition)*: Morgan Kaufmann, 2012.

[16] A. Said and W. A. Pearlman, "A new, fast, and efficient image codec based on set partitioning in hierarchical trees," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, pp. 243-250, 1996.

[17] Z. Lu, D. Y. Kim and W. A. Pearlman, "Wavelet compression of ECG signals by the set partitioning in hierarchical trees algorithm," *IEEE Transactions on Biomedical Engineering*, vol. 47, pp. 849-856, 2000.

[18] E. J. Candes, J. Romberg and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, pp. 489-509, 2006.

[19] E. J. Candes and T. Tao, "Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies?," *IEEE Transactions on Information Theory*, vol. 52, pp. 5406-5425, 2006.

[20] E. J. Candes and M. B. Wakin, "An Introduction To Compressive Sampling," *IEEE Signal Processing Magazine*, vol. 25, pp. 21-30, 2008.

[21] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, pp. 1289-1306, 2006.

[22] M. Davenport, M. Duarte, Y. Eldar and G. Kutyniok, *Introduction to Compressed Sensing, in Compressed Sensing: Theory and Applications*: Cambridge University Press, 2012.

[23] B. Jalil, O. Laligant, E. Fauvet and O. Beya, "Detection of QRS complex in ECG signal based on classification approach," in *IEEE International Conference on Image Processing*, 2010, pp. 345-348.

[24] N. Twomey, N. Walsh, O. Doyle, B. McGinley, M. Glavin, E. Jones, and W. P. Marnane, "The effect of lossy ECG compression on QRS and HRV feature extraction," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2010, pp. 634-637.

[25] C. P. Shen, W. C. Kao, Y. Y. Yang, M. C. Hsu, Y. T. Wu, and F. Lai, "Detection of cardiac arrhythmia in electrocardiograms using adaptive feature extraction and modified support vector machines," *Expert Systems with Applications*, vol. 39, pp. 7845-7852, 2012.

[26] A. Bakhshipour, M. Pooyan and H. Mohammadnejad, "Myocardial ischemia detection with ECG analysis, using Wavelet Transform and Support Vector Machines," in *17th Iranian Conference of Biomedical Engineering (ICBME)*, 2010, pp. 1-4.

[27] K. Harris, D. Edwards and J. Mant, "How can we best detect atrial fibrillation?," *The journal of the Royal College of Physicians of Edinburgh*, vol. 42, pp. 5-22, 2012.

[28] M. S. Manikandan and K. P. Soman, "A novel method for detecting R-peaks in electrocardiogram (ECG) signal," *Biomedical Signal Processing and Control*, vol. 7, pp. 118-128, 2012.

[29] D. Nabil and F. Bereksi Reguig, "Ectopic beats detection and correction methods: A review," *Biomedical Signal Processing and Control,* vol. 18, pp. 228-244, 2015.
[30] G. Antoniol and P. Tonella, "EEG data compression techniques," *IEEE Transactions on Biomedical Engineering,* vol. 44, pp. 105-114, 1997.
[31] T. A. Welch, "A Technique for High-Performance Data Compression," *Computer,* vol. 17, pp. 8-19, 1984.
[32] K. Li, Y. Pan, F. Chen, K. T. Cheng and R. Huan, "Real-time lossless ECG compression for low-power wearable medical devices based on adaptive region prediction," *Electronics Letters,* vol. 50, pp. 1904-1906, 2014.
[33] S. L. Chen, G. A. Luo and T. L. Lin, "Efficient fuzzy-controlled and hybrid entropy coding strategy lossless ECG encoder VLSI design for wireless body sensor networks," *Electronics Letters,* vol. 49, pp. 1058-1060, 2013.
[34] C. J. Deepu, X. Zhang, W. S. Liew, D. L. T. Wong and Y. Lian, "An ECG-on-Chip With 535 nW/Channel Integrated Lossless Data Compressor for Wireless Sensors," *IEEE Journal of Solid-State Circuits,* vol. 49, pp. 2435-2448, 2014.
[35] Z. Arnavut, "ECG Signal Compression Based on Burrows-Wheeler Transformation and Inversion Ranks of Linear Prediction," *IEEE Transactions on Biomedical Engineering,* vol. 54, pp. 410-418, 2007.
[36] J. P. Abenstein and W. J. Tompkins, "A New Data-Reduction Algorithm for Real-Time ECG Analysis," *IEEE Transactions on Biomedical Engineering,* vol. BME-29, pp. 43-48, 1982.
[37] J. R. Cox, F. M. Nolle, H. A. Fozzard and G. C. Oliver, "AZTEC, a Preprocessing Program for Real-Time ECG Rhythm Analysis," *IEEE Transactions on Biomedical Engineering,* vol. BME-15, pp. 128-129, 1968.
[38] P. O. Borjesson, G. Einarsson and O. Pahlm, "Comments on "Compression of the ECG by Prediction or Interpolation and Entropy Encoding"," *IEEE Transactions on Biomedical Engineering,* vol. BME-27, pp. 674-674, 1980.
[39] D. A. DiPersio and R. C. Barr, "Evaluation of the fan method of adaptive sampling on human electrocardiograms," *Medical and Biological Engineering and Computing,* vol. 23, pp. 401-410, 1985.
[40] D. S. Taubman and M. W. Marcellin, *JPEG 2000: Image Compression Fundamentals, Standards and Practice*: Kluwer Academic Publishers, 2001.
[41] A. Bilgin, M. W. Marcellin and M. I. Altbach, "Compression of electrocardiogram signals using JPEG2000," *IEEE Transactions on Consumer Electronics,* vol. 49, pp. 833-840, 2003.
[42] R. Benzid, F. Marir, A. Boussaad, M. Benyoucef and D. Arar, "Fixed percentage of wavelet coefficients to be zeroed for ECG compression," *Electronics Letters,* vol. 39, pp. 830-831, 2003.
[43] H. Imai, N. Kiraura and Y. Yoshlda, "An efficient encoding method for electrocardiography using spline functions," *Systems and Computers in Japan,* vol. 16, pp. 85-94, 1985.
[44] G. Nave and A. Cohen, "ECG compression using long-term prediction," *IEEE Transactions on Biomedical Engineering,* vol. 40, pp. 877-885, 1993.

[45] E. J. Candes, Y. C. Eldar, D. Needell and P. Randall, "Compressed sensing with coherent and redundant dictionaries," *Applied and Computational Harmonic Analysis,* vol. 31, pp. 59-73, 2011.

[46] J. A. Tropp, "Just relax: convex programming methods for identifying sparse signals in noise," *IEEE Transactions on Information Theory,* vol. 52, pp. 1030-1051, 2006.

[47] V. M. Patel and R. Chellappa, *Sparse representations and compressive sensing for imaging and vision*: SpringerBriefs, 2013.

[48] S. S. Chen, D. L. Donoho and M. A. Saunders, "Atomic decomposition by Basis Pursuit," *Journal on Scientific Computing,* vol. 20, pp. 33–61, 1999.

[49] L. Du, R. Wang, W. Wan, X. Q. Yu and S. Yu, "Analysis on greedy reconstruction algorithms based on compressed sensing," in *International Conference on Audio, Language and Image Processing (ICALIP)*, 2012, pp. 783-789.

[50] S. A. Razavi, E. Ollila and V. Koivunen, "Robust greedy algorithms for compressed sensing," in *20th European Signal Processing Conference (EUSIPCO)*, 2012, pp. 969-973.

[51] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing,* vol. 41, pp. 3397-3415, 1993.

[52] J. A. Tropp and A. C. Gilbert, "Signal Recovery From Random Measurements Via Orthogonal Matching Pursuit," *IEEE Transactions on Information Theory,* vol. 53, pp. 4655-4666, 2007.

[53] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from noisy samples," *Applied and Computational Harmonic Analysis,* vol. 26, pp. 301–321, 2009.

[54] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, S. Ting, K. F. Kelly, and R. G. Baraniuk, "Single-Pixel Imaging via Compressive Sampling," *IEEE Signal Processing Magazine,* vol. 25, pp. 83-91, 2008.

[55] J. A. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg and R. G. Baraniuk, "Beyond Nyquist: Efficient Sampling of Sparse Bandlimited Signals," *IEEE Transactions on Information Theory,* vol. 56, pp. 520-544, 2010.

[56] J. N. Laska, S. Kirolos, M. F. Duarte, T. S. Ragheb, R. G. Baraniuk, and Y. Massoud, "Theory and Implementation of an Analog-to-Information Converter using Random Demodulation," in *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2007, pp. 1959-1962.

[57] S. Kirolos, J. Laska, M. Wakin, M. Duarte, D. Baron, T. Ragheb, Y. Massoud, and R. Baraniuk, "Analog-to-Information Conversion via Random Demodulation," in *IEEE Dallas/CAS Workshop on Design, Applications, Integration and Software*, 2006, pp. 71-74.

[58] J. Yoo, S. Becker, M. Loh, M. Monge, E. Candes, and A. Emami-Neyestanak, "A 100MHz-2GHz 12.5x sub-Nyquist rate receiver in 90nm CMOS," in *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, 2012, pp. 31-34.

[59] D. Gangopadhyay, E. G. Allstot, A. M. R. Dixon, K. Natarajan, S. Gupta, and D. J. Allstot, "Compressed Sensing Analog Front-End for Bio-Sensor

Applications," *IEEE Journal of Solid-State Circuits,* vol. 49, pp. 426-438, 2014.

[60] F. Chen, A. P. Chandrakasan and V. M. Stojanovic, "Design and Analysis of a Hardware-Efficient Compressed Sensing Architecture for Data Compression in Wireless Sensors," *IEEE Journal of Solid-State Circuits,* vol. 47, pp. 744-756, 2012.

[61] F. Chen, A. P. Chandrakasan and V. Stojanovic, "A signal-agnostic compressed sensing acquisition system for wireless and implantable sensors," in *IEEE Custom Integrated Circuits Conference (CICC)*, 2010, pp. 1-4.

[62] H. Mamaghanian, N. Khaled, D. Atienza and P. Vandergheynst, "Design and Exploration of Low-Power Analog to Information Conversion Based on Compressed Sensing," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems,* vol. 2, pp. 493-501, 2012.

[63] H. Mamaghanian, N. Khaled, D. Atienza and P. Vandergheynst, "Compressed Sensing for Real-Time Energy-Efficient ECG Compression on Wireless Body Sensor Nodes," *IEEE Transactions on Biomedical Engineering,* vol. 58, pp. 2456-2466, 2011.

[64] L. F. Polania, R. E. Carrillo, M. Blanco-Velasco and K. E. Barner, "Exploiting Prior Knowledge in Compressed Sensing Wireless ECG Systems," *IEEE Journal of Biomedical and Health Informatics,* vol. 19, pp. 508-519, 2015.

[65] D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Energy-efficient Compressed Sensing for ambulatory ECG monitoring," *Computers in Biology and Medicine,* vol. 71, pp. 1-13, 2016.

[66] F. Ansari-Ram and S. Hosseini-Khayat, "ECG signal compression using compressed sensing with nonuniform binary matrices," in *16th CSI International Symposium on Artificial Intelligence and Signal Processing (AISP)*, 2012, pp. 305-309.

[67] A. M. R. Dixon, E. G. Allstot, D. Gangopadhyay and D. J. Allstot, "Compressed Sensing System Considerations for ECG and EMG Wireless Biosensors," *IEEE Transactions on Biomedical Circuits and Systems,* vol. 6, pp. 156-166, 2012.

[68] Z. Zhang, T.-P. Jung, S. Makeig and B. D. Rao, "Compressed Sensing for Energy-Efficient Wireless Telemonitoring of Noninvasive Fetal ECG Via Block Sparse Bayesian Learning," *IEEE Transactions on Biomedical Engineering,* vol. 60, pp. 300-309, 2013.

[69] A. J. Casson and E. Rodriguez-Villegas, "Signal agnostic compressive sensing for Body Area Networks: Comparison of signal reconstructions," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2012, pp. 4497-4500.

[70] D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Compressed Sensing for Bioelectric Signals: A Review," *IEEE Journal of Biomedical and Health Informatics,*, vol. 19, pp. 529-540, 2015.

[71] A. Mishra, F. Thakkar, C. Modi and R. Kher, "ECG signal compression using Compressive Sensing and wavelet transform," in *Annual International*

Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), 2012, pp. 3404-3407.

[72] D. H. Chae, Y. F. Alem, S. Durrani and R. A. Kennedy, "Performance study of compressive sampling for ECG signal compression in noisy and varying sparsity acquisition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2013, pp. 1306-1309.

[73] A. Mishra, F. N. Thakkar, C. Modi and R. Kher, "Selecting the Most Favorable Wavelet for Compressing ECG Signals Using Compressive Sensing Approach," in *International Conference on Communication Systems and Network Technologies (CSNT)*, 2012, pp. 128-132.

[74] A. Mishra, F. Thakkar, C. Modi and R. Kher, "Comparative Analysis of Wavelet Basis Functions for ECG Signal Compression through Compressive Sensing," *International Journal of Computer Science and Telecommunications*, vol. 3, pp. 23-31, 2012.

[75] A. M. R. Dixon, E. G. Allstot, A. Y. Chen, D. Gangopadhyay and D. J. Allstot, "Compressed sensing reconstruction: Comparative study with applications to ECG bio-signals," in *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2011, pp. 805-808.

[76] H. Mamaghanian, N. Khaled, D. Atienza and P. Vandergheynst, "Structured sparsity models for compressively sensed electrocardiogram signals: A comparative study," in *IEEE Biomedical Circuits and Systems Conference (BioCAS)*, 2011, pp. 125-128.

[77] D. Wei and O. Milenkovic, "Subspace Pursuit for Compressive Sensing Signal Reconstruction," *IEEE Transactions on Information Theory*, vol. 55, pp. 2230-2249, 2009.

[78] T. Blumensath, "On the difference between orthogonal matching pursuit and orthogonal least squares," *Technical Report from the University of Edinburgh*, March 2007.

[79] T. Blumensath and M. E. Davies, "Normalized Iterative Hard Thresholding: Guaranteed Stability and Performance," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, pp. 298-309, 2010.

[80] Shimmer Research, [Online]. Available: *http://www.shimmersensing.com/*. Accessed: April 2016.

[81] Y. Zigel, A. Cohen and A. Katz, "The weighted diagnostic distortion (WDD) measure for ECG signal compression," *IEEE Transactions on Biomedical Engineering*, vol. 47, pp. 1422-1430, 2000.

[82] L. F. Polania, R. E. Carrillo, M. Blanco-Velasco and K. E. Barner, "Compressed sensing based method for ECG compression," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 761-764.

[83] J. K. Pant and S. Krishnan, "Compressive Sensing of Electrocardiogram Signals by Promoting Sparsity on the Second-Order Difference and by Using Dictionary Learning," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 8, pp. 293-302, 2014.

[84] L. F. Polania and K. E. Barner, "Multi-scale dictionary learning for compressive sensing ECG," in *IEEE Digital Signal Processing and Signal Processing Education Meeting (DSP/SPE)*, 2013, pp. 36-41.

[85] MIT-BIH Arrhythmia Database, [Online]. Available: *http://www.physionet.org/physiobank/database/mitdb/*. Accessed: September 2013.

[86] MIT-BIH Arrhythmia Database Directory, [Online]. Available: *http://www.physionet.org/physiobank/database/html/mitdbdir/records.htm*. Accessed: September 2013.

[87] MIT-BIH Noise Stress Test Database, [Online]. Available: *http://www.physionet.org/physiobank/database/nstdb/*. Accessed: September 2015.

[88] Z. Qian, C. Chao, L. Zhinan, H. Anpeng, J. Bingli, D. Xiaohui, and X. Linzhen, "A novel multi-resolution SVM (MR-SVM) algorithm to detect ECG signal anomaly in WE-CARE project," in *ISSNIP Biosignals and Biorobotics Conference (BRC)*, 2013, pp. 1-6.

[89] H. ChuDuc, K. NguyenPhan and D. NguyenViet, "A Review of Heart Rate Variability and its Applications," *APCBEE Procedia*, vol. 7, pp. 80-85, 2013.

[90] V. X. Afonso, W. J. Tompkins, T. Q. Nguyen and L. Shen, "Filter bank-based ECG beat detection," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 1996, pp. 1037-1038.

[91] G. Higgins, B. McGinley, S. Faul, R. P. McEvoy, M. Glavin, W. P. Marnane, and E. Jones, "The Effects of Lossy Compression on Diagnostically Relevant Seizure Information in EEG Signals," *IEEE Journal of Biomedical and Health Informatics*, vol. 17, pp. 121-127, 2013.

[92] G. Higgins, B. McGinley, N. Walsh, M. Glavin and E. Jones, "Lossy compression of EEG signals using SPIHT," *Electronics Letters*, vol. 47, pp. 1017-1018, 2011.

[93] R. J. Martis, U. R. Acharya and H. Adeli, "Current methods in electrocardiogram characterization," *Computers in Biology and Medicine*, vol. 48, pp. 133-149, 2014.

[94] R. J. Martis, U. R. Acharya, H. Adeli, H. Prasad, J. H. Tan, K. C. Chua, C. L. Too, S. W. J. Yeo, and L. Tong, "Computer aided diagnosis of atrial arrhythmia using dimensionality reduction methods on transform domain representation," *Biomedical Signal Processing and Control*, vol. 13, pp. 295-305, 2014.

[95] D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Effects of Non-Uniform Quantization on ECG acquired using Compressed Sensing," in *4th International Conference on Wireless Mobile Communication and Healthcare (MobiHealth)*, 2014, pp. 79-82.

[96] K. Kreutz-Delgado, J. Murray, B. Rao, K. Engan, T. Lee, and T. Sejnowski, "Dictionary Learning Algorithms for Sparse Representation," *Neural Computation*, vol. 15, pp. 349-396, 2003.

[97] D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Impact of compressed sensing on clinically relevant metrics for ambulatory ECG monitoring," *Electronics Letters*, vol. 51, pp. 323-325, 2015.

[98] M. Fira, L. Goras, N. Cleju and C. Barabasa, "Results on ECG compressed sensing using specific dictionaries and its validation," in *34th International Conference on Information Technology Interfaces (ITI)*, 2012, pp. 423-428.

[99] M. Aharon, M. Elad and A. Bruckstein, "K-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation," *IEEE Transactions on Signal Processing*, vol. 54, pp. 4311-4322, 2006.

[100] K. Engan, S. O. Aase and J. Hakon Husoy, "Method of optimal directions for frame design," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 1999, pp. 2443-2446.

[101] R. Rubinstein, A. M. Bruckstein and M. Elad, "Dictionaries for Sparse Representation Modeling," *Proceedings of the IEEE*, vol. 98, pp. 1045-1057, 2010.

[102] R. Tibshirani, "Regression Shrinkage and Selection Via the Lasso," *Journal of the Royal Statistical Society, Series B*, vol. 58, pp. 267-288, 1994.

[103] V. M. Patel and R. Chellappa, "Sparse Representations, Compressive Sensing and dictionaries for pattern recognition," in *First Asian Conference on Pattern Recognition (ACPR)*, 2011, pp. 325-329.

[104] M. Elad, R. Rubinstein and M. Zibulevsky, "Efficient Implementation of the K-SVD Algorithm using Batch Orthogonal Matching Pursuit " *Technical Report - CS, Technion*, April 2008.

[105] K-SVD Matlab Toolbox, [Online]. Available: *http://www.cs.technion.ac.il/~ronrubin/software.html*. Accessed: September 2013.

[106] G. R. Tsouri and M. H. Ostertag, "Patient-Specific 12-Lead ECG Reconstruction From Sparse Electrodes Using Independent Component Analysis," *IEEE Journal of Biomedical and Health Informatics*, vol. 18, pp. 476-482, 2014.

[107] T. Ince, S. Kiranyaz and M. Gabbouj, "A Generic and Robust System for Automated Patient-Specific Classification of ECG Signals," *IEEE Transactions on Biomedical Engineering*, vol. 56, pp. 1415-1426, 2009.

[108] A. Khazaee and A. Ebrahimzadeh, "Classification of electrocardiogram signals with support vector machines and genetic algorithms using power spectral features," *Biomedical Signal Processing and Control*, vol. 5, pp. 252-263, 2010.

[109] E. v. d. Berg and M. P. Friedlander, "Probing the Pareto Frontier for Basis Pursuit Solutions," *SIAM Journal on Scientific Computing*, vol. 31, pp. 890-912, 2009.

[110] D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Adaptive Dictionary Reconstruction for Compressed Sensing of ECG Signals," *IEEE Journal of Biomedical and Health Informatics*, accepted for publication, Febuary 2016.

[111] G. Cohen and P. de Chazal, "Automated detection of sleep apnea in infants: A multi-modal approach," *Computers in Biology and Medicine*, vol. 63, pp. 118-123, 2015.

[112] J. Oster, J. Behar, R. Colloca, L. Qichen, L. Qiao, and G. D. Clifford, "Open source Java-based ECG analysis software and Android app for Atrial

Fibrillation screening," in *Computing in Cardiology Conference (CinC)*, 2013, pp. 731-734.

[113] U. Rajendra Acharya, K. Paul Joseph, N. Kannathal, C. M. Lim and J. S. Suri, "Heart rate variability: a review," *Medical and Biological Engineering and Computing*, vol. 44, pp. 1031-1051, 2006.

[114] Task Force of the European Society of Cardiology and The North American Society of Pacing and Electrophysiology, "Heart Rate Variability Standards of Measurement, Physiological Interpretation, and Clinical Use," *Circulation*, vol. 93, pp. 1043–1065, 1996.

[115] P. Dehkordi, A. Garde, W. Karlen, D. Wensley, J. M. Ansermino, and G. A. Dumont, "Pulse rate variability compared with heart rate variability in children with and without sleep disordered breathing," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2013, pp. 6563-6566.

[116] C. Vidaurre, T. H. Sander and A. Schlögl, "BioSig: The Free and Open Source Software Library for Biomedical Signal Processing," *Computational Intelligence and Neuroscience*, vol. 2011, pp. 1-12, 2011.

[117] G. M. Friesen, T. C. Jannett, M. A. Jadallah, S. L. Yates, S. R. Quint, and H. T. Nagle, "A comparison of the noise sensitivity of nine QRS detection algorithms," *IEEE Transactions on Biomedical Engineering*, vol. 37, pp. 85-98, 1990.

[118] W. Zhu, N. Zeng and N. Wang, "Sensitivity, Specificity, Accuracy, associated confidence interval and ROC analysis with practical SAS implementations," in *NESUG: Proceedings: Health Care and Life Sciences*, 2010.

[119] Analog Devices Blackfin DSP - ADSP-BF537, [Online]. Available: *http://www.analog.com/en/processors-dsp/blackfin/adsp-bf537/products/product.html*. Accessed: June 2014.

[120] Texas Instruments Bluetooth Low-Energy System-on-Chip - CC2540, [Online]. Available: *http://www.ti.com/product/cc2540*. Accessed: January 2015.

[121] M. Angelopoulou, K. Masselos, P. K. Cheung and Y. Andreopoulos, "Implementation and Comparison of the 5/3 Lifting 2D Discrete Wavelet Transform Computation Schedules on FPGAs," *Journal of Signal Processing Systems*, vol. 51, pp. 3-21, 2008.

[122] K. Andra, C. Chakrabarti and T. Acharya, "A VLSI architecture for lifting-based forward and inverse wavelet transform," *IEEE Transactions on Signal Processing*, vol. 50, pp. 966-977, 2002.

[123] B. Murmann, "ADC Performance Survey 1997-2014," [Online]. Available: *http://web.stanford.edu/~murmann/adcsurvey.html*. Accessed: December 2014.

[124] "Estimating Power for ADSP-BF534/BF536/BF537 Blackfin Processors," [Online]. Available: *http://www.analog.com/static/imported-files/application_notes/EE-297.Rev.3.11.07.pdf*. Accessed: June 2014.

[125] M. Trakimas, R. D. Angelo, S. Aeron, T. Hancock and S. Sonkusale, "A Compressed Sensing Analog-to-Information Converter With Edge-Triggered

SAR ADC Core," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 60, pp. 1135-1148, 2013.

[126] M. Siekkinen, M. Hiienkari, J. K. Nurminen and J. Nieminen, "How low energy is bluetooth low energy? Comparative measurements with ZigBee/802.15.4," in *IEEE Wireless Communications and Networking Conference Workshops (WCNCW)*, 2012, pp. 232-237.

APPENDIX A

Journal Submissions Arising From Thesis

The following are the journal papers published from the work completed for this thesis:

D. Craven, B. McGinley, L. Kilmartin, M. Glavin, and E Jones, "Compressed Sensing for Bioelectric Signals: A Review," *IEEE Journal of Biomedical and Health Informatics*, vol. 19, pp. 529-540, 2015.

D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Impact of compressed sensing on clinically relevant metrics for ambulatory ECG monitoring," *Electronics Letters*, vol. 51, pp. 323-325, 2015.

D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Adaptive Dictionary Reconstruction for Compressed Sensing of ECG Signals," accepted for publication in *IEEE Journal of Biomedical and Health Informatics*, February 2016.

D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Energy-Efficient Compressed Sensing for Ambulatory ECG Monitoring," *Computers in Biology and Medicine*, vol. 71, pp. 1-16, 2016.

D. Craven, M. O'Halloran, B. McGinley, R. C. Conceicao, L. Kilmartin, E. Jones, and M. Glavin, "Compressive Sampling for Time Critical Microwave Imaging Applications," *Healthcare Technology Letters*, vol. 1, pp. 6-12, 2014.