

DOCTOR OF PHILOSOPHY

Approaches to transmission reduction protocols in low-frequency Wireless Sensor Networks deployed in the field

Wilkins, Ross

Award date:
2015

Awarding institution:
Coventry University

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of this thesis for personal non-commercial research or study
- This thesis cannot be reproduced or quoted extensively from without first obtaining permission from the copyright holder(s)
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Approaches to transmission reduction
protocols in low-frequency
Wireless Sensor Networks
deployed in the field

Ross Ivan Ryan Wilkins

A thesis submitted in partial fulfilment
of the University's requirements for the Degree of
Doctor of Philosophy

September 2014

Coventry University
Faculty of Engineering and Computing

Abstract

A key barrier in the adoption of Wireless Sensor Networks (WSNs) is achieving long-lived and robust real-life deployments. Issues include: reducing the impact of transmission loss, node failure detection, accommodating multiple sensor modalities, and the energy requirement of the WSN network stack. In systems where radio transmissions are the largest energy consumer on a node, it follows that reducing the number of transmissions will, in turn, extend node lifetime. Research in this area has led to the development of the Dual Prediction Scheme (DPS). However, the design of specific DPS algorithms in the literature have not typically considered issues arising in real world deployments. Therefore, this thesis proposes solutions to enable DPSs to function in robust and long-lived real-world WSN deployments. To exemplify the proposed solutions, COGENT-HOUSE, an end-to-end open-source home environmental and energy monitoring system, is considered as a case study. COGENT-HOUSE was deployed in 37 homes generating 235 evaluation data traces, each spanning periods of two weeks to a year.

DPSs presented within the literature are often lacking in the ability to handle several aspects of real world deployments. To address issues in real-life deployments this thesis proposes a novel generalised framework, named Generalised Dual Prediction Scheme (G-DPS). G-DPS provides: i) a multi-modal approach, ii) an acknowledgement scheme, iii) heartbeat messages, and iv) a method to calculate reconstructed data yield. G-DPS's multi-modal approach allows multiple sensor's readings to be combined into a single model, compared to single-modal which uses multiple instances of a DPS. Considering a node sensing temperature, humidity and CO₂, the multi-modal approach transmissions are reduced by up to 27%, signal reconstruction accuracy is improved by up to 65%, and the energy requirement of nodes is reduced by 15% compared to single-modal DPS. In a lossy network use of acknowledgements improves signal reconstruction accuracy by up to 2 \times and increases the data yield of the system up to 7 \times , when compared to an acknowledgement-less scheme, with only up to a 1.13 \times increase in energy consumption. Heartbeat messages allow the detection of faulty nodes, and yet do not significantly impact the energy requirement of functioning nodes. Implementing DPS algorithms within the G-DPS framework enables robust deployments, as well as easier comparison of performance between differing approaches.

DPSs focus on modelling sensed signals, allowing accurate reconstruction of the signal from fewer transmissions. Although transmissions can be reduced in this way, considerable savings are also possible at the application level. Given the information needs of a specific application, raw sensor measurement data is often highly compressible. This thesis proposes the Bare Necessities (BN) algorithm, which exploits on-node analytics by transforming data to information closer to the data source (the sensing device). This approach is evaluated in the context of a household monitoring application that reports the percentage of time a room of the home spends in various environmental conditions. BN can reduce the number of packets transmitted to the sink by 7000 \times compared to a sense-and-send approach.

To support the implementation of the above solutions in achieving long lifetimes, this thesis explores the impact of the network stack on the energy consumption of low transmission sensor nodes. Considering a DPS achieving a 20 \times transmission reduction, the energy reduction of a node is only 1.3 \times when using the TinyOS network stack. This thesis proposes the Backbone Collection Tree Protocol (B-CTP), a networking approach utilising a persistent backbone network of powered nodes. B-CTP coupled with Linear Spanish Inquisition Protocol (L-SIP) decreases the energy requirement for sensing nodes by 13.4 \times compared to sense-and-send nodes using the TinyOS network stack. When B-CTP is coupled with BN an energy reduction of 14.1 \times is achieved.

Finally, this thesis proposes a quadratic spline reconstruction method which improves signal reconstruction accuracy by 1.3 \times compared to commonly used linear interpolation or model prediction based reconstruction approaches. Incorporating sequence numbers into the quadratic spline method allows up to 5 hours of accurate signal imputation during transmission failure.

In summary, the techniques presented in this thesis enable WSNs to be long-lived and robust in real-life deployments. Furthermore, the underlying approaches can be applied to existing techniques and implemented for a wide variety of applications.

Acknowledgements

Completing this thesis and surrounding work has been the most daunting task I have undertaken to date. It is with the support of the group around me that I am now submitting.

First and foremost, my thanks go to my Director of Studies, Professor Elena Gaura. Elena was the person that provided me the opportunity to start my career in research. She has always helped me to see the bigger picture and push me in the right direction. Her support and encouragement, during both the challenging and good times, over the years has been invaluable and I am a better researcher because of it.

My thanks extend to my second supervisor Dr. James Brusey. James has always been there to ask difficult questions and to get me thinking around the subject. He has been a great support during technical discussions and has been a major influence on the way I approach the development of ideas. Furthermore, James has helped me to greatly improve my skills in presenting and writing.

Thanks also go to John Halloran, my third supervisor. John's feedback from a different perspective was a great aid and helped guide this thesis to be seen from a wider view.

Working within the Cogent Labs is not a solitary task. I've had the opportunity to work with a great group of people, both current and past members. I would like to thank each and everyone of them for the support over the years. A special thanks goes to Dr. John Kemp for advice over the years and help in proof reading this thesis throughout.

I would also like to thank the Orbit Housing Group for supporting this work through both funding, expertise, and access to properties for monitoring. I am grateful to all of the participants who opened their homes and allowed me to deploy in their homes. Without access to these homes the work in this thesis would not be possible.

Thanks to all of my close friends who have provided plenty of respite (sometimes too much) from everything when times got tough.

Finally I would like to express my gratitude to my parents. They have always supported me in my goals and aims in life—without their backing over the years I wouldn't be where I am today.

Contents

1	Introduction	1
1.1	Research questions	4
1.2	Method	7
1.3	Research scope	9
1.4	Contributions to knowledge	10
1.5	Publications	11
1.6	Thesis structure	12
1.7	Acknowledgement of contributed work	12
1.8	Acknowledgement of software tools	13
2	Energy efficient Wireless Sensor Networks	15
2.1	Wireless Sensor Networks	15
2.2	Wireless node energy consumption	17
2.3	Approaches to energy saving in WSNs	23
2.4	Data-driven optimisation in WSNs	25
2.5	Use of DPS in WSNs deployed in the field	34
2.6	WSN signal reconstruction	38
2.7	Summary	39
3	Generalised Dual Prediction Scheme	41
3.1	Motivation	42
3.2	Definitions	43
3.3	The Generalised Dual Prediction Scheme	44
3.4	The Linear Spanish Inquisition Protocol	52
3.5	Evaluation of G-DPS	53
3.6	The Backbone Collection Tree Protocol for low transmission WSNs	71
3.7	Summary	76
4	Spline-based data reconstruction in Wireless Sensor Networks	79
4.1	Reconstruction methods	80
4.2	Evaluation	85
4.3	Dual quadratic splines on a lossy network	89
4.4	The effect of reconstruction on statistical summaries	92
4.5	Summary	93
5	Bare Necessities—Knowledge driven design	95
5.1	Motivation	96
5.2	Application-level state: the time-discounted distribution summary	97
5.3	The Bare Necessities algorithm	99
5.4	Evaluation of BN	101
5.5	Summary	112
6	Conclusions	113
6.1	Answers to research questions	114
6.2	Future work	120
6.3	Generalising to other applications	122
6.4	Concluding remarks	123

References	125
A Cogent House: A WSN for the Built Environment	139
A.1 System concept	139
A.2 System aims	140
A.3 System design	141
A.4 Cogent-House overview	142
A.5 System hardware description	143
A.6 System software description	145
A.7 System evaluation	148
A.8 Summary	153
B Deployments and datasets	155
C Publications, presentations and attended conferences	163
C.1 Journal publications	163
C.2 Conference proceedings	163
C.3 Technical reports	164
C.4 Presentations and demos	167
C.5 Conferences attended	167
D Ethical approval	203

List of Figures

2.1	Typical multi-hop WSN architecture	16
2.2	Taxonomy of approaches to energy savings in sensor networks.	24
2.3	A comparison of data reduction techniques applied to the Intel Lab dataset.	32
3.1	Example L-SIP output for a monitored dining room.	55
3.2	Example of a signal where multi-modal would provide no overall benefit compared to single-modal.	58
3.3	Graph showing percentage of transmissions which are heartbeat messages for a given heartbeat period.	60
3.4	Transmission model	62
3.5	The effect of transmission failure on the number of required transmissions for multi-modal L-SIP with and without end-to-end acknowledgements.	63
3.6	The effect of transmission failure on the accuracy of the reconstructed temperature signal for multi-modal L-SIP with and without end-to-end acknowledgements.	64
3.7	The effect of transmission failure on the annual energy requirement of a L-SIP node with and without end-to-end acknowledgements.	65
3.8	Demonstration of failed transmissions in acknowledgement and acknowledgement-less transmissions.	66
3.9	The effect of transmission failure on the reconstructed data yield for multi-modal L-SIP with and without end-to-end acknowledgements and sense-and-send.	67
3.10	Typical battery voltage of nodes running Collection Tree Protocol (CTP) and sense-and-send, and a node using B-CTP and L-SIP.	70
3.11	Annual energy usage compared to achieved transmission reduction when using CTP and Low Power Listening (LPL).	71
3.12	Proposed B-CTP networking infrastructure.	73
3.13	Annual energy usage compared to achieved transmission reduction when using B-CTP and LPL.	75
4.1	Faceted plot showing linear interpolation, predictive model, and dual quadratic spline signal reconstruction methods applied to a sample of 5 minute sense-and-send home air temperature data. The black points indicate where L-SIP transmissions occur.	81
4.2	Combined line plot showing Linear interpolation, predictive model, and dual quadratic spline signal reconstruction methods applied to a sample of 5 minute sense-and-send home air temperature data. The black points indicate where L-SIP transmissions occur.	82
4.3	Large reconstruction error caused by linear interpolation.	82
4.4	Comparison of average temperature reconstruction accuracy of reconstruction methods.	86
4.5	Comparison of average relative humidity reconstruction accuracy of reconstruction methods.	87
4.6	Comparison of average CO ₂ reconstruction accuracy of reconstruction methods.	87
4.7	Example of large error introduced by the two polynomial style spline methods.	88
4.8	Demonstration of the benefit of using the adjusted dual quadratic spline method versus non-adjusted dual quadratic splines.	90
4.9	Reconstruction accuracy over 9 data traces during communication failure simulations.	91
4.10	Comparison of temperature exposure graphs using original data and reconstructed data.	93
5.1	Example temperature time-discounted distribution summary.	98
5.2	BN distribution signal with a decay of one month	103
5.3	Example BN time-discounted temperature distribution summary.	104

5.4	Example post-processed time-discounted temperature distribution summary generated at monthly intervals for a monitored bedroom.	104
5.5	Percentage of annual BN packet transmissions over a range of error thresholds.	106
5.6	RMSE of the BN reconstruction error over a range of thresholds.	106
5.7	Annual energy usage for nodes with BN and L-SIP.	111
A.1	Conceptual design of building monitoring system.	140
A.2	System flow overview.	142
A.3	Node hardware platform	144
A.4	WSN data flow overview	146
A.5	Overview of Passivhaus deployment area.	149
A.6	Typical node deployment in an example two bedroom home.	150
A.7	WSN data yield heatmap.	152

List of Tables

2.1	Power consumption of typical commercial motes.	19
2.2	Reported energy consumption of selected sensors	19
3.1	Summary of datasets used	54
3.2	Example L-SIP performance metrics, for a temperature signal compressed with L-SIP . .	56
3.3	Comparison of multi-modal and single-modal approaches of a node sensing temperature and humidity.	57
3.4	Comparison of multi-modal and single-modal approaches on a node of a node sensing temperature, humidity, and CO ₂	57
3.5	Summary of L-SIP deployments	69
3.6	Baseline micro-benchmark estimates for a TelosB running sense-and-send.	70
3.7	Microbenchmark estimates for a TelosB running SIP with CTP and LPL.	70
3.8	Microbenchmark estimates for a TelosB running L-SIP with B-CTP and LPL.	74
3.9	Microbenchmark estimates for a TelosB running L-SIP with B-CTP and LPL.	74
3.10	Microbenchmark estimates for using sense-and-send and L-SIP on TelosB for various MAC approaches.	75
4.1	Error introduced into exposure graph bands by the reconstruction process.	92
5.1	Defined temperature exposure bands.	97
5.2	Example BN performance metrics	105
5.3	Average number of BN information packet transmissions and compression factor for nine rooms.	107
5.4	Performance of BN with L-SIP for one year of temperature data.	108
5.5	Number of transmissions and reconstruction results for the single-modal and multi-modal approach to BN	109
5.6	Microbenchmark estimates for a TelosB running BN with B-CTP.	110
5.7	Microbenchmark estimates for using sense-and-send, L-SIP, and BN on a TelosB mote. . .	110
A.1	WSN yield at the Passivhaus site.	151
A.2	Baseline micro-benchmark estimates for a TelosB running COGENT-HOUSE.	153

List of Algorithms

2.1	On-node algorithm summary for DPS.	29
2.2	Pseudocode for node and sink for SIP	31
3.1	G-DPS node algorithm.	44
3.2	G-DPS sink functions	48
3.3	L-SIP phrased in terms of G-DPS	52
5.1	BN on-node algorithm in terms of G-DPS.	99

Acronyms

AMS	Adaptive Model Selection
AR	AutoRegressive
ARIMA	AutoRegressive Integrated Moving Average
B-CTP	Backbone Collection Tree Protocol
BN	Bare Necessities
CCARC	Cogent Computing Applied Research Center
CO₂	Carbon Dioxide
CTP	Collection Tree Protocol
dEWMA	Dual Exponentially Weighted Moving Average
DKF	Dual Kalman Filter
DPS	Dual Prediction Scheme
ECG	Electrocardiography
ETX	Expected Transmissions
EWMA	Exponentially Weighted Moving Average
G-DPS	Generalised Dual Prediction Scheme
GPR	Gaussian process regression
GPS	Global Positioning System
IoT	Internet of Things
kB	kilobyte
K-RLE	K-Run Length Encoding
LMS	Least Mean Squares
LPL	Low Power Listening
L-SIP	Linear Spanish Inquisition Protocol
MAC	Media Access Control
MCU	Microcontroller
MEMS	MicroElectroMechanical Systems
NLMS	Normalised Least Mean Squares
PAQ	Probabilistic Adaptable Query
PMC	Poor Man's Compression
PMC-MR	Poor Man's Compression - Midrange

PMC-MEAN	Poor Man's Compression - Mean
RAM	Random Access Memory
RLE	Run Length Encoding
RMSE	Root Mean Squared Error
SAF	Similarity-based Adaptive Framework
SIP	Spanish Inquisition Protocol
SQL	Structured Query Language
VOC	Volatile Organic Compounds
WSN	Wireless Sensor Network

Chapter 1

Introduction

There has always been a need to monitor and understand the world around us. Measurements are used in every aspect of life from understanding the weather through to monitoring fuel consumption in a car. Early measurement relied on either manual measurements or complex mechanical systems. For example, in 1980, Johnson [49] devised a system to monitor the HVAC system in a school. This system was composed of a number of thermocouples and flow meters wired to a central metering station. To collate data over time a camera was set up to take a photograph of the meter every 30 minutes, from which raw data could be extracted. Though it allows understanding of the phenomena, manual or meter based measurement is often labour intensive, prone to error, and limited in scope.

Advances in microcontrollers, low powered radio technology, and MicroElectroMechanical Systems (MEMS) sensor technology have given rise to Wireless Sensor Networks (WSNs). WSNs have paved the way to collect data in a range of fields in a less labour intensive and more accurate and timely manner. A WSN consists of a number of physically distributed, autonomous, usually battery powered devices (nodes) equipped with sensing, processing, and communication capabilities. Periodically sensor nodes will *sense* the environment through interfaced sensors, *process* that sensor data using an on-board microcontroller, and *transmit* the results wirelessly to a central data store (sink) using an on-board radio chip. The sink is responsible for storing and processing WSN node sensor readings for eventual data analysis.

Since WSNs are relatively low cost, scalable, and require no fixed infrastructure, they allow the efficient collection of vast quantities of data about the world around us, allowing us to improve our understanding and knowledge about phenomena. To date, WSNs have been implemented in a range of fields. For example, the monitoring of volcanic eruptions [118], soil moisture tension for irrigation management in vineyards [42], sniper fire localisation in battlefields [59], and ice quake detection on glaciers [72].

WSNs are commonly acknowledged today as proven research instruments for several application domains. However, WSNs still have many open research issues including energy management, fault tolerance, deployment processes, and harsh environments. This thesis addresses the question of how to provide an integrated solution which enables robust and long-lived real-world WSN deployments.

A WSN deployment can be considered robust if:

1. The sensor network meets the lifetime requirements of the project.

WSNs have a common primary design goal—to ensure the longest possible device lifetime with the available power budget while meeting application requirements, (performance is usually measured in terms of accuracy of the reconstructed sensed signal). Deployments can be on the order of days in the case of sniper fire localisation or years for glacial monitoring. Failing to achieve deployment lifetime requirements may require node battery changes which could either be expensive, labour intensive, or even impossible. For example, in the case of glacial monitoring once the sensor nodes are embedded in the ice there is no way of accessing the probe to change batteries without significant intervention. Therefore, a WSN developer should consider why, where, and how the energy is used to maximise node lifetime.

2. The sensor network achieves a high data collection yield.

WSNs are generally considered to be lossy networks. Anastasi *et al.* [8] and Arora *et al.* [10] both report high levels of data loss in WSNs using default settings. For an accurate representation of the environment, the network must be reliable enough to achieve a data yield which allows the end-users to make inferences about the cause and effect of events. A low data collection yield means that important events may be lost or analysis skewed due to not having the complete dataset. For example, in the glacial monitoring application, if transmissions were only successful during the day, a calculation of daily averages would be skewed to be warmer than it was. Therefore, a WSN should be designed for maximum transmission success.

3. The sensor network is able to detect / be reactive to sensing node failures.

Nodes in a network may fail for a number of reasons including human intervention, battery failure, or other external influence. For example, glacial monitoring nodes have been known to fall into lakes [23]. In a WSN each sensing node is generally considered to be of equal importance to the other nodes in the deployment. Node failure will impact the data collection yield and affect the network topology, potentially reducing the number of available routes to the sink. Therefore, a WSN should be designed to detect node failure at the earliest instance.

This thesis therefore focuses on techniques and approaches to i) reduce the energy requirement of WSN sensing nodes and ii) maximise the robustness of deployed WSN nodes.

As previously mentioned, a node's energy budget is one of the most important considerations for WSN design. Generally, a WSN node consists of five hardware sub-components: a sensing unit, a processing

unit, a transceiver unit, a storage unit, and a power unit. The primary energy consumer of a node is, generally, the use of the on-board radio chip¹. Polastre *et al.* [85] show that a TelosB node's power consumption with an active radio is $10\times$ greater than when using the Microcontroller (MCU) and $3900\times$ greater than when idle. Therefore, an approach to extend the lifetime of a WSN node would be to reduce the time the radio is used for.

Numerous algorithms have been described in the literature for the purpose of reducing the number of transmissions a node is required to make. An approach often used in the literature is that of a Dual Prediction Scheme (DPS) type algorithm [22, 34, 46, 101, 110, 111]. DPSs share a model of the data between the node and sink. At each sampling interval the node makes a prediction based on the last state transmitted to the sink. Transmissions are only made if this prediction differs from the current sensor reading by more than a defined threshold. However, DPS algorithms presented within the literature are often lacking in the ability to handle several aspects of real world deployments. The aspects include: transmission loss, node failure detection, accommodating multiple sensor modalities, and reduction of the energy requirement of the WSN network stack. This thesis proposes and evaluates a novel generalised framework, named Generalised Dual Prediction Scheme (G-DPS), for the implementation of DPS algorithms in real life WSNs. The goal is to provide a means by which a sensing system designer can implement DPS algorithms to maximise data yield and signal reconstruction accuracy whilst minimising transmissions when deployed in real-world lossy networks.

DPSs, in general, use a model of the sensed signal to allow an accurate reconstruction of the signal at the sink using fewer transmissions. However, end-users tend to work with high-level knowledge. When considering the high-level knowledge content of each transmission there is redundancy in the signal. Furthermore, for some applications the ability to reconstruct the entire time series is unnecessary and it is only important to know the proportion of time spent in a state, or set of states. For example, in human behaviour monitoring applications, end-users are often only interested in how long is spent in a certain modality (walking, driving, standing) in a given day. This thesis proposes and evaluates the combination of DPS with on-node node analytics to significantly reduce the number of transmissions, whilst delivering the knowledge end-users require.

Multi-hop networks are a common approach to networking in WSNs. This approach allows for WSNs to be deployed over large geographical areas, routing data through intermediary routing nodes to reach the sink. While there are large bodies of literature related to multi-hop networks and to DPS algorithms, few publications attempt to answer the question of how these two technologies interact with each other.

¹Note that this applies to the use of passive sensors, discussed in Section 2.2 on page 17, the use of active sensors can outweigh the energy cost of the radio.

Raza *et al.* [93] show that once the number of transmissions a node is required to make is significantly reduced, the Media Access Control (MAC) and routing layers have the greatest energy requirement. This thesis evaluates the performance of DPSs coupled with the commonly used TinyOS network stack which is composed of the Collection Tree Protocol (CTP) and Low Power Listening (LPL). This thesis also proposes a network topology that minimises listening time to minimise the energy requirement of a sensing node.

DPS algorithms commonly use either basic linear interpolation or the output of the predictive model for the purpose of signal reconstruction. It is expected by the author that spline-based reconstruction methods can provide a higher accuracy than linear interpolation or model prediction. This thesis evaluates three spline-based methods, and compares their performance to both linear interpolation and model prediction.

This rest of this chapter is structured as follows: First, Section 1.1, presents the research questions which this thesis aims to address. Section 1.2 describes the method undertaken for the work in this thesis. Section 1.4 lists the contributions to knowledge this thesis provides. Section 1.5 details peer-reviewed publications resulting from this thesis. Finally, Section 1.6 describes the structure of this thesis, Section 1.7 acknowledges contributed work, and Section 1.8 acknowledges the tools used in this thesis.

1.1 Research questions

The aim of this thesis is to provide generalised methods and protocols for the implementation of DPSs in deployed WSNs. This thesis aims to answer the following overarching research question:

How can WSN nodes be designed to achieve robust and long-lived real-life WSN deployments?

Overall 6 questions are posed to answer the overarching research question. These are as follows:

RQ1: What features can improve the robustness of DPSs implemented in deployed WSNs?

Section 2.5 on page 34 shows that DPS algorithms in the literature have not typically considered issues in real world deployments. The following issues have been identified:

1. DPSs are generally designed with the aim of compressing a single sensing modality, however, generally nodes include more than one sensing modality.
2. DPSs require the node and sink to have identical copies of the sink state. However, DPSs have not

been designed to handle lossy networks, therefore, a node cannot determine if a transmission has failed.

3. The sink is unable to distinguish if a node is suppressing messages as intended or the node has failed.

Chapter 3 proposes a generalised framework to enable the robust deployment of DPSs in real-life WSN deployments. This framework was developed by answering the following three subquestions:

RQ1A: Does combining multiple sensor readings into a single model allow a greater reduction in the number of packets transmitted and improve signal reconstruction accuracy compared to compressing each stream individually?

Sensing nodes generally include multiple sensors of differing types. However, existing DPS algorithms only consider a node with one sensor. This thesis will examine how to integrate multiple sensors in DPSs. In a single-modal approach there are multiple instances of a DPS algorithm and individual sensor model states are transmitted when there is an event. However, when dealing with multiple signals in a given environment, the signals are often highly correlated (for example, temperature and relative humidity). As an alternative a multi-modal approach combines multiple sensor's readings into a single model. Since the multi-modal approach will update the sink state for any given modality more frequently compared to sending individual state updates, the approach is expected to reduce signal reconstruction error. However, the total number of required transmissions is expected to be reduced compared to multiple single-modal DPS instances.

Chapter 3 (specifically Section 3.5.3 on page 56) answers this question by comparing the transmission reduction and reconstruction accuracy of multi-modal to single-modal as a part of the proposed generalised framework.

RQ1B: Can heartbeat messages allow detection of node failure within a user specified time period, without producing a large impact on the energy requirement of a functioning node?

The continuous sequence of packets received from a sense-and-send node at regular, and predictable, intervals provide a simple mechanism to check the health status of a node. However, DPS algorithms transmit at irregular and unpredictable intervals and therefore, a packet might not be received either because the node is functioning but is suppressing messages as intended or because the node has failed. The end-user is unable to distinguish between transmission suppression and node failure. Specifying

a maximum allowed time for a node not to transmit should allow node failure to be detected without impacting on a functioning node's energy requirement. Chapter 3 (specifically Section 3.5.4 on page 59)) evaluates the use of a heartbeat message to detect node failure.

RQ1C: Can the use of end-to-end acknowledgements with DPSs allow for a greater reconstructed data yield compared to an acknowledgement-less schemes?

Multi-hop reliability can be very poor, less than 30% in some cases. DPSs rely on a high packet delivery rate to enable accurate reconstruction of the sensed signal. A mechanism to alleviate the issue with lossy networks is required to maximise reconstructed data yield. However, as shown by Anastasi *et al.* [8], this will come at a cost of a higher energy consumption.

Chapter 3 (specifically Section 3.5.5 on page 61)) answers this question by evaluating an end-to-end acknowledgement based scheme, comparing the reconstructed data yield, transmissions increase, and energy use as a part of the generalised framework.

RQ2: Can the lifetime of a WSN node implementing transmission reduction approaches be increased further by using a persistent backbone network of mains powered routing nodes?

Section 2.5 on page 34 shows that when transmissions are substantially reduced, the primary energy consumer for a WSN node is the energy overheads of networking. Specifically the total energy requirement for listening is greater than any other process. If mains powered nodes can be utilised for routing, battery powered nodes can reduce their radio duty cycle by removing the need to listen. Reducing the radio duty cycle of a battery powered node should significantly decrease the overall energy requirement of the node. Chapter 3 (specifically Section 3.6 on page 71) answers this question by proposing and evaluating a WSN networking technique for low transmission algorithms.

RQ3: Can a spline-based signal reconstruction method improve the accuracy of reconstructed signals compared to piecewise linear methods, for example linear interpolation or model prediction, when using DPS algorithms such as Linear Spanish Inquisition Protocol (L-SIP)?

Section 2.6 on page 38 shows that the selection of the best method to accurately reconstruct the original signal based on the output of DPS algorithms has received little attention. Traditionally, linear interpolation is used to reconstruct the sensed signal. However, in addition to the sensed value Spanish Inquisition Protocol (SIP) provides the signal gradient in the state update. Furthermore, the known

bounds on the suppressed samples can be inferred. This additional information lends itself to the use of spline-based methods to reconstruct a signal. If signal reconstruction accuracy can be increased then a larger error threshold could be used in SIP to give the same final reconstruction accuracy, allowing a further reduction in transmissions. Chapter 4 answers this question by evaluating three spline-based reconstruction methods, along with model predictions and traditional linear interpolation.

RQ4: Can a combination of DPS concepts with the calculation of application-level information on-node significantly reduce the energy requirements of a node further than the current state of the art?

This question asks if designing a DPS which models application-level metrics rather than the sensed signal can provide better performance than the current state of the art, SIP. Compared to the knowledge output required by end-users there is, generally, much redundancy in the sensed signal. Performing on-node processing to generate application-level metrics offers the opportunity to further compress the raw signal. Using application-level metrics as a model for a DPS should significantly reduce the number of transmissions and thus the node energy requirement. Chapter 5 answers this question by proposing and evaluating an algorithm using this approach.

1.2 Method

Raman and Chebrolu [91] critique the WSN literature from a systems perspective. The authors found that there are inconsistencies between i) WSN algorithm and protocol design, and ii) WSNs designed for real-life applications. They conclude, “an application-driven, bottom-up approach is required for meaningful solution of any networking issues in WSNs”. Although deployments in the field are often time consuming and expensive, they offer a real in-situ evaluation of a system revealing issues which would not be detected through simulation alone.

In this thesis I have taken this *bottom-up* approach suggested by Raman and Chebrolu. Cogent-House described in Appendix A is the application for the application-driven approach. COGENT-HOUSE, an end-to-end open-source home environmental and energy monitoring WSN, was developed as a tool to collect data from real-life deployments. The concept and design of COGENT-HOUSE is detailed in Appendix A.

The datasets used to exemplify approaches developed in this thesis and to evaluate performance were drawn from the deployment of COGENT-HOUSE within 37 homes. The deployment homes consist of flats and houses with between one and five bedrooms, between one and seven occupants, and built between the 1940s and 2010. These homes represent a wide variety of builds and occupancy patterns. The

characteristics of these datasets and the performance of the deployment can be found in Appendix B. This approach allowed the performance of algorithms presented in this thesis to be analysed over many datasets in the targeted environment.

Much of the work in this thesis is described with reference to G-DPS. The G-DPS framework, presented in Chapter 3, provides a generic framework to implement DPSs on a WSN node, and provides a basis within which the contributions to knowledge are evaluated. The G-DPS framework builds upon the SIP algorithm developed by Goldsmith and Brusey [34] by accounting for the issues discovered during the deployment of Cogent-House in the field. To evaluate the performance of G-DPS, L-SIP, an implementation of G-DPS for monitoring linear signals based on SIP, was used as an exemplar using the COGENT-HOUSE datasets. The evaluation methodology is presented in detail in Section 3.5 on page 53.

Chapter 3 also presents Backbone Collection Tree Protocol (B-CTP), an extension to the commonly used CTP networking protocol. The TOSSIM simulator developed by Levis *et al.* [64] was used to evaluate the performance of the new approach. To calculate the energy requirement of B-CTP compared to CTP, the microbenchmarking approach (see Section 2.2.3 on page 20) was used. Microbenchmarking is an approach to calculate the energy requirement of a node in a timely and cost-effective way whilst not imposing hardware constraints. This approach is presented in Section 2.2.3 on page 20.

Chapter 4 presents methods to reconstruct signals from DPSs. To evaluate the effectiveness of the reconstruction methods, each method was applied to all of the COGENT-HOUSE datasets to compare and contrast. The evaluation approach is presented in Section 4.2 on page 85.

Chapter 5 presents Bare Necessities (BN), an approach combining DPS concepts with on-node processing to convert data to information. The reports created for Orbit Heart of England included a number of novel “metrics” to summarise building performance in a way which is readily understood by the end-user (surveyors at Orbit Heart of England). The time-discounted distribution summary (see Section 5.2 on page 97) was selected from these as an example and baseline for evaluation. This evaluation makes use of a year long dataset (house 1) from the full set of data collated by COGENT-HOUSE. A full description of the evaluation can be found in Section 5.4 on page 101.

In summary, to develop and evaluate the contributions presented in this thesis, the author has taken an approach of empirical data-driven evaluation backed up with on-node in-situ evaluation.

1.3 Research scope

In the field of WSNs there is a wide variety of potential applications with differing requirements. MacRuairi *et al.* [68] and Römer and Mattern [99] both describe a taxonomy for WSN applications. This section uses this prior work to describe the scope of applications the work in this thesis applies to. This criteria will be broadly split into application, devices, and network.

1.3.1 Application

This criteria describes the class of applications such as the type of phenomena to be sensed, sampling periods and deployment methodologies. The applications considered are that of environmental monitoring with the following attributes:

Sensed Phenomena: The work here applies to applications sensing multiple or single phenomena.

The sensed phenomena can either be distributed (e.g, temperature in a room), or discrete (e.g., cumulative gas consumption).

Temporal Resolution: Low frequency applications are considered, with sample periods between a few minutes and hours.

Spatial-Resolution / Coverage: The approaches presented in this thesis do not consider spatial resolution, however it is assumed each node is in transmission range of at least one other node.

Size: The work in this thesis has been applied to real sensing system with a range of network sizes, from 4 nodes to over 100 nodes.

Deployment: Deployments are manually performed, and can be either iterative or one-time. When deploying an opportunistic nature is assumed (e.g., if deploying in a home nodes are placed where possible).

Mobility The work assumes that all nodes and gateways are immobile, and have no mobility.

Lifetime: No lifetime requirements are assumed. However, the aim of this thesis is to achieve lifetimes on the order of years.

1.3.2 Devices

The class of devices is described by their heterogeneity and energy source.

Heterogeneity: It is assumed all sensor nodes (and repeaters) will have the same underlying hardware, however the sensors interfaced to these nodes may vary between nodes.

Energy source: All nodes are considered to be battery powered. The integration of energy harvesting energy sources would improve the performance of techniques presented here, however, this is beyond the scope of the work.

1.3.3 Network

The final criteria to describe the application class is the nature of the network:

Communication modality: It is assumed sensor nodes will be using radio communication. Other techniques such as IR and ultrasound are not considered.

Infrastructure A single gateway per sensor network is considered which will have an internet connection if available.

Network: A multi-hop network is considered with a single sink node/base station at the trees root.

Connectivity The network is always considered to be connected.

Latency: There are no strict latency requirements, however, it is assumed a message will be received at the base station within a few seconds after transmission from the sensing node.

Bandwidth: The applications considered have small continuous packet transmissions. A typical packet is less than 100 bytes,

1.4 Contributions to knowledge

By answering the research questions set forth in Section 1.1, this thesis provides the following contributions to knowledge:

1. Generalised Dual Prediction Scheme (G-DPS)—A novel, generalised framework for the implementation of DPS in real life deployments. G-DPS is described and evaluated in Chapter 3.
2. Backbone Collection Tree Protocol (B-CTP)—An extension to CTP to utilise a persistent powered backbone network, which reduces the energy requirement for listening in order to extend node lifetime. Chapter 3 presents B-CTP.

3. The use of a dual quadratic spline signal reconstruction method which improves signal reconstruction accuracy of DPSs This approach is described and evaluated in Chapter 4.
4. Bare Necessities—An algorithm utilising on-node processing to deliver information rather than data, significantly reducing node transmissions Chapter 5 describes and evaluates BN.

1.5 Publications

The work in this thesis has resulted in the following peer-reviewed publications:

Journal papers

- Elena I. Gaura, James Brusey, Michael Allen, **Ross Wilkins**, Daniel Goldsmith, and Ramona Rednic. “Edge mining the Internet of things”. In *IEEE Sensors Journal*. vol. 13, no. 10, Oct. 2013, pp. 3816–3825.

Conference proceedings

- Elena I. Gaura, John Halloran, James Brusey, **Ross Wilkins**, and Ramona Rednic. “Sustainable future? Building and life-style assessment”. In *Proceedings 2012 International Conference on Advanced Computer Science and Information Systems*, Dec. 2012, pp. 7–11.
- Elena I. Gaura, James Brusey, **Ross Wilkins**. “Bare necessities—Knowledge-driven WSN design”. In *Proceedings of 10th IEEE Sensors Conference*, Oct. 2011, pp. 66–70.
- Elena I. Gaura, James Brusey, **Ross Wilkins**, and John Barnham. “Wireless Sensing For The Built Environment: Enabling Innovation Towards Greener, Healthier Homes”. In *Proceedings of Clean Technology 2011*, June. 2011, pp. 367–372.
- Elena I. Gaura, James Brusey, **Ross Wilkins**, and John Barnham. “Inferring Knowledge From Building Monitoring Systems: The Case For Wireless Sensing In Residential Buildings”. In *Proceedings of Clean Technology 2011*, June. 2011, pp. 353–358.

Appendix C on page 163 details further outputs resulting from this work, plus includes the full copies of these papers.

1.6 Thesis structure

This chapter has provided an introduction to this thesis, including the motivation for the work, the research approach adopted, and the research questions and contributions to knowledge. The rest of this thesis is organised as follows:

Chapter 2 discusses relevant background literature to the topics introduced throughout this thesis, focusing on power consumption of nodes, techniques to achieve node longevity focusing on data-driven approaches, and the performances of these approaches in real world deployments.

Chapter 3 proposes and evaluates Generalised Dual Prediction Scheme (G-DPS), a novel generalised framework to develop DPS-style algorithms. G-DPS aims to solve issues that effect the robustness of DPS algorithms in real life deployments. G-DPS is evaluated using L-SIP, an implementation of G-DPS for monitoring linear signals, making use of data collated from real world deployments. This G-DPS framework provides the base for the approaches presented in this thesis. Furthermore, this chapter proposes B-CTP a modification to CTP which extends node lifetime using persistent backbone powered nodes. A real world in-situ evaluation of L-SIP coupled with B-CTP is also presented.

Chapter 4 presents and evaluates a number of reconstruction methods, with a focus on dual quadratic splines which incorporates the known gradient of the signal state to improve signal reconstruction accuracy..

Chapter 5 proposes Bare Necessities (BN), an approach to designing WSNs nodes to deliver only the high-level information an end-user is interested in.

1.7 Acknowledgement of contributed work

This section details the contribution made by other researchers which have aided the work presented in this thesis:

- The software development work for the open-source house monitoring system, COGENT-HOUSE presented in Appendix A was developed in collaboration with Dr. James Brusey.
- Dr. Ramona Rednic and Dr. Olukunle Ojetola, fellow researchers at the Cogent Computing Applied Research Center (CCARC), designed and developed the interface boards and additional hardware required for COGENT-HOUSE.

- Dr. Daniel Goldsmith, a fellow researcher at CCARC, has provided both data from his own deployments of COGENT-HOUSE, as well as updating the open-source code where bugs, issues, and enhancements were identified.

1.8 Acknowledgement of software tools

A number of open-source tools have been used to create the work presented in this thesis. This sections briefly details the main tools used:

TinyOS TinyOS is an open source embedded operating system. Code was written in NesC for TinyOS to develop the WSN nodes in the COGENT-HOUSE system described in Appendix A.

Python The Python programming language is used two-fold in this work: i) It is used for the server-side processing in the COGENT-HOUSE system, ii) it is also used for developing offline versions of the contributions presented in this thesis for testing.

R R is a programming language for statistical process. In this thesis R has been used to evaluate the data resulting from running the python scripts for offline testing. These R scripts make heavy use of Hadley Wickham's suite of tools in particular the ggplot2 package to create all graphs in this thesis.

LyX LyX is the document processing tool used to create this thesis document.

Chapter 2

Energy efficient Wireless Sensor Networks

The aim of this chapter is to: inform the work presented in this thesis, provide background and support for the developments proposed, and reveal the gaps in knowledge and practice in the existing work. The literature review will cover:

1. an introduction to Wireless Sensor Networks (WSNs),
2. the energy consumption of WSN nodes,
3. approaches to energy reduction in WSNs, with a focus on Dual Prediction Scheme (DPS) and,
4. WSN signal reconstruction and data imputation,

This chapter is structured as follows: Section 2.1 provides an overview of WSNs. Section 2.2 discusses the energy consumption of WSN nodes. Section 2.3 provides an overview of the categories of power optimisation approaches in WSNs. Section 2.4 describes attempts to reduce energy through data-driven approaches which forms one of the core themes of this thesis. Section 2.5 describes the performance of DPS in WSNs in real life deployments. Section 2.6 outlines methods to reconstruct sensed signals specifically where samples are missing or have been suppressed. Finally, Section 2.7 summarises the literature surveyed, the gaps found, and their relationship to the work in this thesis.

2.1 Wireless Sensor Networks

Advances in microcontrollers, low powered radio technology, and MicroElectroMechanical Systems (MEMS) sensor technology have given rise to Wireless Sensor Networks (WSNs). WSNs have paved the way to collect data in a range of fields in a less labour intensive and more accurate and timely manner. Akyildiz *et al.* [3] describes a WSN as a number of autonomous sensor devices distributed over a geographical

This item has been removed due to third party copyright. The unabridged version of the thesis can be viewed at the Lanchester Library, Coventry University.

Figure 2.1: Typical multi-hop WSN architecture. Reprinted from Reddy [95]

area. This thesis considers the specific case of a static WSN which consists of a (large) number of sensor nodes deployed in a given environment. The sensor nodes periodically *sense* the environment through interfaced sensors, *process* that sensor data using an on-board microcontroller, and *transmit* the results wirelessly using multi-hop communication, to a central sink node, using an on-board radio chip. This sink node collects and forwards data, from all sensor nodes, to a base station that is responsible for the storage and processing of sensor readings. Figure 2.1 shows a typical topology for a WSN deployment.

WSNs have become proven research instruments for several application domains including:

Natural Environment A wide range of natural environments have been monitored using WSNs, from monitoring extreme environments not suitable for human habitation through to more commercial examples such as farmland. Examples of monitoring include: volcanic eruptions [118], soil moisture in vineyards [42], and ice quakes on glaciers [72].

Built Environment Building monitoring has been used for a range of applications, including monitoring the environmental conditions of homes to advise on heating strategies [98], monitoring heritage buildings for structural movement [19], and monitoring environmental conditions when storing museum artefacts [115].

Health In health applications, WSNs have been used to: monitor patients with Parkinson's disease [83], monitor behaviour in the elderly or vulnerable [69], and detect falls [80].

Defence Several battlefield applications have been developed, such as: sniper detection and localisa-

tion [59], battlefield surveillance [13] and monitoring bomb suit operative physiological conditions [30].

WSNs are commonly acknowledged today as proven research instruments for several application domains. However, WSNs still have many open research issues including energy management, fault tolerance, scalability, deployment processes, and harsh environments. This thesis focuses on the issues of i) energy constraints (aiming to maximise node lifetime by minimising the energy requirement of a node) and ii) fault tolerance (providing mechanisms by which the impact of faults can be reduced).

2.2 Wireless node energy consumption

Due to the often inaccessible nature of WSN deployments, the primary power source for a WSN node is in most cases batteries. Since batteries have only a finite supply of energy, long term deployments may require maintenance to replace depleted batteries. Replacing batteries can present problems as nodes may be inaccessible, or replacement may be costly or time consuming, such as when deployed on a glacier, in a jet turbine engine, or within a smart home. Due to the limited energy available, the energy requirement of a node must be minimised in order to maximise node lifetime¹.

To optimise a node's energy use it is useful to start by understanding how the energy is used. This section describes: the design tradeoffs for developing energy efficient WSNs, typical lifetimes of existing WSNs, the different ways in which nodes use their energy.

Römer and Mattern [99] show that the lifetime requirements of WSN vary greatly depending on the application. They evaluate the characteristics of 15 deployments showing that the lifetimes range from a few hours, in the case of furniture assembly monitoring, up to five years for a WSN deployed in the ocean. The average lifetime is shown to be between several months and a year.

Even though the ocean deployment has a lifetime up to five years, it is rare for systems to achieve longer than a 1 year lifetime. For example, the Torre Aquila project, described by Ceriotti *et al.* [19], deployed a WSN to monitor the structural integrity of a heritage building to better plan maintenance. The deployment consisted of a number of node types measuring temperature, relative humidity, light, acceleration and fibre optic sensors all sampling at an interval of 10 minutes. Using two pairs of size C batteries (7600 mAh), Ceriotti *et al.* estimate a node lifetime of one year. SensorScope, described by Schmid *et al.* [102], is a system for an indoor environmental monitoring network, built around the Telosb

¹The phrase node lifetime can refer to several attributes of a node. For example the length of time before a software failure, or before a hardware component fails, or the lifetime in terms of energy. In this thesis the phrase node lifetime refers to a node's battery life given a fixed capacity battery, assuming node hardware/software works indefinitely.

platform (measuring temperature, relative humidity and light) and uses the B-MAC [86] networking protocol (commonly referred to as Low Power Listening). With the nodes sampling at a two minute interval, Schmid and Dubois estimate that the system will run for 61 days on a pair of AA batteries. The WISE-MUSE project, described by Peralta *et al.* [84], resulted in a WSN which monitors temperature, relative humidity and light in an art gallery for the preservation of collections. At a sampling interval of 10 minutes a lifetime of two months was achieved using a pair of AA batteries. These three examples show the limited expected WSN node longevity monitoring simple measurands such as temperature and humidity at relatively low sample frequencies.

2.2.1 The lifetime-accuracy-timeliness trade-off

WSN nodes often have a limited energy supply. While in some cases nodes can be connected to mains power, usually the application demands they are powered by batteries. When developing WSN systems a key design goal is to therefore extend node lifetime by efficiently using this limited power supply. Bajwa *et al.* [12] describes the power-distortion-latency (or lifetime-accuracy-timeliness) trade-off, in which the design of a WSN is a trade off between sensor node lifetime (power), data accuracy (distortion) and timeliness of data updates (latency).

To increase a node's lifetime a designer of a WSN could extend a node's lifetime by:

1. Reducing the accuracy of the sensing node, for example, by reducing the sample rate.
2. Reducing the timeliness of the data being received by, for example, buffering data and sending once per day or even requiring manual data download.

However, the goal of WSN research is to design WSN nodes that use minimal energy to collect and disseminate an accurate representation of the phenomena being observed.

2.2.2 WSN node power requirements

A WSN node usually incorporates five main sub-components: a sensing unit, a processing unit, a transceiver unit, a storage unit, and a power unit [3, 90]. Table 2.1 gives an overview of the energy consumption of three widely used commercial node platforms (TelosB, Mica2, and MicaZ) that will be referred to throughout this section. The remainder of this section describes the sensing, radio, processing, and storage energy requirements of WSN nodes.

The primary aim of a WSN is to sense the environment and transmit readings to a central store and therefore a WSN node is equipped with sensors and associated circuitry to interface to the micro-

Table 2.1: Power consumption of typical commercial nodes [85]

Mode	Mica2 (2002)	MicaZ (2003)	TelosB (2004)
Mote standby (RTC on)	0.06 mW	0.09 mW	0.02 mW
MCU idle (DC0 on)	10.56 mW	10.56 mW	0.18 mW
MCU active	26.40 mW	26.40 mW	5.94 mW
MCU + radio RX	49.83 mW	76.89 mW	71.94 mW
MCU + radio TX	83.82 mW	69.30 mW	64.35 mW
MCU + flash read	31.02 mW	31.02 mW	13.53 mW
MCU + flash write	71.28 mW	71.28 mW	49.83 mW

Table 2.2: Reported energy consumption of selected sensors

Sensor	Active / Passive	Energy consumption
ELT Inc. B-530 CO ₂ sensor module	Active	1560 mW
Applied Sensor IAQ2000 VOC sensor module	Active	150 mW
Sensiron SHT11 Humidity sensor	Passive	1.5 mW
Sensiron SHT11 Temperature sensor	Passive	1.5 mW
Hamamatsu S1087 Light sensor	Passive	4.8 mW

processor's inputs. Raghunathan *et al.* [90] define two categories of sensors: active and passive. For passive sensors, such as temperature, the phenomena acts on the sensor, causing a change in properties, for example, changing its resistance (thermistors) or generating a voltage (thermocouple). Active sensors (such as radar, and Carbon Dioxide (CO₂)) act on the environment to take a measurement.

Table 2.2 provides a breakdown of the energy use of select sensors. Passive sensors draw little current compared to other components of sensor nodes. However, active sensors can be large consumers of power. In the case of the CO₂ sensor module (ELT Inc. B-530), the typical peak current is 130 mA—over five times more than the radio. Over a single sample period the B-530 CO₂ sensor consumes 0.013 Wh over its 30 second “active” time compared to the 9×10^{-6} Wh consumed by the radio over half a second. Therefore for the node the use of the CO₂ sensor will be the main limiting factor on the node lifetime. Passive sensors are much more common in WSN deployments, meaning that the cost of sensing is often minimal.

To communicate data or receive commands, WSN nodes are equipped with a radio chip. From a simplistic view the radio has three possible states: i) sending packets, ii) listening / receiving packets, or iii) switched off. Table 2.1 shows that radio usage has the greatest energy requirement and that to transmit packets uses as much energy as receiving. Polastre, Szewczyk, and Culler [85] showed that the TelosB's radio consumption is ten times greater than processing and 3900 times greater than idle. Pottie, and Kaiser [88] show that the energy used to transmit 1 kilobyte (kB) of data is the equivalent of performing 3 million operations. Since the radio is the greatest energy consumer, reducing radio usage

should increase node lifetime.

Table 2.1 on the preceding page shows that the power consumption for the Microcontroller (MCU) is small compared to other node subsystems such as the radio. A node's lowest energy consumption state is when the node's MCU is idle. Therefore, when the application permits, the MCU should be in an idle state to extend the node's lifetime. Interestingly, the power required for the MCU has decreased over subsequent platform revisions, whereas the energy required for the radio has stayed approximately the same.

Some applications may require the storing of historical sensor values. A node offers two types of storage. A small amount of data (often on the order of tens of kilobytes) can be stored in Random Access Memory (RAM), however this storage is vulnerable to power loss. If the quantity of data to be stored is large, or requires permanent storage, data can be written to integrated flash memory. Table 2.1 shows that while the energy to read from flash is small, writing to flash is close to that of radio usage. Nguyen *et al.* [77] gives an example of caching data before transmission, in this example more energy is required to write the data to flash compared with transmitting every sample. Therefore, any use of on-node storage needs as much consideration as transmission.

2.2.3 Estimating node lifetime

The previous section describe how a node uses its energy. This section describes how the lifetime of a node can be estimated using these measurements.

The simplest method of determining the lifetime of a node is to allow the node to run until it is no longer functional. While this is able to give an accurate estimate of lifetime, the approach is impractical for lifetimes approaching periods of months or years.

Dutta *et al.* [26] describe iCount. iCount measures energy usage by counting the switching cycles of the regulator. iCount provides a benefit over other approaches by being simple in design. However, in terms of accuracy the approach can have errors up to 20% compared to the actual measurement during constant load.

Kopke and Wolisz [58] describe the energy measuring approach for the SANDbed WSN testbed [41]. Kopke and Wolisz developed a circuit for precise energy measurements in-situ. The energy is derived by measuring the current draw of a node over a shunt resistor and the supply voltage. This energy is then reported back to the node. This approach has two disadvantages i) as with iCount measuring the idle current is inaccurate, and ii) it is unclear what additional energy is required by the node to manage the measuring circuit. Milenkovic *et al.* [75] discuss a similar approach, however do not report the accuracy

or performance achieved.

Jiang *et al.* [48] describe SPOT, another hardware based approach. Compared to the other three approaches SPOT achieves a very high resolution (down to 1 μA for constant currents) and is able to react to the rapidly changing current draw of a WSN node. However, compared to the previous approaches the set up is much more complex and the paper does not detail the accuracy of the measurement circuit.

Hardware approaches are generally said to be more accurate than software / calculation based approaches. However, when considering a node using minimal energy, hardware approaches are unable to accurately measure at the microamp level. As discussed later in this chapter, when implementing transmission reduction algorithm nodes are primarily in a sleep state therefore pure hardware approaches are not suitable. An alternative to this approach is microbenchmarking.

2.2.4 Microbenchmarking

Microbenchmarking is an alternative approach that takes measurements of tasks to calculate energy requirements for an application [53, 57, 60, 73]. The energy requirement of a node is derived from the current consumption I over time t , to calculate the total charge Q .

Current I is the rate of charge flow. Therefore, the total charge, Q , that flows through the two points over a period of time t , if the current is constant, is given as:

$$Q = It$$

The potential difference or voltage V across two points is defined as the energy E dissipated or transferred per coulomb of charge, Q , that moves through a single point:

$$V = E/Q$$

Therefore the energy consumed (or dissipated) by an electrical circuit over a period of time and assuming a constant current for that period is:

$$E = VQ$$

In the case of WSN's the voltage is fixed. For example, the Cogent-House system nodes are powered by two AA batteries with a capacity of 2400 mAh. Therefore, for all calculations a fixed voltage of 3 V is assumed (although this method can be used for any supply voltage). Since there is a fixed voltage the consumed charge is proportional to the energy consumed and therefore will be used to show the relative

energy improvement.

The microbenchmarking approach calculates the energy requirements of a node from two measures:

1. The average current draws of sleep states I_{sleep} and operations $I_{\text{operation}}$. Typical operations that a node performs are: sensing, processing, listening, and transmissions. The current draw of an operation is measured by performing an operation continuously and measuring the current draw of the node using a digital multimeter in series.
2. The time $t_{\text{operation}}$ to completion of an operation, recorded from timestamps in the code.

Microbenchmarking uses the two measures of each operation to calculate the energy consumption for a task, $E_{\text{operation}}$,

$$E_{\text{operation}} = I_{\text{operation}} \times t_{\text{operation}}$$

The energy consumed by all operations performed during a sample period E_{sample} is the sum of the energy for all operations performed. For example, for a typical sense-and-send application,

$$E_{\text{sample}} = E_{\text{sensing}} + E_{\text{processing}} + E_{\text{listening}} + E_{\text{send}} + E_{\text{sleep}}$$

Since the discharge curve of batteries can vary between manufacturers, this thesis reports on the annual energy requirement of a node rather than a predicted lifetime. The energy requirement for a node over a year E_{year} is calculated from the energy for a single sample period E_{sample} and the number of samples performed in a year n

$$E_{\text{year}} = E_{\text{sample}} \times n$$

Compared to the hardware approaches microbenchmarking is: less time consuming, cost-effective, and has no node platform constraints. However, some accuracy may be lost due to the consumption not being measured in-situ and not accommodating the battery model. The approach is accurate enough to give an annual energy requirement. Microbenchmarking will be used in this thesis to measure the energy performance of algorithms coupled with in-situ performance analysis.

2.2.4.1 Microbenchmarking method

This section gives a step-by-step method for performing microbenchmarking measurements and evaluation:

1. For each operation (an operation is defined here as being a contiguous set of processing steps that use a similar amount of power, such as performing calculation, sensing temperature, or transmitting wirelessly) in a WSN application:
 - (a) Create a program that repeats an operation N times (for some large value of N) to estimate the average running time for that operation.
 - (b) Place a multimeter in series between the battery pack and WSN node, and record the average current over a number of iterations
 - (c) Calculate the charge for an operation using $Q = It$
2. Calculate the energy required for one sample period by summing all operations from step 2. The time required for when a node is idle can be calculated by subtracting the total time of all operations in a sample from the sample period:

$$t_{\text{sleep}} = t_{\text{sample}} - (t_{\text{processing}} + t_{\text{listening}} + t_{\text{send}} + t_{\text{sense}})$$

3. Estimate energy use for a year. For example, in a sense-and-send application, with energy requirements of 0.011 mAh per sample and sensing at 5 minutes, use the formula stated in the previous section:

$$E_{\text{year}} = E_{\text{sample}} \times n$$

$$E_{\text{year}} = 0.011 \times 105120$$

$$E_{\text{year}} = 1,156 \text{ mAh/year}$$

2.3 Approaches to energy saving in WSNs

Raghunathan *et al.* [90] state that a key design constraint for many WSN applications is the energy consumption of the node. Furthermore it has been identified that power optimisation is a key roadblock to the adoption of the Internet of Things (IoT) [114]. This section examines approaches to energy conservation in WSNs.

Anastasi *et al.* [9] defines an energy saving taxonomy for WSNs as shown in Figure 2.2. This taxonomy defines techniques to be based on either the networking subsystem (the design of efficient networking protocols, or duty cycling), or the sensing and processing subsystem (data-driven reduction). Mobile

Figure 2.2: Taxonomy of approaches to energy savings in sensor networks. Reprinted from Anastasi *et al.* [9].

This item has been removed due to third party copyright. The unabridged version of the thesis can be viewed at the Lanchester Library, Coventry University.

sensing is also discussed, however this thesis assumes fixed position sensing nodes. Two approaches not considered by Anastasi *et al.*'s taxonomy is that of hardware optimisation or energy harvesting.

Hardware Optimisation Refers to either i) the design of energy efficient components (e.g., low power radios), or ii) hardware resource management through energy aware operating systems such as PixieOS [67]. Promising avenues include future low power processors, low power communication units, and energy harvesting to produce embedded systems that can function with significantly less energy [74]. Technology in this area is making relatively slow progress, for example while Table 2.1 on page 19 demonstrated that the energy consumption of the MCU was reduced between three sensor nodes, the power requirements of the radio have remained relatively static.

Energy harvesting extends node lifetime through the conversion of external environment sources (e.g., thermal energy [14], vibration [7], or solar power [65]) into electrical energy.

Networking Optimisation The network can be optimised through either i) the design of energy efficient network protocols [2], or ii) energy-efficient MAC Layers [27, 86]. Though many differing network protocols have been suggested, these approaches have rarely been implemented or evaluated in real deployments [31].

Data-driven Optimisation Software-based approaches can reduce energy consumption through analysis of the sensed signal. Such techniques may take the form of data compression, data aggregation, or data prediction. The primary benefit of these approaches are that they are not tied to specific hardware platforms. The main limiting factor of these approaches is the characteristics of the signal and the efficiency of the underlying hardware.

These approaches can be implemented individually, however it is often the case that a combination of these approaches will be used. The approach of data-driven optimisation is the most generic one and can be applied to a wide range of applications. The following section provides details of the current state of the art around data-driven optimisation in WSNs.

2.4 Data-driven optimisation in WSNs

The work in this thesis primarily investigates data-driven optimisation in WSNs. The taxonomy shown in Figure 2.2 splits the data-driven optimisation approach into two sub-approaches:

Energy efficient acquisition In some cases the energy cost of sensing may be higher than the energy requirement of the radio, for example the CO₂ sensor module discussed in Section 2.2. Therefore

an approach is required that reduces the number of acquisitions (i.e., data samples). Examples of these approaches include adaptive sampling [5], hierarchical sampling [105], and model-based active sampling [82].

Transmission reduction This approach reduces the energy requirement of a node by exploiting characteristics of the sensed signal in order to minimise the number of transmissions required.

As previously shown in Section 2.2, in many applications, the energy requirement for sensing is low compared to the energy requirement for radio use. Therefore, the following sections will look at transmission reduction techniques.

2.4.1 Event detection

One approach to reduce the number of transmissions to the sink is to only transmit when a pre-defined event occurs, for example a temperature exceeding a set value. Event detection algorithms have been used in applications such as sniper detection [59] or detection of animal calls for localisation [6], in these cases a node only reports when a target has been located.

Event detection can reduce the number of transmissions significantly, however, Kapitanova, Son, and Kang [20] show that specifying event thresholds can be difficult. The context of the event must be considered, for example, if a node is deployed to measure air temperature in the summer a temperature of 0 °C would be considered a significant event, whereas in the winter it would not be unusual. Since WSNs are generally expected to be deployed in environments that change over time (for example, seasonal variation), the relevance to the application of the predefined event may change. This may result in missing potentially interesting events that would give additional understanding of the environment.

In an event detection approach the number of transmissions is proportional to the number of events that occur. Therefore, in applications which experience infrequent events, for example, detecting ice quakes on glaciers, lifetimes can be substantially increased. However, event-based transmissions are a destructive approach, only reporting when an event has occurred. WSNs are deployed to gain an understanding of a environment therefore it is often important to have a view of the signal leading up to the event to determine how and why an event occurs, for example, is an ice quake due to a faster rate of melting of the glacier.

Werner-Allen *et al.* attempt to solve this issue through Lance [119]. Lance calculates a summary of a window of data on the sensor node. The full data window is stored on the node, and the summary transmitted to the sink. Through a user specified rule based policy the sink determines if the summary

indicates data is interesting. If so the full raw data stream is requested from the sensor node. Although Lance attempts to reduce the number of transmissions to extend node lifetime, Section 2.2 has shown that storing data in flash is nearly as energy expensive as transmissions.

2.4.2 Data compression

Rather than just reporting events, data compression allows the original signal to be reconstructed. Data compression is a technique which encodes information using fewer bits than the original representation. There are two main classes of compression algorithm: lossless compression can exactly reconstruct the original signal from the compressed signal, whereas lossy compression removes some “less important” information to achieve greater compression ratios. Lossy compression therefore only allows an approximation of the original signal.

Run Length Encoding (RLE) is a lossless data compression algorithm, which has been used in the compression of images, for example palette-based bitmapped images such as computer icons. RLE replaces repeating sequences of values with a single copy of the sequence and a count of occurrences. For example the string:

BBBBBBBBBBBBBBBBBAAAAAAAAACCCCCCOOONN

would be represented as 15B8A6C3O2N, from which you can reconstruct the original string. A disadvantage to RLE is that it relies on the signal containing repeated sequences. Natural phenomena, such as temperature, however, do not tend to produce signals with simple repeating patterns. Capo-Chichi, Guyennet, and Friedt [18] attempted to solve the issue of variability in signals with K-Run Length Encoding (K-RLE). K-RLE is a lossy data compression algorithm which introduces an error budget, allowing similar readings to be grouped and thus reducing the effect of the variability in the signal. Similar approaches to event detection have been proposed by Lazarus and Mehrotra [62] termed Poor Man’s Compression (PMC):

Poor Man’s Compression - Midrange (PMC-MR) only transmits when the midpoint of a data window exceeds twice a defined error threshold. When the segment range exceeds the error threshold, the midpoint of the range and a count of samples since last transmission is transmitted.

Poor Man’s Compression - Mean (PMC-MEAN) on the other hand reports the mean value of a data window when it exceeds a threshold rather than the midpoint. This approach reduces the mean error of the reconstructed data, however it produces a higher number of transmissions than the PMC-MR.

Olsten, Loo, and Widom [81] propose the approximate caching algorithm. Instead of transmitting an actual value the algorithm reports an interval approximation of the sensed value, for example, the value is guaranteed to lie between 10 and 20. When a sensed value falls outside the range of the previous approximated interval then the interval is updated to take account of the new value, termed a value-initiated refresh. On the sink, Structured Query Language (SQL) like queries are registered with a central stream processor, along with the maximum error permissible for that query. If the transmitted approximated interval exceeds the threshold for the query then the sink executes a query-initiated refresh which returns the exact sensor value and an update range. The disadvantage of this approach is that only a value interval is reported, therefore the original data stream cannot be accurately reconstructed. This may affect any analysis of data gathered, for example, when calculating an average both an upper and lower bound of that value would need to be calculated and reported.

Rather than compress data at an individual node level, in-network aggregation processes data at an intermediary node (usually a cluster-head) to attempt to reduce transmissions. Fasolo, Rossi, and Widmer [28] provide a review of in-network processing techniques. The review defines the process of in-network aggregation as the process of gathering and routing information through a multi-hop network, processing data at intermediate nodes with the objective of reducing resource consumption (in particular energy) and thereby increasing network lifetime.

Heinzelman, Chandrakasan, and Balakrishnan [40] developed LEACH. LEACH is a cluster-based approach that reduces communication by aggregating data at a intermediate node. All nodes in a cluster transmit their sensor readings to the cluster head, once all reading are received they are aggregated into a single packet and forwarded to the sink. This approach will reduce the overall traffic of the network, however, each node will still be required to transmit at each sample interval.

Intanagonwiwat *et al.* [44] describes the directed diffusion approach. In this approach a user specifies an *interest*, which comprises of an event definition, the area in which this event is expected, and the interval (data rate) required. The algorithm disseminates the interest through the network, forming gradients (or paths) from which the data matching the interest can be drawn back to the sink. Direct diffusion includes support for data aggregation to happen at neighbouring nodes. Intanagonwiwat *et al.* show that when matching reports from neighbouring nodes are suppressed, this approach was shown to save five times the energy compared to no suppression.

TAG [71], as described by Madden *et al.*, aims to reduce the number of transmissions by producing user specified summary statistics (such as minimum, maximum or average) at intermediary nodes in the network tree. Using SQL group by style commands allow for an overall smaller packet, at a cost of

Algorithm 2.1 On-node algorithm summary for DPS.

At each sensing cycle:

1. obtain sensor reading
 2. estimate new state
 3. predict state based on last transmitted state
 4. if new state estimate is significantly different from predicted state
 - (a) transmit new state to sink
-

accuracy, to be transmitted. Although TAG reduces packet size through summarisation Chu *et al.* [22] show, that when given the option to produce results using in-network aggregation the approach has been unpopular, with end-users opting to obtain the raw signal from the nodes. However, the end-users of the application were domain scientists who are used to working with raw signals, in many applications non-specialist end-users are more interested in a summary.

Though in-network aggregation looks promising, Gaura *et al.* [31] observe that WSNs rarely collaboratively work together at the application-level—only the Media Access Control (MAC) layers require cooperation between nodes. Furthermore these approaches mostly reduce transmissions in routing and so each sensing node will still need to transmit every sensed sample to an intermediate node, which is the largest drain on energy reserves.

A disadvantage of these compression approaches, generally, is that they do not consider how a signal evolves over time. Since many environmental phenomena follow a predictable trend (such as outside temperature) there is often redundancy in the data that could be exploited to reduce the number of transmissions required. Many approaches have considered modelling the signal to reduce transmissions further.

2.4.3 Data modelling approaches

Even though sensor data is often highly compressible, the benefit of transmitting smaller data packets is often outweighed by the energy consumption overheads associated with transmission. For example, Kim [56] demonstrates that, with the CC2530 radio using Z-Stack, around 60% of the energy to transmit is the overheads related to switching on and configuring the radio. Therefore reducing the number of transmissions will reduce energy usage much more than reducing the quantity of data per transmission.

One approach to reduce the number of transmissions is based on the node and sink having a shared model of the signal, with predictions made using state updates transmitted by the node. Transmissions

are only made when the predicted state of the signal differs from the current sensor reading by a set condition. This approach, named Dual Prediction Scheme (DPS) [63], is summarised in Algorithm 2.1. The rest of this section focuses on this approach.

Jain, Chang, and Wang [46] suggest a Dual Kalman Filter (DKF) approach. The DKF approach makes use of identical Kalman filter models at the node and sink. The estimation stage of the Kalman filter estimates the next sensor reading. If the error between the actual sensor reading and the prediction exceeds some given threshold then an update of Kalman filter parameters are sent to the sink. A Kalman filter may seem a suitable choice for modelling a signal but can introduce significant latency [16], requires additional computation, and is difficult to tune [51], often needing extensive historical data from the intended deployment environment. The need for prior data streams and the difficulty in tuning the Kalman filter makes it unsuitable for a generic approach to transmission reduction.

Santini and Römer [101] present an alternative approach using a Least Mean Squares (LMS) adaptive filter for the state estimate. Adaptive filters, such as LMS, are generally used where there is no knowledge of the phenomena signal (such as the noise-level) available. Therefore, this approach improves upon DKF by removing the need for a-priori knowledge of the observed phenomena. In addition, LMS is an adaptive filter, therefore, it reduces the need for significant training compared to filters such as the Kalman filter in the DKF method. Though this approach provides a large reduction in packets, it has been shown the LMS adaptive filter has a convergence period where the reported values do not match the sensed signal. During this convergence period the node will require extra transmissions and have a higher reconstruction error during this period.

The Ken approach, described by Chu *et al.* [22], is a DPS approach which uses a pair of probabilistic models at the node and sink. In this paper Chu *et al.* uses a model of the signal, an example given is a linear prediction model, to calculate a probability distribution function, and a transition model to forward predict. As with other DPS approaches the accuracy of the model is compared against the current model state—if the error in the model exceeds a user specified threshold then an update of the model state is transmitted to the sink. The paper also shows that the model can take into account correlation between neighbouring nodes to improve performance, however, this requires significant intra-node communication. Furthermore, the KEN approach requires a significant period of training data which ranges from 4–15 days (11520–43200 samples) on the Intel data set, requiring a large amount of additional energy at the start of the deployment.

Liu, Wu, and Tsao [66] suggest the use of AutoRegressive Integrated Moving Average (ARIMA) [15] models. Though able to accurately model a signal, this approach requires a long training phase which

Algorithm 2.2 Pseudocode for node and sink for SIP**Node:**

```

 $s \leftarrow \text{query sensor}()$ 
 $\mathbf{x}' \leftarrow \text{estimate new state}(s, \mathbf{x}_{\text{old}}, t_{\text{old}})$ 
 $\mathbf{x}_{\text{sink}} \leftarrow \text{predict sink state}(\mathbf{x}_{\text{sink}}, t_{\text{sink}})$ 
if  $|\mathbf{x}' - \mathbf{x}_{\text{sink}}| > \varepsilon$  :
    transmit( $\mathbf{x}'$ )
     $\mathbf{x}_{\text{sink}} \leftarrow \mathbf{x}'$ 
     $t_{\text{sink}} \leftarrow t$ 
 $\mathbf{x}_{\text{old}} \leftarrow \mathbf{x}'$ 
 $t_{\text{old}} \leftarrow t$ 

```

Sink:

```

[On receipt of new state estimate( $\mathbf{x}$ )]
 $\mathbf{x}_{\text{sink}}(t) \leftarrow \mathbf{x}; t_{\text{last}} \leftarrow t$ 

[Estimate value for time( $t$ )]
if  $t \geq t_{\text{last}}$ 
    predict from  $\mathbf{x}_{\text{sink}}(t_{\text{last}})$ 
else
    interpolate from neighbouring  $\mathbf{x}_{\text{sink}}$ 

```

requires additional computation and storage of historical values. The ARIMA approach is also computationally expensive. Probabilistic Adaptable Query (PAQ) [111] and Similarity-based Adaptive Framework (SAF) [110] reduce the complexity and computation of ARIMA models by focusing on an Autoregressive (AR) model. PAQ improves upon SAF, by incorporating a trend component which reduces the training set required. Both of these approaches allow correlations between nodes to be exploited, however, they still require a long period of training to derive an accurate model. Borgne, Santini, and Bontempi [63] present Adaptive Model Selection (AMS), which attempts to reduce the training time for AR models. While in the training phase, several models are run simultaneously. A racing algorithm is implemented to select the best performing model for data prediction. This approach, however, still requires approximately 1000 readings for training.

Rather than make use of complex filtering, ARIMA models, or probabilistic models, Goldsmith and Brusey propose the Spanish Inquisition Protocol (SIP) [34] (shown in Algorithm 2.2). SIP extends prior work on DPS by transmitting a state vector rather than individual readings. This state vector consists of the current reading and rate of change of the phenomena. Accounting for the trends in the data stream allows the use of a simple piecewise linear model to predict future sensor readings.

Figure 2.3 shows the transmission reduction performance of selected DPS techniques, as reported by Goldsmith [33]. These protocols were all tested offline using the Intel Lab dataset [70] using an error threshold of 0.5 °C. A threshold of 0.5 °C is commonly for indoor environment monitoring applications. However, in applications such as human body monitoring tighter thresholds must be used, however, this will come at the cost of additional transmissions.

KEN was the worst performer requiring 65% of samples to be transmitted, this value does not consider the training period required therefore the number of transmission may be significantly higher. Filter-based approaches such as DKF and LMS improve on this by a factor of 6 and are able to reconstruct the sensed signal using approximately 10% of the collected samples. Out of these two approaches the

This item has been removed due to third party copyright. The unabridged version of the thesis can be viewed at the Lanchester Library, Coventry University.

Figure 2.3: A comparison of data reduction techniques, as applied to the publicly available Intel Lab dataset [70] using an error threshold of 0.5°C. Reprinted with permission from Goldsmith [33].

LMS method is favourable as it does not require extensive tuning or prior data. The SAF approach has a similar transmission reduction performance to SIP requiring less than 3% of transmissions. However, for SAF, the Intel dataset was preprocessed to adjust the sampling rate and infer missing values [110] and thus the results may not be indicative of actual deployed performance.

The work in this thesis builds upon the current work in DPSs. DPS techniques promise significantly extended lifetimes by reducing the number of required transmissions. However, design and evaluation of DPS algorithms to respond to the needs of real world deployments has received little attention (see Section 2.5 on page 34). In Chapter 3, this thesis examines the issues of deploying SIP in real life deployments and proposes a generalised framework, named Generalised Dual Prediction Scheme (G-DPS), for the development of DPS algorithms targeted towards real world deployments. Since SIP [34] reduces transmissions further than any other method within its allowed threshold it has been shown to be the best performer. Therefore, SIP is used as an exemplifier for implementing a DPS algorithm in terms of the G-DPS framework.

2.4.4 On-node processing

As an alternative to modelling the sensed signal there may be benefits to processing on the sensing node itself. On-node processing algorithms use application specific algorithms to transform or summarise data and provide a reduced transmission size or number of transmissions to the sink. While this has parallels to the concepts such as data compression, the focus here is the transformation of data into another form of information.

Pottie and Kaiser [88] suggest that if application and infrastructure permit, it is beneficial to process the data locally to reduce traffic volume and reduce energy costs. End-users are typically interested in converting raw signals into high-level knowledge [43, 121]. This means that traditional sense-and-send systems waste energy on the transmission of data with little informational relevance.

Albu, Lukkien, and Verhoeven [4] propose a method of on-node processing of Electrocardiography (ECG) signals. Rather than transmit a raw ECG signal the algorithm calculates and transmits R-Peaks (the peak amplitude of a typical ECG signal) which allows for the detection of simple arrhythmia (e.g., high/low heart rates or missing beats). This on-node processing leads to a node lifetime increase by a factor of 3–5 \times compared to a sense-and-send approach, however, this approach does not take account of the changes between R-Peaks over time. Incorporating an event detection type approach, such as only transmitting when changes occur in the detected heartbeat, could additionally increase the lifetime. PEAR [106] takes a similar approach of calculating R-Peaks on-node. Using the interval between peaks, PEAR derives a person's activity-level through a decision tree-based activity classifier. The energy conservation of the approach is dependent on how active the users are. Assuming users are active 18%–28% of the monitoring period, the approach is able to increase node lifetime up to 2.5 \times compared with conventional ECG sensing approaches.

The approach proposed by Kasi *et al.* [54] annotates detected events on sensor nodes and executes application rules on these events. By performing data processing locally, the sensor nodes makes decisions quickly and remotely without the need for instructions from gateway nodes. This approach achieved an energy saving of 33% for a transmission reduction of 60%. This approach was evaluated through simulation and only for a period of one minute (at 1Hz sampling rate). It could be that over longer periods the algorithm performs better, but no long-term evaluation has been undertaken.

Histogram-like summaries of time spent in certain states are useful in a variety of applications, and can be calculated on-node. Reddy *et al.* [96], uses Global Positioning System (GPS) and accelerometer data from a mobile phone to determine the percentage of time an individual spends in particular transportation states (e.g., bike, walk, motor). From the same research group this concept has been used by Ryder *et*

al. [96, 100] to monitor the mobility patterns of individuals using mobile phone GPS data. However, these approaches are often computed from raw data traces at the sink and thus do not explicitly target node-level data reduction.

When considering a specific application, the knowledge end-users are trying to obtain is a refined form of the sensed signal [17]. The combination of on-node processing with DPS techniques, should be able to significantly decrease the number of transmissions required by a node and therefore increase node lifetime whilst delivering the information required by the end user. This thesis examines this technique in Chapter 5.

2.5 Use of DPS in WSNs deployed in the field

In this chapter, DPSs have been shown to significantly reduce the number of transmissions required by a WSN node with the aim of reducing the energy requirement of a node. However, the design and evaluation of DPSs often does not consider the case of multi-hop networks or lossy networks. This section explores the literature relating to these two areas.

2.5.1 DPS on multi-hop networks

While there are large bodies of literature related to multi-hop networks and to DPS algorithms, few publications attempt to answer the question of how these two technologies interact with each other.

Raza *et al.* [93] investigate the effect of a DPS algorithm, Derivative-Based Prediction (DBP), deployed in an operational road tunnel. This deployed WSN is implemented using a combination of Collection Tree Protocol (CTP) and Low Power Listening (LPL) (also named Box-MAC [76]). In their experiments they show that the average radio duty cycle (the fraction of the sensing cycle in which the radio is active) with a DPS algorithm and good compression is essentially the same as when using CTP with no transmissions. They conclude that CTP beacons become dominant in terms of transmissions once you can achieve good compression and therefore, further compression provides little gain unless the MAC and routing layers can be improved.

Building on their previous work, Raza *et al.* [94] proposed increasing the wakeup interval for LPL and decreasing the number of CTP beacons as a solution to reduce the overall power consumption of a multi-hop network where DPS is used. In a sense-and-send system this would result in longer transmission times (due to the reduced listening periods), and thus have a much higher energy requirement. However, due to the compression offered by their DPS style algorithm, the listening and the beacons are the overwhelming

power consumers and therefore the savings gained with their changes outweigh the increased cost of transmitting.

In conclusion, to further reduce the energy requirement of a node when transmissions have been sufficiently reduced, network maintenance related to radio usage should be kept to a minimum to reduce the radio duty cycle.

2.5.2 DPS in real-life deployments

Anastasi *et al.* [8] looked into reliability and energy efficiency for 802.15.4 networks as a function of the sleep/wake cycling and the MAC level settings. They showed that multi-hop reliability can be very poor, providing less than 60% delivery rate, when using the default settings, regardless of the duty cycling method (around 5% delivery for the fully synchronised scheme up to around 60% for a ZigBee compliant scheme). Anastasi *et al.* showed that by changing the MAC settings, the reliability of all duty cycling methods they examined can be significantly improved (85% to 100% depending on the scheme), though with the trade-off of higher energy consumption and end-to-end latency.

Arora *et al.* [10], discovered the same result as Anastasi *et al.* for a different type of network. Using the distance-vector routing and queue management protocols provided as the default in TinyOS, only 33.7% of packets were delivered across their network. This figure was increased to 81% by implementing a new routing scheme making use of knowledge about the node deployment.

The Intel Lab dataset [70] is often used to evaluate DPS offline. Environmental data was collected from 54 Mica2Dot nodes deployed in the Intel Berkeley Research lab over a period of 30 days. Data was collected using the TinyDB in-network query processing system, built on the TinyOS platform. An analysis of data yield shows that only $30\% \pm 0.1$ of packets were successfully transmitted in this network—matching the results by Arora *et al.* and Anastasi *et al.*.

In the case of DPSs, a high packet delivery rate is important to enable an accurate reconstruction of the sensed signal, and therefore a reliable delivery mechanism is required to alleviate the issue with lossy networks. However, as shown by Anastasi *et al.*, this might only be achieved as a trade-off with higher energy consumption.

Raza *et al.* [93] point out that additional measures (beyond that offered by the network stack) are required to ensure reliable delivery to provide correct operation of the DPS algorithm. Furthermore, techniques are required to detect permanent node failure as early as possible.

The rest of this section investigates node health monitoring in WSNs. Node health monitoring has a wealth of research and a variety of systems have been proposed. However, these approaches generally

apply to WSNs implementing a sense-and-send approach. Node health monitoring for WSNs implementing event-driven or transmission suppression approaches are less covered. In DPS algorithms, detecting node failure is more crucial since the approach relies on both the node and the sink having a synchronised state of the model for accurate signal reconstruction.

Node failures can be either transient, intermittent, or permanent. Transient issues include packet loss due to a collision of packets in a network, and are usually short term. Intermittent failures include a node losing a route to the sink, which are usually longer in duration, but can fix themselves. Permanent failures are unrecoverable without intervention, for example, the exhaustion of a node's energy supply, theft of a node, or damage to the hardware. In event-based systems, such as DPS algorithms, transmissions are intermittent compared to sense-and send where packets are received on a regular basis. This leads to two issues i) it is not possible to detect when transmissions have been lost, and ii) it is not possible to detect that a node has failed. The rest of this section focuses on the use of acknowledgements, sequence numbers, and heartbeat messages to detect node failures.

DPS algorithms rely on both the sink and node having identical copies of the sink state. Therefore a node must be able to detect when the sink fails to receive and store a state update transmission. In WSNs acknowledgements have been used in sense-and send systems to detect transmission failure. Dam and Langendoen [113] present T-Mac a contention-based Medium Access Control protocol for wireless sensor networks. In their design they include the use of acknowledgements to confirm the delivery of packets but do not evaluate their use. In the same vain WiseMAC (Wireless Sensor MAC) presented by El-Hoiydi and Decotignie [27], use acknowledgements to confirm delivery of packets. The Monnit Corporation² use acknowledgements in their system to confirm the reliable transmission of sensor readings—if an acknowledgement is not received from the sink the node attempts to transmit a further three times.

When considering event and DPS algorithms acknowledgements have received little attention. Tezcan *et al.* [107] present end-to-end reliable event transfer schemes for WSNs. Tezcan *et al.* show that the use of acknowledgements improves delivery of events by 20%, trading off against additional computational requirements. However there is no discussion of the increase in energy required for this improvement. End-to-end acknowledgements enable a node and sink to keep their state models synchronised, meaning a better data yield and improved reconstruction accuracy. However, this is a traded off against an increase in required transmissions.

This thesis investigates the use of acknowledgements with DPSs and their impact on energy requirements, data yield, and signal reconstruction accuracy.

²<http://www.monnit.com/support/faqs.php>.

Further to a node being able to detect when one of its transmissions have failed, it is important for the sink to detect transmission failure. The sink must be able to detect this to know when it is unable to trust the model state for a node. The use of sequence numbers to tag packets is a common feature of WSN design. The continuity of sequence numbers received from a node can provide a means of detecting packet loss. Zhao and Govindan [122] used sequence numbers to investigate packet delivery performance in a dense sense-and-send WSN. Detecting gaps in a monotonically increasing sequence number tagged to a packet allowed for the calculation of packet loss. Iyer *et al.* [45], Wan *et al.* [116], Wang *et al.* [117] among others all used sequence numbers in the same way to detect transmission loss. Since WSNs are often lossy, using a sequence number in model state updates should allow for the sink to detect where transmissions have failed. In turn this will allow for the calculation of data yield, and show where reconstructed data may not meet accuracy requirements.

The use of a heartbeat message is another common feature in WSNs and distributed systems to detect failure. A heartbeat message indicates that a node is still functional when the node has not reported for a defined period. Athanassoulis *et al.* [11], Cunah *et al.* [25], Gobriel *et al.* [32], and Sharaf *et al.* [103] all describe the use of a uniform heartbeat message to detect failures. If a heartbeat message (or a number of heartbeat messages) is not received by the sink within a user specified time period then the node is deemed to be faulty.

Haas *et al.* [37] and Yadav and Khilar [78] describe a method where nodes inter-communicate to check neighbours are still functional. A node requests a heartbeat message from its neighbours, if there is no response from the target node then it is deemed to be faulty. Once a node has queried its neighbours the information about the state of the network is disseminated to every node. Since this approach will require many more transmissions than the heartbeat approaches described in the previous paragraph this approach will not be suited for DPS algorithms.

Ramanathan *et al.* [92] describe Sympathy, a heartbeat (or metric period) based tool for detecting and debugging node failures devised for sense-and-send monitoring applications. However rather than just send a keep-alive message, Sympathy includes debug information. When a heartbeat message is sent the node gathers information such as the number of packets transmitted, and transmits it to the sink. The sink determines failure if insufficient data has been received from the sensing node on receiving a heartbeat packet. If a node is deemed to have failed the information in the heartbeat packet is used in a decision tree to decide the nature of the failure.

Heartbeats are also used in commercial wireless sensing devices, for example, the Monnit Corporation include a heartbeat period in their sensing devices. Monnit defined two heartbeat types depending on

if the sensor node is event-based or continuous. The event-based heartbeat is a keep alive mechanism. Monnit Corporation recommend a heartbeat of an hour for event-based heartbeats. However, in Section 2.4 it is shown that DPSs can achieve a reduction of 98.2% in the case of the Intel lab dataset. Over a 24 hour deployment this equates a packet around every 5 hours at a five minute sample interval, and therefore heartbeats will be received more often than state updates themselves.

Heartbeats are a well established tool for detecting node failure in WSNs, however their use has so far been limited to pure event-based applications or sense-and-send based WSNs. This thesis will evaluate the use of heartbeats within DPS algorithms, evaluating their impact on the energy requirement of nodes.

Finally, in the instances of the DPSs that have been discussed in this chapter [22, 34, 46, 101, 110, 120], all only consider the case of a single sensing modality. However, sensing nodes generally include multiple sensors of differing types. Furthermore, when dealing with multiple signals in a given environment, the signals are often highly correlated (for example, temperature and relative humidity).

Transmitting the packet payload is typically a small proportion of each packet. For example, assuming a single sensor state per transmission, the packet size for a CTP packet would be 48 bytes—composed of 40 bytes of management data and eight bytes of payload sensor state data. In a multi-modal approach if another sensor state is added (doubling the payload size) there is only a 17% increase in packet size. Kim [56] demonstrates that, for the CC2530 radio using Z-Stack, approximately 60% of the energy to transmit is related to switching on and configuring the radio. Therefore, the energy cost of sending a larger packet with additional data values in the payload is often minimal compared to the overheads of sending a packet at all. The additional energy required to transmit a larger packet will be offset by a reduction in the overall number of transmitted packets.

2.6 WSN signal reconstruction

DPS algorithms provide a significant reduction in the number of transmissions from a node. However, selection of the best method to accurately reconstruct the original signal based on the output of these algorithms has received little attention in the literature.

A common technique to interpolate missing values in a signal is to use linear interpolation. However, this linear reconstructions does not consider the signal gradient and therefore can result in high reconstruction errors. Splines improves on linear interpolation reconstruction by including this trend component. Splines have been used in existing work to provide smooth estimates of signals based on noisy input samples [50, 112] and for samples gathered at a nonuniform rate based on the characteristics

of the signal [36]. Spline based methods take into account the values at a start and end points along with the gradient at both points.

Aggarwal and Parthasarathy [1] present the Conceptual Reconstruction approach. The Conceptual Reconstruction approach calculates missing data from correlated signal streams to create a model of the signal. The evaluation of this approach assumes that missing data is lost randomly. The remaining data is said to be representative of the original signal, however this may not be true for DPSs where the missing data is somewhat engineered based on the event detection rules [104]. Furthermore, this approach relies on the sensor readings in a network being correlated.

Tropp and Gilbert [108] describe an approach of reconstructing signals using the orthogonal matching pursuit (OMP) reconstruction algorithm. OMP is a greedy reconstruction algorithm that computes the best nonlinear approximation to a signal. However, Ji, Sun, and Shen [47] found the approach is computationally expensive when considering a large scale sensor network.

Silberstein *et al.* [104] propose a Bayesian approach to infer the value of suppressed or missing datapoints. This approach may be useful in some applications as it allows estimation of uncertainty in addition to the probabilistic inference of the suppressed / missing datapoints. However, for applications not requiring this capability, a splines based approach is simpler in terms of implementation and computation and takes advantage of the information already provided by DPSs to provide accurate reconstruction of the sensed signal.

Spline-based methods are well suited to DPSs as many approaches transmit the value of the sensed reading, and the additional energy cost of computing and transmitting the signal gradient is slight. This thesis investigates the use of splines, to move beyond the existing literature to improve signal reconstruction by taking advantage of additional information inherent in the use of SIP—signal gradient estimates, bounds on the suppressed samples, and sequence numbers to detect failed transmissions.

2.7 Summary

Numerous algorithms in the literature provide methods to reduce the number of transmissions a node is required to make. An approach often used in the literature is that of a DPS type algorithm. DPSs share a model of the data between the node and sink. The node makes a prediction of the sink's state based on the last state transmitted to the sink. Transmissions are only made if this predicted state differs from the current sensor reading by more than a defined threshold. Based on the literature, SIP (a DPS making use of a linear model) is shown to have the best performance in terms of the accuracy of the reconstructed

signal and the transmission reduction offered. The work presented within this thesis will therefore build upon the prior work in developing SIP by Goldsmith and Brusey [34].

Within the literature there is little evaluation of the performance of DPSs on-node in real life deployments, and therefore have not been designed to handle issues such as: the energy requirement of the WSN network stack, lossy networks, node failure, and accommodating the use of multiple sensing types on a single node. Furthermore, the selection of the best method to accurately reconstruct the original signal based on the output of these algorithms has received little attention. Finally transformation of data into knowledge on the node coupled with DPSs techniques should significantly reduce the number of transmissions a node is required to perform, and thus reduce the energy requirement of a node.

This thesis focuses on the following areas:

1. There is no general framework to implement DPS on WSNs deployed in the field that react to transmission failure, node failure, and the need for multiple sensor modalities on a single node. This discussed is discussed and solved in Chapter 3.
2. An investigation into an approach which allows for an energy usage reduction which is proportional to the transmission reduction is required. This investigation is also presented in Chapter 3.
3. An investigation into methods of DPS signal reconstruction techniques. This is presented in Chapter 4.
4. An evaluation of an approach which combines DPS concepts with the calculation of application-level state (knowledge generation) on-node. This approach is presented and evaluated in Chapter 5.

The next chapter proposes and evaluates G-DPS, a novel generalised framework to implement DPS-style algorithms on-node for use in real life deployments.

Chapter 3

G-DPS: A generalised framework for Dual Prediction Schemes

The previous chapter described Dual Prediction Schemes (DPSs), a type of data reduction algorithm where a model of the data is shared between the node and sink. Transmissions are only made when the predicted state of the model (based on the last state transmitted to the sink) differs from the current state by more than a defined threshold. From these approaches the Spanish Inquisition Protocol (SIP) was found to be the best performing algorithm, achieving a high reconstruction accuracy while reducing transmissions by a factor of $20\times$ in environmental datasets. However, the design of DPS algorithms presented in the literature has not considered issues with real life deployments. Therefore, DPS algorithms as an off-the-shelf software component are unable to handle several aspects of real world deployments including: the energy requirement of the node's wireless network stack, lossy networks, node failure, and accommodating the use of multiple sensing modalities.

This chapter proposes the Generalised Dual Prediction Scheme (G-DPS), a novel generalised framework to develop DPS-style algorithms as described in Chapter 2. The G-DPS framework includes approaches to i) implement DPS with multiple sensors, ii) detect node and transmission failure, iii) calculate reconstructed data yield, and iv) use acknowledgements to maximise model synchronisation time when transmission failure occurs. G-DPS enables a reduction in transmissions, an increase in reconstruction accuracy, and improves reconstructed data yield. The G-DPS framework serves as the base from which the algorithms presented in the remainder of this thesis are derived.

This chapter also presents the Backbone Collection Tree Protocol (B-CTP). Since few nodes are responsible for routing, most nodes do not need to listen for incoming packets. B-CTP reduces the number of nodes involved in listening for packets by utilising a persistent backbone network of permanently powered nodes. Battery powered nodes are only required to transmit a packet to their closest backbone node. Coupling Linear Spanish Inquisition Protocol (L-SIP)—G-DPS implemented with a linear model—with B-CTP reduces the time the radio is used and as a result decreases the annual energy requirement

for a TelosB node by a factor of $13.4\times$.

The evaluation of G-DPS allows the following research questions to be answered: **RQ1:** *What features can improve the robustness of DPSs implemented in deployed Wireless Sensor Networks (WSNs)?*

Furthermore, the evaluation of B-CTP allows the following research question to be answered: **RQ2:** *Can the lifetime of a WSN node implementing transmission reduction approaches be increased further by using a persistent backbone network of mains powered routing nodes?*

This rest of this chapter is structured as follows: Section 3.1 motivates the work presented in this chapter, followed by definitions of terms in Section 3.2. Section 3.3 presents G-DPS, a novel, generalised framework for the implementation of DPSs on-node. Section 3.4 provides an example implementation of G-DPS in the form of L-SIP. Section 3.5 evaluates the G-DPS framework. Finally, Section 3.6 proposes and evaluates B-CTP, a low-power networking approach.

3.1 Motivation

When applied to data-traces, DPSs promise to significantly extended node lifetimes by reducing the number of required transmissions. In the previous chapter, SIP was shown to be the best performer in terms of transmission reduction and signal reconstruction accuracy. However, the design and evaluation of SIP and other DPS algorithms in the literature has not considered issues with real life deployments. When SIP was implemented on hardware and deployed in real-life scenarios a range of issues were revealed and these were: lossy networks, node failure, the need to integrate multiple sensors, excessive energy consumption. This chapter explores solutions to inform robust deployments of DPS algorithms in deployments.

WSNs deployed in the field are often lossy. That is, one cannot guarantee the delivery of each transmitted packet. Anastasi *et al.* [8] looked into reliability and energy efficiency for 802.15.4 networks as a function of the sleep/wake cycling and the MAC level settings. They showed that multi-hop reliability can be very poor, with between 5%–60% delivery rate for a multi-hop ZigBee network. Arora *et al.* [10] found similar results with less than 35% delivery rate for a multi-hop network using the Mica2 platform with TinyOS. The Intel Lab Data [70] is commonly used to evaluate DPS algorithms [22, 34, 46, 101, 110, 120], however, the deployment of 54 nodes for a period of 37 days achieved average yield of only $30\% \pm 0.1$.

DPSs require that both the node and the sink have identical states of the model to reconstruct the signal accurately. However, there is no way for the sink to detect when transmission failure has occurred,

or the node to detect when a state update packet has failed to be received and stored.

DPS algorithms transmit at irregular and unpredictable frequencies. Therefore, a node not reporting data may be either because the node is functioning but is suppressing messages as intended or because the node has failed. The end-user is unable to distinguish between transmission suppression and node failure. Therefore a method is required to distinguish between failed nodes and nodes that are suppressing messages as intended.

DPSs are generally designed with the aim of compressing one sensing modality. However, sensing nodes generally include multiple sensors of differing types. When dealing with multi-modal signals in a given environment, the signals are often highly correlated (for example, temperature and relative humidity). Therefore, a method which combines all sensors into a single model could reduce the number of required transmissions and increase reconstruction accuracy.

Further power savings could be obtained by considering the interaction between the compression algorithm and the network stack. SIP promises significant energy savings by reducing the number of transmissions. While there is a large body of literature related to multi-hop networks and to DPS algorithms, few publications attempt to answer the question of how these two technologies interact with each other. SIP is able to reduce the number of transmission of node by a factor of $20\times$. However, SIP implemented on a TelosB only decreases the energy requirement by a factor of $1.3\times$ due to the networking-related overheads. Therefore the Media Access Control (MAC) layer and communication protocols must also be considered to maximise potential energy savings. Section 3.6 proposes an alternative network topology to significantly decrease the energy requirement of nodes.

The next section provides the definitions for the algorithms, SIP, G-DPS, L-SIP, and B-CTP that are central to this chapter.

3.2 Definitions

The following definitions will be used in the rest of this thesis:

Spanish Inquisition Protocol (SIP) This is the algorithm presented by Goldsmith and Brusey [34] and used here as a starting point for implementing a DPS algorithm within the proposed framework. SIP is described in Algorithm 2.2 on page 31.

Generalised Dual Prediction Scheme (G-DPS) This is a novel generalised framework for implementing DPS style algorithms and forms a contribution of the work in this thesis. This framework is described in Section 3.3.

Algorithm 3.1 G-DPS node algorithm. This is a generalised framework within which DPS algorithms, such as the previously published SIP algorithm, can be implemented [34].

function G-DPS

Step 1. $\mathbf{z} \leftarrow$ obtain vector of sensor readings

Step 2. $t \leftarrow$ current time

Step 3. $\mathbf{x}_{\text{new}} \leftarrow$ estimate new state $(\mathbf{z}, \mathbf{x}_{\text{old}}, t, t_{\text{old}})$

Step 4. $\mathbf{y}_{\text{new}} \leftarrow$ simplify $(\mathbf{x}_{\text{new}})$

Step 5. $\mathbf{y}_s \leftarrow$ predict sink state $(\mathbf{y}_{\text{sink}}, t_{\text{sink}}, t)$

Step 6. if eventful $(\mathbf{y}_{\text{new}}, \mathbf{y}_s)$ or $t - t_{\text{sink}} \geq t_{\text{heartbeat}}$
 (if the state is eventful or if time since the last transmission exceeds a threshold)

(a) transmit $(\mathbf{y}_{\text{new}}, n, t)$

(b) $n \leftarrow n + 1$

(c) when acknowledgement received:

i. $\mathbf{y}_{\text{sink}} \leftarrow \mathbf{y}_{\text{new}}$

ii. $t_{\text{sink}} \leftarrow t$

Step 7. $\mathbf{x}_{\text{old}} \leftarrow \mathbf{x}_{\text{new}}$

Step 8. $t_{\text{old}} \leftarrow t$

Linear Spanish Inquisition Protocol (L-SIP) This is an implementation of G-DPS with a linear model similar to SIP. Section 3.4 presents this implementation of G-DPS.

Backbone Collection Tree Protocol (B-CTP) This is an extension to Collection Tree Protocol (CTP) to utilise a persistent powered backbone network, which reduces sensing node energy requirements for listening. This networking approach is described in Section 3.6.

3.3 The Generalised Dual Prediction Scheme

G-DPS, presented in Algorithm 3.1, is a generalised framework to implement DPS algorithms to reduce the number of transmissions a node is required to make.

1. Sense the environment

In the G-DPS algorithm the node takes a reading from its integrated sensors. The algorithm uses a vector of sensor readings \mathbf{z} rather than an individual reading.

Step 1. $\mathbf{z} \leftarrow$ obtain vector of sensor readings

DPS algorithms in the literature are often defined for use with one sensor signal, whereas sensing nodes generally include multiple sensors of differing types. The G-DPS algorithm integrates all sensor readings into a single vector. Since the sink state for any given sensor will be updated more frequently compared to sending individual state updates, the approach will reduce signal reconstruction error. However, the total number of required transmissions will be reduced compared to multiple single-modal DPS instances.

Once the sensor readings have been taken, the node also records the current time t .

Step 2. $t \leftarrow$ current time

Estimate the new state

The current state is estimated using the vector of readings \mathbf{z} , the previous state \mathbf{x}_{old} , and the previous time t_{old} .

Step 3. $\mathbf{x}_{\text{new}} \leftarrow$ estimate new state ($\mathbf{z}, \mathbf{x}_{\text{old}}, t, t_{\text{old}}$)

This component estimates the current state of the phenomena, transforming the vector of sensor readings into an estimate of the state \mathbf{x}_{new} . The component relies on the selection of appropriate models and state. In the SIP algorithm, for example, the new state is estimated using an Exponentially Weighted Moving Average (EWMA) filter.

The model describes the evolution of the phenomena over time. Usually simple models such as constant or linear models are appropriate choices for a wide variety of phenomena, for example, temperature or, the state of a light switch. More sophisticated models may be required for other signals, for example, Electrocardiography (ECG) signals. In this case, it may make more sense to derive a summary (such as heart rate rather than the full ECG signal).

The state estimate should provide enough information to predict future state. Since a node has limited storage space, a full history of sensor readings cannot be stored¹. Therefore, based on current hardware capabilities, a state estimation method should be selected which has the Markov property (future state only depends on present state) or only requires a small subset of historical data.

3. Simplify the new state

When producing the state estimate, additional information may be required for the calculation that is not necessarily required by the end-user. Simplify is an optional transformation component used to remove unnecessary information from the state estimate before transmission. The estimated state vector \mathbf{x}_{new} may contain more information than must be transmitted. Thus it is often useful to generate a simplified form \mathbf{y}_{new} .

¹Clearly this will not be true for all possible WSN systems, but is a common constraint reported in the literature.

Step 4. $\mathbf{y}_{\text{new}} \leftarrow \text{simplify}(\mathbf{x}_{\text{new}})$

The main limitation of the simplify step is that it can be a destructive process—it is not possible to reconstruct the original signal. Using the example of human posture, the system will only report on the posture the monitored person was classified as being in and not the raw accelerometer data. Therefore, it is often the case the application will need to be well defined to use this simplify step. However, when the simplify step transforms data into narrowly-specified information, the number of potential ways this information may be further mined is reduced, thus improving privacy.

4. Predict the sink state

The sink state \mathbf{y}_s is predicted using the last transmitted state vector \mathbf{y}_{sink} , the current time t , and the time of the last successful state transmission t_{sink} . SIP, for example, achieves this through linear extrapolation.

Step 5. $\mathbf{y}_s \leftarrow \text{predict sink state}(\mathbf{y}_{\text{sink}}, t_{\text{sink}}, t)$

This estimated state is used as an input to the event detection component.

5. Detect events

If either the state is eventful, or if the time since last transmission exceeds a threshold, or if no acknowledgement was received on the last transmission then a state transmission is triggered.

Step 6. if eventful($\mathbf{y}_{\text{new}}, \mathbf{y}_s$) or $t - t_{\text{sink}} \geq t_{\text{heartbeat}}$

Event detection is based on a comparison of the current state estimate and the predicted state. If the comparison satisfies the conditions for an event to be detected, then a state update containing the current state is transmitted.

In addition to detected events, transmissions may be triggered via expiration of a heartbeat period. The predictability of data collection from a sense-and-send node provides a simple mechanism to check the health status of a node. However, DPS algorithms transmit at irregular and unpredictable frequencies. For example, in the case of Bare Necessities (BN), presented in Chapter 5, a node may suppress messages for months at a time. The end-user is unable to distinguish between transmission suppression and node failure.

The G-DPS framework defines a maximum time period, $t_{\text{heartbeat}}$, allowed without a transmission. If the node has not transmitted in this defined period a state transmission is triggered. This transmission is called a heartbeat, and indicates the node is still functional. The sink checks for failure by checking that

a state update is received within the allotted heartbeat time. Any nodes that do not report in the time period are reported to be faulty. The heartbeat approach is to prevent long-term failures. Intermittent failures are treated by sequence numbers and end-to-end acknowledgements described in the next step.

6. Transmit a state update

If a transmission is required then all sensor states \mathbf{y}_{new} are transmitted along with a sequence number n and the current time t .

Step 6 (a.) transmit $(\mathbf{y}_{\text{new}}, n, t)$

The sequence number is incremented for each transmission.

Step 6 (b.) $n \leftarrow n + 1$

Practicalities of implementation dictate that the sequence number will wrap around to some zero at some value N . The sequence numbers are therefore in the range $[0, N)$ where $N = 2^\kappa$ for a sequence number of κ bits. The use of sequence numbers to calculate yield is discussed in Section 3.3.3 on page 49.

When a node transmits a packet, the delivery of the packet to the sink cannot be guaranteed. To accurately reconstruct the sensed signal when using a DPS, both the sensing node and the sink require identical copies of the state estimate. If the state estimate differs between the sensing node and the sink then the reconstruction will not accurately represent the node's true readings. Therefore the reconstruction could exceed the defined error threshold until the next state estimate is received.

To verify that both the sink and nodes are using identical state estimates, software-level end-to-end acknowledgements, (referred to simply as acknowledgements from here) are used to indicate when transmissions have failed. The sensing node's copy of the sink state \mathbf{y}_{sink} will only be updated if an acknowledgement is received from the sink. The approach for acknowledgements is end-to-end as an acknowledgement will only be transmitted when the state estimate has been successfully stored. If the acknowledgement is received before some timeout then the transmission is successful and the state is updated, otherwise the transmission has failed and the node should not update its copy of the sink state. If the state update transmission fails the node will transmit the latest state on subsequent sample periods until a packet has successfully been acknowledged.

Step 6 (c.) when acknowledgement received:

Step 6 (c.i.) $\mathbf{y}_{\text{sink}} \leftarrow \mathbf{y}_{\text{new}}$

Step 6 (c.ii.) $t_{\text{sink}} \leftarrow t$

Algorithm 3.2 G-DPS sink functions

function OnReceive(i, y_t, n) $y_{i,t} \leftarrow y_t$ (store received state for node i at time t) $l_i \leftarrow t$ (update last received time for node i)acknowledge(i, n)**function** Estimate(i, t)**if** $t \geq l_i$ **extrapolate** from y_{i,l_i} to $y_{i,t}$ **else** t_{prev} = last state received before t t_{next} = next sample received after t **interpolate** between $y_{i,t_{prev}}$ and $y_{i,t_{next}}$ to calculate $y_{i,t}$ **function** IsFunctional(i, t)**return** $t - l_i \geq t_{heartbeat}$

7. Update previous state estimate with current state

Finally, the previous stored state vector and time is updated with the current state vector and current time to be used in the next sample period,

Step 7. $\mathbf{x}_{old} \leftarrow \mathbf{x}_{new}$ Step 8. $t_{old} \leftarrow t$ **3.3.1 G-DPS Sink Algorithm**

Algorithm 3.2 describes the G-DPS sink algorithm. The sink provides three functions:

1. OnReceive(i, y_t, n)

The sink maintains a set of application-level state y and timestamp t pairs for each node i . Upon receipt of a new packet from a node, the sink stores the application-level state along with its associated timestamp. If the application-level state is successfully stored then the sink transmits an acknowledgement to the sensing node including the received packet's sequence number n .

2. Estimate(i, t)

The estimate function is used to estimate sensor readings from the stored node states. An end user would use this function in two cases:

- (a) If the end user requires the current state of a node, then the state is extrapolated from the last received state update for that node.
- (b) If the end user requires a past value the sensor value is estimated using interpolation between the two neighbouring states of that time.

3. IsFunctional(i, t)

This function checks that a node i is still functional at the current time t . If the time since the last received application-level state l does not exceed the heartbeat period $t_{heartbeat}$ then the function returns true indicating the node is still functional. Otherwise false is reported indicating a potential fault. As an example this function is used in COGENT-HOUSE in automated reporting of performance.

3.3.2 G-DPS Assumptions

G-DPS is intended to be a generalised approach to DPSs. However, the principals on which it is based require the following assumptions to be made:

The energy cost of sensing and processing the phenomena is less than that of radio transmissions Common to DPS techniques, G-DPS assumes that the energy cost of sensing and processing is small compared to use of the radio. It was shown in Section 2.2 that this is generally true. However, when a node is integrated with active sensors such as Carbon Dioxide (CO₂) it can be the case that sensing can be the greatest energy cost. When the sensing cost is greater than transmission cost techniques that reduce the sample frequency, for example compressive sensing [21], will provide greater savings in transmissions/energy.

Events in the phenomena are not independent Section 3.5.3 shows that a multi-modal approach can be advantageous to DPSs. However, if considering a phenomena where events for different modalities have no relationship then multi-modal will provide no benefit over single-modal.

3.3.3 Detecting transmission loss and calculating data yield

Calculating the yield of a node implementing DPS techniques, such as G-DPS, is not as straight forward as with sense-and-send. With sense-and-send nodes, each sample period has a corresponding transmission and therefore the data yield can be calculated from the number of packets received and the number of

expected packets for a deployment period:

$$y|_a^b = \frac{\text{samples received}}{t_b - t_a}$$

However, as previously discussed, DPSs remove this regular period. Therefore, there is no way of knowing when packets have been lost. Furthermore, DPSs have a one-to-many relationship between a packet and the number of data points that can be reconstructed from that packet. Therefore the yield cannot be calculated using the formulation above. By incorporating a sequence number into packets, the sink can detect transmission loss and infer times where data can be accurately reconstructed.

Consider the case where a node implementing a DPS algorithm sends no sequence numbers with a transmission. If transmissions are received at time t_1 and t_{10} there is no possible way to know if any transmissions were attempted from t_2 to t_9 . Including sequence numbers allows for the calculation of the number of attempted transmissions between sample periods.

If the difference between the sequence number received at t_a , and sequence number received at t_{a+1} is greater than one, this indicates the number of transmissions that have failed. With the inclusion of end-to-end acknowledgements, discussed in the next section, these failures occur one after another until successful transmission. The difference in sequence numbers directly translates to the number of data points that cannot be accurately reconstructed. Therefore sequence numbers provide a simple mechanism to help calculate the data yield of a node implementing a DPS algorithm.

One complication is that sequence numbers have a limited size and can wrap-around. For example, an 8-bit sequence number has a maximum value of 255. When wrap around occurs it can be ambiguous how many packets have been lost—wrap-around may have occurred more than once. The remainder of this section describes a new method to calculate the data yield of a WSN node employing DPS techniques. To limit ambiguity, the yield calculation calculates the possible upper and lower bounds on the number of wrap-arounds.

For a finite period for which a yield is to be calculated, packets with sequence numbers s_a, s_{a+1}, \dots, s_b are received at time t_a, t_{a+1}, \dots, t_b . A packet received at time t_i can be weighted based on the number of data points the packet can reconstruct. This weighting w_i of a packet i is,

$$w_i = q_i - m_i$$

Where q is the number of sample periods between two subsequent transmissions, and m is the number of transmissions that are missing. The number of missed transmissions is less than the number of sample

periods $m_i < q_i$, and therefore the weighting is always non-zero $w_i > 0$. The number of sample periods q is the difference in sample times between the current and subsequent packet,

$$q_i = t_{i+1} - t_i$$

When a node transmits a state update that is not acknowledged the node continues to transmit state updates on subsequent sample periods until an acknowledgement is received. Each transmission increments the sequence number until a transmission is successfully acknowledged. Therefore, the number of data points that cannot be reconstructed m is ordinarily the difference in sequence numbers between between the current and subsequent packet. Therefore the number of missed transmissions is,

$$m_i = s_{i+1} - s_i$$

However, the sequence numbers wrap-around and are thus in the range $[0, N)$ where $N = 2^\kappa$ for a sequence number of κ bits. Furthermore, the number of wrap-arounds $j \in \mathbb{Z}$ may be greater than one during long periods of failure, if the period is sufficiently long to allow it. Since the number of elapsed sample periods is known, and the number of missed data points is less than the number of sample periods, the number of wrap-arounds has an upper bound,

$$j \leq \left\lfloor \frac{q_i - s_{i+1} + s_i + 1}{N} \right\rfloor$$

Furthermore, the lower bound must be 1 if the ending sequence number is lower than or equal to the starting one,

$$j \geq [s_{i+1} \leq s_i]$$

Iverson brackets are used in the above formulation², this gives a result of 1 if the condition is true or 0 otherwise.

The number of wrap-arounds j may be more than one possible value, for example, if the lower bound is zero and the upper bound is 4 the node could have wrapped around between 0–4 times ($j \in \{0, 1, 2, 3, 4\}$), and therefore is ambiguous. In the case where the number of wrap-arounds is ambiguous the number of missed packets should be calculated conservatively using the upper bound with a flag reported to indicate this is a conservative estimate.

²http://en.wikipedia.org/wiki/Iverson_bracket

Algorithm 3.3 L-SIP phrased in terms of G-DPS (See Algorithm 3.1).

estimate new state

dEWMA filtering:

$$\begin{aligned} x'_1 &\leftarrow \alpha z + (1 - \alpha)(x_1 + x_2 \Delta t) \\ x'_2 &\leftarrow \beta(x'_1 - x_1) / \Delta t + (1 - \beta)x_2 \end{aligned}$$

(Update filtered estimates of value x_1 and rate of change x_2 . The time interval between samples is denoted by Δt .)

simplify

$\mathbf{y}_{\text{new}} \leftarrow \mathbf{x}$ (no simplification)

predict sink state

$$\mathbf{y}_s \leftarrow \begin{pmatrix} 1 & t - t_{\text{sink}} \\ 0 & 1 \end{pmatrix} \mathbf{y}_{\text{sink}} \text{ (linear extrapolation)}$$

detect events

$$|\mathbf{y}_{\text{new}} - \mathbf{y}_s| > \varepsilon$$

(The measurement is eventful if the value estimate y_{new} differs from the prediction y_s by at least some threshold ε .)

The data yield percentage between times t_a and t_b can thus be calculated:

$$y|_a^b = \frac{1}{t_b - t_a} \sum_{a \leq i < b} w_i$$

3.4 The Linear Spanish Inquisition Protocol

Linear Spanish Inquisition Protocol (L-SIP), shown in Algorithm 3.3, is an enhanced implementation of the original SIP algorithm [34] using the G-DPS template, encoding the state as a point in time value and rate of change ($\mathbf{y} = (x, \dot{x})^T$). Throughout the rest of this chapter L-SIP is used to exemplify and evaluate the benefit of each G-DPS feature.

When implementing an algorithm based on G-DPS for a given application, four components must be defined: estimate new state, simplify, predict sink state, and detect events. The following describes the implementation of L-SIP, shown in Algorithm 3.3, using G-DPS as a framework.

Estimate new state

Many methods can be used to estimate the current state of a sensed phenomena such as: EWMA, Normalised Least Mean Squares (NLMS), or a Kalman filter. Kalman filters are a popular choice in the literature for modelling data from WSNs. However, Brusey *et al.* [16] show that the use of a Kalman

filter introduces latency into systems because it is slow to process. In addition Julier and Uhlmann [51] show Kalman filters are more difficult to tune than the other methods listed, often requiring a large quantity of data from the intended deployment environment. Furthermore, the Kalman filter requires more computational power compared to simpler methods such as EWMA, whereas in WSNs the amount of computational power is often limited.

SIP as originally designed uses an EWMA filter. However, due to the diurnal nature of typical environmental data, the Dual Exponentially Weighted Moving Average (dEWMA) [52], the second order form of EWMA, is used for state estimation in L-SIP. dEWMA is an extended version of EWMA which accounts for trends in the data. It is sufficiently simple to be implemented efficiently on low powered nodes and often provides a comparable result to a Kalman filter without the added complexity or computational overhead.

Simplify

L-SIP requires the transmission of both the point in time value and rate of change. Therefore L-SIP requires no simplification for the state vector.

Predict sink state

The *predict sink state* component predicts the current state through linear extrapolation of the last transmitted state estimate.

Detect events

The *detect events* component calculates the absolute difference of each component in the current state vector and the predicted state vector. If any difference in a component exceeds an error threshold ε then an event is said to have occurred.

3.5 Evaluation of G-DPS

This section evaluates the benefit of each feature provided by G-DPS. The following features are evaluated: i) multi-modal processing, ii) heartbeat messages, iii) end-to-end acknowledgements.

A new feature is deemed to be a benefit to the G-DPS framework if there is an improvement in any of the following:

1. the number of transmissions is decreased,
2. the accuracy of the reconstructed signal is increased,

Table 3.1: Summary of datasets used

Dataset period	Number of datasets	Equivalent trace-days
Two-weeks	235	3290
One-months	170	5100
Six-months	40	7300
Year	9	3240

3. the yield of the system is increased, or
4. the energy requirement of a node is decreased.

3.5.1 G-DPS evaluation method

The evaluation makes use of data traces from a total of 235 sensors nodes, sensing air temperature and relative humidity, deployed in 38 homes. The homes monitored include a mix of flats and houses, with between 1 and 5 bedrooms, between 1 and 7 occupants, and built between the 1940s and the 2010s. These homes therefore represent a wide variety of builds and occupancy patterns. These datasets consist of data from deployments from two-weeks to two years. For evaluation over differing durations, example periods of two-weeks, one-month, six-months, and a year were extracted from each dataset where deployment periods allowed. The number of datasets for each duration is summarised in Table 3.1, and details of their properties can be found in Appendix B. The data collected from the COGENT-HOUSE (Appendix A) deployments is used instead of the commonly used Intel Lab Data due to larger quantity of available data, the nature of the deployments within a real life (non-laboratory) application, and the availability of datasets with 100% yield allowing for an accurate baseline.

To evaluate a feature, selected datasets were compressed using L-SIP configured with thresholds of 0.5 °C for air temperature and 2% for relative humidity. A dEWMA (see Section 3.4) filtered the sensed signal to estimate the state using the parameter values $\alpha = 0.2$ and $\beta = 0.2$. To reconstruct the sensed signal, the suppressed values between state estimates transmitted by L-SIP were derived through linear interpolation.

The following measures are used to evaluate the features of G-DPS:

1. **Transmission reduction (or compression ratio)**—The percentage of suppressed transmissions.
2. **Reconstructed signal accuracy**—The accuracy of the reconstructed signal compared to the sensed signal measured by Root Mean Squared Error (RMSE).
3. **Transmission yield**—The percentage of successful node transmissions.

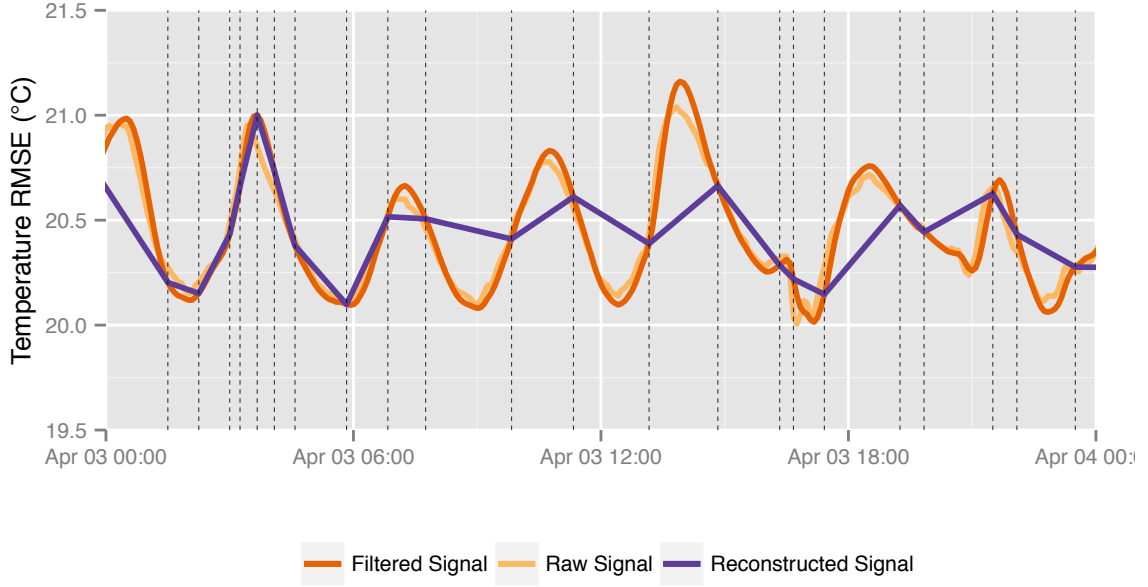


Figure 3.1: Example L-SIP ($\alpha = 0.2$, $\beta = 0.2$, $\varepsilon = 0.5$ °C) output for temperature of a House 1's (See Appendix B) dining room over 24 hours. Changes of temperatures are due to occupants trying to maintain a temperature using a radiator thermostat. Transmissions are indicated by the dotted vertical lines.

4. **Reconstruction data yield**—The percentage of data points that can be reconstructed based on the successful transmissions.
5. **Node energy annual requirement**—This measures the annual energy requirement of a TelosB node using the microbenchmarking approach. It is calculated based on the basis of the nodes using B-CTP as described in Section 3.6 on page 71. B-CTP is a modification to CTP to use a persistent powered backbone network. B-CTP is shown to significantly reduce the energy requirement of a leaf node compared to CTP.

3.5.2 Example L-SIP algorithm output

To explain the terms used in this evaluation, this section examines the output of the L-SIP algorithm over an example day of data.

Figure 3.1 shows the output of L-SIP for a node deployed in a dining room considering a single day's output from the year long dataset. The following terms are used to describe this output:

Raw signal represents the values reported by the sensor. In this case 288 samples were taken during the 24 hour period.

Table 3.2: Example L-SIP performance metrics, for a temperature signal compressed with L-SIP

Packets transmitted	Transmission reduction (ratio)	Reconstruction RMSE	Transmission yield	Reconstruction yield	Estimated energy requirement
19	93.4% ($\times 15$)	0.14 °C	100%	100%	87 mAh/year

Filtered signal is the filtered signal from the dEWMA filter estimating the new state.

Reconstructed signal is the reconstruction of the filtered signal, from the transmitted state updates from the node. The reconstructed signal is compared to the filtered signal using RMSE. In this example, the signal is reconstructed with an RMSE error of 0.14 °C compared to original raw signal.

Transmission represent points where the node transmits a state update to the sink. In this example, a total of 19 state updates are transmitted, a transmission reduction of 93.4% ($15\times$ compression ratio in terms of number of packets). No transmissions are lost and therefore transmission yield and reconstructed data yield are both 100%.

The performance metrics for the example in this section are summarised in Table 3.2.

3.5.3 Multiple sensor compression

G-DPS allows multiple sensor's readings to be combined into a single model, termed a multi-modal approach. Each transmission contains a model state for each monitored sensor. Single-modal is an alternative where there are multiple instances of the DPS algorithm and individual sensor model states are transmitted when there is an event. This section answers **RQ1A** by evaluating the multi-modal approach compared to a single-modal approach, considering the number of packet transmissions, signal reconstruction accuracy, and the energy requirement of the node. It is hypothesised that:

H3.1: *A multi-modal approach (in which all model states are transmitted) will significantly reduce the number of transmissions compared to treating the sensors individually in a single-modal approach.*

H3.2: *A multi-modal approach (in which all model states are transmitted) will maintain or improve the accuracy of the reconstructed signals when compared to a single-modal approach.*

To evaluate the multi-modal approach both single-modal and multi-modal L-SIP were used to compress datasets in accordance with the methodology laid out in Section 3.5.1 on page 54. Initially the approach was tested with a node sensing temperature and humidity. A further evaluation with CO₂ was also per-

Table 3.3: Comparison of multi-modal and single-modal approaches for a node sensing temperature and humidity performance statistics. The results show multi-modal improves performance in all measures. Results are given with \pm a standard deviation it is used to show the standard deviation. This format will be used in subsequent tables in this thesis.

Dataset Period		Transmissions \pm s.d	State updates	Temperature RMSE ($^{\circ}$ C)	Humidity RMSE (%)	Estimated energy requirement (mAh)
Two weeks	Single	260 ± 20	280 ± 20	0.19 ± 0.006	0.82 ± 0.02	5.9 ± 0.2
	Multi	210 ± 20	420 ± 30	0.11 ± 0.005	0.62 ± 0.02	5.4 ± 0.2
	Change	-19%	50%	-42%	-24%	-8.4%
One month	Single	480 ± 40	520 ± 40	0.19 ± 0.007	0.8 ± 0.03	12 ± 0.4
	Multi	400 ± 30	790 ± 60	0.11 ± 0.006	0.61 ± 0.03	11 ± 0.3
	Change	-17%	52%	-42%	-24%	-8.3%
Six months	Single	3300 ± 600	3600 ± 600	0.18 ± 0.02	0.79 ± 0.08	75 ± 6
	Multi	2700 ± 400	5300 ± 900	0.1 ± 0.01	0.64 ± 0.07	68 ± 5
	Change	-18%	47%	-44%	-19%	-9.3%
One year	Single	7700 ± 2000	8000 ± 2000	0.22 ± 0.006	0.9 ± 0.04	160 ± 20
	Multi	6100 ± 1000	12000 ± 2000	0.14 ± 0.01	0.62 ± 0.05	140 ± 10
	Change	-21%	50%	-36%	-31%	-12.5%

formed. The approaches of single-modal and multi-modal L-SIP were compared in terms of reconstruction accuracy, the number of transmissions and the energy requirement of a node.

Table 3.3 shows the results of using both a single-modal and multi-modal approach with L-SIP and the factor improvement of the multi-modal approach. Over each measure and all dataset periods the multi-modal approach improves the performance of L-SIP. Overall, the multi-modal approach reduces transmissions, increases state updates which, in turn, reduces signal reconstruction error, and reduces the energy requirement of a node compared to single-modal.

Table 3.4: Comparison of multi-modal and single-modal approaches on a node sensing temperature, humidity, and CO₂. The results show that the multi-modal approach improves performance in all measures. Due to occupants tending to switch CO₂ nodes off, only dataset periods of two-weeks and one-month are considered

Dataset Period		Transmissions	State updates	Temperature RMSE ($^{\circ}$ C)	Humidity RMSE (%)	CO ₂ RMSE (ppm)	Estimated energy requirement (mAh)
Two weeks	Single	410 ± 50	450 ± 60	0.22 ± 0.01	0.89 ± 0.04	46 ± 2	7.4 ± 0.5
	Multi	300 ± 30	590 ± 70	0.083 ± 0.007	0.36 ± 0.03	35 ± 2	6.2 ± 0.4
	Change	-27%	31%	-62%	-59%	-24%	-16%
One month	Single	750 ± 80	810 ± 80	0.22 ± 0.01	0.86 ± 0.05	45 ± 2	15 ± 0.8
	Multi	560 ± 50	1100 ± 100	0.078 ± 0.007	0.35 ± 0.03	35 ± 2	13 ± 0.6
	Change	-25%	36%	-65%	-60%	-22%	-13%

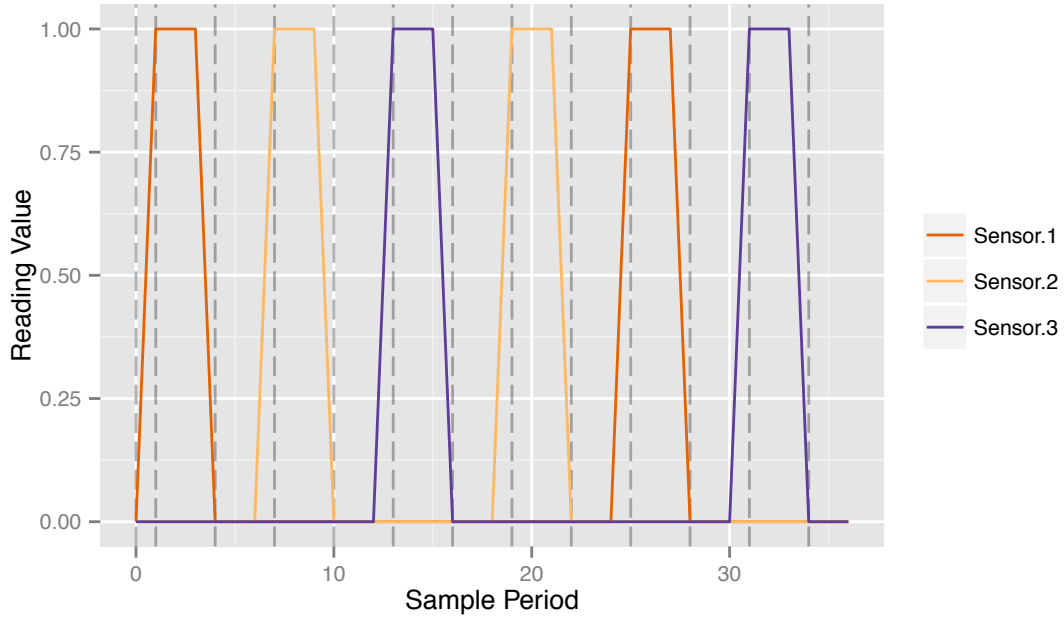


Figure 3.2: Example of a signal where multi-modal would provide no overall benefit compared to single-modal.

The same evaluation was repeated with 80 nodes from the two week datasets, and 60 nodes from the one-month datasets. The selected nodes all sense temperature, relative humidity, and CO₂. The parameters used for CO₂ were threshold $\varepsilon_c = 100\text{ppm}$, and filter parameters $\alpha = 0.2$, $\beta = 0.2$. Since the CO₂ nodes used were prone to be switched off by residents, only dataset periods of two-weeks and one-month were available. As in the previous case Table 3.4 shows that the multi-modal approach reduces transmissions, signal reconstruction error and the energy requirement of a node. Since CO₂ requires more transmissions than either of the other parameters, the inclusion of CO₂ improves the signal reconstruction accuracy of temperature and humidity by a larger factor.

Table 3.4 shows **H3.1** and **H3.2** to be true—when multi-modal is used with a node sensing temperature and humidity and CO₂ transmissions are reduced by up to 27%, signal reconstruction accuracy is improved by up to 65%, and the energy requirement of nodes is reduced by 15% compared to single-modal. However, if considering a phenomena where events for different modalities have no relationship, then multi-modal will provide no benefit over single-modal. For example, see Figure 3.2—in this case both a single-modal and multi-modal approach would transmit at the same time. Therefore multi-modal will not reduce transmissions but will still need to send a larger packet.

To answer **RQ1A**—Yes, when signals are likely to change at the same time (e.g, temperature and humidity) combining multiple sensor modalities into a single model allows for a greater reduction in the

number of packet transmissions and improves signal reconstruction accuracy compared to compressing each stream individually.

3.5.4 Node health check—Heartbeats

This section evaluates the effect of heartbeat messages upon the transmission, energy, and reconstruction performance of G-DPS, and answers **RQ1B**. The evaluation examines i) whether or not use of a heartbeat significantly increases the number of transmissions for a correctly functioning system, and ii) whether the model state should be included in the heartbeat message.

Three cases may occur when a heartbeat packet is transmitted:

1. The node is functioning, however, the state model is able to predict accurately and therefore no state update is required. The heartbeat message indicates the node is still functioning.
2. The heartbeat message fails to transmit. The node will continue to transmit on subsequent sample periods until a successful transmission and acknowledgement. From the sink's point of view no heartbeat message has been received and therefore the node is experiencing problems
3. A transmission has failed sometime between the last transmission to the sink t_{sink} and when a heartbeat should be triggered $t_{\text{sink}} + t_{\text{heartbeat}}$. In this case, a node will continue to transmit on subsequent sample periods until an acknowledgement is received. The sink is only aware that a heartbeat has not been received and therefore the node is experiencing problems.

Only the first case is considered in this evaluation of heartbeat messages, as the heartbeat messages are considered a “still alive” message. The other two cases primarily involve the end-to-end acknowledgements protocol and are discussed in Section 3.5.5. In all failure cases, the sink does not receive a message and thus assumes that failure has occurred. In essence, the absence of a heartbeat message is important for the sink to indicate node failure. The goal of the evaluation is to demonstrate that heartbeats have minimal impact on a functioning node.

The heartbeat approach is evaluated considering a lossless network. Multi-modal L-SIP was first used to compress the datasets as described in the evaluation method in Section 3.5.1. The heartbeat transmitted is of the same form as a state update with the exception that a flag is set to indicate the packet is a heartbeat.

In an ideal situation, a heartbeat period should be selected that produces a minimal number of extra packets but allows the detection of a failed node in a timely manner. Section 3.5.3 shows that multi-modal L-SIP transmits approximately 5% of the total number of samples. Therefore, a packet can be

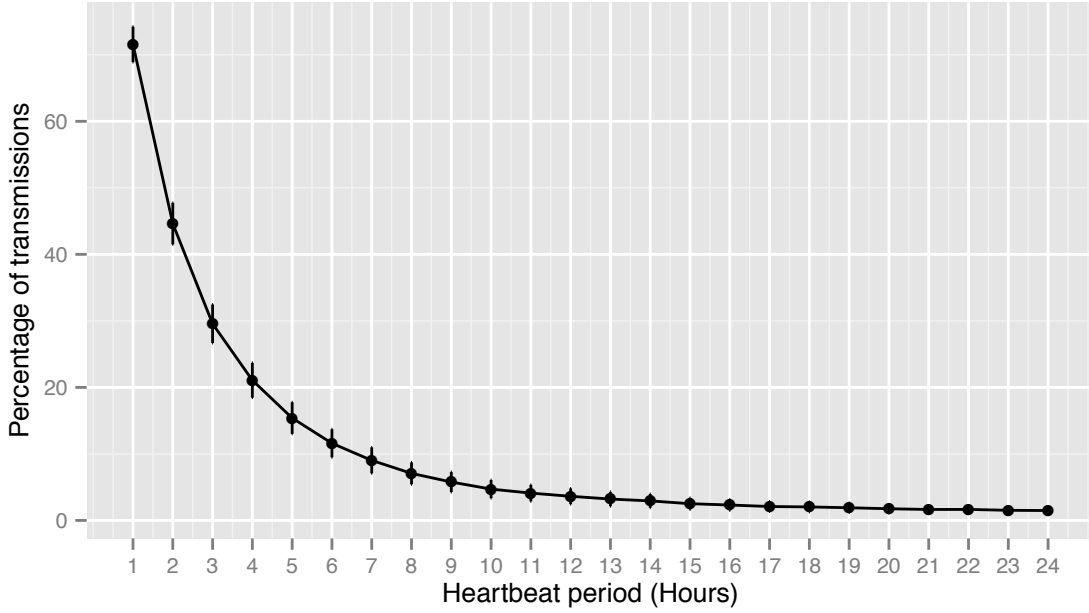


Figure 3.3: Graph showing percentage of transmissions which are heartbeat messages for a given heartbeat period.

expected to be received on average once every 1.7 hours given a five minute sampling period. Therefore it would be expected that 1.7 hours would provide the lower limit of a heartbeat period, setting below this would cause heartbeat transmissions to become a significant percentage of all transmissions sent. From personal experience of deploying these systems, the state of node might get checked at most once a day through an email status update. Furthermore, when implementing L-SIP it is envisaged deployments will be long-term and therefore missing a day of data is not informationally expensive. Based on this experience of deploying and managing WSNs, the remainder of this evaluation will use a heartbeat period of $t_{heartbeat} = 12$ hours. This evaluation tests the following hypotheses.

H3.3 *The use of heartbeats will not increase the number of transmissions of a functioning node in a lossless network when considering application requirements (i.e., the allowable data loss for a system).*

Figure 3.3 shows that even when setting at this lower limit 40% of packet transmissions are that for heartbeats. Therefore the heartbeat period should be carefully considered to minimise data loss and and transmissions. Figure 3.3 also shows that where $t_{heartbeat} = 12$ hoursonly 3.6% of transmissions are that for heartbeat, equating to an extra 3 transmission in a 2 week period.

. L-SIP was evaluated over all datasets with a heartbeat of 12 hours, and compared against using

no heartbeat. These experiments show that in all cases, other than the one month duration, the use of heartbeats do not significantly increase the number of transmissions or reduce the reconstruction accuracy when considering lossless transmissions.

Therefore **H3.3** is true when considering application requirements, balancing possible data loss and the expected number of transmissions the use of a heartbeat will not significantly increase the number of transmissions, or significantly increase the energy consumption of a node. In the literature review it was shown that a number of existing heartbeat approaches include the state information in the heartbeat message. However in a lossless system, such as presented here, including the state in the heartbeat message will not improve reconstruction accuracy since no additional transmissions occur. Therefore, in response to **RQ1B**—Yes, heartbeat messages can detect node failure within a user specified time period, without significantly impacting the energy requirement of a functioning node.

In the case of a lossy network, heartbeats will have much less impact on the node than the need to retransmit. This is evaluated in the next section in the context of acknowledgements.

3.5.5 End-to-end Acknowledgements

This section evaluates the use of end-to-end acknowledgements as a part of the G-DPS framework, and answers **RQ1C**. The evaluation of end-to-end acknowledgements is evaluated in terms of two networks:

A lossless network A WSN in which no transmissions are lost.

A lossy network A WSN in which transmissions fail. In a deployed WSN it is likely that some transmissions fail. As previously discussed in Chapter 2, several deployments achieved transmission yields as low as 35%.

3.5.5.1 Transmission Model

Lossy networks can be modelled by a transmission model that determines the likelihood of transmission success. In the evaluation here, this is based on a two-state model:

1. The previous transmission succeeded,
2. the previous transmission failed.

These two cases have been considered since the success of transmissions is, generally, related to the the last transmissions state. For example, a node that has transmitted successfully on the last sampled period is likely to succeed in the next transmissions, and the reverse for failure. This allows simulation of a larger number of transmission failure types, compared to a randomised failure model.

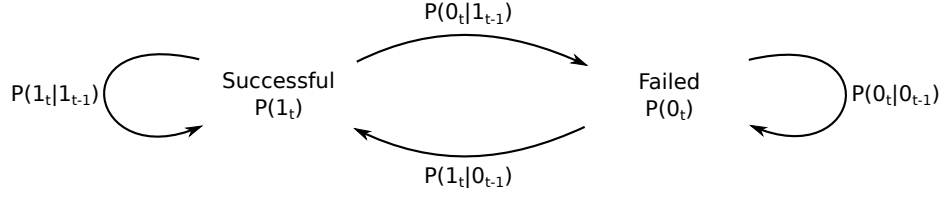


Figure 3.4: Transmission model

Figure 3.4 illustrates the transmission model. The probability space is described in terms of:

$P(1_t|1_{t-1})$ The probability of transmission succeeding given that the previous transmission was successful.

$P(0_t|0_{t-1})$ The probability of transmission failing given that the previous transmission failed.

$P(1_t|0_{t-1})$ and $P(0_t|1_{t-1})$ were not explicitly calculated, but can be determined based on the other two probabilities. To evaluate the whole range of the probability space state, combinations of $P(1_t|1_{t-1})$ and $P(0_t|0_{t-1})$ were selected at 5% intervals. The evaluation in this section is presented in terms of the resulting yields.

The state of the model is updated on each time-step rather than each transmission attempt. The reason for this is that transmissions will be affected by the radio environment which is independent of the need to send a L-SIP state update.

3.5.5.2 Effect of end-to-end acknowledgements on a lossy network

This section evaluates the effect of acknowledgements on transmission reduction, reconstruction accuracy, data yield, and power consumption when a network is lossy.

To evaluate the effect of acknowledgements in the case of transmission failure, one year³ of data from nine sensors with 99.98%⁴ yield (9 traces, 3285 trace-days) was compressed using L-SIP with the parameters defined in Section 3.5.1 ($\alpha = 0.2$, $\beta = 0.2$, $\varepsilon_{temp} = 0.5$ °C, $\varepsilon_{hum} = 2\%$). The traces were required to be of a long duration and to have close to 100% yield such that transmission failure could be meaningfully simulated with a variety of durations. When a transmission is required, the transmission model determines whether the transmission is successful or not. In addition to L-SIP the transmission model was also applied to a sense-and-send approach. The results shown in this section are for one dataset (house 1, master bedroom) which is a representation of the overall performance of the tested datasets.

³Actually 360 days but for brevity will be stated as a year. Taken from House 1 described in Appendix B

⁴Missing values (average 1200 per sensor) are imputed via linear interpolation to avoid discontinuities: the majority of data loss instances were infrequent failures of short duration (e.g, one missing sample)

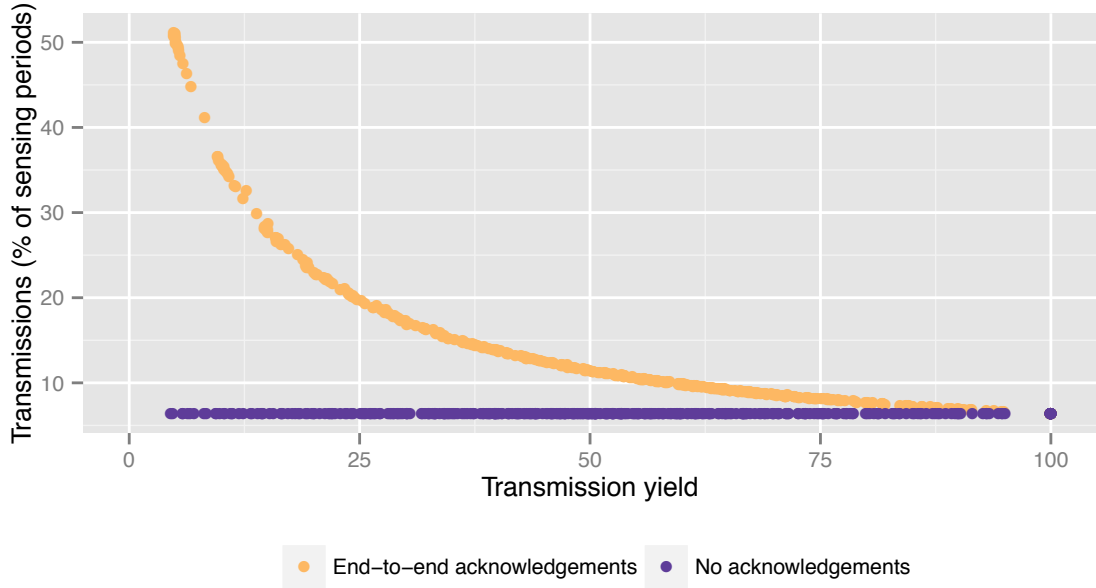


Figure 3.5: The effect of transmission failure on the number of required transmissions for multi-modal L-SIP with and without end-to-end acknowledgements.

When using end-to-end acknowledgements the number of required transmissions will increase as the number of successful transmissions decrease. This ensures that an up-to-date state message is sent as soon as possible, which minimises the error of the reconstructed signal. It is hypothesised:

H3.4 *End-to-end acknowledgements will increase the number of transmissions as failure rates increase (for a fixed sampling frequency). However, they will result in an improvement to the accuracy of the reconstructed signal compared with using no acknowledgements.*

Figure 3.5 shows that the number of transmissions increase as the transmission yield decreases. For example, 20% of state updates are transmitted when the transmission yield is 25% but over 50% transmission are required when the transmission yield is approximately 5%. In the no acknowledgement approach the number of transmissions do not change as the node is unaware of failure and does not attempt to retransmit.

Figure 3.6 shows that with no acknowledgements the reconstruction RMSE increases as the transmission yield decreases. However, when using end-to-end acknowledgements the reconstruction RMSE degrades at a much slower rate, due to state updates being received earlier compared to using no acknowledgements. Considering a system where yield is less than 35% (in line with Anastasi *et al.* [8], Arora

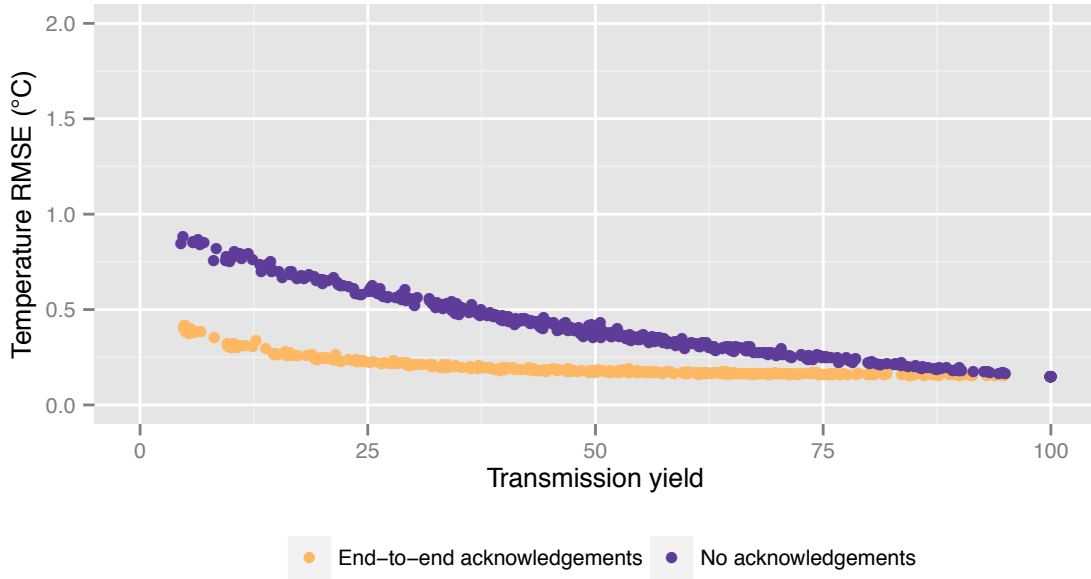


Figure 3.6: The effect of transmission failure on the accuracy of the reconstructed temperature signal for multi-modal L-SIP with and without end-to-end acknowledgements.

et al. [10], and the Intel Lab Data [70]) the use of end-to-end acknowledgements approximately halves the RMSE. Furthermore, when 65% of packets fail (35% transmission yield) the reconstruction RMSE exceeds the set error threshold when using no acknowledgements, whereas the threshold is never exceeded when using acknowledgements. This same effect has been found with both relative humidity and CO₂ signals.

Figure 3.5 and Figure 3.6 show **H3.4** to be true: acknowledgements cause transmissions to increase as failure rates increase, however, this results in improved signal reconstruction accuracy compared to using no acknowledgements

Since using end-to-end acknowledgements increases the number of transmissions in a lossy network it can be expected that the energy requirement of a node is increased. It is hypothesised:

H3.5 *End-to-end acknowledgements have an increased energy requirement compared to using no acknowledgements as transmission failure increases, however both approaches still show a significant decrease in energy consumption compared to sense-and-send.*

Figure 3.7 shows that the node energy requirement, calculated using microbenchmarking, increases as the transmission yield decreases (based on results in Figure 3.5). This is due to the increased number of transmissions required. In the no acknowledgement approach the energy requirement does not change

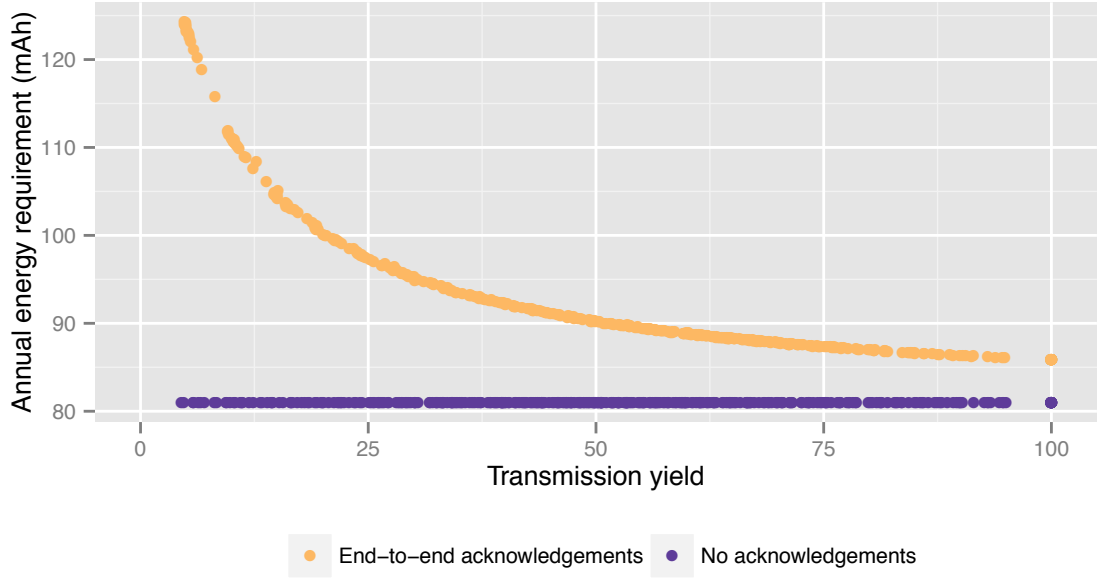


Figure 3.7: The effect of transmission failure on the annual energy requirement of a L-SIP node with and without end-to-end acknowledgements. A nominal radio duty cycle of 0.05% (160 ms) with acknowledgements and 0.005% (16 ms) without acknowledgements is assumed.

as the node is unaware of failure and does not attempt to retransmit, therefore no additional energy is required. Considering the case where yield is approximately 35%, when using acknowledgements (90 mAh) the node requires a factor of $1.13\times$ more energy than when using no acknowledgements (80 mAh). In the worse case where transmission yield is in the region of 5%, L-SIP with acknowledgements uses a factor of $1.5\times$ more energy than no acknowledgements (120 mAh vs 80 mAh). However, this is still an energy decrease by a factor $10\times$ compared to CTP sense-and-send (1200 mAh, see Table 3.6 on page 70) and a factor $1.4\times$ compared to B-CTP sense-and-send (170 mAh). The small decrease compared to B-CTP sense-and-send is due to the idle state now accounting for 90% of the energy. If the idle energy requirement is disregarded, assuming future hardware improvements, the energy decrease would be a factor of $6\times$.

Figure 3.7 shows **H3.5** to be true: the energy requirement for a L-SIP node increases as transmission failure increases when using acknowledgements, however in the worse case this energy requirement is still $10\times$ less than CTP sense-and-send.

When transmission failure occurs end-to-end acknowledgements coupled with sequence numbers can identify which data points cannot be reconstructed accurately. However, when not using acknowledgements it is not possible to determine where the failure occurred and therefore all data points between the

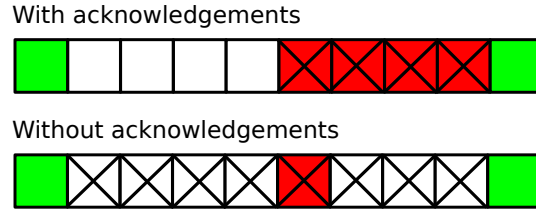


Figure 3.8: Demonstration of failed transmissions in acknowledgement and acknowledgement-less cases. The red boxes indicate transmission failure. The boxes with crosses in are data points which must be discarded, as they cannot be guaranteed to be within the error threshold.

last successful and the next successful packet following failure must be discarded. Since the failed state update could have occurred on the sensing cycle following the previous successful update, no points can be guaranteed to be within the error threshold. Figure 3.8 demonstrates this issue.

Since the use of acknowledgements allows the detection of points where the accuracy of the data cannot be guaranteed, and the low transmissions rates of DPS reduce the chance of transmissions failing it can be hypothesised:

H3.6 *Using end-to-end acknowledgements coupled with sequence numbers improves the number of data points that can be reconstructed compared to using no acknowledgements.*

H3.7 *Using end-to-end acknowledgements and sequence numbers with L-SIP results in less susceptibility to data loss compared to sense-and-send.*

Figure 3.9 shows the reconstructed data yield, as calculated using the method in Section 3.3.3, versus the transmissions yield for L-SIP with and without acknowledgements and for sense-and-send.

Since sense-and-send has a one-to-one relationship between transmissions and data points, the data yield is equal to the transmission yield. L-SIP without acknowledgements does not perform favourably compared to the other two options. The issue with using no acknowledgements is the one-to-many relationship—if a transmission fails it must be assumed that all data between the next received state update and the previous potentially exceeds the defined error threshold and thus must be discarded. However, when using end-to-end acknowledgements the data yield is still high even when transmission yield is low. Considering the case of a node achieving a 35% yield using no acknowledgements results in a data reconstruction yield of 12%, however, when using acknowledgements 85% of data is still recoverable—this is an improvement by a factor of $7\times$. This approach of using L-SIP with acknowledgements can be seen to have two advantages:

1. The infrequent transmissions mean L-SIP may simply not need to transmit during periods of radio

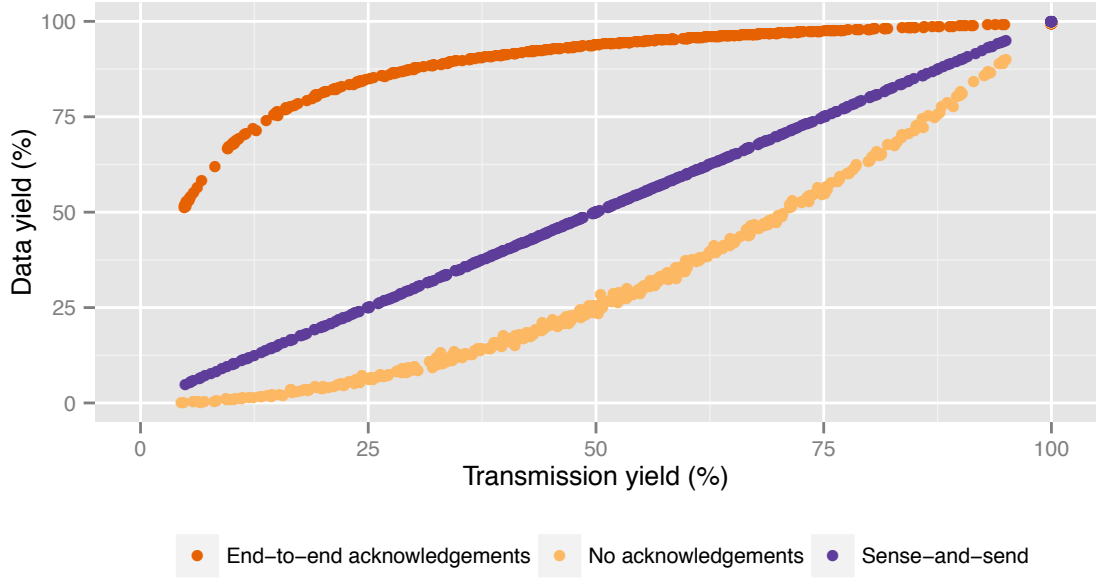


Figure 3.9: The effect of transmission failure on the reconstructed data yield for multi-modal L-SIP with and without end-to-end acknowledgements and sense-and-send.

interruption.

2. When transmission failure occurs L-SIP with acknowledgements uses a pseudo sense-and-send mode, therefore only loses as many packets as sense-and-send would during that period.

Figure 3.9 shows **H3.6** and **H3.7** to be true: using acknowledgements means L-SIP is less susceptible to data loss and, when coupled with sequence numbers, can improve the amount of data that can be accurately reconstructed.

3.5.5.3 Effect of end-to-end acknowledgements on a lossless network

In a lossless network, the difference between using and not using acknowledgements is the radio duty time. When using acknowledgements it is expected the radio will be on for a longer period. To evaluate the effect of acknowledgements on the radio duty cycle, two nodes were programmed using L-SIP and B-CTP with one node using acknowledgements and the other without. Both nodes made use of multi-modal sensing, heartbeats, and sequence numbers. These nodes were deployed next to each other in an office environment for a week with a sample period of one minute. The error threshold ε for both temperature and humidity was set to -1 to force a transmission every sample period.

When using acknowledgements, a node has a median radio duty cycle of $0.05\% \pm 0.006$ (160 ms), which

is an increase of $10\times$ compared to using no acknowledgements, which gave $0.005\% \pm 0.0002$ (16 ms). The increase in time for acknowledgements results from the time taken for the sink to receive a packet, decode and store the data, and send the resulting acknowledgement. As previously shown in Table 3.9, using the microbenchmarking approach a node uses 87 mAh/year when using acknowledgements using the duty cycle calculated here. However, when considering the median duty cycle for no acknowledgements the node requires 83 mAh/year, a decrease of only a factor of $1.05\times$.

3.5.5.4 End-to-end acknowledgements summary

This section has shown that end-to-end acknowledgements slightly increase the node energy requirement. However, this increase is offset by an increase in data yield, and improved signal reconstruction accuracy. It should be noted, however, that there are a few system design considerations when implementing acknowledgements:

1. Acknowledgements should not be used when the round trip time of sending a packet and receiving an acknowledgement is greater than the defined sample period.
2. In applications where the frequency of transmissions is high, end-to-end acknowledgements will have a higher energy requirement due to increased radio duty cycle.

This section has shown that the use of acknowledgements in a lossy network can bring several improvements. Assuming a node achieving a 35% yield, as reported in the literature, the use of acknowledgements improves signal reconstruction accuracy by a factor of $2\times$ and increases the data yield of the system up to a factor of $7\times$, when compared to acknowledgement-less L-SIP. In a lossless system, acknowledgements only increase a node's annual energy requirement by a factor of $1.05\times$. In a lossy network acknowledgements only increase a node's annual energy requirement by a factor of $1.16\times$ when the node's yield is 35%. In the worst case of 5% yield, L-SIP with acknowledgements uses a factor of $1.5\times$ more energy than no acknowledgements.

Therefore to answer **RQ1C**—Yes, end-to-end acknowledgements increase reconstructed data yields compared to an acknowledgement-less schemes.

3.5.6 On-node evaluation of L-SIP

Since simulation approaches do not necessarily show the real performance of algorithms, this section evaluates two deployments of nodes implementing L-SIP. The first deployment (D1) was a four bedroom

Table 3.5: Summary of L-SIP deployments

	Transmission reduction	Transmission yield	Reconstruction yield
D1	85% \pm 4%	74% \pm 5%	95% \pm 1%
D2	97% \pm 2%	77% \pm 19%	98.9% \pm 2%

detached home⁵, gathering data for 38 weeks. The second deployment comprised of an unoccupied home (D2) for a period of 3 months. The nodes in these deployments were all programmed with L-SIP, using the same parameters as defined in section 3.5.1 on page 54 ($\alpha = 0.2$, $\beta = 0.2$, $\varepsilon_{temp} = 0.5^\circ\text{C}$, $\varepsilon_{hum} = 2\%$).

The evaluation of on-node L-SIP is based on four measures of performance: i) transmission reduction, ii) transmission yield, iii) Reconstructed data yield, and iv) battery consumption.

Table 3.5 shows a summary of the evaluated L-SIP deployments. Deployment D1 achieved an average transmission reduction per node of 85% \pm 4% and an average per node transmission yield of 74% \pm 5% (the transmission yield was affected by sink failure in week 10 of the deployment). From these successful transmissions an average of 95% \pm 1% of the data per node could be accurately reconstructed within the defined error thresholds of 0.5 $^\circ\text{C}$ for temperature and 2% for relative humidity—using the data yield calculation presented in Section 3.3.3.

In D2 the conditions in the unoccupied home were much more stable than for D1 with a per node average transmission reduction of 97% \pm 2%. The transmission yield was marginally better with an average of 77% \pm 19% per node, and the average per node reconstruction yield was 98.9% \pm 2%.

Figure 3.10 shows the battery discharge curve of a sense-and-send node, and node implementing L-SIP. It is shown that the battery discharges significantly faster when using sense-and-send with CTP and Low Power Listening (LPL) compared to using L-SIP and B-CTP. This demonstrates the significantly longer lifetime achievable with L-SIP.

This section shows that, when deployed, L-SIP performs as expected—reducing energy usage and the number of required transmissions. However, occupants tended to switch backbone nodes off, which shows that a battery backup or hard-wired power connection is required for AC powered nodes.

⁵House 38 described in Appendix B

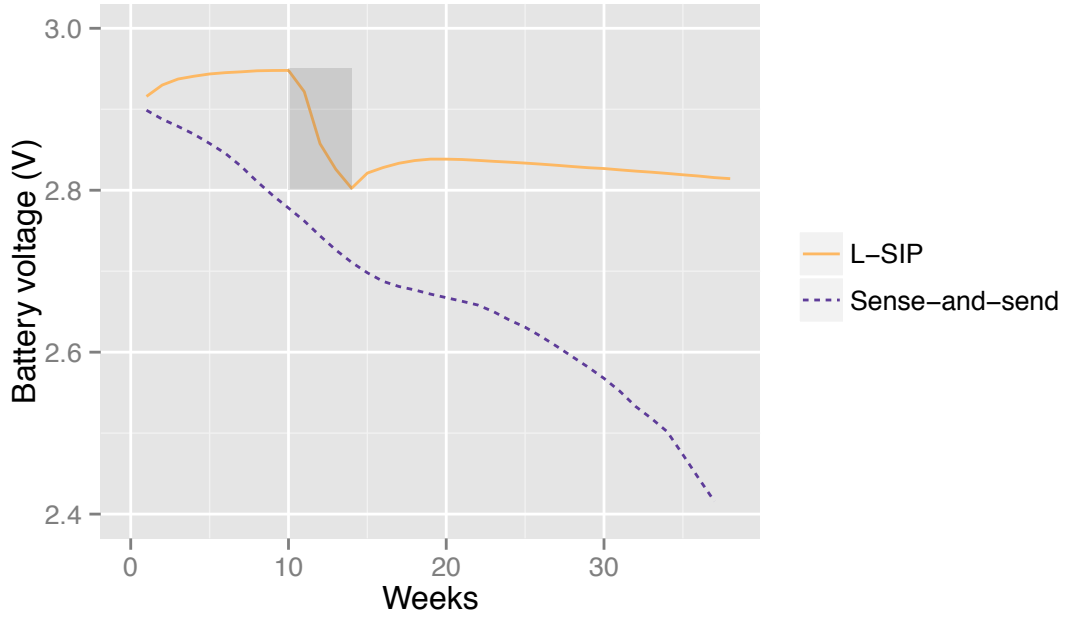


Figure 3.10: Typical battery voltage of a node running CTP with sense-and-send, and a node using L-SIP with B-CTP. The shaded grey area is where sink failure occurred during the deployment.

Table 3.6: Baseline microbenchmark estimates for a sense-and-send approach on TelosB node with a five minute sampling cycle. CTP send time is based on logs from a 200+ node network and include retries.

Process	Annual samples		Time (ms)		mA		mAh/year
Sense	105120	×	295	×	0.458	=	3.9
Processing	105120	×	1	×	0.182	=	0.01
CTP send	105120	×	473	×	18.920	=	260
LPL listen	105120	×	1,500	×	18.920	=	830
Idle	105120	×	297,732	×	0.009	=	78
Totals							1171.9

Table 3.7: Annual energy consumption of a TelosB node implementing L-SIP with CTP and LPL on a five minute sampling cycle. CTP send time is based on logs from a 200+ node network and includes retries.

Process	Annual samples		Time (ms)		mA		mAh/year
Sense	105120	×	295	×	0.458	=	3.9
Processing	105120	×	44	×	0.182	=	0.2
CTP send	5256	×	473	×	18.920	=	13
LPL listen	105120	×	1,500	×	18.920	=	830
Idle	105120	×	297,732	×	0.009	=	78
Totals							925.1

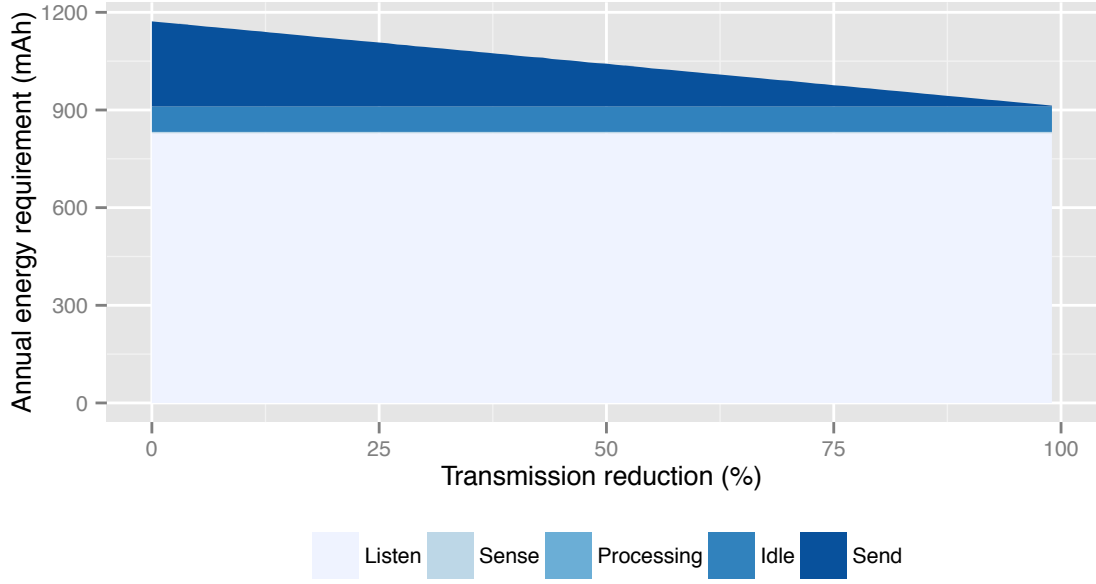


Figure 3.11: Annual energy usage compared to achieved transmission reduction when using CTP and LPL. The primary energy use is for listening.

3.6 The Backbone Collection Tree Protocol for low transmission WSNs

The aim of DPS algorithms is to increase the lifetime of a node. Since L-SIP reduces transmissions by 95% compared to sense-and-send it can be hypothesised:

H3.8: *A node implementing L-SIP will have significantly lower energy requirement than a node implementing sense-and-send using a network stack comprised of CTP and LPL.*

The microbenchmarking method has been used to estimate the annual energy use of TelosB nodes in order to compare the use of L-SIP versus sense-and-send. The TelosB nodes are assumed to make use of the commonly used network stack comprised of CTP and LPL implemented within TinyOS. This experiment assumes a sampling interval of five minutes for both algorithms. Using microbenchmarking the TelosB node energy use for L-SIP was calculated assuming a factor of $20\times$ (or 95%) packet reduction. Table 3.6 shows the results of microbenchmarking for sense-and-send while Table 3.7 shows the results for L-SIP. The annual energy requirement of a TelosB node implementing L-SIP and CTP (925.1 mAh/year) is reduced by a factor of $1.3\times$ compared to a sense-and-send approach (1171.9 mAh/year).

Figure 3.11 shows the annual energy requirement for a node using CTP and LPL over a range of transmission reductions. Since the node has to always listen, sense, and sleep during each sample period the energy requirement for these processes are constant regardless of the number of transmissions that are suppressed. It is clear to see from Figure 3.11 that the energy requirement for listening is greater than any other process. When transmissions are significantly reduced, listening accounts for over 90% of the annual energy requirement. Therefore, in addition to reducing the number of required transmissions the energy requirement of the MAC layer must be considered. Raza *et al.* [93] evaluated a similar algorithm and also concluded the MAC layer is a major consideration.

Figure 3.11 shows **H3.8** to be false. When transmissions have been substantially decreased, focus must be on reducing the energy overhead of networking. The largest energy consumer is the radio listening for packets in LPL. Therefore, removing the need to listen where possible will increase node lifetime.

3.6.1 The Backbone Collection Tree Protocol

In a tree-based multi-hop network, such as CTP, nodes act as either leaf nodes or routing nodes. These classifications are typically assigned when the network is formed based on factors such as signal strength. The routing nodes are required to have their radio turned on for longer periods to perform the routing duties. However, if the state of the network changes nodes may change from leaf to routing nodes and vice-versa. Therefore, all nodes are required to listen periodically for incoming packets.

A typical deployment usually consists of more leaf nodes than routing nodes. Since most nodes do not need to be involved in the routing process they do not necessarily need to listen for incoming packets. Furthermore, especially within the built environment, a large percentage of nodes can be powered from the mains electricity or other high capacity energy source. In the case of the deployment, described in Appendix A, 54% of nodes are mains powered. These mains powered nodes can leave their radios on permanently, with no penalty to the lifetime, to form a dense permanent network backbone.

Using the concept of a backbone network, a network infrastructure named Backbone Collection Tree Protocol (B-CTP) is proposed. B-CTP is based on the combination of a backbone of powered nodes with CTP. Figure 3.12 shows this approach. Two types of node are defined:

Leaf nodes are sense-and-send battery powered nodes, which forward their data to a neighbouring backbone node. These leaf nodes only switch their radio on long enough to transmit a packet and receive the subsequent acknowledgement packet.

Backbone nodes are mains powered nodes and have their radios on continuously. Additionally to their

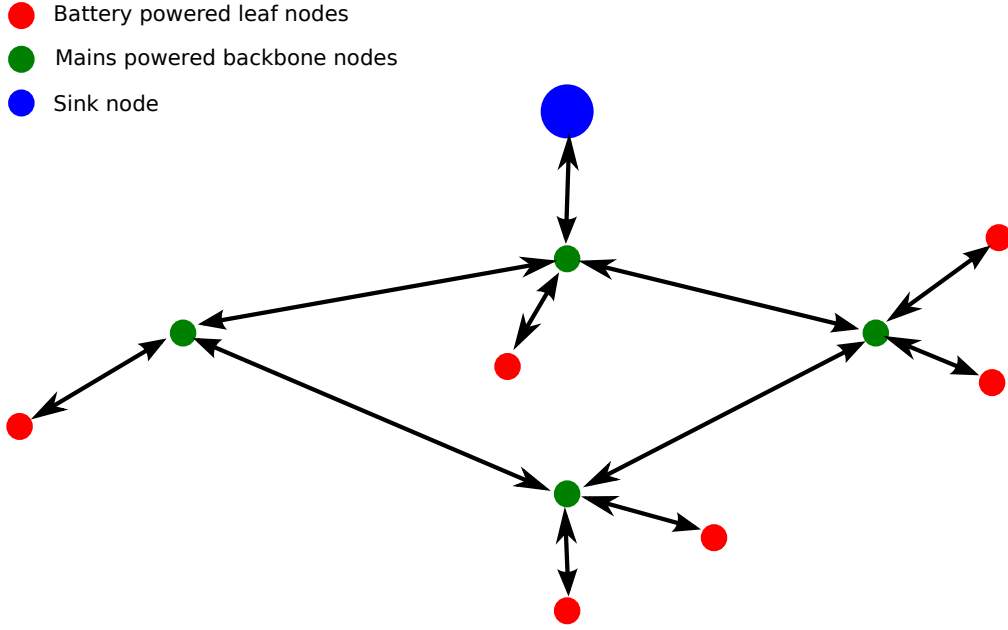


Figure 3.12: Proposed B-CTP network infrastructure: Battery powered leaf nodes send data packets to powered backbone nodes which are responsible for forwarding data to sink nodes.

own sense-and-send cycle, they have the added responsibility of routing packets using CTP, along with associated tasks such as determining link quality.

To implement B-CTP, CTP has been modified to only use backbone nodes for routing. CTP uses a distance vector routing protocol with a four-bit link estimator [29] to calculate the Expected Transmissions (ETX) link quality between nodes. The route cost is the sum of all link costs in a particular route. The four-bit link estimator has been modified for B-CTP to avoid using leaf-to-leaf links and give priority to backbone nodes. When a link quality is requested, and the link is leaf-to-leaf (determined in the current implementation by the node id), the maximum (worst) ETX value is returned to avoid the link being used. This resembles the way collection tree roots advertise a cost of zero ETX. The change means the leaf nodes only need to listen for packets when awaiting acknowledgements. Thus radio is only switched on when a transmission is required, then switched off when either an acknowledgement is received or a predetermined time out time is triggered.

The number of backbone nodes required, as with the deployment of any WSN, depends upon the environment and the application needs. For example, deploying in an open-space would require less than deploying in a construction site with no direct line of site. In the home deployments it was found that conditions such as, Wi-Fi and Zigbee channels overlapping can degrade a signal, therefore possibly

Table 3.8: Annual energy consumption of a TelosB B-CTP leaf node node with a five minute sampling cycle implementing sense-and-send. B-CTP send time is based on experimentation described in Section 3.5.5.

Process	Annual samples		Time (ms)		mA		mAh/year
Sensing	105120	×	295	×	0.458	=	3.9
Processing	105120	×	44	×	0.182	=	0.2
B-CTP send	105120	×	160	×	18.920	=	88.4
Listening	0	×	0	×	0	=	0.0
Idle	105120	×	297,732	×	0.009	=	79.0
Totals							171.5

Table 3.9: Annual energy consumption of a TelosB B-CTP leaf node node with a five minute sampling cycle implementing L-SIP. B-CTP send time is based on experimentation described in Section 3.5.5.

Process	Annual samples		Time (ms)		mA		mAh/year
Sensing	105120	×	295	×	0.458	=	3.9
Processing	105120	×	44	×	0.182	=	0.2
B-CTP send	5256	×	160	×	18.920	=	4.4
Listening	0	×	0	×	0	=	0.0
Idle	105120	×	297,732	×	0.009	=	79.0
Totals							87.5

warranting the need for more back bone nodes. From my experience in deploying nodes implementing this protocol in homes, typically one or two backbone nodes are deployed, depending on the properties size⁶. WSN building deployments are often opportunistic, therefore where permitted a backbone node should be deployed in place of a leaf node.

Table 3.8 shows the annual energy requirement of a TelosB node implementing sense-and-send and B-CTP is 171.5 mAh/year, this a reduction by 6.7× compared to using CTP alone.

Since mains powered nodes are prioritised for routing and leaf nodes can switch their radios off for much longer periods it can be hypothesised that:

H3.9: *B-CTP will significantly lower the energy requirement of a node when coupled with L-SIP, compared to CTP with L-SIP.*

Table 3.9 shows the annual energy requirement of a TelosB node implementing L-SIP and B-CTP is 87.5 mAh/year, assuming a sample rate of 5 minutes and a packet reduction of 95%. When compared to sense-and-send with CTP, the combination of B-CTP and L-SIP reduces the annual energy requirement by a factor of 13.4×.

⁶2 backbones nodes are currently being used in a seven bedroom home which has a footprint of 364.7 m², a single backbone node has been used in homes up to 120 m²

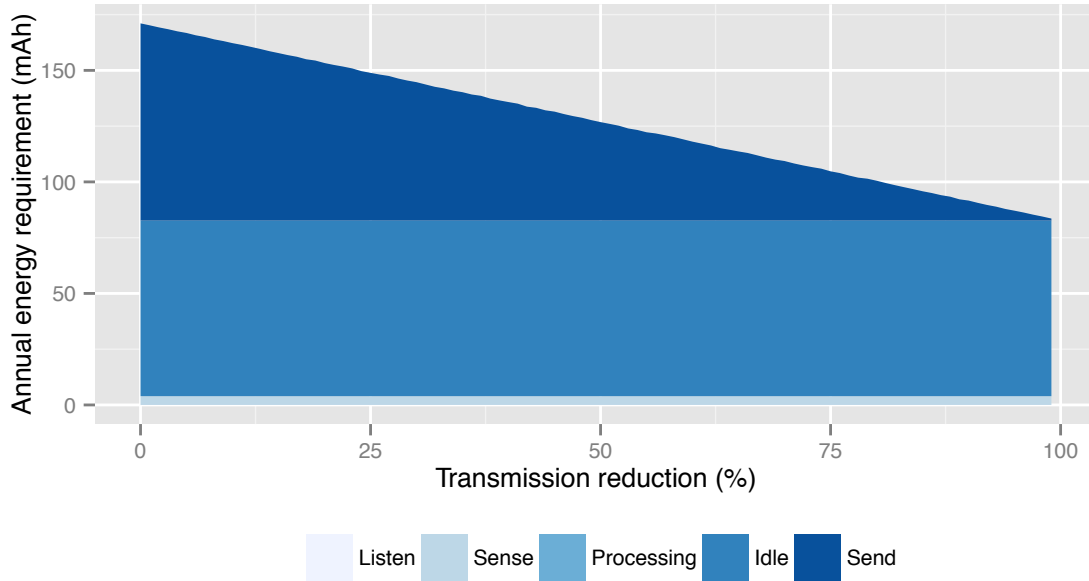


Figure 3.13: Annual energy usage compared to achieved transmission reduction when using B-CTP and LPL.

Table 3.10: Summary of microbenchmark estimates for using sense-and-send and L-SIP on TelosB for the MAC approaches of LPL and B-CTP.

Protocol	Estimated energy consumption (mAh/year)	Energy reduction factor relative to LPL sense-and-send
LPL sense-and-send	1171.9	1.0
LPL L-SIP	925.1	1.3
B-CTP sense-and-send	171.5	6.8
B-CTP L-SIP	87.5	13.4

Figure 3.13 shows the annual energy requirement for a leaf node using B-CTP over a range of transmission reductions. As with CTP, B-CTP still has a constant energy draw for a node being idle, and for sensing. However, there is no longer a requirement for listening—reducing the overall energy requirement considerably. Considering a L-SIP node, with a factor of $20\times$ reduction in transmissions, the node being in an idle state now accounts for 90% of the energy requirement. To reduce this requirement, improvements to the node hardware are required (for example, more efficient voltage conversion or parts with lower power sleep modes).

Table 3.10 summarises the estimated relative improvement in power use of sense-and-send and L-SIP using the two MAC approaches of LPL and B-CTP. With a sense-and-send approach, B-CTP networking can reduce the annual energy requirement by a factor of $6.8\times$, while implementing L-SIP with B-CTP

decreases the annual energy requirement for a leaf node by a factor of $13.4\times$.

Table 3.10 shows **H3.9** to be true: coupling L-SIP with B-CTP to reduce the time the radio is used, for both listening and transmissions, decreases the annual energy requirement for a node by a factor of $13.4\times$. From Tables 3.8 and 3.9 it can be seen that these significant energy reductions are due to the reduced use of the radio (both transmissions, and listening).

In response to **RQ2**—Yes, the lifetime of a WSN node implementing transmission reduction approaches can be increased further by using a persistent backbone network of mains powered routing nodes.

Compared to CTP, however, B-CTP is unable to react to changes in the network, for example, introducing a new backbone node. CTP broadcasts a control beacon with routing information at an adaptive interval. When a node is first switched on this control beacon is sent at regular frequency, however, as the network becomes stable it is eventually backed off to be sent at a maximum of once every 8 minutes. Since leaf nodes have their radios off for long durations, if a new backbone node is introduced to the network it is unlikely a leaf node will detect control beacon. Therefore when introducing a new node in a formed network, leaf nodes will be required to be reset to detect the new backbone node.

This section has therefore shown that B-CTP can significantly reduce the annual energy requirement for a node.

3.7 Summary

This chapter has presented two contributions to knowledge:

1. G-SIP: a novel, generalised framework for the implementation of DPSs.
2. B-CTP—An extension to CTP to utilise a persistent powered backbone network, which reduces the energy requirement for listening in order to extend node lifetime.

3.7.1 Generalised Dual Prediction Scheme summary

The G-DPS framework provides solutions to enable DPS algorithms to function in real life deployments. This will lead to deployments which are more robust to transmission failures, can detect node failures, and support multiple sensing modalities. This results in deployed WSNs being easier to maintain, providing lower maintenance costs, and producing a higher data yield compared to what is currently possible with DPSs in real-life networks. The G-DPS framework includes:

1. A multi-modal approach that allows DPSs to be implemented with multiple sensors. Considering a node sensing temperature, humidity and CO₂, the multi-modal approach transmissions are reduced by up to 27%, signal reconstruction accuracy is improved by up to 65%, and the energy requirement of TelosB nodes is reduced by 15% compared to single-modal DPS.
2. The use of acknowledgements in a lossy network, assuming a node achieving a 35% yield, improves signal reconstruction accuracy by a factor of 2 \times , increases the data yield of the system a factor of 7 \times , and only increases a node's annual energy requirement by a factor of 1.13 \times . when compared to acknowledgement-less L-SIP. In a lossless system, acknowledgements only increase a TelosB node's annual energy requirement by a factor of 1.05 \times .
3. Using a heartbeat period of 12 hours, heartbeat messages are shown to only increase the number of transmissions by a factor of up to 1.02 \times on a functioning node compared to using no heartbeat messages. Heartbeats allow the detection of faulty nodes.
4. The use of sequence numbers in state update transmissions, allows for the calculation of reconstructed data yield.

3.7.2 Backbone Collection Tree Protocol summary

To support the G-DPS framework a network topology to increase the lifetime of sensing nodes was proposed. B-CTP makes use of a persistent powered backbone network to significantly extend node lifetime. The proposed B-CTP coupled with L-SIP was shown to decrease the annual energy requirement for a TelosB node up to a factor of 13.4 \times . This networking approach will allow for much longer lived deployments, reducing the need to change batteries resulting in unobtrusive WSNs.

As shown in Section 2.6 the process of accurately reconstructing the original signal based on the output of DPSs has received little attention. The next chapter investigates techniques to provide an accurate reconstruction of signals from DPS algorithms.

Chapter 4

Spline-based data reconstruction in Wireless Sensor Networks

The previous chapter described the Generalised Dual Prediction Scheme (G-DPS) which was evaluated in terms of an implementation named Linear Spanish Inquisition Protocol (L-SIP). L-SIP is a Dual Prediction Scheme (DPS) algorithm which utilises a linear model that encodes the state as a filtered estimate of the value and rate of change. L-SIP significantly reduces the number of transmissions over sense-and-send. However, selection of the best method to accurately reconstruct the original signal based on the output of DPS algorithms has received little attention in the literature.

The following research question is therefore answered in this chapter: **RQ3:** *Can a spline-based signal reconstruction method improve the accuracy of reconstructed signals compared to piecewise linear methods, for example linear interpolation or model prediction, when using DPS algorithms such as L-SIP?*

This chapter shows that considering the signal gradient and the known bounds on the suppressed samples allows the sensed signal to be more accurately reconstructed. Five reconstruction methods are evaluated—three spline-based methods proposed here, along with model predictions and traditional linear interpolation as baselines (the latter of which does not consider the gradient and bounds). The three spline-based reconstruction methods are:

1. a cubic polynomial spline,
2. a quartic polynomial spline, and
3. a pair of quadratic splines with a discontinuity at the join.

The evaluation of the five reconstruction methods (implemented with L-SIP in a home environment monitoring application) shows that dual quadratic splines provide lower Root Mean Squared Error (RMSE) than the other methods. The chapter also shows that when transmission failures occur, an extension to the dual quadratic spline based method (“adjusted quadratic splines”) allows reconstruction with significantly lower error than would otherwise be possible. Overall, dual quadratic splines are shown to be

the preferable method for reconstructing a signal based on the output of a DPS algorithm.

This chapter is structured as follows: Section 4.1 describes the reconstruction methods considered in this chapter. Section 4.2 provides the results of the reconstruction method analysis, following this Section 4.3 evaluates the best performing method on a lossy network. Section 4.4 investigates the effect of spline-based reconstruction on further analysis of gathered data (using the case study of exposure graphs). Finally, Section 4.5 summarises the results.

4.1 Reconstruction methods

This chapter provides an evaluation of whether spline-based reconstruction techniques can provide higher accuracy than linear interpolation or model prediction. Specifically, it is expected that splines can improve on linear interpolation by incorporating the following information:

1. the estimated gradient of the signal, and
2. the known bounds on predicted data values.

Figures 4.1 and 4.2 show an example of the linear interpolation, predictive model, and dual quadratic spline methods applied to a sample of home air temperature data.

The three spline-based methods have the requirement that the state estimate consists of a timestamp, value, and gradient. Additionally, in common with DPS generally, there is the requirement of no loss of state updates (due, for example, to sensor or network failures), otherwise the reconstruction may not be accurate. Relaxing this lossless network requirement is considered in Section 4.3.

4.1.1 Linear interpolation

Linear interpolation is a simple well-known reconstruction method, using only the data value provided in each update. A disadvantage of linear interpolation is demonstrated in Figure 4.3—a period of suppressed samples followed by a rapid change in signal trend can result in high reconstruction errors as the bounds for the suppressed samples are not considered. This problem is also examined by Silberstein *et al.* [104].

4.1.2 Predictive model

A natural choice for reconstruction when using a DPS algorithm is to use the output of the predictive model implemented at the sink. The reconstruction will therefore follow the predictions from each state update. This has the disadvantage that it introduces discontinuities into the reconstruction as

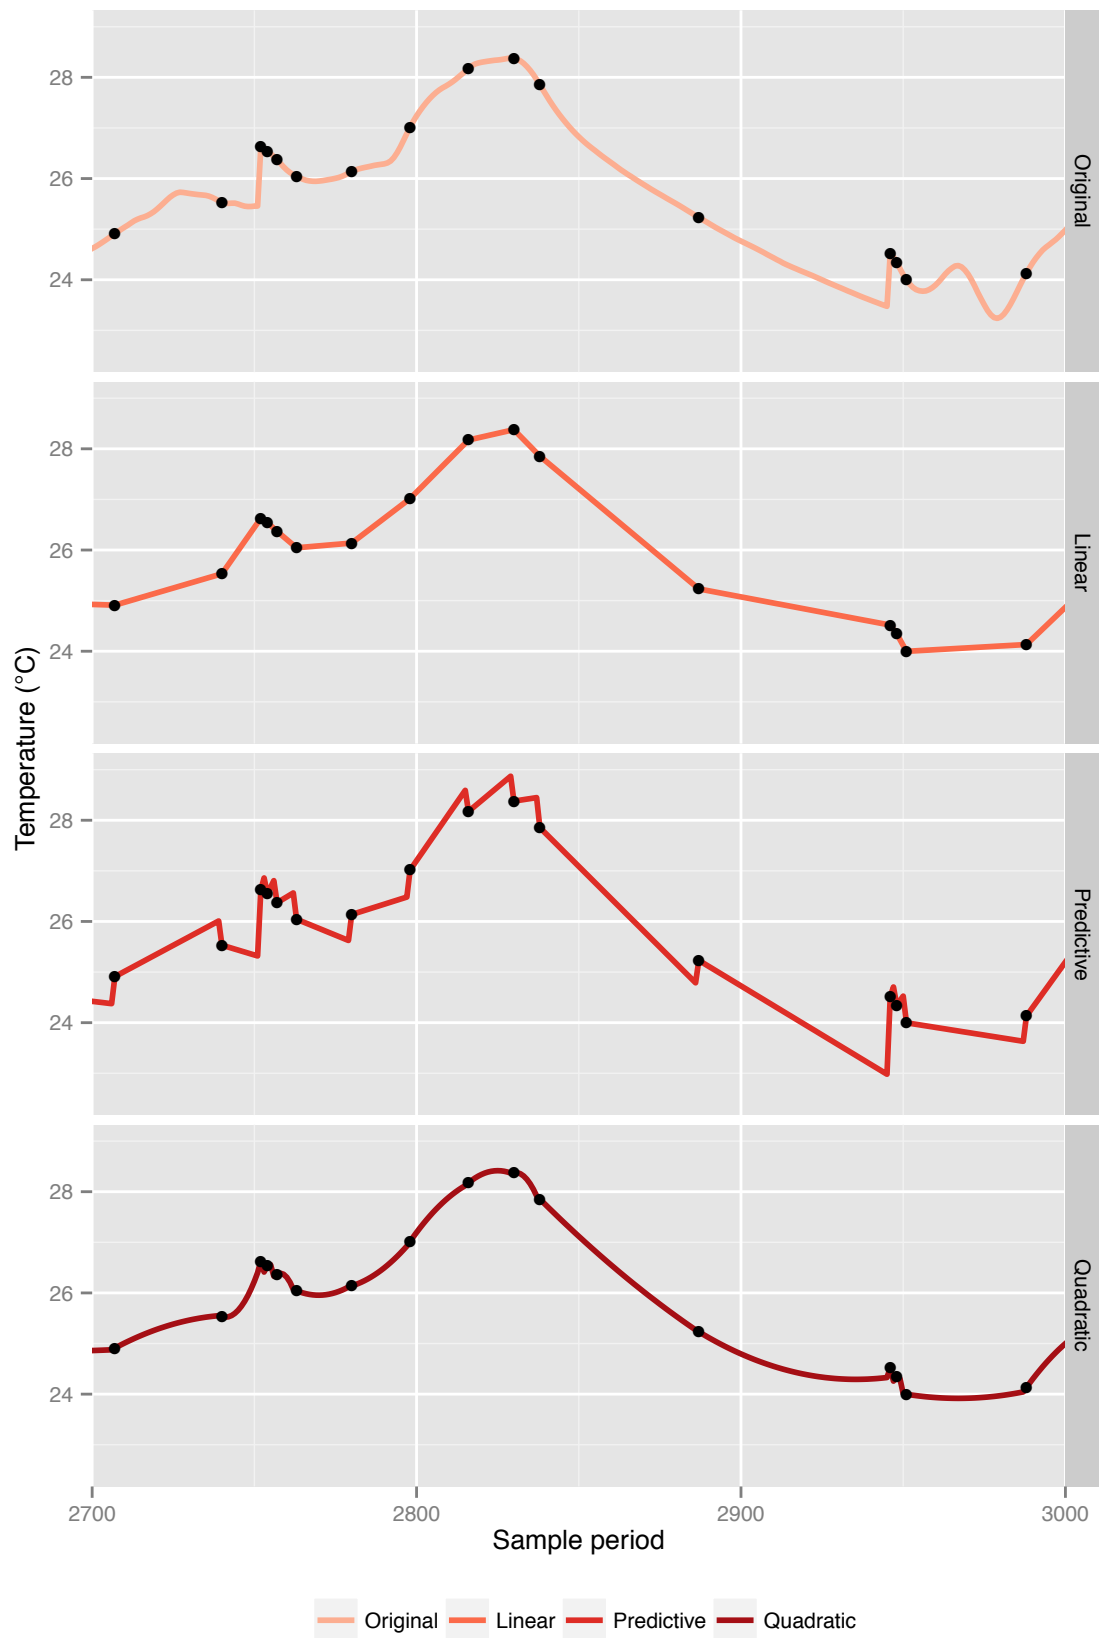


Figure 4.1: Faceted plot showing linear interpolation, predictive model, and dual quadratic spline signal reconstruction methods applied to a sample of 5 minute sense-and-send home air temperature data. The black points indicate where L-SIP transmissions occur.

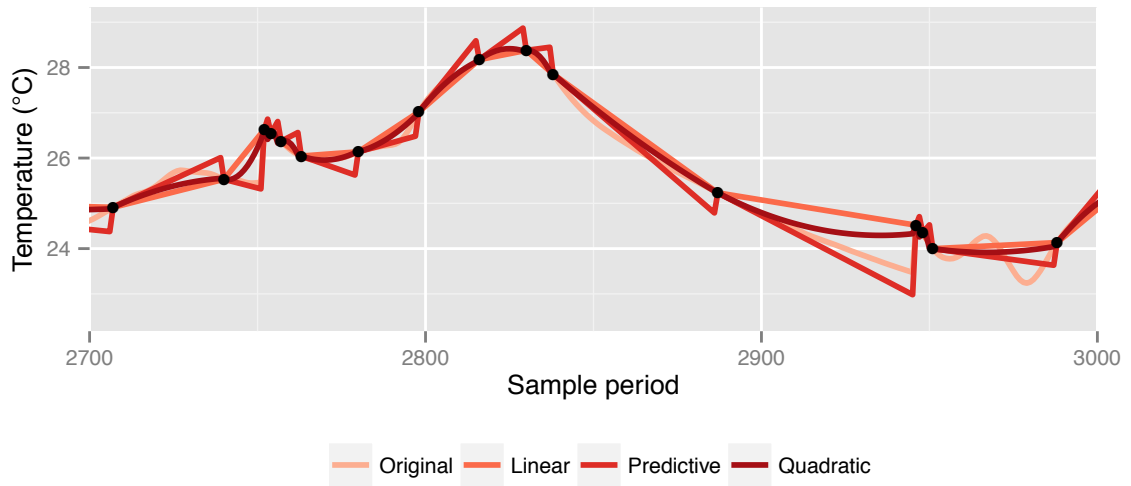


Figure 4.2: Combined line plot showing Linear interpolation, predictive model, and dual quadratic spline signal reconstruction methods applied to a sample of 5 minute sense-and-send home air temperature data. The black points indicate where L-SIP transmissions occur.

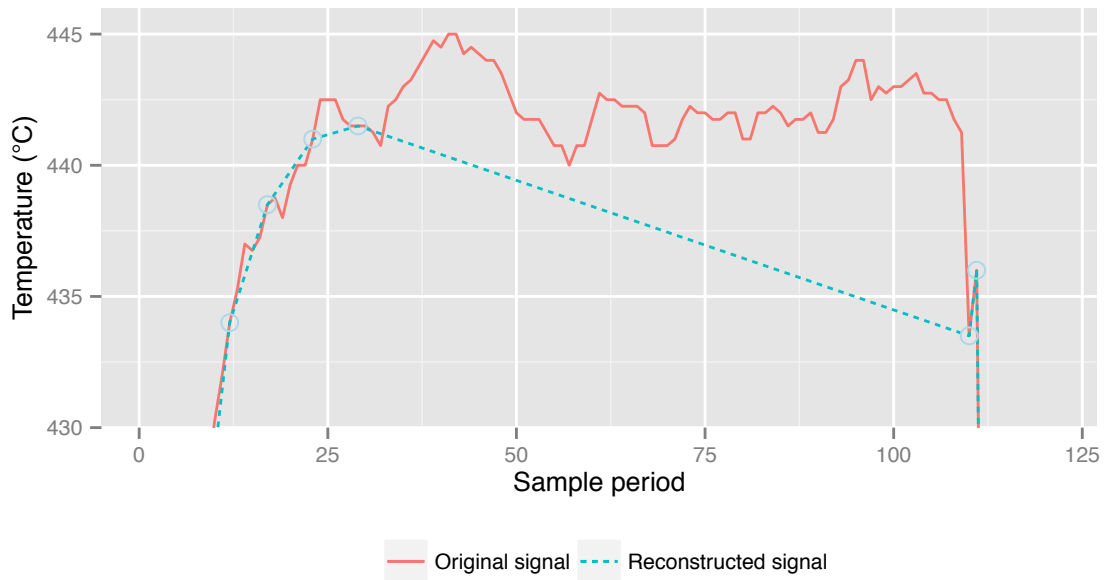


Figure 4.3: Large reconstruction error caused by linear interpolation given a change in signal trend following a period of acceptably predicted values. Hollow circles represent state update transmissions. This example is based on gas turbine engine temperatures, with a threshold of 5 °C.

the predictions diverge out to the error threshold and then “snap back” when an update occurs, as demonstrated in Figure 4.2. This effect results in a visually unappealing reconstruction and introduces more error than other methods. This method does, however, bound the error of every datapoint to be within the defined threshold—a property that the other methods do not have. In certain applications this may be a requirement.

4.1.3 Cubic splines

The lowest order polynomial that accommodates the demand of matching the gradients at two arbitrary points is a cubic of the form,

$$f(t) = At^3 + Bt^2 + Ct + D$$

The corresponding function of the signal x is $h(x) = f(t(x))$. It is desirable that the signal $h(x)$ is defined for all points $[x_j, x_k]$ and that the first point x_j corresponds to the initial state $t(x_j) = 0$, while the second point x_k corresponds to the next transmitted state $t(x_k) = 1$. This can be ensured by making $t(x)$ the linear function:

$$t(x) = \frac{(x - x_j)}{(x_k - x_j)}$$

where $x_k \neq x_j$, $t'(x) = 1/(x_k - x_j)$, and $x'(t) = x_k - x_j$. Given the two points defined by tuples $(a, b) = (f(t_j), f'(t_j))$ and $(c, d) = (f(t_k), f'(t_k))$, solutions can be found:

$$A = 2a + b - 2c + d$$

$$B = -3a - 2b + 3c - d$$

$$C = b$$

$$D = a.$$

We typically know the gradient of the signal with respect to real time x (or $\frac{df}{dx}$) instead of with respect to parametric time t (or $\frac{df}{dt}$), so it is helpful to apply the chain rule:

$$\frac{df}{dt} = \frac{df}{dx} \cdot \frac{dx}{dt} = (x_k - x_j) \frac{df}{dx}$$

4.1.4 Quartic splines

When sensing cycles occur between x_j and x_k , the reconstruction might be improved by incorporating the knowledge that (assuming that messages were not lost) the previous sensing cycle to x_k (denoted

x_{k-1}) produced a state estimate that could be linearly predicted from the state at x_j plus or minus an allowed threshold ε . That is,

$$|r(x_{k-1}) - h(x_{k-1})| \leq \varepsilon$$

where $r(x_{k-1}) = a + b(x_{k-1} - x_j)$ is the extrapolated estimate based on the state at x_j .

If the basic form $h(x_{k-1})$ given by the spline meets this criteria, then no adjustment is needed. However, if it does not meet the criteria, then an additional constraint can be added to require the curve to pass through either the maximum or minimum error possible, depending on which one is closer,

$$h_2(x_{k-1}) = \begin{cases} r(x_{k-1}) + \varepsilon & \text{if } h(x_{k-1}) > r(x_{k-1}) + \varepsilon \\ r(x_{k-1}) - \varepsilon & \text{if } h(x_{k-1}) < r(x_{k-1}) - \varepsilon \\ h(x_{k-1}) & \text{otherwise.} \end{cases}$$

A higher order polynomial is needed to incorporate this constraint,

$$f_2(t) = At^4 + Bt^3 + Ct^2 + Dt + E$$

As with the cubic spline in Section 4.1.3, the spline can be derived by solving for A,B,C,D,E given $(a, b) = (f_2(t_j), f'_2(t_j))$, $(c, d) = (f_2(t_k), f'_2(t_k))$ and the additional constraint $e = f_2(t_{k-1})$.

4.1.5 Dual quadratic splines

In Section 4.1.4, a spline control point is calculated at x_{k-1} that is linearly predicted from the state at x_j plus or minus an allowed threshold ε . Rather than producing a curve that fits through this point two quadratic splines are calculated, one that goes through $h(x_j)$, $h(x_{k-1})$, with the gradient $h'(x_j)$, and one that goes through $h(x_{k-1})$, $h(x_k)$ with the gradient $h'(x_k)$ at x_k . These two new splines take the form:

$$f_3(x) = Ax^2 + Bx + C$$

$$f_4(x) = Dx^2 + Ex + F$$

The first quadratic spline is calculated with parameters $a = f_3(t_j)$, $b = f_3(t_k)$, and $c = f'_3(t_k)$. Similarly, once $h(x_{k-1})$ has been estimated, a quadratic spline with parameters $d = f_4(t_j)$, $e = f_4(t_k)$, and $f = f'_4(t_j)$ helps provide a smooth line.

4.2 Evaluation

Due to the large number of different applications and sensing modalities that DPS algorithms may be used for, the context of the evaluation must be clearly defined. Particularly, the evaluation presented here focuses on periodic low frequency signals that would be encountered when monitoring air temperature in occupied homes. Data composed of high frequency waveforms, such as vibration and acoustic data, is not within the scope. When analysing these types of data it is often important that frequency and phase information can be correctly extracted. Furthermore, the linear model used by L-SIP is not suited to high frequency cyclic data, providing a relatively low data reduction.

The data used for the analysis in this chapter consists of two sets drawn from the same pool. The complete pool of data is made up of traces from a total of 235 air temperature sensors deployed in 37 homes. The homes consist of flats and houses with between 1 and 5 bedrooms, between 1 and 7 occupants, and built between the 1940s and the 2010s. These homes therefore represent a wide variety of builds and occupancy patterns. The two sets of data extracted from these are:

Dataset 1

Two weeks of data from each sensor with 100% yield (235 traces, 3290 trace-days). The duration was selected to allow the same period to be extracted from each sensor trace given the limits of differing deployment durations and yields over time. The evaluation of the five reconstruction methods considered here is based on these traces.

Dataset 2

One year¹ of data from nine nodes in a home with 99.98%² yield (9 traces, 3240 trace-days). This was selected to allow demonstration of the adjusted splines method (see Section 4.3). The traces were required to be of a long duration such that transmission failure could be meaningfully simulated with a variety of durations. The specific sensors selected were located in an end-terrace 4 bedroom house with 5 occupants. The data contains variation on a number of time scales—seasonal cycle, daily cycle, and occupant driven transients.

This section evaluates the following hypothesis:

H4.1: *Dual quadratic spline-based reconstruction for L-SIP will increase the accuracy of the reconstructed signal compared to linear interpolation or model prediction reconstruction.*

¹Actually 360 days but for brevity will be stated as a year. Taken from House 1 described in Appendix B

²Missing values (average 1200 per sensor) are imputed via linear interpolation to avoid discontinuities: the majority of data loss instances were infrequent failures of short duration (e.g, one missing sample)

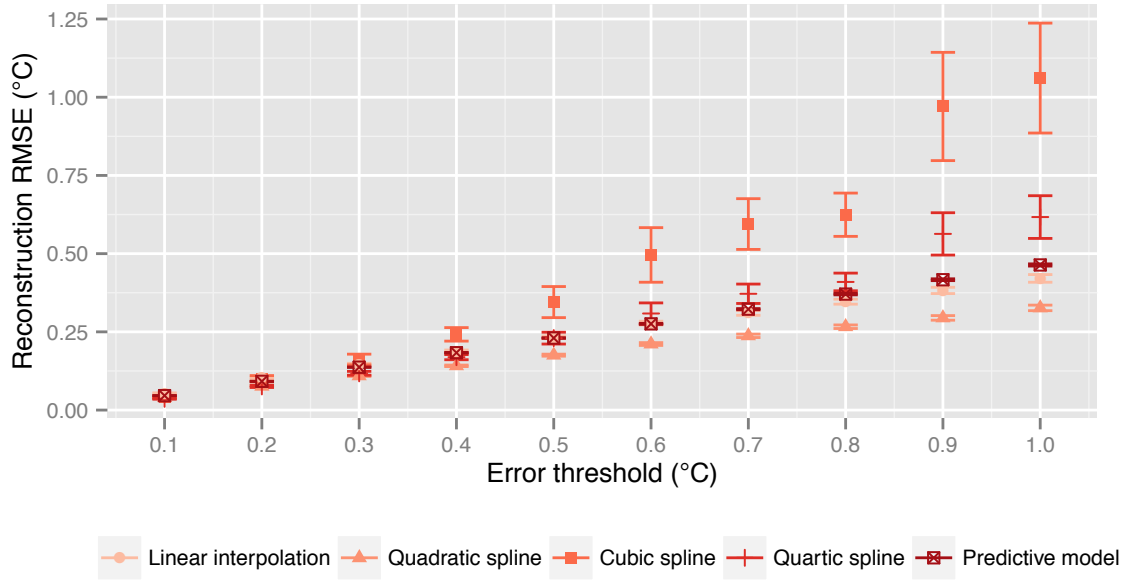


Figure 4.4: Comparison of average temperature reconstruction accuracy (with 95% confidence interval) of linear interpolation, predictive model, and the three spline-based reconstruction methods.

To evaluate the proposed reconstruction methods, L-SIP (with a range of error thresholds from 0.1 °C to 1 °C) was applied to Dataset 1. Each trace was reconstructed using each of linear interpolation, predictive model, cubic splines, quartic splines, and dual quadratic splines. The performance of the reconstruction methods were compared using the RMSE between the original (smoothed) signal and the reconstructed signal.

Figure 4.4 shows the average reconstruction error (along with the 95% confidence interval) for each method over a range of error threshold values. The results show that cubic splines provide the least accurate reconstruction in all cases and that the confidence interval for this method becomes larger as the error threshold increases (meaning that the results for any given trace become less predictable the greater the error threshold). At thresholds of 0.9 °C and 1 °C, the RMSE for cubic splines can exceed the error threshold—a result not seen for any other reconstruction method.

The dual quadratic spline approach provides the lowest reconstruction error for all thresholds considered—approximately 25% lower than linear interpolation and approximately 15% to 30% lower than the predictive model. Over all traces, the performance of linear interpolation, predictive model, and dual quadratic splines was highly predictable—indicated by the small confidence intervals. The increase in RMSE for dual quadratic splines is approximately 0.03 °C per 0.1 °C increase in threshold, compared to 0.04 °C for linear interpolations. The largest confidence interval for dual quadratic splines was 0.015 °C, compared

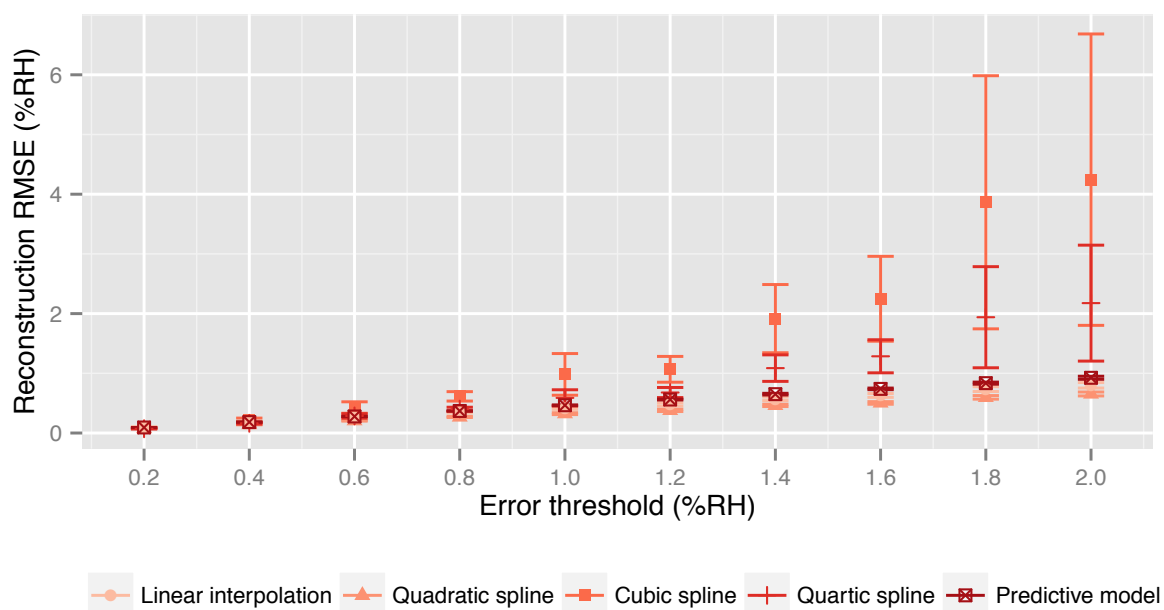


Figure 4.5: Comparison of average relative humidity reconstruction accuracy (with 95% confidence interval) of linear interpolation, predictive model, and the three spline-based reconstruction methods.

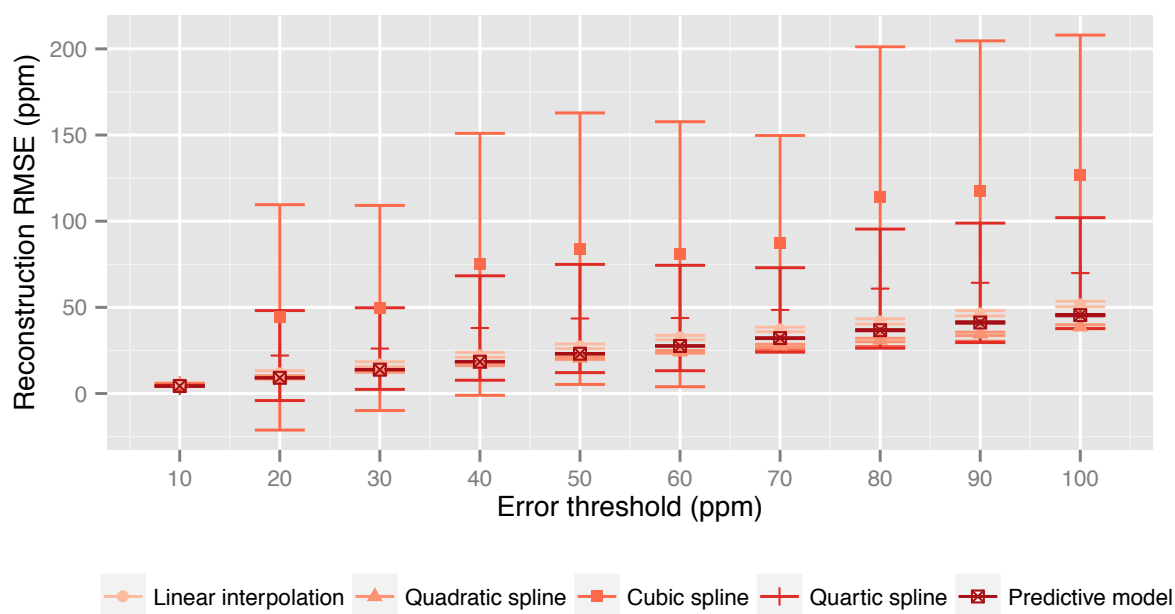


Figure 4.6: Comparison of average CO₂ reconstruction accuracy (with 95% confidence interval) of linear interpolation, predictive model, and the three spline-based reconstruction methods.

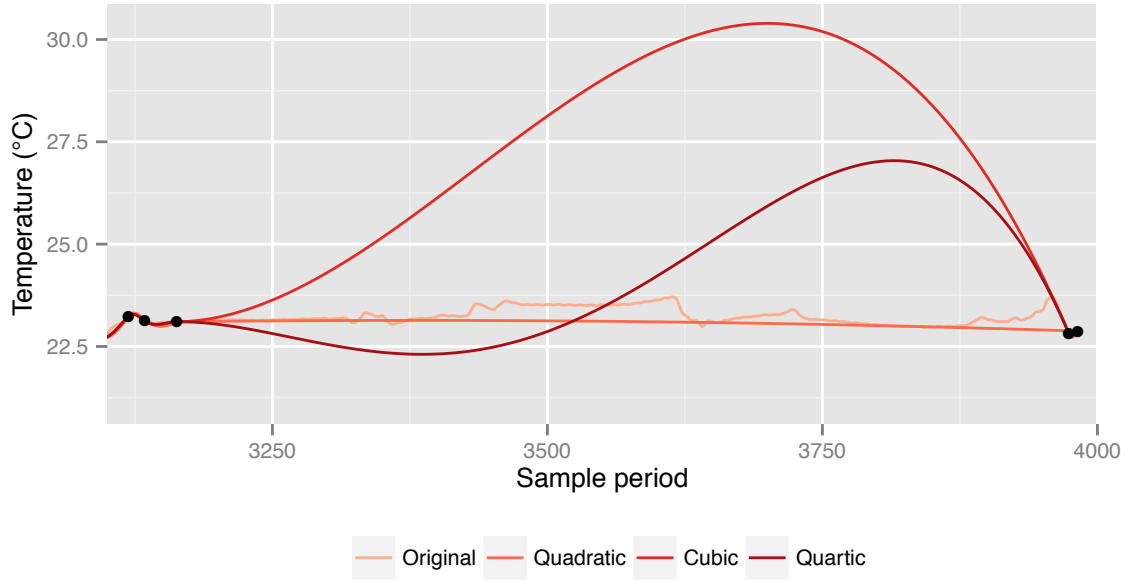


Figure 4.7: Example of large curves created by the cubic spline and quartic spline methods. The black dots represent when state transmissions have occurred.

to 0.021 °C for linear interpolation. This indicates that, overall, dual quadratic splines provide a better reconstruction than linear interpolation for air temperature data of the type used here. The predictive model provided the smallest confidence interval of the reconstruction methods but a higher RMSE than linear interpolation at higher thresholds.

The results shown do not appear to be particular to air temperature. Figure 4.5 shows that the results for reconstruction for relative humidity match that of air temperature. Figure 4.6 shows the results for reconstruction of Carbon Dioxide (CO₂) data from the homes (80 nodes, of the full 235 node set, that included CO₂ sensors). The same ordering of methods from most to least accurate can be seen for CO₂ as with air temperature with the exception of the predictive model, which performed better than linear interpolation for this data type. Overall, dual quadratic splines continue to provide the lowest reconstruction RMSE of the methods considered.

A problem observed with both of the higher order polynomial-based splines (cubic and quartic) is that if the state at either end of the spline has a large gradient then the polynomial will create a large looping curve that may be drastically different to the original signal. Figure 4.7 demonstrates this effect; the large gradient in the state estimate at sample period 3974 creates a large loop in the cubic and quartic reconstruction method. However, when using the quadratic splines, clamping at x_{k-1} , is not subject to this problem. . The dual quadratic spline method is therefore more accurate in signal reconstruction.

The combination of low and predictable reconstruction RMSE with avoidance of the “large loop” problem exhibited by the other spline-based methods indicates that dual quadratic splines are the best option of those considered here for reconstructing data such as interior air temperatures.

This evaluation has shown **H4.1** to be true, the dual quadratic splines reconstruction method provides the most accurate reconstruction of signals (versus linear interpolation, predictive model, cubic splines, and quartic splines) resulting in a RMSE approximately 25% lower than linear interpolation and approximately 15%–30% lower than the predictive model.

4.3 Dual quadratic splines on a lossy network

Section 4.1 listed “no loss of state updates” as a requirement of the spline-based reconstruction methods presented in this chapter and of DPS signal reconstruction more generally. In a real-world scenario this is often unlikely to be the case. The loss of transmitted packets will impact the accuracy of the reconstruction. This section investigates the effect of failure to transmit updates and proposes an extension to the dual quadratic spline reconstruction method to minimise the effect of transmission loss or reconstruction error. Note that the technique described here has the requirement that the original sensing period is fixed—it does not apply to systems with a variable sensing period.

One of the requirements imposed by DPS, is that the node and sink share the same model of the sensed data. Specifically, Spanish Inquisition Protocol (SIP) requires the sink to acknowledge receipt of state updates and only performs prediction using the new state when such an acknowledgement has been received. This means that when packet loss occurs, state updates (including an incrementing sequence number) will continue to be transmitted following every sensing cycle until an acknowledgement is received from the sink (Note: there is a special case where the state reverts and thus the sink state is again considered accurate. This case is ignored for the purpose of this analysis.) . It is hypothesised:

H4.2: *Using sequence numbers in the dual quadratic spline reconstruction algorithm to detect transmission failure will increase the accuracy of the reconstructed signal following a network failure*

State predictions are accurate (within the defined error threshold) until the point at which the subsequent state update should have been transmitted. Taking a state update u_j received at x_j with sequence number s_j and a second state update u_k received at timestamp x_k with sequence s_k , the timestamp up to which u_j is valid is $x_{k-(s_k-s_j)}$, referred to hereafter as x_m for clarity. The point at x_m with predicted value v_p can therefore be used as the join between the two splines in place of the point at x_{k-1} . In the implementation used here, the predicted value at the join is calculated based on a weighted combination of a forward

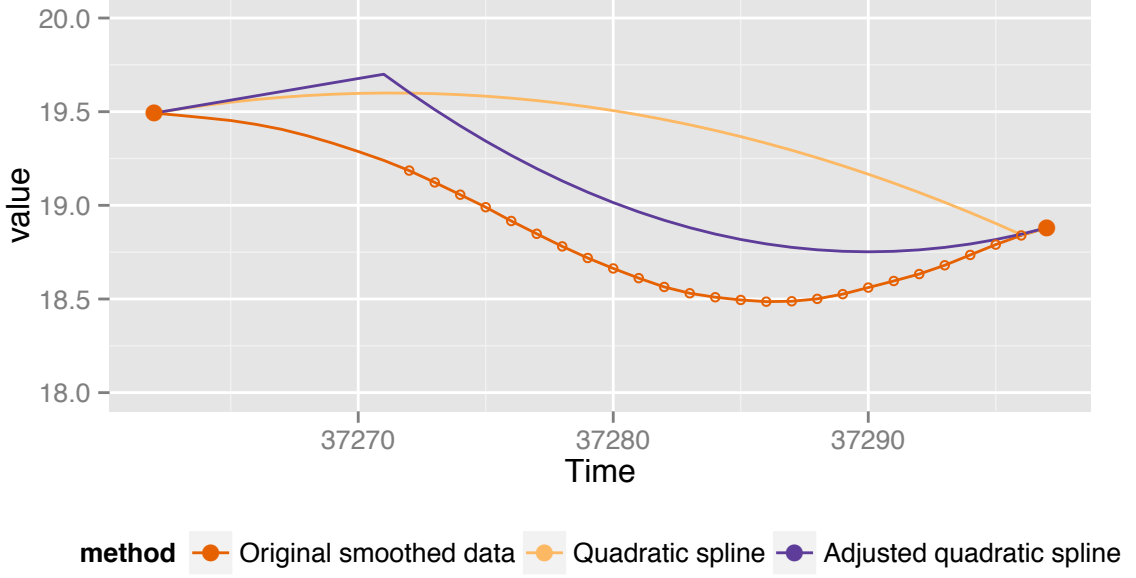


Figure 4.8: Demonstration of the benefit of using the adjusted dual quadratic spline method (RMSE of 0.279 °C) versus non-adjusted dual quadratic splines (RMSE of 0.569 °C). Smaller open points indicate failed transmissions.

prediction, v_f , and a backward prediction, v_b :

$$\begin{aligned}
 l &= x_k - x_j \\
 x_m &= x_{k-(s_k-s_j)} \\
 v_p &= v_f \frac{x_j - x_m}{l} + v_b \frac{x_m - x_k}{l}
 \end{aligned}$$

The predicted values v_f and v_b are linear extrapolations based on u_j and u_k .

Figure 4.8 shows the effect of packet loss based on a trace where state updates failed to transmit for 25 sample periods, with success on the 26th transmission attempt. Using the adjusted dual quadratic splines method reduces reconstruction RMSE for this period by a factor of 2.

Figure 4.9 shows a comparison of the results for the original and adjusted dual quadratic spline methods applied to temperature data in Dataset 2 (9 traces of one year each). In total, splines were computed for 500 simulated transmission failure periods for each of 15 possible failure durations (7500 total simulations). The RMSE for each method was calculated over the failure period. Overall, the adjusted dual quadratic splines method provides a lower RMSE than the original dual quadratic splines method during periods of transmission loss. For example, for failure durations between 10 to 25 samples, the adjusted

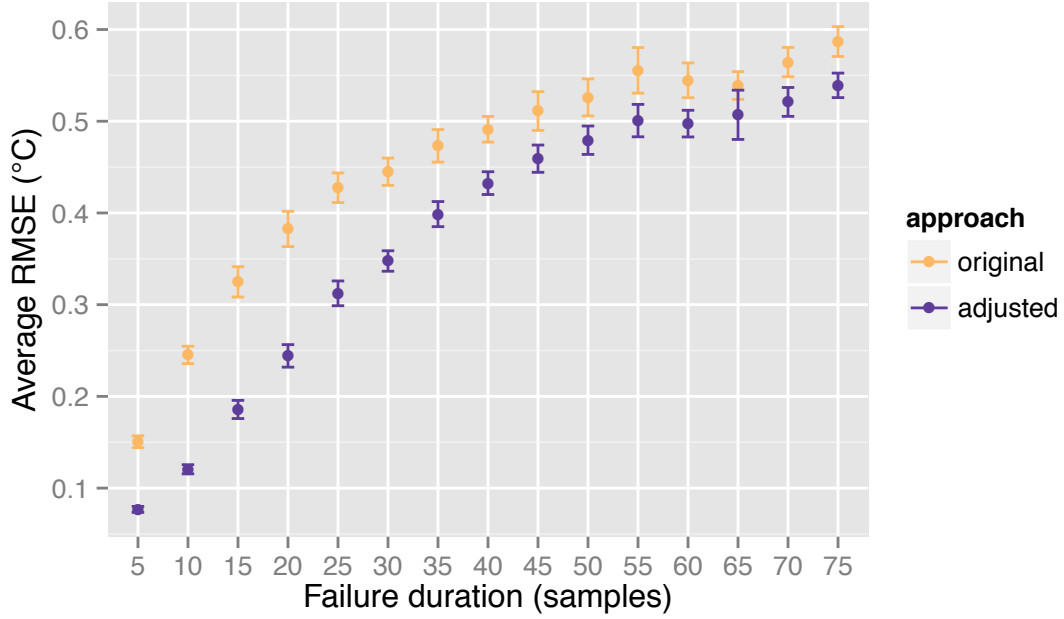


Figure 4.9: Reconstruction accuracy over 9 data traces (500 communication failure simulations per trace per failure duration).

dual quadratic splines method improves reconstruction by 0.13 °C. The original dual quadratic spline method can provide a reconstruction within the error threshold (0.5 °C) for up to 40 failed transmissions (3 hours, 20 minutes), whereas the adjusted dual quadratic spline provides an accurate reconstruction for up to 60 failed transmissions (5 hours)—an increase by a factor of 1.5×. Beyond this duration, the two methods cannot guarantee accuracy within the allowable threshold limits. This is because neither method is able to extract sufficient information from the received state updates to allow more accurate reconstruction over such long periods.

The evaluation has shown **H4.2** to be true—integrating sequence numbers in the dual quadratic spline reconstruction algorithm can improve signal reconstruction accuracy. The adjusted method allows an accurate reconstruction of data when transmissions fail for up to 5 hours (for the system and environment studied here). Beyond this duration it provides similar reconstruction accuracy as the original method, although is consistently better. The conclusion therefore is that the adjusted dual quadratic spline method can provide a lower RMSE for short duration failures and performs no worse than the dual quadratic spline method for longer duration failures.

Table 4.1: Error introduced into exposure graph bands by the reconstruction process. Values not given where no datapoints (original or reconstructed) fell within the given band. The two largest errors are indicated in bold.

	Error per band due to reconstruction process (%)				
	Health issues	Cold	Comfort	Warm	Overheating
Bathroom 1	0.56	0.14	0.54	0.16	-
Dining room	-	0.17	0.23	0.06	-
Bedroom 1	-	0.40	0.32	0.06	0.01
Bedroom 2	0.04	0.25	0.10	0.23	0.03
Bedroom 3	0.22	0.07	0.07	0.34	0.01
Kitchen	0.01	0.71	0.81	0.09	-
Living room	0.01	0.22	0.39	0.16	0.00
Sitting room	0.00	0.13	0.22	0.08	-
Spare room	-	0.09	0.45	0.44	0.08

4.4 The effect of reconstruction on statistical summaries

An important question for monitoring systems utilising data reduction algorithms is how the reconstruction process affects the analysis performed by the end-user. Any error introduced by the reconstruction process will propagate through to the analysis, potentially changing the conclusions reached or prompting incorrect decision making. However, the effect of reconstruction error for individual points is likely to be lessened by any summarisation process applied to the data, for example when generating exposure graphs (showing the percentage of time the sensed values fell within particular ranges).

To analyse the effect of the reconstruction on statistical summaries, temperature time-discounted distribution summaries (See Section 5.2 on page 97) created from each trace of Dataset 2 (9 traces, one year duration each) were compared based on the original traces and reconstructed traces using dual quadratic splines. The temperature exposure graphs show the percentage of data points that fell into each of five ranges (or bands), labelled “overheating”, “warm”, “comfort”, “cold”, and “health issues”. These ranges are important for home monitoring purposes as they show the effectiveness of a given heating strategy for the home and allow problematic rooms to be easily identified.

Table 4.1 provides the results of the comparison, showing the error per band for each room. The errors are relatively low as demonstrated by Figure 4.10, which compares the exposure graphs generated for data from the kitchen sensor. The reconstruction of this data resulted in the highest band error of any of the traces—0.8% for the comfort band. Even for this highest error, the difference in the band values is visually almost imperceptible. This indicates that spline based reconstruction does not significantly impact the further analysis of home monitoring data.

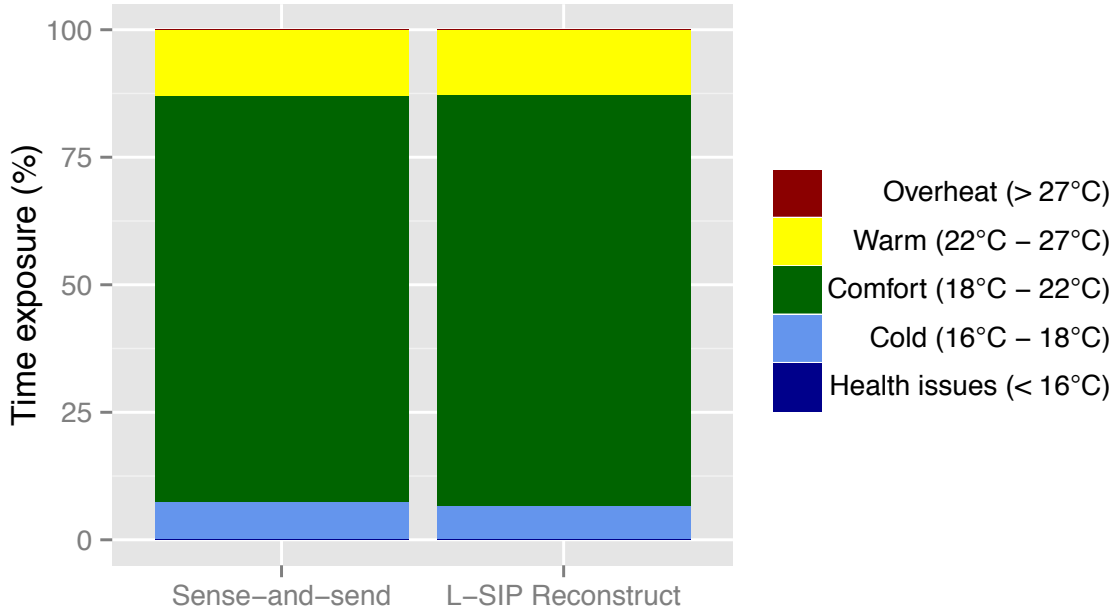


Figure 4.10: Comparison of temperature exposure graphs using original data and reconstructed data. Trace used was the kitchen sensor, giving the largest band error (0.8%).

4.5 Summary

This chapter has demonstrated the benefit of using splines for signal reconstruction, and specifically the use of dual quadratic splines. These make use of two pieces information provided by L-SIP: i) gradient estimates and ii) bounds on suppressed samples. The results show that the dual quadratic splines reconstruction method provides more accurate reconstruction of signals than linear interpolation, predictive model, cubic splines, and quartic splines, resulting in an RMSE approximately 25% lower than linear interpolation and approximately 15% to 30% lower than the predictive model. Indoor air temperature monitoring was used as the application for evaluation, however, the reconstruction method is suitable for a range of data types.

In the case of transmission failure it was shown that the dual quadratic spline reconstruction method can be extended to use sequence numbers to determine the limits of prediction for the last received state. The original dual quadratic spline method can provide a reconstruction within the error threshold for up to 40 failed transmissions (3 hours, 20 minutes), whereas the adjusted dual quadratic spline provides an accurate reconstruction up to 60 failed transmissions (5 hours)—an increase by a factor of 1.5 \times . The conclusion therefore is that the adjusted dual quadratic method provides a more accurate reconstruction

in a lossy network.

Finally, it is important to understand the effect of reconstruction on the further analysis of gathered data. The case study of exposure graphs was considered here to compare knowledge generated from Sense-and-send and L-SIP reconstruction. It was demonstrated that even for this highest error, the difference in the band values is visually imperceptible.

This chapter aimed to answer **RQ3**. To answer—yes, dual quadratic splines improves the accuracy of reconstructed signals compared to the output of simpler methods, for example linear interpolation or model prediction, when using data suppression algorithms such as SIP.

The next chapter describes how the data can be transformed on-node to produce information rich metrics which reduce the need for transmissions, thus extending node lifetime.

Chapter 5

Bare Necessities—Knowledge driven design

The previous chapters in this thesis have focused on Dual Prediction Schemes (DPSs), which model the sensed signal to allow an accurate reconstruction of the signal at the sink using fewer transmissions. Although transmissions can be reduced through modelling the signal in this way, considerable savings are also possible by considering the context of a specific application.

This chapter presents and evaluates Bare Necessities (BN), an approach which pushes the calculation of application-level state on-node. BN is so called as it only transmits the “bare necessities”—the application-level information required by the end-user.

The description and evaluation of BN in this chapter allows the following research question to be answered: **RQ4:** *Can a combination of DPS concepts with the calculation of application-level information on-node significantly reduce the energy requirements of a node further than the current state of the art?*

BN is evaluated in the context of a household monitoring application that reports the percentage of time a room spends in various environmental conditions. BN provides packet reduction by a factor of $7000\times$ compared to a sense-and-send approach, and a factor of $190\times$ compared to Linear Spanish Inquisition Protocol (L-SIP). When implemented on-node with the Backbone Collection Tree Protocol (B-CTP) (as described in Chapter 3), BN reduces a TelosB node’s annual energy requirement by a factor of $14.1\times$.

This chapter is structured as follows: The next section motivates the work presented in this chapter. Section 5.2 then introduces the concept of application state, and describes the example used for evaluation of BN. Section 5.3 describes the BN algorithm, which is then evaluated in Section 5.4.

5.1 Motivation

Given the information needs of a specific application, raw sensor measurement data is often highly compressible. For some applications the ability to reconstruct the entire time series is unnecessary and it is only important to know the proportion of time spent in a state, or set of states. For example, human behaviour monitoring applications usually contain high data rate sensors such as accelerometers and gyroscopes with sampling rates from tens to hundreds of hertz per sensor [80, 97]. However, to understand the subject's general behaviour the end-user is often only interested in how long is spent in a certain state (walking, driving, standing) in a given day. This information is often much smaller in terms of number of bits (or more compressible) than the raw signal used to generate it.

In addition to the size of the data, application-level information tends to be more stable over time than the raw signal. For example, in building monitoring applications the ratio of energy consumption to degree days (ratio which bears a rough correspondence to the building heat loss) is expected to remain the same season after season, year after year. As a result, any significant change in the value may be of importance. For example, a change may indicate: a refurbishment improved the insulation, or new tenants have adjusted the heating system. Identifying when key metrics change can be insightful. Therefore on-node analytics to generate application-level state lend themselves to DPS-based techniques. The transmission reduction allowed by DPS algorithms should, theoretically, considerably extend the network life while having minimal effect on the usefulness of the information gathered.

Dasu and Johnson [24] state that 80% of data analysis effort is spent on the process of cleaning and preparing the data. Therefore, if the calculation of application-level information is performed on-node, then the only processing required by the end-user is to visualise the received information in a meaningful format. This minimises the amount of time required for post-processing.

An additional benefit achieved by transmitting only application-level information is that of privacy. Privacy issues are often not considered in the design of Wireless Sensor Network (WSN) systems. Collected data may be misappropriated for other uses beyond the original monitoring specification. For example, when monitoring humidity in a bathroom to assess mould risk, the raw signal could then be used for unauthorised purposes such as identifying when and for how long showers or baths are used. Langheinrich [61] states that use limitation should be placed on information that is clearly not part of the original intent of monitoring. Therefore, in the case of BN, transforming the data into application-level information at the first opportunity makes it difficult to misuse data, and aids in providing a measure of privacy.

Table 5.1: Temperature exposure bands, derived from the work of Hacker *et al.* [38], and Nurse *et al.* [79]

Range (°C)	Description
$T \leq 16$	Room presents a health risk to occupants
$16 < T \leq 18$	Room is too cold (slight health risk)
$18 < T \leq 22$	Optimal thermal comfort
$22 < T \leq 27$	Room is warm (wasting energy on heating)
$T > 27$	Room is overheated (wasting energy on heating)

5.2 Application-level state: the time-discounted distribution summary

The previous section motivated pushing the calculation of application-level information (or “metrics”) closer to the data source. In this thesis the application-level metric is termed **application-level state**. Application-level state is so called, since it is both at an application-level (in terms of context) and referring to the condition of the environment at a point in time. This section describes one such application-level state: the time-discounted distribution summary.

As part of the case study used in this thesis, home performance reports created for Orbit Heart of England required a number of novel metrics to summarise building performance in a way which can be readily understood by the end-user. One proposed metric, the time-discounted distribution summary, is used as an exemplar for BN. The time-discounted distribution summary will be used throughout this chapter to evaluate the performance of the BN approach.

Generally, time-series plots are used to graphically represent environmental data. However, the end-users for the case study used throughout this thesis (surveyors at Orbit Heart of England) were more interested in the frequency of conditions in the home. For example, for what proportion of the time was the living room in a thermally comfortable state? Tufte [109] showed that time-series plots are ideal for showing trends in the data, however, when considering the time apportioned to a condition, Tufte recommends bar charts over time-series plots. The time-discounted distribution (see Figure 5.1 on the next page) uses the concept of bar-charts to group readings into similar bands of readings. For example, a cold band which represents all temperatures below 16 °C.

The time-discounted distribution summarises the relative amount of time that a set of observations fall into a number of value ranges (bands), and is visualised using a stacked bar chart. In the case of environmental monitoring, measurements are split by each well-defined space in the home (rooms, corridors, etc.), and the conditions within that room (cold, comfortable, warm, etc.).

While it is expected that a home monitoring system will sense several types of data, for demonstration

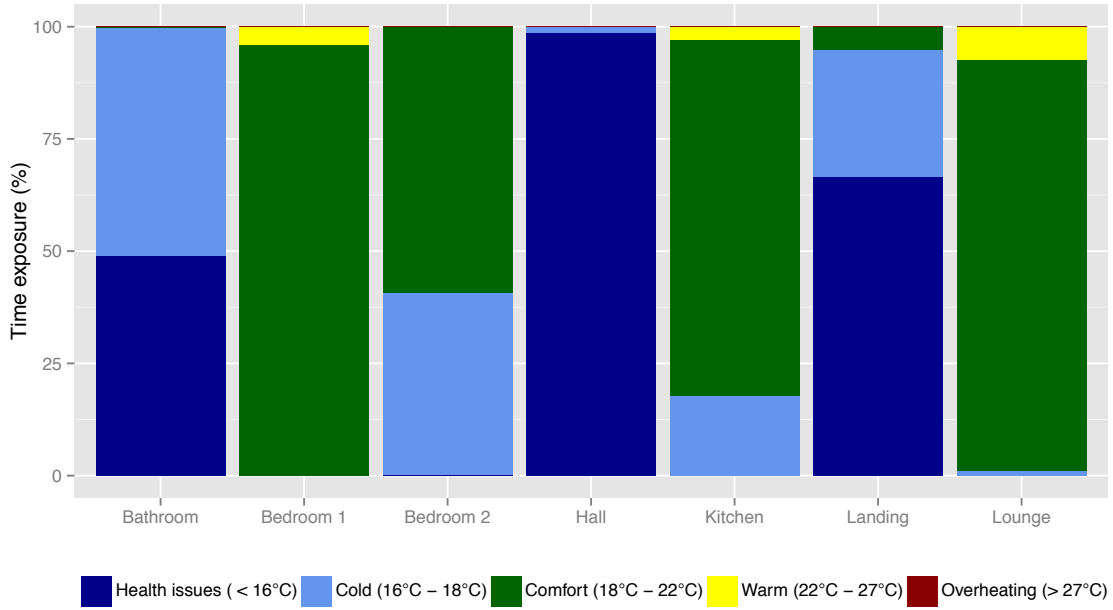


Figure 5.1: Temperature time-discounted distribution summary indicating comfortable living areas but cold utility and transition areas.

purposes this example will consider only the case of temperature sensing. The temperature band ranges have been derived considering literature in the application domain, such as the work by Hacker *et al.* [38] and Nurse *et al.* [79], and are presented in Table 5.1.

The time-discounted distribution summary is usually presented in the form of stacked bar charts with each band coloured appropriately. For example, using red for hot and blue for cold. Intuitive colouring allows a quick understanding of the environmental conditions within the home by the end-user.

Figure 5.1 presents an example of the time-discounted distribution output for temperature in a home. The figure shows that the main living areas (living room, bedrooms and kitchen) are mostly in thermally acceptable conditions. The two transition rooms (hall, landing), and the bathroom are areas for concern since the rooms are cold or in a state where potential health risks can occur. The graph suggests that the heating strategy employed by the occupants favours the living areas and provides minimal heating to the remaining rooms.

In the motivation of this chapter it was stated that raw sensor measurement data is often highly compressible in terms of application-level information. As an example, consider a WSN node monitoring temperature and humidity in a room at a five minute sensing interval. A sense-and-send approach would require 6.8×10^6 bits to be transmitted to the sink over a period of a year. However, when using time-discounted distribution summaries the number of bits required for temperature and humidity monitoring

Algorithm 5.1 Online time-discounted histogram encoding algorithm for estimating the exposure distribution phrased in terms of G-DPS (described in Algorithm 3.1).

estimate new state

$x_i \leftarrow \gamma x_i + b(i, z)$, (*update band count*)
for all $i \in \mathcal{B}$.

The predicate function $b(i, z)$ gives 1 if the reading z is in band i and zero otherwise. The update decays the current count estimate by decay constant γ and then increments the active band.

simplify

$y_i \leftarrow x_i / \sum_{i \in \mathcal{B}} x_i$, (*update distribution*)
for all $i \in \mathcal{B}$.

This normalises the distribution such that the band counts sum to 1.

predict sink state

$\mathbf{y}' \leftarrow \mathbf{y}_{\text{sink}}$ (static distribution assumption)

detect events

yes if $\exists i \in \mathcal{B} : |y_i - y'_i| > \varepsilon$

The distribution is eventful if at least one component has changed by at least some threshold ε

is reduced to $(5 \text{ bands} \times 32 \text{ bits}) \times 2 \text{ parameters} = 320 \text{ bits per update}$. Assuming summary information is transferred only once per month¹ (3840 bits per year), the information reduction would be of the order of $1750\times$. Given that transmission of bits is the main energy cost for wireless nodes this analysis suggests that performing the processing of the time-discounted distribution summary on-node will extend node lifetime.

5.3 The Bare Necessities algorithm

BN is an evolution of the DPS approach. BN takes the concept of DPS and adds on-node analytics. Compared to previous approaches discussed in this thesis (L-SIP and sense-and-send), BN is closely tied to the application and thus is able to make stronger assumptions about the informational content. This enables the majority of raw data to be summarised and thus reduce the number of packets sent by a node. BN imposes some penalties, such as the need for calculation to occur on the node and a slight loss in the accuracy of the resulting transmitted information, compared to post-processing sense-and-send data. However, in the context of the application these penalties are slight.

Algorithm 5.1 shows the BN node algorithm using time-discounted distribution summaries as an implementation of Generalised Dual Prediction Scheme (G-DPS). Each node senses the environment and converts the measured values into application-level state. The application-level state is then normalised

¹Evaluation of BN (see Section 5.4) shows monthly transmissions to be a reasonable assumption.

such that the per-band proportions (or probabilities) sum to 1. The event detection component checks for changes in each element of the state beyond some threshold compared to the predicted sink state. When an event is detected, a packet is transmitted to the database. A more detailed description of each step within the G-DPS framework follows:

Estimate new state

The *estimate new state* component converts data into application-level state, which, in the case of BN, is the time-discounted distribution summary. Time-discounted distribution bands can be thought of as a discrete form of a probability distribution or histogram. For each node, when a measurement of a sensed phenomena is made, the current band distribution is updated to account for new reading.

Let $b(i, z) \rightarrow \{0, 1\}$ be a predicate function that yields 1 if the measurement $z \in \mathcal{Z}$ is in band $i \in \mathcal{B}$ where \mathcal{B} is the set of bands and \mathcal{Z} is the set of possible sensor values. It is assumed at least one band is always applicable and bands do not overlap. For a finite deployment period involving k time intervals the probability that a band i is applicable is the average,

$$\frac{1}{k} \sum_{0 \leq t < k} b_t(i, z)$$

Generally, environments in which WSNs are deployed are not static, and are influenced by a number of factors. For example, the built environment is affected by the number of occupants, heating schedules, and the weather. Since the time-discounted distribution summaries are a cumulative measure, older measurements are discounted to give more importance to newer measurements based on an exponential decay constant $0 < \gamma < 1$ (this resembles the use of α in Exponentially Weighted Moving Average (EWMA)), giving,

$$x_k(i) = \frac{1}{\alpha_k} \sum_{0 \leq t \leq k} b_t(i, z) \gamma^{k-t}$$

where the normalising value α_k is chosen such that $\sum_i x_k(i) = 1$. The decay half-life² is,

$$t_{1/2} = \frac{T \ln 2}{(1 - \gamma)}$$

where T is the sensing period. Therefore, to achieve a specific half-life, the decay factor can be set to:

$$\gamma = 1 - \frac{T \ln 2}{t_{1/2}}$$

²Half-life ($t_{1/2}$) is the amount of time required for a quantity to fall to half its value as measured at the beginning of the time period

Selection of a half-life depends on the variability of the model used and the requirements of the application. If the application-level state experiences frequent meaningful changes then a short half-life is required to react to changes in a timely manner, though this will generate additional packet transmissions. If the application-level state is expected to be relatively stable over time then a longer half-life can be used, reducing the number of required transmissions.

Simplify

Rather than report band counts, the *simplify* component converts the band counts to a distribution by normalising their sum to 1,

$$y_i \leftarrow \frac{x_i}{\sum_{i \in \mathcal{B}} x_i}$$

Predict sink state

Since the distribution is slow changing, the *predict sink state* component assumes the distribution is constant over time. Therefore the last transmission is used as a comparison in the event detection component.

Detect events

Since BN is an implementation of G-DPS, transmissions are only required when a significant change is detected in the distribution of bands. The simplified vector is used as the application-level state. When compared to the predicted sink state, the node state is considered to be eventful if any element changes by some threshold ε (for example, 10%).

5.3.1 BN Assumptions

The following assumption is made for implementing BN:

The sensing frequency is fixed A variable sensing frequency would require each sample to be weighted proportionally to the period that it is applied to.

5.4 Evaluation of BN

This section evaluates the suitability of BN, using the home environment as a case study. The BN algorithm is evaluated using offline datasets with regard to three key performance measures:

1. the percentage of state updates transmitted (the compression ratio),
2. the accuracy of the reconstructed information compared to the state information generated by post-processing sense-and-send data (measured by average Root Mean Squared Error (RMSE) in a band), and
3. the node energy consumption.

5.4.1 Method

BN is envisaged to be deployed in long duration deployments, therefore it has been evaluated considering data collected from a home for a period of a year. The home is an end-terrace 4 bedroom house with 5 occupants built in the early 1900's. The properties of the home and occupants ensured that there was variation on a number of time scales—seasonal cycle, daily cycle, and occupant driven transients. Within the home nine nodes were deployed monitoring temperature, and humidity. These nine deployed nodes had an average yield of 99.98% and equate to data for 3240 trace-days. The nodes were set to sample the environment with a 5 minute sensing period (288 samples per day).

To evaluate BN, the year long datasets were compressed using BN configured with a band error threshold of $\varepsilon = 10\%$. To reconstruct the sensed signal, the suppressed values between state estimates transmitted by BN were derived through linear interpolation. The following measures were used to evaluate BN:

1. **Transmission reduction (or compression ratio)**—The percentage of state update packets transmitted compared to a sense-and-send approach.
2. **Reconstructed signal accuracy**—The accuracy of the reconstructed BN distribution compared to the BN distribution calculated by the node. Accuracy is measured by average RMSE of the bands.
3. **Post-processing accuracy**—The accuracy of the reconstructed distribution signal compared to the distributions created from post-processing sense-and-send data every month. Accuracy is measured by average RMSE of the bands.
4. **Node energy annual requirement**—The annual energy requirement of a node measured using the microbenchmarking approach. The node is assumed to be implementing B-CTP as previously described in Section 3.6 on page 71.

The next section shows an example of the BN algorithm output and how these measures are calculated.

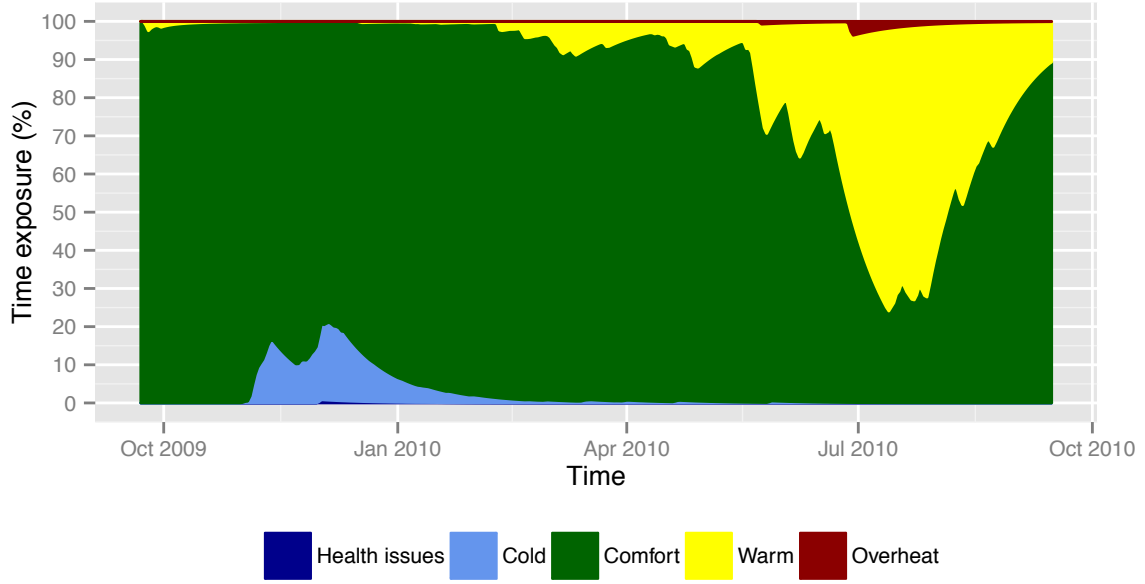


Figure 5.2: BN distribution signal with a decay of $t_{1/2} = 1$ month.

5.4.2 Example BN algorithm output and post-processing summary

To illustrate the terms used in the rest of this chapter, this section demonstrates and discusses the output of the BN algorithm over an example dataset. The following terms are used to describe this output:

Raw signal is the underlying signal reported by the sensor before any processing is undertaken (not shown in Figure 5.3). In this example 103680 samples were taken during the dataset period.

Distribution signal represents the processed sensor readings. In this case the raw signal is summarised as a time-discounted distribution of temperature with a half-life of one month applied. An example was previously shown in Figure 5.2.

Reconstructed distribution signal is the reconstruction of the distributed signal, from the transmitted state updates from a node. In Figure 5.3 the distribution signal is reconstructed using linear interpolation between transmitted state updates. In this evaluation the reconstructed distribution signal is compared to the original distribution signal. At each time-step the distribution signal has 5 values, one for each band, therefore the RMSE of the signal is reported as an average RMSE error of the 5 values. In this example, the distribution signal is reconstructed with an average RMSE error of 2.6% compared to the original distribution signal.

Transmissions represents points where the node transmits an update to the application-level state to

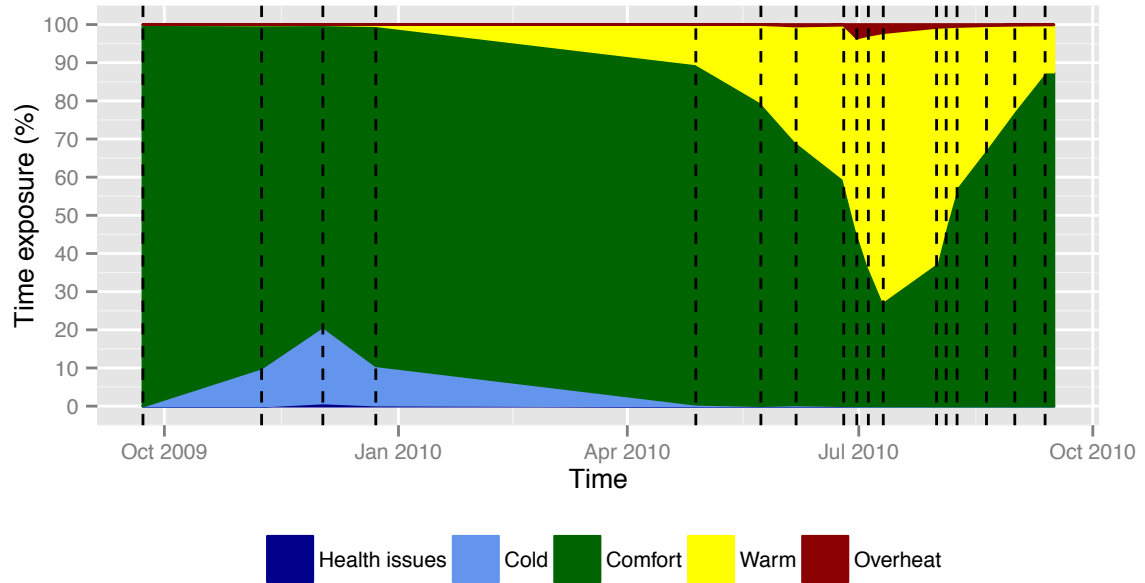


Figure 5.3: Example BN ($t_{1/2} = 1$ month, $\varepsilon = 10\%$) time-discounted temperature distribution summary over time for a monitored bedroom. Transmissions are indicated by the dotted vertical lines.

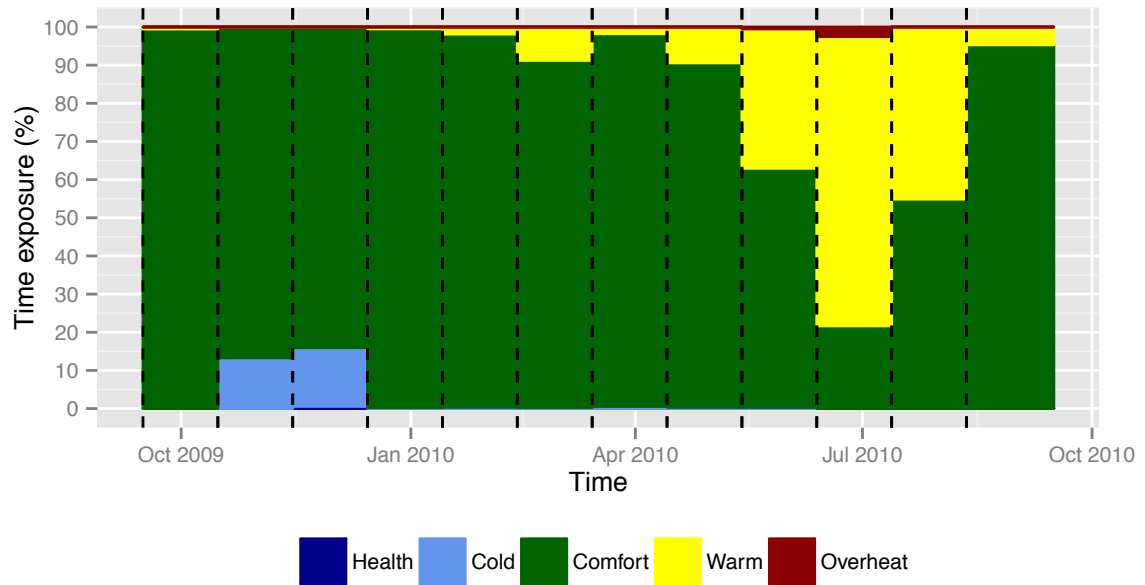


Figure 5.4: Example post-processed time-discounted temperature distribution summary generated at monthly intervals for a monitored bedroom. The dotted vertical lines indicate when post-processing takes place.

Table 5.2: Example BN performance metrics

Samples	Transmission statistics		Reconstructed RMSE	
	Packets transmitted	Transmission reduction (ratio)	vs. distribution signal	vs. post-processing
103680	17	99.98% (6100 \times)	2.6%	12.3%

the sink. In this example a total of 17 state updates would be transmitted to the sink, a transmission reduction of 99.98% (6100 \times compression ration)

Post-processed distribution summary is the time-discounted distribution summary created from post-processing sense-and-send data. Sense-and-send data is processed into time-discounted distributions with non-overlapping periods of two weeks, one month, three months, or six months. Figure 5.4 shows an example of post-processed distribution summary generated at monthly intervals for the same dataset used in Figure 5.3. BN is compared against the post-processed distribution summary since BN is intended to replace the need for any post-processing of data. Each post-processed distribution summary is compared to the output of the reconstructed distribution signal at that time to give an average RMSE error. In this example the average RMSE error of the bands is 12.3%.

The performance metrics for BN described in this section are summarised in Table 5.2.

5.4.3 Effect of error threshold

BN can be tuned for a specific application by setting the error threshold allowed between the predicted sink state and the current estimated state. Setting a higher error threshold will sacrifice some accuracy for fewer transmissions. Furthermore, since larger error thresholds are expected to lower the number of transmissions, it is expected that this will also reduce the reconstruction accuracy.

BN was used to compress the raw temperature signals from the nine nodes in the year long dataset. A range of thresholds $\varepsilon = \{1\%, 2\%, \dots, 20\%\}$ and half-life values $\lambda = \{\text{two weeks, one month, three months, six months}\}$ were used and the percentage of state update packets that require transmission was recorded. To derive the reconstruction accuracy, the RMSE is calculated between the reconstructed distribution signal and the distribution signal calculated on-node.

Figure 5.5 shows that the number of transmissions decrease as the error threshold increases, with less than 0.4% transmissions required when $\varepsilon = 1\%$ and an average of 0.014% when $\varepsilon = 10\%$. Figure 5.6 shows that the reconstruction error increases linearly with the size of the error threshold. Furthermore,



Figure 5.5: Percentage of annual BN packet transmissions over a range of error thresholds.



Figure 5.6: RMSE of the BN reconstruction error over a range of thresholds.

Table 5.3: Average number of BN temperature information packet transmissions and compression ratio, for nine rooms ($\varepsilon = 10\%$). The transmission percentage and compression ratio is calculated using the average number of transmissions over all datasets. This method is used in subsequent tables in this chapter.

Half-life	Transmissions and compression by deployment period				
	One week	Two weeks	One month	Six months	Year
One day	3.3 ± 3	8.3 ± 5	13 ± 6	63 ± 13	208 ± 12.5
	$610\times$	$480\times$	$650\times$	$830\times$	$510\times$
One week	3.3 ± 3	6.2 ± 4	7.0 ± 4.5	13 ± 6	37 ± 10.5
	$610\times$	$650\times$	$1300\times$	$3900\times$	$2800\times$
One month	3.3 ± 3	5.7 ± 4	6.2 ± 4	8.3 ± 5	15 ± 6.5
	$610\times$	$700\times$	$1400\times$	$6300\times$	$7000\times$
Six months	3.3 ± 3	5.7 ± 4	6.1 ± 4	7.5 ± 4.5	9.4 ± 5
	$610\times$	$700\times$	$1400\times$	$7000\times$	$11000\times$

the reconstruction error never exceeds the defined error threshold, remaining at approximately a quarter of the threshold. When increasing the error threshold beyond a certain point, little to no additional transmission reduction will be achieved (as shown in Figure 5.5), yet more error will be introduced into the reconstructed distribution signal.

To conclude, this evaluation has demonstrated that small values for the error threshold will provide a more accurate representation of the data, but require more transmissions to be made. Selecting a higher error threshold will reduce the number of transmissions, however this will also reduce the accuracy of the reconstructed signal.

5.4.4 Effect of half-life on transmissions

In addition to tuning the error threshold ε , the half-life λ is also a tuneable parameter for BN. The resulting decay constant has the effect of weighting current behaviour more highly than past behaviour. Using a larger half-life means that short term transient events will be filtered out, therefore reducing transmissions. The following hypothesis is tested:

H5.1: *The choice of half-life affects the number of required transmissions when using the BN algorithm.*

With a shorter half-life, a large percentage of sampling cycles cause transmission. As the half-life increases, the percentage of transmissions required decreases (for a fixed sampling interval).

BN was used to compress the raw temperature signal from the nine nodes in the year long dataset. This dataset was split into overlapping deployment periods to check how the half-life affects transmission reduction for differing deployment lengths.

Table 5.4: Comparison of the performance of BN ($t_{1/2} = 1$ month, $\varepsilon = 10\%$) with L-SIP ($\varepsilon = 0.5$ °C) for one year of temperature data.

	Transmissions	% of raw (compression ratio)	RMSE in band %
Sense-and-send	102236	100% (1 \times)	n/a
L-SIP	2900 \pm 700	2.80% (36 \times)	0.9% \pm 0.2%
BN	15 \pm 7	0.015% (7000 \times)	7.8% \pm 1%

As shown in Table 5.3, the transmission performance of BN depends on the half-life parameter and the deployment period. When the half-life is longer than the deployment period, little decay in the signal occurs and, therefore, has no effect on the number of transmissions. However, when a short half-life is used on long term deployments, a larger number of transmissions are required as the state reacts more to transient changes.

This experiment has shown **H5.1** to be true: larger half-life values result in a lower number of transmissions as long as the half-life does not exceed the deployment period.

Since the threshold, rather than half-life, affects the accuracy of the reconstructed distribution signal compared to the original distribution signal the accuracy was not considered in this evaluation. The half-life will primarily affect the accuracy of the reconstructed signal when compared to the post-processed distribution summaries. The next section evaluates this.

5.4.5 BN compared to post-processing summaries

BN is intended to replace the need for any post-processing of data other than visualisation. This evaluation tested how the reconstructed distribution signal from BN compares to post-processed distribution summaries calculated from L-SIP and sense-and-send. BN has the most specific application requirements that enables most of the data to be discarded and thus greatly reduces the number of transmissions. To compare BN and L-SIP, the state updates from L-SIP are reconstructed to form the raw signal, from which post-processed distribution summaries are created. This is then compared to the sense-and-send post-processed raw summaries in the same manner as BN. It is expected that BN transmits significantly less than L-SIP and sense-and-send but at a cost of some accuracy.

Table 5.4 compares the message reduction between sense-and-send, L-SIP and BN. In the case of the one year dataset L-SIP reduces the number of transmissions by 97.2% (36 \times) and BN reduces the number of transmissions by 99.985% (7000 \times). Since a node implementing BN transmits significantly less compared to L-SIP, the node is less likely to suffer from transmission failure. Therefore, considering G-DPS the penalty of an increased energy requirements for end-to-end acknowledgements will also be

Table 5.5: Number of transmissions and reconstruction results for the single-modal and multi-modal approach to BN

	Transmissions	Temperature RMSE	Humidity RMSE
Single-modal	35 ± 6	8.6 ± 0.7	13 ± 2
Multi-modal	32 ± 5	8.5 ± 1	8.6 ± 2
Improvement	-8.6%	-1.2%	-33.9%

reduced. Though BN reduces the number of transmissions to a much larger degree the average RMSE in a band is a factor of 7 greater than that of L-SIP and sense-and-send. However, the average RMSE in a band is within the allowed threshold limit of BN, and therefore would have minimal effect on the usefulness of the information gathered.

5.4.6 Multi-modal BN

A natural extension to BN is to support additional sensing modalities alongside temperature, such as relative humidity. A multi-modal approach was shown to be a benefit to G-DPS in Chapter 3. However, it is unclear how multi-modal will effect BN since the measures will often use a differing number of bands which will behave differently. The following hypothesis is tested:

H5.2: *Combining BN distributions into a single state vector will reduce transmissions and improve the accuracy of the distributions.*

To evaluate this approach both single-modal and multi-modal BN were used to compress the raw temperature and relative humidity signals of the year long dataset. BN was configured with an error threshold of $\varepsilon = 10\%$ and a half-life of $\lambda = 1$ month. This evaluation only considers a half-life of a month since it provided the lowest error when compared with one-month post-processed summaries. The approaches of single-modal and multi-modal BN were compared in terms of reconstruction accuracy and the number of transmissions. Recall that in the multi-modal approach each state update transmission includes all monitored parameters.

Table 5.5 shows that the multi-modal approach reduces BN transmissions by compared to single-modal. In terms of reconstruction accuracy, Table 5.5 shows an improvement to reconstruction accuracy for relative humidity by when using a multi-modal approach. However, the temperature reconstruction accuracy receives no significant improvement. Since humidity produces more events than temperature, indicating it is less stable, the additional state updates generated from temperature state updates improves the overall reconstruction accuracy. A paired single-tail Student's t-tests, shows that at the 95%

Table 5.6: Microbenchmark annual energy requirement estimates for a TelosB node with a five minute sampling cycle running BN with B-CTP. Transmission time is based on logs from a 200+ node network and includes retries.

Process	Annual samples		Time (ms)		mA		mAh/year
Sense	105120	×	295	×	0.458	=	3.9
Processing	105120	×	44	×	0.182	=	0.2
Transmissions	45	×	160	×	18.920	=	0.038
Idle	105120	×	299, 231	×	0.009	=	79
Totals							83.1

Table 5.7: Microbenchmark estimates for using sense-and-send, L-SIP, and BN on a TelosB mote.

Algorithm	Percentage of transmissions (%)	Estimated energy consumption (mAh/year)	Energy reduction factor relative to LPL sense-and-send
LPL sense-and-send	100% (1×	1171.9	1.0
B-CTP L-SIP	2.80% (36×	87.5	13.4
B-CTP BN	0.02% (7000×	83.1	14.1

confidence there is only a significant improvement to the humidity accuracy ($p=0.0009$). However, it is expected multi-modal will perform no worse than single-modal.

Therefore **H5.2** is shown to be true: a multi-modal approach to BN is advantageous, reducing transmissions and improving reconstruction accuracy compared to single-modal BN.

5.4.7 Annual energy usage

The aim of the BN algorithm is to increase the lifetime of a node. Since BN reduces the number of required transmissions compared to L-SIP and sense-and-send the following hypothesis can be formed:

H5.3: *BN will have a lower energy requirement than both L-SIP and sense-and-send.*

The microbenchmarking method has been used to estimate the annual energy use of a node. This evaluation assumes a sampling interval of five minutes for all algorithms. Using microbenchmarking the TelosB node energy use for L-SIP was calculated assuming a factor of 20× packet reduction. The annual energy consumption of a TelosB node implementing BN ($t_{1/2} = 1$ month, $\varepsilon = 10\%$) is evaluated assuming BN achieves a factor of 2300× packet reduction.

Table 5.6 shows the microbenchmark calculation for a BN node using the B-CTP approach, while Table 5.7 compares the number of transmissions and energy requirement for sense-and-send, L-SIP, and BN. BN coupled with B-CTP networking provides an annual energy decrease by a factor of 14.1× compared to sense-and-send. However, BN only reduces the energy requirement of a node by a factor of

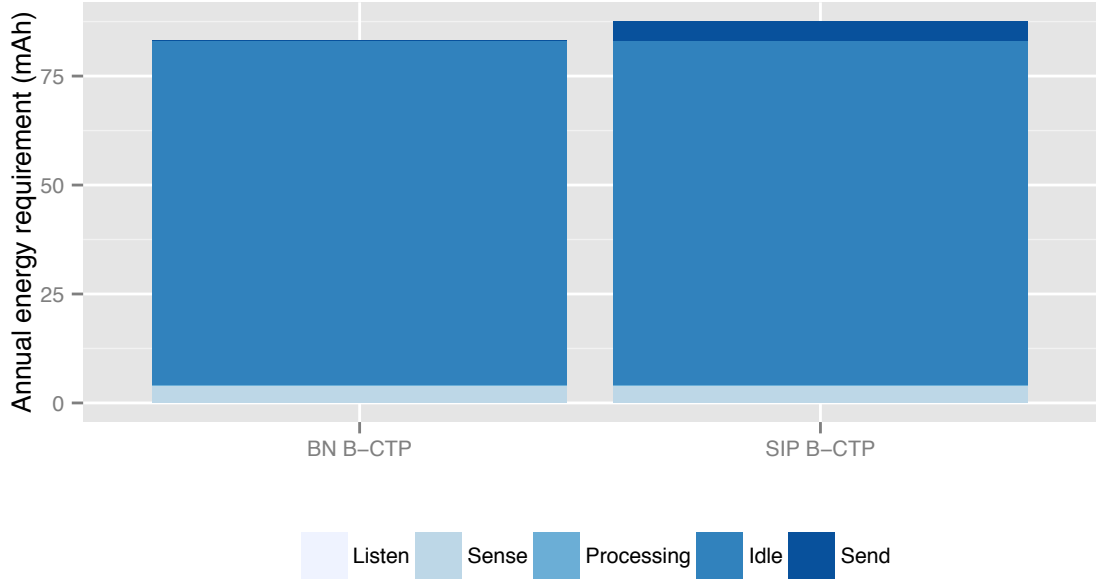


Figure 5.7: Annual energy usage for BN ($t_{1/2} = 1$ month, $\varepsilon = 10\%$) and L-SIP. It can be noted that the dominant energy use is the idle consumption of the node.

$1.05\times$ compared to L-SIP. As discussed in Section 3.6 on page 71, when transmissions are substantially reduced the limiting factor to further energy improvements is the node's hardware. In this case, Figure 5.7 shows that 95% of the node's energy requirement is for the node being in an idle state. As work towards reducing hardware energy requirements progresses, the gains by BN will have a greater effect on the node lifetime achieved.

Table 5.6 has shown **H5.3** to be true. Compared to sense-and-send, BN reduces the energy requirement of a node by a factor of $14.1\times$. However, when compared to L-SIP the decrease in annual energy use is small, a factor of $1.05\times$, despite a factor of $190\times$ reduction in transmissions. The limiting factor to the improvement of BN over L-SIP is baseline hardware energy requirements such as the idle consumption.

5.5 Summary

This chapter presented an evolution of DPSs that focuses the WSN system on transmitting the “bare necessities”—the context-specific information that end-users require to gain an understanding of the phenomena under study.

Understanding the application domain can lead to the establishment of knowledge generation techniques. The Bare Necessities algorithm combines the calculation of this knowledge on-node with event detection to drastically reduce the number of transmissions required. Though, BN imposes some penalties, such as the need for calculation to occur on the node and a slight loss in the accuracy of the resulting information compared to post-processing sense-and-send data, these penalties are slight in the context of the application considered here. BN is able to reduce the number of transmissions by a factor up to $7000\times$. The transmission reduction achieved by BN is $190\times$ greater than existing techniques such as Spanish Inquisition Protocol (SIP). Though BN reduces the number of transmissions to a much larger degree, the average RMSE in a band is a factor of $7\times$ greater than that of L-SIP and sense-and-send approaches followed by post-processing.

Combined with the B-CTP networking approach, BN leads to a reduction of annual energy use compared to sense-and-send by a factor of $14.1\times$. Compared to L-SIP, BN only reduces the energy requirement of a TelosB node by a factor of $1.05\times$. Improvements to the node hardware are required (for example, more efficient voltage conversion or, parts with lower sleep modes) to maximise the savings possible by reducing transmissions as much as BN.

This chapter has demonstrated that designing DPSs with application requirements in mind, rather than modelling the raw signal, can significantly reduce the number of transmissions required by a node. When combined with G-DPS and B-CTP node lifetimes can be significantly increased, with nodes delivering information the end-user requires, and being robust to issues which arise in real-life WSN deployments.

This chapter aimed to answer **RQ4**. To answer—yes, combining of DPS concepts with the calculation of application-level information on-node can reduce the energy requirement of a node further than the current state of the art.

The next chapter concludes this thesis, summarising the work presented and providing answers to the research questions posed at the start of this thesis.

Chapter 6

Conclusions

This thesis has investigated approaches to the design of long-lived and robust Wireless Sensor Networks (WSNs) for use in real-life deployments. The work presented in this thesis has therefore focused on the development of generalised methods and protocols for the implementation of Dual Prediction Schemes (DPSs) in WSNs deployed in the field.

WSNs are useful for a variety of applications including the monitoring of volcanic eruptions [118], soil moisture tension for irrigation management in vineyards [42], sniper fire localisation in battlefields [59], and ice quake detection on glaciers [72]. When designing WSNs, the node's energy budget, and thus lifetime, is one of the most important considerations. DPSs promise significantly extended node lifetimes by reducing the number of required transmissions. However, there remain problems with achieving robust and long-lived real-world deployments of nodes implementing DPSs. These problems include: the energy requirement of the WSN network stack, lossy networks, node failure, and accommodating the use of multiple sensing modalities. This thesis set out to answer the overarching research question: *How can WSNs nodes be designed to achieve robust and long-lived WSN deployments?*

The work presented in this thesis resulted in the following contributions to knowledge:

1. Generalised Dual Prediction Scheme (G-DPS)—a novel, generalised framework for the implementation of DPSs in real-life deployments.
2. The Backbone Collection Tree Protocol (B-CTP)—an extension to the Collection Tree Protocol (CTP) to significantly extend node lifetime via a persistent powered backbone.
3. A dual quadratic spline method to reconstruct signals which uses the gradient of the signal, and known error bounds to increase the accuracy of the reconstructed signal compared to linear interpolation, predictive model, cubic spline, and quartic spline reconstruction techniques.
4. Bare Necessities (BN)—an implementation of G-DPS that uses on-node analytics to deliver information (rather than data) to significantly reduce transmissions.

At the start of this thesis it was identified that a key barrier to the adoption of WSNs is achieving

long-lived and robust real-life deployments. Issues include: reducing the impact of transmission loss, node failure detection, accommodating multiple sensor modalities, and the energy requirement of the WSN network stack. The generalised solutions presented in this thesis enable the design and real-life deployment of a WSN which has an extended lifetime (potentially $14.1\times$ times greater than using a sense-and-send approach), is able to limit the effect of transmission failure to maximise data yield, can detect node failure, and allows the integration of multiple sensors. This is accomplished while still providing timely delivery of the information the end users are interested in.

COGENT-HOUSE, a full end-to-end open-source home environmental and energy monitoring system (see Appendix A), was developed to exemplify and evaluate the proposed solutions presented in this thesis. A total of 37 real life deployments, performed by the author and colleagues, of the COGENT-HOUSE system were performed to generate 235 evaluation datasets of periods between two week and a year. To further demonstrate the benefits of these proposed solutions the COGENT-HOUSE system was deployed and evaluated in-situ in two deployments.

This chapter is structured as follows: Section 6.1, provides answers to the research questions posed in Chapter 1. Section 6.2 proposes future research topics. Finally, Section 6.4 concludes this thesis.

6.1 Answers to research questions

This thesis has answered the following research questions:

RQ1: What features can improve the robustness of DPSs implemented in deployed WSNs?

When designing WSNs a sensing node's energy budget is one of the most important considerations. The largest consumer of energy in a WSN node is often¹ the radio, therefore reducing the use of the radio should improve node lifetime. A common approach found in the literature is that of a DPS algorithms. DPSs model the sensed signal to allow an accurate reconstruction of the signal at the sink using fewer transmissions. However, DPS algorithms presented within the literature are often lacking in the ability to handle several aspects of real world deployments. The aspects include: transmission loss, node failure detection, accommodating multiple sensor modalities, and reduction of the energy requirement of the WSN network stack.

To address the issue of robustness of DPSs in real life deployments, this thesis proposes a novel generalised framework named G-DPS. G-DPS provides: i) a multi-modal approach, ii) an acknowledgement scheme, iii) heartbeat messages, and iv) a method to calculate reconstructed data yield. G-DPS was the

¹Note that power consumption of active sensors swamps radio usage.

result of the following sub-questions:

RQ1A: Does combining multiple sensor modalities into a single model allow a greater reduction in the number of packets transmitted and improve signal reconstruction accuracy compared to compressing each stream individually?

Yes, if sensor signals are likely to change at the same time. For example, temperature and humidity.

Section 2.5 on page 34 shows that DPSs are generally designed with the aim of compressing one sensing modality. However, when designing a WSN for deployment, sensing nodes generally include multiple sensors of differing types. G-DPS allows multiple sensor's readings to be combined into a single model, named a multi-modal approach. Each transmission contains a model state for each monitored sensor. Single-modal is an existing technique where there are multiple instances of the DPS and individual sensor model states are monitored separately.

The multi-modal approach was evaluated in Section 3.5.3 on page 56. The multi-modal and single-modal approaches were applied to data traces collected from 80 nodes (sensing temperature, relative humidity, and Carbon Dioxide (CO₂)) deployed in 37 homes. The evaluation showed that when a multi-modal approach is used transmissions are reduced by a factor of 1.4 \times , signal reconstruction accuracy is improved by a factor of 2 \times , and the node energy requirement of nodes is reduced by a factor of 1.2 \times compared to single-modal.

The limitation with this approach is when events for different modalities have no relationship, then multi-modal will provide no benefit over single-modal. Larger packets will be transmitted with no saving in the number of transmissions that are required.

RQ1B: Can heartbeat messages allow detection of node failure within a user specified time period, without producing a large impact on the energy requirement of a functioning node?

Yes, it is possible to balance acceptable data loss with the expected number of additional transmissions from heartbeats, to minimise the impact on a nodes energy requirement.

DPS-based algorithms transmit at irregular and unpredictable frequencies. Therefore, a node not reporting data may be either because the node is functioning but is suppressing messages as intended or because the node has failed. The end-user is unable to distinguish between transmission suppression and node failure. The G-DPS framework therefore defines a maximum time period allowed without a transmission. If the node has not transmitted in this defined period an update transmission is forced. This transmission is called a heartbeat, and indicates the node is still functional.

The use of a heartbeat message was evaluated using Linear Spanish Inquisition Protocol (L-SIP) in Section 3.5.5 on page 61. Since heartbeat messages are considered a “still alive” message, the evaluation only considers the case of the impact of heartbeats on a functioning node. All dataset periods extracted from the 235 datasets were compressed with L-SIP using heartbeats and without heartbeats. Using a heartbeat period of 12 hours, derived from the experience of deploying and managing long term WSNs, the use of heartbeat messages was shown to increase the number of transmissions by a maximum of $1.02\times$ on a functioning node compared to using no heartbeat messages. Therefore, since heartbeats do not increase the number of transmissions, node energy requirements are not significantly increased.

The limitation with this approach is that the method requires data to already be collected from a node, without this heartbeats may become a significant energy expenditure,

RQ1C: Can the use of end-to-end acknowledgements with DPSs allow for a greater reconstructed data yield compared to an acknowledgement-less schemes?

Yes, end-to-end acknowledgements works in applications where the sample period is significantly greater than the round trip time of a transmissions

Section 2.5 on page 34 shows that multi-hop networking reliability can be very poor—evaluations of system performance, in the literature, have shown yields less than 35% [8, 10, 70]. In the case of DPSs, reliable packet delivery is important to enable an accurate reconstruction of the sensed signal. Therefore, an approach is required to allow nodes to detect when a transmission fails. To verify that both the sink and nodes are producing identical state predictions, software-level end-to-end acknowledgements are included in G-DPS to indicate when transmission and storage of a state update has failed. The sensing node’s copy of the sink state will only be updated if an acknowledgement is received from the sink.

The end-to-end acknowledgement approach was evaluated using L-SIP in Section 3.5.5 on page 61. The evaluation of end-to-end acknowledgements considered two types of networks: i) lossless and ii) lossy.

A lossless network was considered to evaluate if acknowledgements had an adverse effect on the energy consumption of a functioning system. The only difference between acknowledgement and acknowledgement-less schemes is the radio duty cycle, which translates to a potential usage energy increase. When using acknowledgements on-node evaluation shows a median radio duty cycle of $0.05\% \pm 0.006$ (160 ms) an increase of $10\times$ compared to using no acknowledgements, which gave $0.005\% \pm 0.0002$ (16 ms). Using this information the microbenchmarking approach shows a node not implementing acknowledgements uses 83 mAh/year, compared to a node implementing acknowledgements requiring 87 mAh/year—an increase of only a factor of $1.05\times$.

In a deployed network it is very unlikely that every transmission will be successful. Therefore, the use of acknowledgements were also evaluated in a lossy network. To evaluate the effect of acknowledgements in the case of transmission failure, one year of data from nine sensors with 99.98% yield (house 1, 9 traces, 3240 trace-days) was compressed using L-SIP, with the success of state update transmissions being decided by a probabilistic model. In a lossy network, for a node achieving a 35% transmission yield (as reported in the literature), the acknowledgement scheme improves signal reconstruction accuracy by a factor of $2\times$ and increases the data yield of the system up to a factor of $7\times$ when compared to an acknowledgement-less scheme. Furthermore, node annual energy requirements are only increased by a factor of $1.13\times$ when using acknowledgements.

The approach of using end-to-end acknowledgements do not work when considering applications with high-frequency transmissions. When the round trip time of sending a packet and receiving an acknowledgement is greater than the defined sample period the sensing nodes knowledge of the sink state will not be updated in time. An alternative approach to solve this is discussed in Section 6.2.

RQ2: Can the lifetime of a WSN node implementing transmission reduction approaches be increased further by using a persistent backbone network of mains powered routing nodes?

Yes, when there is a persistent or large capacity power source (e.g., a car battery) for some of the nodes are available.

This thesis evaluated the energy consumption of a node implementing DPSs along with the commonly used TinyOS network stack which is composed of CTP and Low Power Listening (LPL). DPSs promise a significant reduction in transmissions which should reduce the energy requirement of a node. However, microbenchmarking shows that when a factor of 20 transmission reduction is achieved there is only a factor of 1.05 reduction in the energy requirement. An evaluation of the energy use per node process shows that listening for packets using LPL accounts for 90% of a nodes energy requirement. Therefore an approach was required to reduce the radio duty cycle of nodes implementing DPSs. Section 3.6 on page 71 proposed the B-CTP networking topology, an extension to CTP which reduces a node's energy requirement for listening. This is achieved by reducing the number of nodes required for routing and listening for packets via the use of a persistent backbone network of mains powered nodes.

To evaluate the energy requirement for a node implementing B-CTP the annual energy requirement for a leaf node using B-CTP over a range of transmission reductions was calculated using microbenchmarking. Compared to sense-and-send using the CTP and LPL network stack, use of B-CTP reduces the energy requirement of a node by a factor of $7\times$. Considering a L-SIP node, with a factor of $20\times$ reduction

in transmissions, use of B-CTP decreases the energy requirement of a node by a factor of $13.4\times$. The microbenchmarking approach shows that with B-CTP, when transmissions are significantly reduced, 90% of the energy requirement is for the sleep state. Therefore, to reduce the energy requirement any further, improvements to the node hardware are required (for example, more efficient voltage conversion or parts with lower power sleep modes). A limitation of the B-CTP approach is that since leaf nodes perform no listening they do not receive control beacons from their neighbours therefore are unable to react to changes in the network, such as the introduction of new backbone nodes. Therefore, when introducing a new node in a formed network, leaf nodes will be required to be reset to detect the new backbone node.

B-CTP has two limitations i) it is unable to react to changes in the network, for example, introducing a new backbone node. CTP broadcasts a control beacon with routing information at an adaptive interval. ii) It relies on a multi-tiered network with the network level requiring a large capacity or persistent power source, ideally in a WSNs all nodes should have the same responsibilities.

RQ3: Can a spline-based signal reconstruction method improve the accuracy of reconstructed signals compared to piecewise linear methods, for example linear interpolation or model prediction, when using DPS algorithms such as L-SIP?

A dual quadratic spline-based signal reconstruction method outperforms both linear interpolation and model prediction.

Selection of the best method to accurately reconstruct the original signal based on the output of DPS algorithms has received little attention in the literature. Chapter 4 on page 79 examined five reconstruction methods—three spline-based methods along with model predictions and traditional linear interpolation as baselines.

To evaluate the proposed reconstruction methods, L-SIP (with a range of error thresholds from $0.1\text{ }^{\circ}\text{C}$ to $1\text{ }^{\circ}\text{C}$) was applied to 235 temperature data traces. Each trace was reconstructed using each of the methods, from which the Root Mean Squared Error (RMSE) between the original (smoothed) signal and the reconstructed signal was recorded. The evaluation of the reconstruction methods presented in Section 4.2 shows dual quadratic splines to be the most accurate in reconstructing a signal. The RMSE for dual quadratic splines is a factor of $1.3\times$ lower than the RMSE for linear interpolation and model predictions. The dual quadratic spline approach provides the lowest reconstruction error for all thresholds considered—around 25% lower than linear interpolation and around 15%–30% lower than the predictive model. In the case of transmission failure an extension to the dual quadratic spline was proposed which takes into account the number of failed transmissions. This adjusted dual quadratic spline method was shown to provide an accurate reconstruction for up to 60 failed transmissions (5 hours)—an increase by

a factor of $1.5\times$ compared to the standard dual quadratic splines method.

The limitation with this approach is that the approach does not work when significant transmission loss occurs. In the example presented in this thesis, transmission loss over 5 hours still cannot be reconstructed accurately. Considering other measures such as correlation between nodes may improve accuracy further.

RQ4: Can a combination of DPS concepts with the calculation of application-level information on-node reduce the energy requirements of a node further than the current state of the art?

Yes, when the information requirements of an application area are well known, and raw data is not required, calculation of application-level information on-node reduces the energy requirements of a node

DPSs, in general, use a model of the sensed signal to allow an accurate reconstruction of the signal at the sink using fewer transmissions. Although transmissions can be reduced through modelling the signal in this way, considerable savings are also possible at the application level. Chapter 5 on page 95 an implementation of G-DPSs which transmits the “bare necessities”—the context-specific information that end-users require to gain an understanding of the phenomena under study. The proposed BN algorithm is an implementation of the G-DPS framework that utilises on-node processing to deliver information rather than data, significantly reducing the number of transmissions a node is required to make.

Comparing the performance of BN ($t_{1/2} = 1$ month, $\varepsilon = 10\%$) with L-SIP ($\varepsilon = 0.5$ °C), and sense-and-send for one year of temperature data, BN is able to reduce the number of transmissions by a factor of up to $7000\times$ compared to sense-and-send and up to $190\times$ greater than existing techniques such as the L-SIP. Though BN reduces the number of transmissions to a much larger degree, the average RMSE in a band is a factor of $7\times$ greater than the distribution summaries calculated using L-SIP or sense-and-send data. Compared to L-SIP, BN only reduces the energy requirement of nodes by a factor of $1.05\times$, this is due to the idle consumption accounting for 90% of the energy consumption.

This thesis has also shown how the BN approach can be generalised to other domain areas, for example, Kemp *et al.* [55] have already explored the use of BN to monitor an elderly person’s activity at home using wearable devices.

The limitation with this approach is that even though the number of transmissions is reduced significantly, the relative energy reduction is only small when compared to L-SIP. Therefore, for this approach to be more beneficial, improvements to the node hardware are required to maximise the savings possible by reducing transmissions.

6.2 Future work

There are several areas of future work that can be investigated to expand on the work presented in this thesis. This section presents such areas related to improving the performance of DPS-based algorithms such as L-SIP and BN.

Buffering DPS-based algorithms during transmission failure

An alternative approach to the acknowledgement approach described in Section 3.3 on page 44, is to buffer readings when transmissions fail. One of the major benefits of the Spanish Inquisition Protocol (SIP) is the reduced storage space required to record the sensed phenomena (assuming reconstruction occurs when analysing data). The TelosB node is equipped with 1 MB of flash memory for applications which require local storage. By using SIP, a node could potentially store an estimated 1.4 years of data for five parameters (assuming five minute sampling, 90% message reduction and 70 bytes per record). If a transmission failure occurs, rather than retransmit as per the end-to-end acknowledgement scheme, a node could buffer the state until the next successful state transmission. Once a successful transmission occurs all buffered states would be sent. This potentially would increase the reconstructed data yield, however, the time synchronisation of transmitted packets and the impact of energy consumption would require investigation.

Solving the issue would allow signal events to be recorded when they actually occurred, removing any lag from end-to-end acknowledgements, and allow for a greater data yield.

Transmission loss with high-frequency signals

To verify that both the sink and nodes are using identical state estimates, software-level end-to-end acknowledgements are included in the G-DPS framework to indicate when transmissions have failed. However, the acknowledgements have only been evaluated with low-frequency signals. When using an end-to-end acknowledgement approach in an application with high-frequency signals, such as jet engine temperature monitoring with a sample rate of 10 Hz², an acknowledgement will only be received after the next sensing cycle starts. Therefore, the node has no knowledge whether or not the transmission was successful and therefore the current state of the sink. Furthermore in applications where the frequency of transmissions is high, end-to-end acknowledgements will have a higher energy requirement due to increased radio duty cycle. Therefore, an investigation is required into i) the impact of not using

²Though the temperature of a jet engine is usually stable in flight, when there is a change in engine speed jet engines have large and rapid changes in temperature, warranting the need for a high sampling frequency

acknowledgements in high-frequency applications, and, ii) alternative acknowledgement approaches. One alternative would be to make use of the buffering approach described previously in this section, effectively using flash storage as a “temporary sink”. This temporary sink would have a separate management process for the transmission of states (either individually, or as a window). The issues raised in the previous discussion of buffering (time synchronisation and the impact of additional energy consumption) require consideration in this case.

Solving this issue will allow the G-DPS framework to be expanded for additional high-frequency signal applications, such as the jet engine monitoring example. Furthermore, in applications with high frequency transmissions, the number of transmissions could potentially be reduced further than using a DPS alone, reducing the load on the network and chance of transmission failure.

Support beaconing update in B-CTP

Chapter 3 shows that B-CTP can substantially decrease the energy requirement of a node by making use of a persistent powered network. However, since leaf nodes do not perform any listening they are unable to alter routing tables when new powered nodes are introduced into the network. Therefore, to detect a new backbone node the sensing nodes will need to be reset. A method is therefore required to allow B-CTP to react to any changes to the network. One solution would be for the leaf nodes to issue a beacon request when a node attempts a transmission. This may increase the radio duty cycle, however it would allow for the node to react to changes in the network and radio environment, creating stronger links to the sink resulting in improved data yield and less failed transmissions. However, this increased radio duty cycle will come at the cost of additional energy.

Solving the issue would allow for reactive networks similar to CTP, it would all the maintainer of the deployment to include additional backbone nodes to enhance the network in problem ares, leading to increased data yields, and less node failure.

Data imputation with Gaussian Process Regression

In Chapter 4 the use of a dual quadratic spline reconstruction method to reconstruct a signal resulting from L-SIP transmissions was described. Section 4.3 considered the case of signal reconstruction during transmission failure. Evaluation shows that when coupled with sequence numbers, the dual quadratic spline method was shown to provide an accurate reconstruction of for up to 60 failed transmissions (5 hours). Beyond this 5 hour transmission loss duration, however, the accuracy of the reconstructed signal cannot be guaranteed to be within allowable error threshold limits. As an alternative, Gaussian

process regression (GPR) could be used to impute missing data. Goldsmith *et al.* [35] have demonstrate the implementation of a Virtual Sensor implemented using GPR, which combines the historical data collected by a temporarily deployed node with correlated data from a subset of permanent sensor nodes. While the approach has been demonstrated for a sense-and-send approach, the approach may work well in conjunction with L-SIP during long periods of node failure. If a GPR approach could impute missing data in a signal this could remove the need for acknowledgements in the G-DPS design, and removing the need for two way software-level communication in the network.

6.3 Generalising to other applications

This thesis has presented techniques to implement long-lived indoor environment monitoring WSNs by generating application-level information on-node. This section shows how these techniques can be implemented for other applications. The following are examples of potential feasible applications:

Human behaviour Kemp *et al.* [55] have explored the use of BN in monitoring an elderly person at home using wearable devices. The example given is to detect behavioural changes that might indicate that assistance is required. A distribution of time spent performing daily activities is generated. If at any point this distribution changes by more than 10% this indicates a change in behaviour. The number of transmissions made is equivalent to approximately one transmission per month following the initial settling period during which the subject's routine is characterised.

GPS dwell regions Rather than tracing activities, another possibility is to track the time spent at specific locations. Madan *et al.* [69] show that the more time spent at home the more likely it is that the subject is ill. Monitoring the distribution of dwell regions for a change in behaviour may suggest the subject has a problem. for example, the subject spending more time at home.

Room occupancy In the domain of built environment monitoring, the pattern of room occupancy can be useful for building control systems. Tracking a distribution of when a room is occupied, based on time and day of week, can lead to the development of effective heating strategies leading to a potential reduction in the energy requirements of a building. Tracking this distribution will only require a transmission of an update when the use of the room changes significantly.

Energy per degree day Another application for building monitoring, compares the overall (externally provided) energy usage (combining gas and electricity) with the number of heating (or cooling) degree days for the same period. Heating degree days are defined as the integral over time of the

difference between the external temperature and a base temperature. Typically $15.5\text{ }^{\circ}\text{C}$ is used as the base temperature. Roughly speaking, the ratio of heat energy used to heating degree days is proportional to the specific heat loss of the building. Generally, energy meters are situated outside of a property therefore a single node would be required to monitor external temperature and energy consumption. As previously mentioned in this chapter, the Energy per degree day metric tends to remain the same season after season, year after year. Therefore using the BN approach would significantly reduce transmissions.

In Chapter 3 the B-CTP networking approach is introduced to significantly extend the lifetime of a node. For the built environment case study presented in this thesis, a large percentage of nodes can be powered from the mains electricity. However, many WSN applications required outdoor deployments [cite, volcano, glacier, redwood, fire]. If B-CTP is to be used in these applications an approach would be required where nodes with ease of reach would be powered by a larger capacity battery. For example, in the case of detecting forest fires [39] nodes are placed at different locations within a forest, nodes deployed near fire tracks are easier to reach than those deeper in the tree line. Therefore nodes on the fire-track could be installed with a high-capacity car battery (topped-up with solar energy harvesting) which could be changed more often, whilst nodes deeper in the tree line would have 2 AA batteries as standard but can be left alone for years (depending on sensing modalities and application requirements).

6.4 Concluding remarks

This thesis set out to answer the overarching research question: *How can WSN nodes be designed to achieve robust and long-lived real-life WSN deployments?*

The G-DPS framework provides solutions to enable robust deployments of DPS algorithms in real life deployments, compensating for transmission-loss, detecting node failure, providing data yield calculations, and allowing DPS to be implemented with multiple sensors. To maximise the lifetime of a node the proposed B-CTP utilises a persistent powered backbone network to reduce a node's energy requirement for listening. On the server side, spline based reconstruction is shown to improve reconstruction accuracy compared to the commonly used linear interpolation. Finally, to further reduce transmissions and extend node lifetime the BN algorithm demonstrates that designing DPSs with application requirements in mind, rather than modelling the raw signal, can significantly reduce the number of transmissions required by a node. These techniques enable WSNs to be long-lived and robust in real-life deployments. Furthermore, the underlying approaches can be applied to existing techniques and implemented for a wide variety of

applications. Application-driven WSN design, such as BN, supported by implementation frameworks, as presented in this thesis, will not only improve performance for existing applications but also enable many new ones that were previously impractical.

References

- [1] Charu C. Aggarwal and Srinivasan Parthasarathy. Mining Massively Incomplete Data Sets by Conceptual Reconstruction. In *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '01, pages 227–232, 2001.
- [2] Kemal Akkaya and Mohamed Younis. A survey on routing protocols for wireless sensor networks. *Ad Hoc Networks*, 3(3):325–349, 2005.
- [3] I.F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. Wireless sensor networks: a survey. *Computer Networks*, 38(4):393–422, 2002.
- [4] D. Albu, J. Lukkien, and R. Verhoeven. On-Node Processing of ECG Signals. In *Proceedings of the 7th IEEE Conference on Consumer Communications and Networking*, CCNC, pages 1–5, 2010.
- [5] C. Alippi, G. Anastasi, C. Galperti, F. Mancini, and M. Roveri. Adaptive Sampling for Energy Conservation in Wireless Sensor Networks for Snow Monitoring Applications. In *Proceedings of the IEEE International Conference on Mobile Adhoc and Sensor Systems*, MASS '07, pages 1–6, Oct 2007.
- [6] M. Allen, L. Girod, R. Newton, S. Madden, D.T. Blumstein, and D. Estrin. VoxNet: An Interactive, Rapidly-Deployable Acoustic Monitoring Platform. In *Proceedings of the International Conference on Information Processing in Sensor Networks*, IPSN '08, pages 371–382, 2008.
- [7] Yasser Ammar, Aurélien Buhrig, Marcin Marzencki, Benoît Charlot, Skandar Basrour, Karine Matou, and Marc Renaudin. Wireless sensor network node with asynchronous architecture and vibration harvesting micro power generator. In *Proceedings of the 2005 joint conference on Smart objects and ambient intelligence: innovative context-aware services: usages and technologies*, pages 287–292, 2005.
- [8] G. Anastasi, M. Conti, M. Di Francesco, and V. Neri. Reliability and energy efficiency in multi-hop IEEE 802.15.4/ZigBee Wireless Sensor Networks. In *Proceedings of the IEEE Symposium on Computers and Communications 2010*, ISCC '10, pages 336–341, June 2010.
- [9] Giuseppe Anastasi, Marco Conti, Mario Di Francesco, and Andrea Passarella. Energy conservation in wireless sensor networks: A survey. *Ad Hoc Networks*, 7(3):537–568, 2009.

- [10] A. Arora, R. Ramnath, E. Ertin, P. Sinha, S. Bapat, V. Naik, V. Kulathumani, Hongwei Zhang, Hui Cao, M. Sridharan, S. Kumar, N. Seddon, C. Anderson, T. Herman, N. Trivedi, M. Nesterenko, R. Shah, S. Kulkarni, M. Aramugam, Limin Wang, M. Gouda, Young ri Choi, D. Culler, P. Dutta, C. Sharp, G. Tolle, M. Grimmer, B. Ferreira, and K. Parker. ExScal: elements of an extreme scale wireless sensor network. In *Proceedings of the 11th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications, 2005.*, pages 102–108, Aug 2005.
- [11] Manos Athanassoulis, Ioannis Alagiannis, and Stathes Hadjiefthymiades. A Multi-Criteria Message Forwarding Architecture for Wireless Sensor Networks. In *Proceedings of the 10th Pan-Hellenic Conference on Informatics*, 2005.
- [12] W. Bajwa, J. Haupt, A. Sayeed, and R. Nowak. Compressive wireless sensing. In *Proceedings of the 5th International Conference on Information Processing in Sensor Networks*, IPSN '06, pages 134–142, 2006.
- [13] Tatiana Bokareva, Wen Hu, Salil Kanhere, Branko Ristic, Travis Bessell, Mark Rutten, and Sanjay Jha. Wireless sensor networks for battlefield surveillance. In *Proceedings of the Land Warfare Conference 2006*, 2006.
- [14] Harald Bottner, J. Nurnus, A. Gavrikov, G. Kuhner, M. Jagle, C. Kunzel, D. Eberhard, G. Plescher, A. Schubert, and K. H Schlereth. New thermoelectric components using microsystem technologies. *Journal of Microelectromechanical Systems*, 13(3):414–420, 2004.
- [15] George E. P. Box, Gwilym M. Jenkins, and Gregory C. Reinsel. *Time Series Analysis: Forecasting and Control*. Wiley Series in Probability and Statistics. Wiley-Blackwell, 2008.
- [16] James Brusey, Elena Gaura, Daniel Goldsmith, and James Shuttleworth. FieldMAP: A Spatiotemporal Field Monitoring Application Prototyping Framework. *IEEE Sensors Journal*, 9(11):1378–1390, 2009.
- [17] David Campbell. Is It Still "Big Data" If It Fits In My Pocket? *Proceedings of the 37th International conference on Very Large Data Bases*, 4(11):694, 2011.
- [18] Eugene Pamba Capo-Chichi, Herve Guyennet, and Jean-Michel Friedt. K-RLE: A New Data Compression Algorithm for Wireless Sensor Network. In *Proceedings of the International conference on Sensor Technologies and Applications*, pages 502–507, 2009.

- [19] Matteo Ceriotti, Luca Mottola, Gian Pietro Picco, Amy L. Murphy, Stefan Guna, Michele Corra, Matteo Pozzi, Daniele Zonta, and Paolo Zanon. Monitoring heritage buildings with wireless sensor networks: The Torre Aquila deployment. In *Proceedings of the 2009 International Conference on Information Processing in Sensor Networks*, IPSN '09, pages 277–288, 2009.
- [20] MaggieX. Cheng, Xuan Gong, and Peng-Jun Wan. Minimum Delay Routing in Multihop Wireless Networks. In Yu Cheng, DoYoung Eun, Zhiguang Qin, Min Song, and Kai Xing, editors, *Wireless Algorithms, Systems, and Applications*, volume 6843 of *Lecture Notes in Computer Science*, pages 146–156. Springer Berlin Heidelberg, 2011.
- [21] Chun Tung Chou, R. Rana, and Wen Hu. Energy efficient information collection in wireless sensor networks using adaptive compressive sensing. In *Proceedings of the 34th IEEE Conference on Local Computer Networks*, LCN '09, pages 443–450, 2009.
- [22] D. Chu, A. Deshpande, J.M. Hellerstein, and Wei Hong. Approximate Data Collection in Sensor Networks using Probabilistic Models. In *Proceedings of the 22nd International Conference on Data Engineering*, ICDE '06, 2006.
- [23] Alex Clayton. Lonely Glaciologist: Prepping for fieldwork. Online: <http://lonelyglaciologist.wordpress.com/>.
- [24] T. Dasu and T. Johnson. *Exploratory Data Mining and Data Cleaning*. Wiley Series in Probability and Statistics. Wiley, 2003.
- [25] D.O. de Cunha, O.C.M.B. Duarte, and G. Pujolle. An efficient data transport protocol for event-driven field-estimation on sensor networks. In *Proceedings of the 9th IFIP International Conference on Mobile Wireless Communications Networks*, pages 31–35, sept. 2007.
- [26] P. Dutta, M. Feldmeier, J. Paradiso, and D. Culler. Energy Metering for Free: Augmenting Switching Regulators for Real-Time Monitoring. In *Proceedings of the International Conference on Information Processing in Sensor Networks*, IPSN '08, pages 283–294, 2008.
- [27] A. El-Hoiydi and J-D Decotignie. WiseMAC: an ultra low power MAC protocol for the downlink of infrastructure wireless sensor networks. In *Proceedings of the Ninth International Symposium on Computers and Communications*, ISCC '04, pages 244–251, 2004.
- [28] E. Fasolo, M. Rossi, J. Widmer, and M. Zorzi. In-network aggregation techniques for wireless sensor networks: a survey. *IEEE Wireless Communications*, 14(2):70–87, 2007.

- [29] Rodrigo Fonseca, Omprakash Gnawali, Kyle Jamieson, and Philip Levis. Four-Bit Wireless Link Estimation. In *Proceedings of the Sixth Workshop on Hot Topics in Networks*, HotNets, 2007.
- [30] E.I. Gaura, J. Brusey, J. Kemp, and C.D. Thake. Increasing Safety of Bomb Disposal Missions: A Body Sensor Network Approach. *Journal of IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 39(6):621–636, Nov 2009.
- [31] Elena Gaura, James Brusey, Michael Allen, Lewis Girod, and Geoffrey Challen. *Wireless Sensor Networks: Deployments and Design Frameworks*. Springer, 2010.
- [32] S. Gabriel, S. Khattab, D. Mosse, and R. Melhem. GroupBeat: Wireless sensor networks made reliable. In *Proceedings of the 5th IEEE International Conference on Mobile Ad Hoc and Sensor Systems*, MASS '08, pages 477–484, 2008.
- [33] Daniel Goldsmith. *Model-based Transmission Reduction and Virtual Sensing in Wireless Sensor Networks*. PhD thesis, Coventry University, 2013.
- [34] Daniel Goldsmith and James Brusey. The Spanish Inquisition Protocol: Model based transmission reduction for wireless sensor networks. In *IEEE Sensors*, pages 2043–2048, November 2010.
- [35] Daniel Goldsmith, Ami Preis, Michael Allen, and Andrew J. Whittle. Virtual Sensors to Improve On-Line Hydraulic Model Calibration. In *Proceedings of the 12th Water Distribution Systems Analysis Conference (WDSA)*, 2010.
- [36] U. Grunde. Non-Stationary Signal Reconstruction from Level-Crossing Samples using Akima Spline. *Electronics and Electrical Engineering*, 117(1):9–12, 2012.
- [37] C. Haas, D. Krüger, D. Pfisterer, S. Fischer, D. Dudek, P. Rothenpieler, A. Kuntz, and M. Zitterbart. FleGSens - secure area monitoring using wireless sensor networks. In *Proceedings of the International Conference on Sensor Networks, Information, and Ubiquitous Computing, ICSNIUC 2009*, 2009.
- [38] Jake Hacker, Stephen Belcher, and Richenda Connell. Beating the Heat: keeping UK buildings cool in a warming climate. UKCIP Briefing Report. Online: http://www.ukcip.org.uk/wordpress/wp-content/PDFs/Beating_heat.pdf, UKCIP, Oxford, 2005.
- [39] M. Hefeeda and M. Bagheri. Wireless sensor networks for early detection of forest fires. In *IEEE International Conference on Mobile Adhoc and Sensor Systems, 2007. MASS 2007.*, pages 1–6, Oct 2007.

- [40] W.R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-efficient communication protocol for wireless microsensor networks. In *Proceedings of the 33rd Annual Hawaii International Conference on System Sciences*, 2000.
- [41] A. Hergenröder, J. Horneber, and J. Wilke. SANDbed: A WSN Testbed for Network Management and Energy Monitoring. In *Proceedings of the 8. GI/ITG KuVS Fachgespräch "Drahtlose Sensornetze"*, pages 71–73, 2009.
- [42] Mark Holler. High Density, Multiple Depth Wireless Soil Moisture Tension Measurements for Irrigation Management. Technical report, Camalie Vineyards, 2008.
- [43] V. Huang and M.K. Javed. Semantic Sensor Information Description and Processing. In *Proceedings of the Second International Conference on Sensor Technologies and Applications, SENSORCOMM '08*, pages 456–461, 2008.
- [44] Chalermek Intanagonwiwat, Ramesh Govindan, Deborah Estrin, John Heidemann, and Fabio Silva. Directed diffusion for wireless sensor networking. *IEEE/ACM Transactions on Networking*, 11(1):2–16, 2003.
- [45] Y.G. Iyer, S. Gandham, and S. Venkatesan. Stcp: a generic transport layer protocol for wireless sensor networks. In *Proceedings of the 14th International Conference on Computer Communications and Networks, ICCCN 2005*, pages 449–454, Oct 2005.
- [46] Ankur Jain, Edward Y. Chang, and Yuan-Fang Wang. Adaptive stream resource management using Kalman Filters. In *Proceedings of the ACM SIGMOD international conference on Management of data, SIGMOD '04*, pages 11–22, 2004.
- [47] Sai Ji, Yajie Sun, and Jian Shen. A method of data recovery based on compressive sensing in wireless structural health monitoring. *Mathematical Problems in Engineering*, pages 1–9, January 2014.
- [48] Xiaofan Jiang, Prabal Dutta, David Culler, and Ion Stoica. Micro power meter for energy monitoring of wireless sensor networks at scale. In *Proceedings of the 6th International Conference on Information Processing in Sensor Networks, IPSN '07*, pages 186–195, 2007.
- [49] G.L. Johnson. *A Low Cost Data-logging System*. Building research notes. Division of Building Research, National Research Council of Canada, 1980.

- [50] Michael J. Johnson, Zuowei Shen, and Yuhong Xu. Scattered data reconstruction by regularization in b-spline and associated wavelet spaces. *Journal of approximation theory*, 159(2):197–223, 2009.
- [51] S.J. Julier and J.K. Uhlmann. Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3):401–422, March.
- [52] Prajakta S. Kalekar. Time series Forecasting using Holt-Winters Exponential Smoothing. Online: http://www.it.iitb.ac.in/~praj/acads/seminar/04329008_ExponentialSmoothing.pdf, December 2004.
- [53] A. Karagiannis, D. Vouyioukas, and P. Constantinou. Energy consumption measurement and analysis in wireless sensor networks for biomedical applications. In *Proceedings of the 4th International Conference on Pervasive Technologies Related to Assistive Environments*, PETRA '11, 2011.
- [54] Mumraiz Khan Kasi, Annika Hinze, Catherine Legg, and Steve Jones. SEPSen: semantic event processing at the sensor nodes for energy efficient wireless sensor networks. In *Proceedings of the 6th ACM International Conference on Distributed Event-Based Systems*, DEBS '12, pages 119–122, 2012.
- [55] John Kemp, Elena Gaura, Ramona Rednic, and James Brusey. Long-Term Behavioural Change Detection through Pervasive Sensing. In *Proceedings of the 14th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing*, SNPD, pages 629–634, 2013.
- [56] Chester Kim. Measuring Power Consumption of CC2530 With Z-Stack. Technical report, Texas Instruments, 2012.
- [57] Kevin Klues, Vlado Handziski, Chenyang Lu, Adam Wolisz, David Culler, David Gay, and Philip Levis. Integrating concurrency control and energy management in device drivers. In *Proceedings of Twenty-first ACM SIGOPS Symposium on Operating Systems Principles*, SOSP '07, pages 251–264, 2007.
- [58] A. Kopke and A. Wolisz. Measuring the Node Energy Consumption in USB Based WSN Testbeds. In *Proceedings of the 28th International Conference on Distributed Computing Systems Workshops*, ICDCS '08., pages 333–338, 2008.

- [59] Patrick Kuckertz, J. Ansari, J. Riihijarvi, and Petri Mahonen. Sniper Fire Localization using Wireless Sensor Networks and Genetic Algorithm based Data Fusion. In *Proceedings of the IEEE Military Communications Conference, MILCOM 2007.*, 2007.
- [60] O. Landsiedel, K. Wehrle, and S. Götz. Accurate prediction of power consumption in sensor networks. In *Proceedings of the The Second IEEE Workshop on Embedded Networked Sensors, 2005, EmNetS-II*, pages 37–44, May 2005.
- [61] Marc Langheinrich. Privacy by Design—Principles of Privacy-Aware Ubiquitous Systems. In Gregory D. Abowd, Barry Brumitt, and Steven Shafer, editors, *UbiComp 2001: Ubiquitous Computing*, volume 2201 of *Lecture Notes in Computer Science*, pages 273–291. Springer Berlin Heidelberg, 2001.
- [62] I. Lazaridis and S. Mehrotra. Capturing sensor-generated time series with quality guarantees. In *Proceedings of the 19th International Conference on Data Engineering*, pages 429–440, 2003.
- [63] Yann-Aël Le Borgne, Silvia Santini, and Gianluca Bontempi. Adaptive model selection for time series prediction in wireless sensor networks. *ACM Signal Processing*, 87(12):3010–3020, December 2007.
- [64] Philip Levis, Nelson Lee, Matt Welsh, and David Culler. TOSSIM: Accurate and Scalable Simulation of Entire TinyOS Applications. In *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems, SenSys '03*, pages 126–137, 2003.
- [65] Kris Lin, Jennifer Yu, Jason Hsu, Sadaf Zahedi, David Lee, Jonathan Friedman, Aman Kansal, Vijay Raghunathan, and Mani Srivastava. Helimote: enabling long-lived sensor networks through solar energy harvesting. In *Proceedings of the 3rd international conference on Embedded networked sensor systems, SenSys '05*, pages 309–309, 2005.
- [66] Chong Liu, Kui Wu, and Min Tsao. Energy efficient information collection with the ARIMA model in wireless sensor networks. In *Proceedings of the IEEE Global Telecommunications Conference, GLOBECOM '05*, 2005.
- [67] Konrad Lorincz, Bor-rong Chen, Jason Waterman, Geoff Werner-Allen, and Matt Welsh. Resource aware programming in the Pixie OS. In *Proceedings of the 6th ACM conference on Embedded network sensor systems, SenSys '08*, pages 211–224, 2008.

- [68] R. MacRuairi, M.T. Keane, and G. Coleman. A wireless sensor network application requirements taxonomy. In *Second International Conference on Sensor Technologies and Applications, 2008.*, pages 209–216, Aug 2008.
- [69] Anmol Madan, Manuel Cebrian, David Lazer, and Alex Pentland. Social sensing for epidemiological behavior change. In *Proceedings of the 12th ACM international conference on Ubiquitous computing, Ubicomp '10*, pages 291–300, 2010.
- [70] Samuel Madden. Intel Lab Data. Online: <http://db.lcs.mit.edu/labdata/labdata.html>, 2010.
- [71] Samuel Madden, Michael J. Franklin, Joseph M. Hellerstein, and Wei Hong. TAG: a Tiny AGgregation service for ad-hoc sensor networks. In *Proceedings of twenty-first ACM Specialist Interest Group On Operating Systems symposium on operating systems principles*, pages 131–146, 2002.
- [72] Kirk Martinez, Royan Ong, and Jane Hart. Glacsweb: a sensor network for hostile environments. In *Proceedings of the First IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks*, 2004.
- [73] Gaurav Mathur, Peter Desnoyers, Paul Chukiu, Deepak Ganesan, and Prashant Shenoy. Ultra-low power data storage for sensor networks. *ACM Transactions on Sensor Networks*, 5(4), 2009.
- [74] F. Mattern and C. Floerkemeier. From the Internet of Computers to the Internet of Things. *From Active Data Management to Event-Based Systems and More: Papers in Honor of Alejandro Buchmann on the Occasion of His 60th Birthday*, 1:242–259, 2010.
- [75] A. Milenkovic, M. Milenkovic, E. Jovanov, D. Hite, and D. Raskovic. An environment for runtime power monitoring of wireless sensor network platforms. In *Proceedings of the Thirty-Seventh Southeastern Symposium on System Theory, SSST '05.*, pages 406–410, 2005.
- [76] David Moss and Philip Levis. BoX-MACs: Exploiting Physical and Link Layer Boundaries in Low-Power Networking. Technical report, Stanford, 2008.
- [77] Hoang Anh Nguyen, A. Forster, D. Puccinelli, and S. Giordano. Sensor node lifetime: An experimental study. In *Proceedings of the 2011 IEEE International Conference on Pervasive Computing and Communications Workshops, PERCOM '11*, pages 202–207, 2011.
- [78] P.M. Khilar Nishi Yadav. Cluster based distributed diagnosis in manet. *International Journal of Computer and Information Technology (IJCIT)*, 1:54 – 60, September 2012.

- [79] Jo Nurse, Ben Smith, Margaret Hanson, Rebecca Wilkinson, Corrine Thompson, and Julia Mason. Fuel Poverty and Health Strategy. Online: <http://www.southamptonhealth.nhs.uk/EasySiteWeb/getresource.axd?AssetID=2474&type=full&servicetype=Attachment>, May 2004.
- [80] O. Ojetola, E.I. Gaura, and J. Brusey. Fall detection with wearable sensors—safe (smart fall detection). In *Proceedings of the 7th International Conference on Intelligent Environments*, IE '07, pages 318–321, July 2011.
- [81] Chris Olston, Boon Thau Loo, and Jennifer Widom. Adaptive precision setting for cached approximate values. In *Proceedings of the 2001 ACM SIGMOD international conference on Management of data*, SIGMOD '01, pages 355–366, 2001.
- [82] Paritosh Padhy, Rajdeep K. Dash, Kirk Martinez, and Nicholas R. Jennings. A utility-based sensing and communication model for a glacial sensor network. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, AAMAS '06, pages 1353–1360, 2006.
- [83] S. Patel, Bor rong Chen, T. Buckley, R. Rednic, D. McClure, D. Tarsy, L. Shih, J. Dy, M. Welsh, and P. Bonato. Home monitoring of patients with parkinson's disease via wearable technology and a web-based application. In *Proceedings of the 2010 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, EMBC '10, pages 4411–4414, Aug 2010.
- [84] L. M. Rodríguez Peralta, L. M. Pestana Leão de Brito, B. A. Teixeira Gouveia, D. J. Gomes de Sousa, and C. Da Silva Alves. The WISE-MUSE Project: Environmental Monitoring and Controlling of Museums based on Wireless Sensors Networks. *EJSE Special Issue: Sensor Network on Building Monitoring: from theory to real application*, 1:46–57, 2009.
- [85] J. Polastre, R. Szewczyk, and D. Culler. Telos: enabling ultra-low power wireless research. In *Proceedings of the Fourth International Symposium on Information Processing in Sensor Networks*, IPSN '05, pages 364–369, April 2005.
- [86] Joseph Polastre, Jason Hill, and David Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the 2nd international conference on Embedded networked sensor systems*, SenSys '04, pages 95–107, 2004.
- [87] Joseph Polastre, Robert Szewczyk, Alan Mainwaring, David Culler, and John Anderson. Analysis of wireless sensor networks for habitat monitoring. In *Wireless sensor networks*, chapter Analysis

- of wireless sensor networks for habitat monitoring, pages 399–423. Kluwer Academic Publishers, 2004.
- [88] G. J. Pottie and W. J. Kaiser. Wireless integrated network sensors. *Communications of the ACM*, 43(5):51–58, May 2000.
- [89] W.B. Pottner, L. Wolf, J. Cecilio, P. Furtado, R. Silva, J.S. Silva, A. Santos, P. Gil, A. Cardoso, Z. Zinonos, J. do O, B. McCarthy, J. Brown, U. Roedig, T. O'Donovan, C.J. Sreenan, Z. He, T. Voigt, and A. Jugel. WSN evaluation in industrial environments first results and lessons learned. In *Proceedings of the 2011 International Conference on Distributed Computing in Sensor Systems and Workshops*, DCOSS '11, pages 1–8, June 2011.
- [90] V. Raghunathan, C. Schurgers, Sung Park, and M.B. Srivastava. Energy-aware wireless microsensor networks. *IEEE Signal Processing Magazine*, 19(2):40–50, March 2002.
- [91] Bhaskaran Raman and Kameswari Chebrolu. Sensor networks: A critique of "sensor networks" from a systems perspective. *SIGCOMM Comput. Commun. Rev.*, 38(3):75–78, July 2008.
- [92] Nithya Ramanathan, Kevin Chang, Rahul Kapur, Lewis Girod, Eddie Kohler, and Deborah Estrin. Sympathy for the sensor network debugger. In *Proceedings of the 3rd international conference on Embedded networked sensor systems*, SenSys '05, pages 255–267, New York, NY, USA, 2005. ACM.
- [93] U. Raza, A. Camerra, A.L. Murphy, T. Palpanas, and G.P. Picco. What does model-driven data acquisition really achieve in wireless sensor networks? In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications*, PERCOM '12, pages 85–94, 2012.
- [94] Usman Raza, Amy L. Murphy, and Gian Pietro Picco. Pushing a standard wireless sensor network stack for ultra-low data rates. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, SenSys '13, pages 49:1–49:2, 2013.
- [95] Adi Mallikarjuna Reddy. Wsn.svg. Online: <http://commons.wikimedia.org/wiki/File:WSN.svg>.
- [96] Sasank Reddy, Min Mun, Jeff Burke, Deborah Estrin, Mark Hansen, and Mani Srivastava. Using Mobile Phones to Determine Transportation Modes. *ACM Transactions on Sensor Networks*, 6(2):13:1–13:27, 2010.
- [97] Ramona Rednic, Elena Gaura, and James Brusey. Classact: Accelerometer-based real-time activity classifier. In *Proceedings of the 2nd Wireless Sensing Interest Group Showcase*, pages 21–25, 2009.

- [98] Alex Rogers, Siddhartha Ghosh, Reuben Wilcock, and Nicholas R. Jennings. A scalable low-cost solution to provide personalised home heating advice to households. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 1:1–1:8, 2013.
- [99] K. Römer and F. Mattern. The design space of wireless sensor networks. *Wireless Communications, IEEE*, 11(6):54–61, Dec 2004.
- [100] J. Ryder, B. Longstaff, S. Reddy, and D. Estrin. Ambulation: A Tool for Monitoring Mobility Patterns over Time Using Mobile Phones. In *Proceedings of the International Conference on Computational Science and Engineering*, volume 4 of *CSE '09*, pages 927–931, 2009.
- [101] Silvia Santini and Kay Römer. An Adaptive Strategy for Quality-Based Data Reduction in Wireless Sensor Networks. In *Proceedings of the 3rd IEEE International Conference on Networked Sensing Systems*, INSS '06, 2006.
- [102] Thomas Schmid, Henri Dubois-ferrière, and Martin Vetterli. Sensorscope: Experiences with a wireless building monitoring sensor network. In *Proceedings of the First Workshop on Real-World Wireless Sensor Networks*, REALWSN '05, 2005.
- [103] A. Sharaf, Jonathan Beaver, Alexandros Labrinidis, and K. Chrysanthis. Balancing energy efficiency and quality of aggregate data in sensor networks. *The Very Large DataBase Journal*, 13(4):384–403, December 2004.
- [104] Adam Silberstein, Gavino Puggioni, Alan Gelfand, Kamesh Munagala, and Jun Yang. Suppression and failures in sensor networks: a bayesian approach. In *Proceedings of the 33rd international conference on Very large data bases*, VLDB '07, pages 842–853, 2007.
- [105] A. Singh, D. Budzik, W. Chen, M.A. Batalin, M. Stealey, H. Borgstrom, and W.J. Kaiser. Multiscale sensing: A new paradigm for actuated sensing of high frequency dynamic phenomena. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 328–335, Oct 2006.
- [106] Feng-Tso Sun, C. Kuo, and M. Griss. PEAR: Power efficiency through activity recognition (for ECG-based sensing). In *Proceedings of the 5th International Conference on Pervasive Computing Technologies for Healthcare*, PervasiveHealth '11, pages 115–122, 2011.

- [107] N. Tezcan, E. Cayirci, and M.U. Caglayan. End-to-end reliable event transfer in wireless sensor networks. In *Proceedings of the 15th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, volume 2 of *PIMRC 2004*, pages 989–994, 2004.
- [108] J.A Tropp and AC. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on Information Theory*, 53(12):4655–4666, Dec 2007.
- [109] Edward R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 1983.
- [110] Daniela Tulone and Samuel Madden. An energy-efficient querying framework in sensor networks for detecting node similarities. In *Proceedings of the 9th ACM international symposium on Modeling analysis and simulation of wireless and mobile systems*, MSWiM '06, pages 191–300, 2006.
- [111] Daniela Tulone and Samuel Madden. PAQ: Time Series Forecasting for Approximate Query Answering in Sensor Networks. In Kay Römer, Holger Karl, and Friedemann Mattern, editors, *Wireless Sensor Networks*, volume 3868 of *Lecture Notes in Computer Science*, pages 21–37. Springer Berlin Heidelberg, 2006.
- [112] Michael Unser and Thierry Blu. Generalized Smoothing Splines and the Optimal Discretization of the Wiener Filter. *IEEE Transactions on Signal Processing*, 53(6):2146–2159, June 2005.
- [113] Tijs van Dam and Koen Langendoen. An adaptive energy-efficient mac protocol for wireless sensor networks. In *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems*, SenSys '03, pages 171–180, 2003.
- [114] Ovidiu Vermesan, Peter Friess, Patrick Guillemin, Sergio Gusmeroli, Harald Sundmaeker, Alessandro Bassi, Ignacio Soler Jubert, Margaretha Mazura, Mark Harrison, Markus Eisenhauer, and Pat Doody. Internet of Things Strategic Research Roadmap. Online: http://www.internet-of-things-research.eu/pdf/IoT_Cluster_Strategic_Research_Agenda_2011.pdf, June 2011. Cluster of European Research Projects on the Internet of Things.
- [115] F. Viani, M. Salucci, P. Rocca, G. Oliveri, and A. Massa. A multi-sensor wsn backbone for museum monitoring and surveillance. In *Proceedings of the 6th European Conference on Antennas and Propagation*, EUCAP '12, pages 51–52, March 2012.
- [116] Chieh-Yih Wan, Andrew T. Campbell, and Lakshman Krishnamurthy. PSFQ: A Reliable Transport Protocol for Wireless Sensor Networks. In *Proceedings of the 1st ACM International Workshop on Wireless Sensor Networks and Applications*, WSNA '02. ACM, 2002.

- [117] Chuang Wang, Taiming Feng, Jinsook Kim, Guiling Wang, and Wensheng Zhang. Catching packet droppers and modifiers in wireless sensor networks. *IEEE Transactions on Parallel and Distributed Systems*, 23(5):835–843, May 2012.
- [118] Geoff Werner-Allen, Jeffrey B. Johnson, Mario Ruiz, Jonathan Lees, and Matt Welsh. Monitoring Volcanic Eruptions with a Wireless Sensor Network. In *Proceedings of the Second European Workshop on Wireless Sensor Networks*, 2005.
- [119] Geoffrey Werner-Allen, Stephen Dawson-Haggerty, and Matt Welsh. Lance: optimizing high-resolution signal collection in wireless sensor networks. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, SenSys '08, pages 169–182, 2008.
- [120] B. Widrow, J.M. McCool, M. Larimore, and C.R. Johnson. Stationary and nonstationary learning characteristics of the LMS adaptive filter. *Proceedings of the IEEE*, 64(8):1151–1162, 1976.
- [121] Alex Wun and Milenko Petrovic. A system for semantic data fusion in sensor networks. In *Proceedings of the inaugural international conference on Distributed event-based systems*, DEBS 2007, pages 75–79. ACM, 2007.
- [122] Jerry Zhao and Ramesh Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems*, SenSys '03, pages 1–13, New York, NY, USA, 2003. ACM.

Appendix A

Cogent House: A WSN for the Built Environment

This appendix details the concept and design of a Wireless Sensor Network (WSN) to monitor a building's environment conditions and energy consumption. COGENT-HOUSE is a full end-to-end open-source home environmental and energy WSN monitoring system which aims to meet this purpose. Cogent-House is designed to gather sensor data (such as temperature, humidity, electricity usage, gas usage, heat metering, CO₂, VOC, etc) from all types of buildings and to transmit that data to a central database where it can be viewed with a web-browser.

This appendix is structured as follows: Section A.1 provides a high level conceptual description of the system. Section A.2 describes the aims and requirements of the system. Section A.3 describes the design choices of the system. Section A.4 gives a high-level overview of the system, which is presented in more detail regarding the hardware in Section A.5, and software in Section A.6. Finally, Section A.7 provides an evaluation of the system.

A.1 System concept

Figure A.1 on the next page shows a conceptional view of the system. The following stages are performed:

Sense Environmental and energy data such as temperature, relative humidity, Carbon Dioxide (CO₂), Volatile Organic Compounds (VOC), air quality, electricity consumption, and heat metering is gathered.

Send Data from all sensing points is wireless sent to a central store. Due to the ad-hoc and short-term nature of the deployments wired transmissions cannot be used.

Store The data in the sense component is stored on-site for collection post the deployment period.

Off-Site Transfer Datasets are transferred off site for storage and data analysis. This is through either

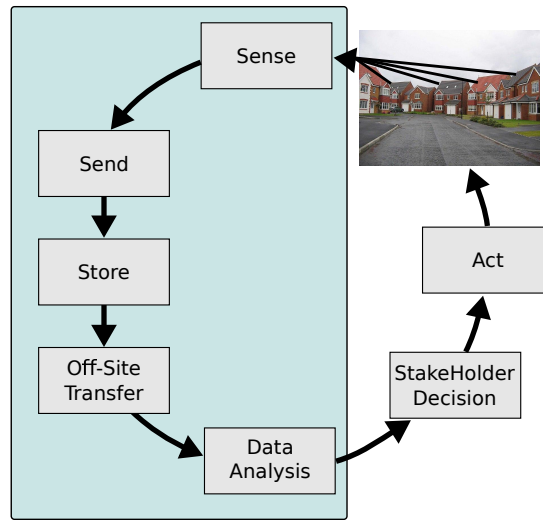


Figure A.1: Conceptual design of building monitoring system. The blue square indicates the autonomous system components. Outside the blue square are tasks undertaken by the stakeholders.

physically collecting systems, or data transfer over 3G.

Data Analysis The data analysis stage converts the raw sensor data into information metrics. During the data analysis, data cleansing techniques are applied to the raw data to create a clean dataset, from this the data samples are ran through various processing techniques to create the metrics. The analyses produced at this stage are presented to the stakeholder for review.

Stakeholder Decision From the reported analysis the stakeholder makes a decision on if and how a property should be altered. For example, if the stakeholder knows mould is likely to be present in the bathroom and the data backs this up, the stakeholder may decide to install better ventilation.

Act At this stage all decisions have been made and any remedial changes are performed on the property. The properties are then either monitored again to check whether the changes have made a positive change.

A.2 System aims

From the system concept, a number of high level requirements for the system have been identified:

1. Loss of individual sensing devices (due to power or communication reasons) should not impact other devices within the network.
2. The system should be portable and reusable, so that it can be flexibly redeployed.

3. The system should be robust to failure such that the system can be left unattended while in deployment.
4. The lifetime of the node should be a minimum of six months.

From these requirements a list of main features have been identified as key components of the system:

1. Integration of sensors to monitor the following parameters: temperature, relative humidity, air quality (CO₂ and VOC), electricity usage, and gas consumption.
2. Ability to form ad-hoc wireless networks for the transportation of the sensed data for central storage.
3. The sensing devices should be configurable in terms of sensing modalities and sample frequency.
4. The system needs sufficient reports, error checking and fault tolerance to allow faults to be detected.
5. The energy requirement of the sensing device should be minimised by limiting the radio usage.

Given the nature of the built environment and the requirements as outlined, WSNs are a suitable technology to be used as part of the system design for several reasons:

1. Protected, indoors deployment environment.
2. Communication ranges are short, as servers and router / gateways can be situated in close proximity of the deployed sensing nodes within buildings and networks tend to be dense.
3. Mains power proximity to ensure long-lived deployments.
4. Data rates are low, given the slow changing nature of most environmental parameters, leading to low network traffic and the use of low power techniques.
5. Wide availability, at low cost, of appropriate micro sensors for physical phenomena of interest.

A.3 System design

Figure A.2 shows an overview of how the system would be used. First, based on business needs a site for monitoring would be identified. The systems is deployed for a minimum of two weeks up to a number of years, depending on application needs. A WSN collects environmental and energy data which is transmitted to and stored on a local server. This local server is connected to the Internet

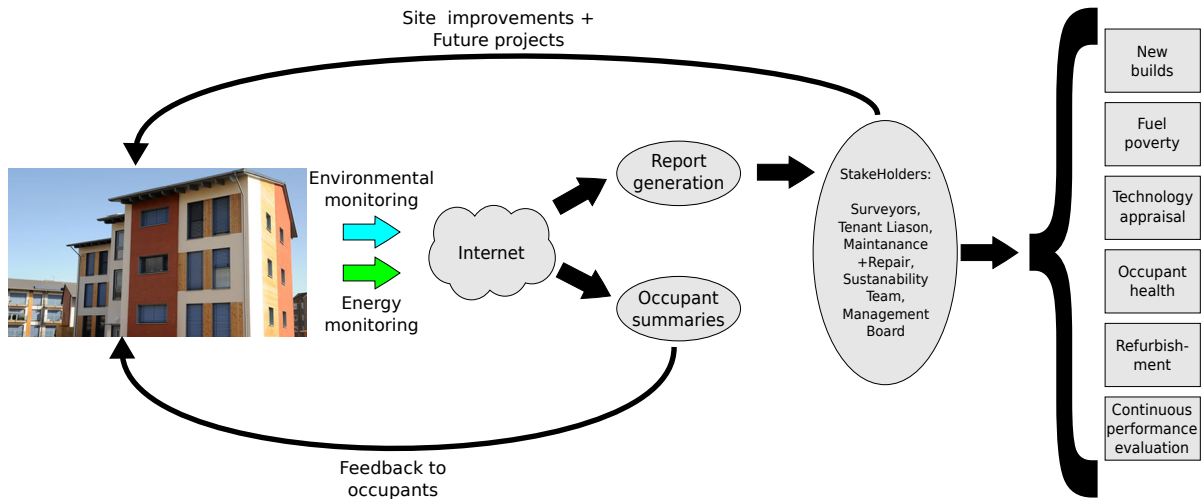


Figure A.2: System flow overview.

and the collected data is pushed to a remote server¹. Periodically the data is analysed to produce two performance reports, which are:

Site Performance This report is passed on to a variety of differing stakeholders, who use the report two-fold 1) take remedial action on the monitored site, 2) help direct future policies and business decisions based on the results of the report.

Occupant Summary This reports to occupants how well they are controlling the environment in their property and how much energy they are consuming.

A.4 Cogent-House overview

COGENT-HOUSE is the result of development work to meet the system design and concept. COGENT-HOUSE is an open-source, low-cost, wireless building monitoring solution that supports a wide variety of sensor types, provides a web-interface, e-mail alerts, and can push data to the cloud. Unlike many other building monitoring solutions, the COGENT-HOUSE solution is open source. All source code is available from: <https://code.google.com/p/cogent-house>. This has a number of benefits to the end-user:

1. the software and any upgrades are freely available;
2. any part of the solution can be extended in non-standard ways (although the resulting system must remain 'open-source');

¹If an Internet connection is not available, data is transmitted to the remote server when the deployment is ended and the local server collected

3. the end-user is not ‘locked-in’ to the solution and can readily export / convert data to other solutions;
4. multiple hardware providers exist guaranteeing a low-cost solution.

COGENT-HOUSE employs WSN mesh networking, which means that the network is not limited by the distance to the base station and repeater stations are, generally, not required. Instead, each node can act as a repeater. Furthermore, the network automatically adapts its routing tree to best suit changing environmental conditions, node movement and node loss. Since COGENT-HOUSE is wireless, there are no wiring costs and ‘live’ sensor data is always available. The ‘live’ aspect has three key benefits:

1. The ‘health’ and data yield of the system can be monitored during the deployment;
2. It is not necessary to recover the system before making use of the data;
3. It is possible to automate processes based on the sensor measurements (for example, to send alerts when sensor measurement values exceed certain bounds).

COGENT-HOUSE has been proven to easy and quick to deploy. Tests with end-users demonstrate that the system can be deployed in a typical residential home within one hour.

The next two sections provide an overview of the hardware (Section A.5) and software (Section A.6) components.

A.5 System hardware description

This section provides an overview of the hardware of the COGENT-HOUSE system.

A.5.1 Local Server

Either a Raspberry Pi or Ubuntu PC can be used as a local server. The local server is responsible for the storage of data received from the WSN.

This server can optionally be connected to the Internet (via ethernet or 3G dongle) to support ‘pushing’ the data to the cloud (see Section A.6.4). The key advantage of this approach is to enable aggregation of data from a number of parallel deployments. Furthermore, it means that the data can be more easily accessed over the Internet. Providing an Internet connection allows for remote access allowing debugging in-situ.

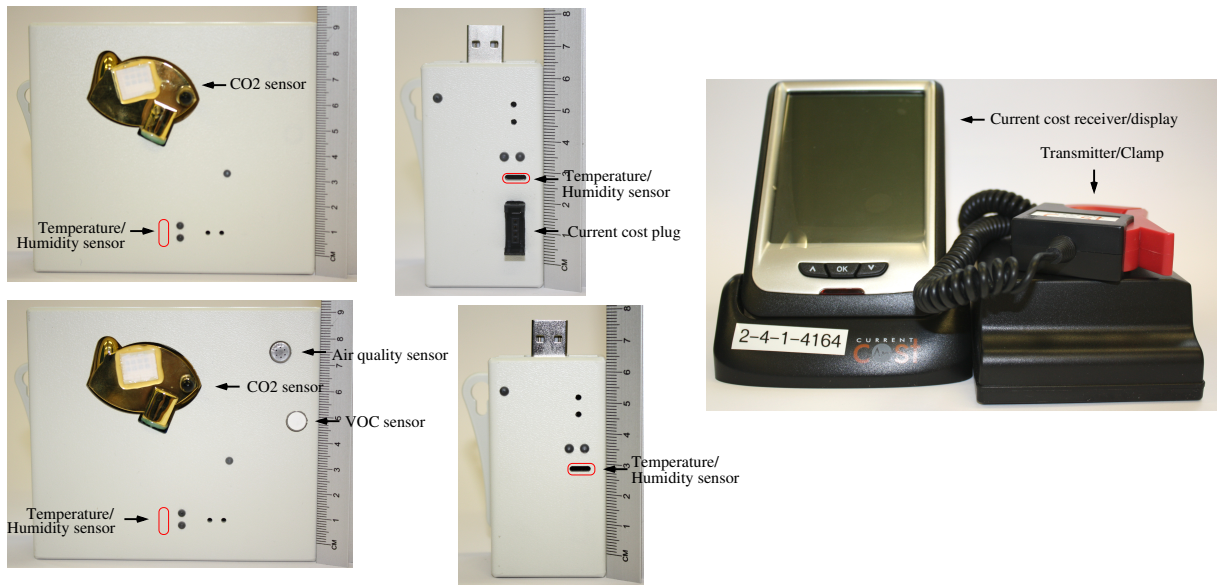


Figure A.3: Node hardware platform (From 1st Row L-R: CO₂ Node, Extension Node, Current Cost Unit. 2nd Row L-R: Air Quality Node, Base Node)

A.5.2 Remote Server

Currently the Cogent Computing Applied Research Centre provide ‘cloud’ aggregation services that allow any number of deployments to be monitored simultaneously. Data is transferred from the Local Server through the use of a developed push synchronisation / aggregation software through as Internet connection.

A.5.3 Base node

The system comprises of a number of differing node types. All nodes are built around the TelosB hardware platform, which is an all in one communication (CC2420 Radio) and processing platform (MSP430 F1611 Microcontroller) which includes integrated temperature and relative humidity sensors (Sensirion SHT11). The TelosB node was selected as the basis of the system due to its low energy consumption, small form factor, ease of integration, and proven use in research.

A.5.4 Sink node

The sink node is attached to the Local Server, this node is a base node programmed with the Sink Node Software (Appendix A.6.2)

A.5.5 Sensor Extensions

The base node can be extended to implement a variety of additional sensing modalities. A number of in-house sensor-boards have been developed to interface additional sensors to the TelosB's microcontroller. Additional sensing includes:

A.5.5.1 Electrical power sensing

To understand occupant electricity usage and consumption an extension to the basic node is provided to interface with the Current Cost EnviR electricity clamp display unit. This sensor requires AC mains power (power converter is not shown in the diagram). A further variant supports interfacing to the OptiSmart pulse reader. This pulse reader can be attached to electricity meters with optical pulse output. This approach provides greater accuracy.

A.5.5.2 Air Quality

To monitor the air quality of a room, CO₂, and VOC sensors are interfaced to the base node. These sensors require additional power and thus must be AC mains powered to operate their full sensor set. The sensor system duty cycles the sensor to minimise the effect of drift and reduce overall power consumption. In the case of AC mains power loss, the node can continue to operate (without CO₂, VOC or AQ sensing) using a battery backup.

A.5.5.3 Gas

Occupant gas usage and consumption sensing is provided by integration with the Magpeye Opto / Ferro. These devices provide an ATEX compliant solution that is suitable for most commonly available gas meters.

A.5.5.4 Heat metering

Heat metering can be used to disaggregating heating energy consumption between thermal environment heating and hot water heating. Interfacing is provided for heat meters that provide an electrical pulse output. has been tested with the Zenner Zelsius heat meter.

A.6 System software description

This section describes the software components of the COGENT-HOUSE system.

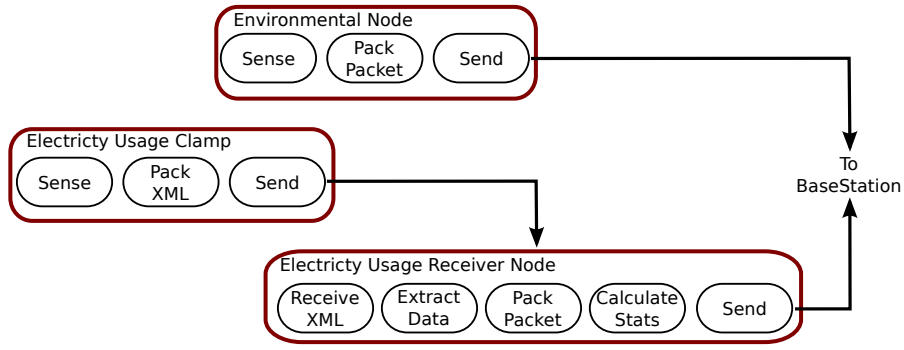


Figure A.4: WSN data flow overview

A.6.1 Node Software

Figure A.4 shows the data flow of the node software at each sampling interval a node will i) sample it's interfaced sensors, ii) process sensor data (For example, parse the xml received from the current cost unit) iii) pack sensor samples into a data packet, and iv) transmit the data packet wirelessly to the sink.

The COGENT-HOUSE nodes have been implemented using TinyOS. TinyOS has been selected as the development environment due to the ease of use, proven track record, readily available software components, and the community support available. Mesh networking is achieved through commonly used network stack comprised of Collection Tree Protocol (CTP) and Low Power Listening (LPL) provided by TinyOS

A.6.2 Sink node Software

The Sink Node has the role of wirelessly receiving packets from each individual sensor node, and forwarding these messages through serial-over-USB to the Local Server which is read by the base logging software.

A.6.3 Base logging

The Base Logging component receives sensor data from the sink node, extracts the relevant data, and 'logs' it to a MySQL database.

A.6.4 Push synchronisation / aggregation

The 'push synchronisation' system provides a way to synchronise between a local and remote data to 'push' logged data from the local to the remote. This system can be used in one of two ways:

1. if the local system is connected to the Internet, ‘push’ can be run periodically to keep the remote up-to-date, or
2. if the local system is not connected to the Internet during deployment, ‘push’ can be run when the system is returned from the deployment to upload the gathered data.

A key benefit of the ‘push’ system is to bring logged data from multiple deployments into a single database. Furthermore, pushing logged data to a database ‘in the cloud’ simplifies access to that data (it need no longer be hidden behind a firewall, for example).

A.6.5 Web interface

The web interface to the server provides a way to interact with the monitoring system and to see live sensor readings as they are gathered. The key features of the web interface to the server are:

- The web interface supports deployment allowing discovery of nodes as they are installed in the house and configuring of them.
- The health of the monitoring system can be displayed, including data yield, battery levels, network tree, etc.
- Main data types supported (i.e. temperature, humidity, CO₂, VOC, electricity) can be displayed live. The graph display supports zooming in or out (from 1 hour to 3 months).
- Logged data can be exported to the CSV (comma separated values) file format which can then be loaded by a statistical package such as Microsoft Excel or R.

A.6.6 Calibration support

The system supports per sensor calibration post-collection. This means that it is possible to supply linear calibration constants for each sensor on each node either before, during or after the data is collected. Note that most sensors (including the TelosB temperature and humidity sensors) are factory calibrated. The additional calibration support can be used to provide enhance accuracy over the default factory calibration.

A.6.7 Data Analysis

Alongside the Cogent-House Building Monitoring System a set of analysis tools have been developed. The key features include:

- Exposure analysis (for temperature, humidity, CO₂, etc)
- Energy versus degree days
- Comfort assessment
- Energy benchmarking
- Tornado plots

A.7 System evaluation

This section provides an evaluation of the COGENT-HOUSE system described in this appendix. A number of existing projects [19, 87, 89] define the metrics to evaluate WSN systems as: i) the useful data yield of the system and ii) the energy consumption of a node.

The system is primarily evaluated in terms of the deployment at a Passivhaus housing estate comprising of five homes and 18 flats. As illustrated in Figure A.5 the site comprises of three blocks:

1. Block one to the right of the site comprising of two houses and twelve flats (Houses 15–28),
2. Block two in the centre, comprises of three houses (Houses 29–31),
3. Block three to the left of the site, comprises of six flats (Houses 32–37).

Details of the homes can be found in Appendix B

Three servers, with Internet access for remote data collection, were deployed in the communal areas of Blocks one and three, and in one of the houses in Block two. A total of 176 sensing nodes were deployed with each server supporting between 27 and 107 sensing nodes.

A.7.1 Data yield

The analysis presented here assigns data loss according to the following categories:

Overall Yield Total yield percentage of data successfully received.

Theft/Server Fault Loss It is assumed if the server is in place it will receive packets. Therefore if zero nodes are reporting for a period of time this loss is down to a faulty server or the server was stolen.

Powered off Loss If the server is functioning but a powered node has not reported it is assumed the node has been switched off at the socket.



Figure A.5: Overview of deployment area: (A) Wireless Sensor Nodes, (B) Data collection server with Internet access, (C) A zoomed in version of a property showing the deployment strategy can be seen in Figure A.6

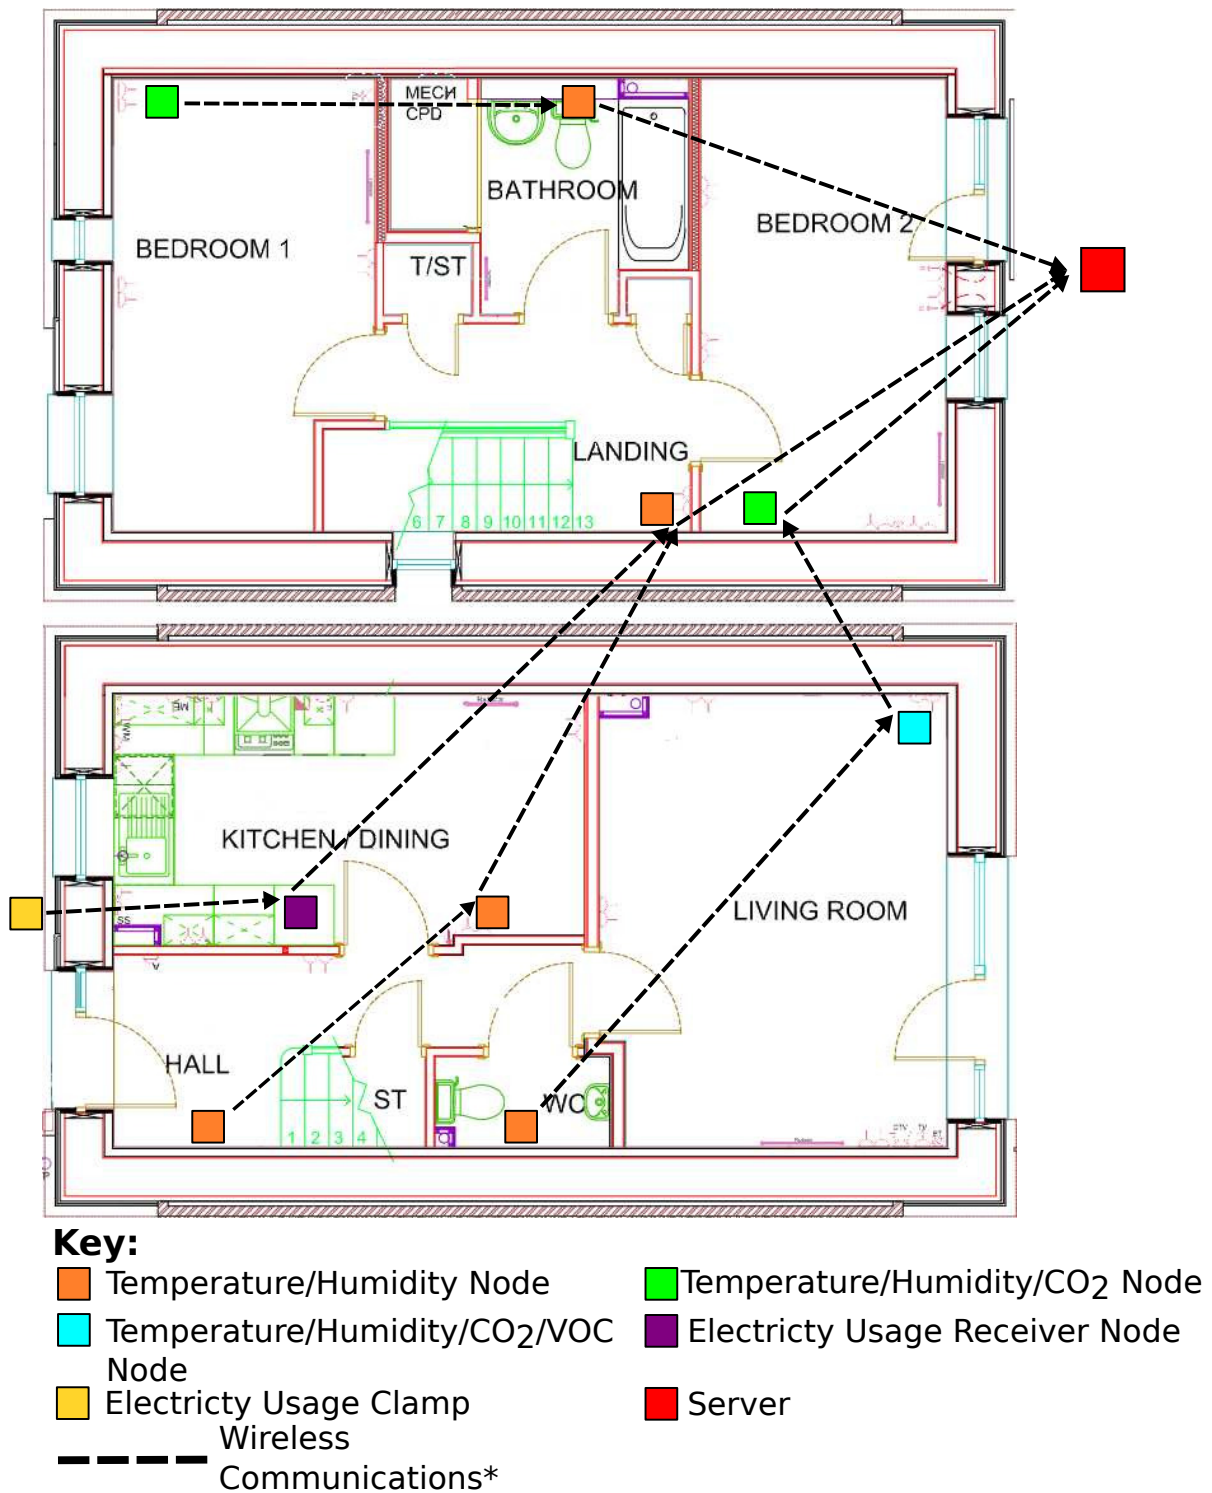


Figure A.6: Typical node deployment in an example two bedroom home. (* Multi-hop networking with suitable tree structure – note: multiple homes are served by the shown server)

Table A.1: Summary of the yield achieved and source of packet loss by the three servers deployed at the Passivhaus site

	Server 1 (14 properties)	Server 2 (3 properties)	Server 3 (6 properties)
Overall Yield	61%	55%	50%
Server Theft Loss	19%	11%	2%
Powered-off Loss	11%	29%	43%
Low battery Loss	3%	0%	4%
Network Loss	6%	5%	1%

Low battery Loss The percentage of packets sent where the battery voltage is less than 2.3V.

Network Loss The remaining loss is due to dropped/corrupt packets.

The main source of data loss resulted from 2 causes: i) physical access to server (one occurrence of theft, 2 occurrences of vandalism), ii) residents being able to switch the AIR-QUALITY nodes off at the socket (this issue has been further rectified by a revised version of the air quality interface board, which switches to battery power if switched off at the socket). Table A.1 gives a break down of the yield of the site split by the three deployed servers between July 2011–October 2012.

Rather than using tables, another form of analysing server performance in terms of yield is the use of heat-maps, which graphically represents daily yield percentages. From Figure A.7 it is clear to see before the 21st September 2011 there was an issue affecting the entire network, which was later found to be the flooding of corrupted dissemination packets. As a further example between May 9th 2012 to July 3rd 2012 no data was recorded at all, due to an instance of theft resulting in a major loss of data. In August the server’s hard-drive became corrupt and had to be replaced. Since there was a large amount of down time, between May–August 2012, the batteries in the nodes failed much earlier than anticipated, leading to a lower yield until they were replaced on 14th September 2012. Further vandalism occurred on the 21st September 2012 with the server being switched off.

A.7.2 Node lifetime

When the node’s battery level falls below 2.3 V the ability to transmit uncorrupted packets becomes unreliable. Table A.2 shows the annual energy requirement of a COGENT-HOUSE node (1200 mAH) using microbenchmarking. Assuming a node is powered by two standard AA batteries (3000 mAh) the estimated lifetime of a COGENT-HOUSE node is 2.5 years. However, from an analysis of six nodes which

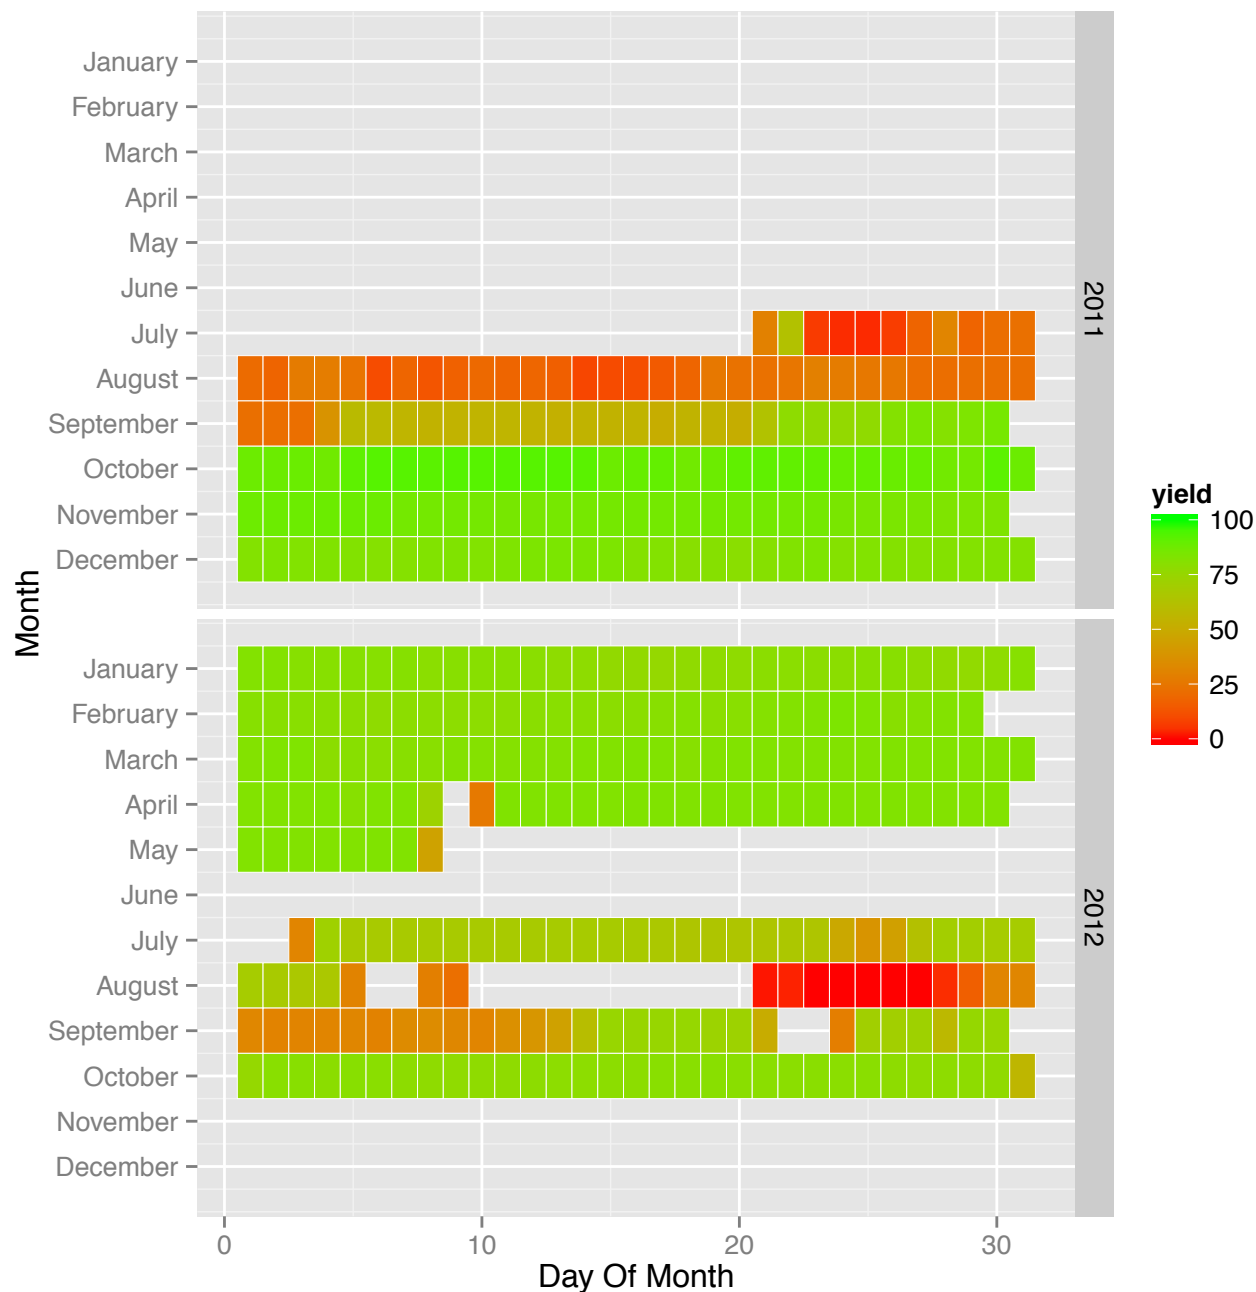


Figure A.7: Yield HeatMap for Server 1 at the Passivhaus site showing daily yields for a 15 month period

Table A.2: Baseline microbenchmark estimates for a COGENT-HOUSE TelosB node with a five minute sampling cycle. CTP send time is based on logs from a 200+ node network and include retries.

Process	Annual samples		Time (ms)		mA		mAh/year
Temperature	105120	×	220	×	0.458	=	2.9
Relative humidity	105120	×	75	×	0.458	=	1
Voltage	105120	×	0.017	×	0.536	=	0.00027
CTP send	105120	×	473	×	18.920	=	260
LPL check	105120	×	1,500	×	18.920	=	830
Idle	105120	×	297,732	×	0.009	=	78
Totals							1200

missed their scheduled battery change, batteries fell below 2.3 V in an average time of 300 days with a standard deviation of 10 days. This is a performance of 33% as compare to the microbenchmarking theoretical estimate, the reason for the performance difference is due to the microbenchmarking approach assuming the node will function until it is fully drained, 0 V, rather than the functional range of 2.3 V. This is why the measure of annual energy requirement has been used throughout this thesis. Battery life was acceptable for the application, as batteries could be changed every six months in accordance with scheduled system maintenance and tenant audits.

A.8 Summary

This appendix has described the design and development of COGENT-HOUSE a full end-to-end open-source home environmental and energy monitoring system. COGENT-HOUSE, was developed to exemplify and evaluate algorithms presented in this thesis.

The next appendix provides detail on the 37 Cogent-House deployments from which data extracted from 235 nodes have been used as evaluation datasets in this thesis.

Appendix B

Deployments and datasets

The proceeding table shows information regarding the deployments which produced the datasets used for evaluation throughout this thesis. The table includes the following details:

1. House Id—A unique identifier for the deployment.
2. House type—An identifier for the house type formed of <HouseType><No. Bedrooms><Heating Source>. A house type can be either a house (H), flat (F), or bungalow (B). Heating source can either be electricity (E), gas (G), ground source heat pump (GS), Air source heat pump (A), or district heating system (D). For example, considering a 4 bedroom bungalow heated by a ground source heat pump the house type would be-B4GS.
3. Monitoring duration—This is the number of days the deployment was performed over.
4. Data yield—The yield of data usable for analysis.
5. Sensor type—Which phenomena were sensed in the home.
6. Min—Minimum value sensed by the phenomena.
7. Mean—Average value, and standard deviation, sensed by the phenomena.
8. Max—Maximum value sensed by the phenomena.
9. Evaluation dataset periods—The evaluation dataset periods the deployment is used in.

House Id	House type	Monitoring duration (days)	No. of nodes	Data yield (%)	Sensor type	Min.	Mean	Max.	Evaluation dataset duration (No. of nodes)
1	H4G	360	14	99.98	Temp. (°C)	14	20 ± 1	34	Two week (14)
					Hum. (%)	28	64 ± 5	84	One month (14)
					CO2 (ppm)	460	890 ± 400	3600	Six month (9)
									One year (9)
2	H3E	33	8	91.6	Temp. (°C)	17	22 ± 1	31	Two week (8)
					Hum. (%)	25	45 ± 5	84	One month (8)
					CO2 (ppm)	620	900 ± 200	1800	
3	F2E	64	6	95.7	Temp. (°C)	15	21 ± 2	29	Two week (6)
					Hum. (%)	31	47 ± 7	94	One month (6)
					CO2 (ppm)	600	760 ± 70	1100	
4	F2E	20	6	95.1	Temp. (°C)	7.2	20 ± 2	28	Two week (6)
					Hum. (%)	31	50 ± 9	97	
5	F1E	33	5	98.5	Temp. (°C)	15	19 ± 2	23	Two week (5)
					Hum. (%)	44	53 ± 3	68	One month (5)
6	B2GS	53	6	99.2	Temp. (°C)	13	19 ± 3	32	Two week (5)
					Hum. (%)	22	54 ± 10	82	One month (5)
					CO2 (ppm)	490	1600 ± 700	3900	
7	B2GS	53	6	95.1	Temp. (°C)	17	26 ± 5	37	Two week (5)
					Hum. (%)	16	35 ± 10	78	One month (5)
					CO2 (ppm)	550	1100 ± 250	2200	

House Id	House type	Monitoring duration (days)	No. of nodes	Data yield (%)	Sensor type	Min.	Mean	Max.	Evaluation dataset duration (No. of nodes)
8	B2GS	22	6	94.9	Temp. (°C)	17	24 ± 2	32	Two week (2)
					Hum. (%)	23	44 ± 6	66	
9	H2G	16	9	91.1	Temp. (°C)	10	18 ± 3	26	Two week (9)
					Hum. (%)	32	53 ± 11	97	
10	H4E	20	11	97.3	Temp. (°C)	15	21 ± 2	30	Two week (11)
					Hum. (%)	28	50 ± 4	82	
11	H3E	27	8	72.4	Temp. (°C)	10	19 ± 2	29	Two week (8)
					Hum. (%)	29	54 ± 6	95	
12	H3E	20	9	99.6	Temp. (°C)	14	20 ± 2	50	Two week (9)
					Hum. (%)	23	55 ± 4	94	
13	B2A	16	8	99.6	Temp. (°C)	13	18 ± 2	26	Two week (9)
					Hum. (%)	31	58 ± 9	90	
14	B2E	19	8	100	Temp. (°C)	10	17 ± 2	29	Two week (8)
					Hum. (%)	31	59 ± 10	97	
15	H2D	470	9	50.4	Temp. (°C)	20	22 ± 0.85	26	Two week (8) One month (7)
					Hum. (%)	37	48 ± 3.9	60	
					CO2 (ppm)	540	790 ± 280	1900	
16	H3D	470	10	46.8	Temp. (°C)	21	24 ± 0.96	27	Two week (9) One month (9)
					Hum. (%)	31	47 ± 5	85	
					CO2 (ppm)	570	840 ± 160	1600	

House Id	House type	Monitoring duration (days)	No. of nodes	Data yield (%)	Sensor type	Min.	Mean	Max.	Evaluation dataset duration (No. of nodes)
17	F2D	470	6	32.6	Temp. (°C)	22	25 ± 0.98	27	Two week (4)
					Hum. (%)	33	45 ± 4.8	73	One month (4)
					CO2 (ppm)	550	860 ± 210	1700	
18	F2D	470	6	43.8	Temp. (°C)	23	26 ± 1.1	30	Two week (6)
					Hum. (%)	37	49 ± 5.9	85	One month (6)
					CO2 (ppm)	540	770 ± 230	3500	
19	F2D	470	6	44.2	Temp. (°C)	22	27 ± 1	30	Two week (6)
					Hum. (%)	29	38 ± 3.5	60	One month (6)
					CO2 (ppm)	530	720 ± 150	2300	
20	F2D	470	6	47.7	Temp. (°C)	22	25 ± 0.9	28	Two week (6)
					Hum. (%)	33	44 ± 4	84	One month (6)
					CO2 (ppm)	520	800 ± 200	2500	
21	F2D	470	6	47.3	Temp. (°C)	21	25 ± 1	29	Two week (6)
					Hum. (%)	34	46 ± 4	77	One month (6)
					CO2 (ppm)	490	850 ± 200	3400	
22	F2D	470	6	48.6	Temp. (°C)	21	25 ± 1	31	Two week (6)
					Hum. (%)	22	40 ± 7	77	One month (6)
					CO2 (ppm)	530	780 ± 200	1200	

House Id	House type	Monitoring duration (days)	No. of nodes	Data yield (%)	Sensor type	Min.	Mean	Max.	Evaluation dataset duration (No. of nodes)
23	F2D	470	6	44.7	Temp. (°C)	20	25 ± 2	28	Two week (6)
					Hum. (%)	28	43 ± 5	80	One month (6)
					CO2 (ppm)	530	800 ± 200	1900	
24	F2D	470	6	45.1	Temp. (°C)	17	23 ± 2	27	Two week (6)
					Hum. (%)	35	48 ± 6	82	One month (6)
					CO2 (ppm)	520	830 ± 200	1900	
25	F2D	470	6	29.8	Temp. (°C)	18	23 ± 1	28	Two week (5)
					Hum. (%)	32	47 ± 6	76	One month (4)
					CO2 (ppm)	520	890 ± 300	2700	
26	F2D	470	6	47.4	Temp. (°C)	18	23 ± 1	29	Two week (6)
					Hum. (%)	34	49 ± 4	76	One month (6)
					CO2 (ppm)	520	860 ± 200	3100	
27	F2D	470	6	26.4	Temp. (°C)	23	25 ± 1	30	Two week (3)
					Hum. (%)	33	42 ± 4	67	One month (3)
28	F2D	470	6	37.3	Temp. (°C)	22	26 ± 1	29	Two week (6)
					Hum. (%)	31	42 ± 5	52	One month (6)
					CO2 (ppm)	570	810 ± 100	1700	Six month (0)
29	H2D	470	9	61.0	Temp. (°C)	19	22 ± 1	24	Two week (8)
					Hum. (%)	27	48 ± 8	60	One month (8)
					CO2 (ppm)	510	700 ± 300	3100	Six month (8)

House Id	House type	Monitoring duration (days)	No. of nodes	Data yield (%)	Sensor type	Min.	Mean	Max.	Evaluation dataset duration (No. of nodes)
36	F2D	470	6	36.2	Temp. (°C) Hum. (%)	23 39	25 ± 0.7 49 ± 6	29 86	Two week (3) One month (3) Six month (3)
37	F2D	470	6	19.5	Temp. (°C) Hum. (%) CO2 (ppm)	17 26 560	23 ± 1 44 ± 8 900 ± 200	29 78 3300	Two week (6) One month (6) Six month (3)
38	H4G	244	18	95	Temp. (°C) Hum. (%)	14 31	19 ± 2 65 ± 9	27 89	Deployment was used as a L-SIP live deployment exemplar

Appendix C

Publications, presentations and attended conferences

The work in this thesis has resulted in the following peer-reviewed publications, presentations, and technical reports.

C.1 Journal publications

Edge mining the Internet of Things

Elena I. Gaura, James Brusey, Michael Allen, **Ross Wilkins**, Daniel Goldsmith, and Ramona Rednic. “Edge mining the Internet of things”. In *IEEE Sensors Journal*. vol. 13, no. 10, Oct. 2013, pp. 3816–3825.

C.2 Conference proceedings

Sustainable future? Building and lifestyle assessment.

Elena I. Gaura, John Halloran, James Brusey, **Ross Wilkins**, and Ramona Rednic. “Sustainable future? Building and life-style assessment”. In *Proceedings 2012 International Conference on Advanced Computer Science and Information Systems*, Dec. 2012, pp. 7–11.

Bare necessities—Knowledge-driven WSN design.

Elena I. Gaura, James Brusey, **Ross Wilkins**. “Bare necessities—Knowledge-driven WSN design”. In *Proceedings of 10th IEEE Sensors Conference*, Oct. 2011, pp. 66–70.

Wireless Sensing For The Built Environment: Enabling Innovation Towards Greener, Healthier Homes.

Elena I. Gaura, James Brusey, **Ross Wilkins**, and John Barnham. “Wireless Sensing For The Built Environment: Enabling Innovation Towards Greener, Healthier Homes”. In Proceedings of Clean Technology 2011, June. 2011, pp. 367–372.

Inferring knowledge from building monitoring systems: The case for wireless sensing in residential buildings.

Elena I. Gaura, James Brusey, **Ross Wilkins**, and John Barnham. “Inferring Knowledge From Building Monitoring Systems: The Case For Wireless Sensing In Residential Buildings”. In Proceedings of Clean Technology 2011, June. 2011, pp. 353–358.

C.3 Technical reports

COGENT-ETI.002: Holistic Assessment of Building Thermal Insulation—Post-Refurbishment In Use Test

Ross Wilkins, James Brusey. “Holistic Assessment of Building Thermal Insulation—Post-Refurbishment In Use Test”. Technical Report COGENT-ETI.002, Coventry University, 2014

COGENT-ETI.001: Holistic Assessment of Building Thermal Insulation—In Use Test

Ross Wilkins, James Brusey. “Holistic Assessment of Building Thermal Insulation—In Use Test”. Technical Report COGENT-ETI.001, Coventry University, 2014

COGENT-ORBIT-11: Sampson Close Monitoring Report: Final Overview

Elena Gaura, John Kemp, Ramona Rednic, **Ross Wilkins**, James Brusey, and John Halloran. “Sampson Close Monitoring Report: March 2013 to June 2013”. Technical Report COGENT-ORBIT.11, Coventry University, 2013

COGENT-ORBIT-10: Sampson Close Monitoring Report: June 2013 to September 2013

Elena Gaura, John Kemp, Ramona Rednic, **Ross Wilkins**, James Brusey, and John Halloran. “Sampson Close Monitoring Report: June 2013 to September 2013”. Technical Report COGENT-ORBIT.10, Coventry University, 2013

COGENT-ORBIT-08: Sampson Close Monitoring Report: March 2013 to June 2013

Elena Gaura, John Kemp, Ramona Rednic, **Ross Wilkins**, James Brusey, and John Halloran. “Sampson Close Monitoring Report: March 2013 to June 2013”. Technical Report COGENT-ORBIT.08, Coventry University, 2013

COGENT-ORBIT-07: Sampson Close Monitoring Report: September 2012 to March 2013

Elena Gaura, John Kemp, Ramona Rednic, **Ross Wilkins**, James Brusey, and John Halloran. “Sampson Close Monitoring Report: September 2012 to March 2013”. Technical Report COGENT-ORBIT.07, Coventry University, 2013

COGENT-ORBIT-06: Sampson Close Monitoring: Annual Report (September 2011 – September 2012)

Elena Gaura, John Kemp, Ramona Rednic, **Ross Wilkins**, James Brusey, and John Halloran. “Sampson Close Monitoring: Annual Report (September 2011 – September 2012)”. Technical Report COGENT-ORBIT.06, Coventry University, 2012

COGENT-ORBIT-05: Sampson Close Monitoring: May to September 2012

Elena Gaura, John Kemp, Ramona Rednic, **Ross Wilkins**, James Brusey, and John Halloran. “Sampson Close Monitoring: May to September 2012”. Technical Report COGENT-ORBIT.05, Coventry University, 2012

COGENT-ORBIT-04: Sampson Close Monitoring: May 2012

Elena Gaura, John Kemp, Ramona Rednic, **Ross Wilkins**, and James Brusey. “Sampson Close Monitoring: May 2012”. Technical Report COGENT-ORBIT.04, Coventry University, 2012

COGENT-ORBIT-03: Sampson Close Monitoring: Winter Report and Tenant Survey

Elena Gaura, James Brusey, John Halloran, **Ross Wilkins**, Ramona Rednic, and John Kemp. “Sampson Close Monitoring: Winter Report and Tenant Survey. Technical Report”. Technical Report COGENT-ORBIT.03, Coventry University, 2012

COGENT-SC.001: Baseline Occupant Survey and Energy Consumption Analysis

Elena Gaura, James Brusey, John Halloran, **Ross Wilkins**, Ramona Rednic, and John Kemp. “Sampson Close Monitoring: Baseline Occupant Survey and Energy Consumption Analysis”. Technical Report COGENT-SC.01, Coventry University, 2011

COGENT-ORBIT-02: Sampson Close Monitoring: Benefits and First Results

Elena Gaura, James Brusey, John Halloran, **Ross Wilkins**, Ramona Rednic, and John Kemp. “Sampson Close Monitoring: Benefits and First Results”. Technical Report COGENT-ORBIT.02, Coventry University, 2011

COGENT-ORBIT-01: Orbit Deployments Winter and Summer 2010: Temperature, Humidity and Comfort Exposure Report

Ross Wilkins. “Orbit Deployments Winter and Summer 2009: Temperature, Humidity, and Comfort Exposure Report”. Technical Report COGENT-ORBIT.01, Coventry University, 2010

COGENT.008: The Need For Home Environment Monitoring For The Increase In Energy Efficiency And Indoor Environment Quality

Ross Wilkins. “The Need For Home Environment Monitoring For The Increase In Energy Efficiency And Indoor Environment Quality”. Coventry University, 2009

C.4 Presentations and demos

- **Buildings Performance Evaluations Using Wireless Sensor Networks**, NTU – Coventry Research Exchange – Nottingham Trent University, UK
- **DECENT Homes: Distributed Evaluation of Carbon Emissions through NeTworked sensing in Homes** (September 2012), Southampton – Coventry Research Exchange – Coventry University, UK
- **Implementing the Passivhaus concept: Towards Lower Carbon and Increased Comfort Housing** (September 2011), Southampton – Coventry Research Exchange – Coventry University, UK
- **Introducing the Passivhaus Project** (June 2010), Southampton – Coventry Research Exchange – Coventry University, UK
- **Poster: Wireless Sensor Networks For Building Monitoring** (May 2010), Coventry University Research Symposium 2010 – Coventry University, UK

C.5 Conferences attended

- **The Building Performance Gap closing it through better measurement** (December 2012) – London, UK
- **WiSIG: Advances in Wireless Sensor Networks for Hostile Environments** (May 2012) – Derby, UK
- **M&E – The Building Services Event 2011** (October 2011) – London, UK
- **Clean Technology 2011** (June 2011) – Boston, MA, USA
- **WiSIG: Wireless Sensing for Smart Buildings** (February 2011) – Coventry University, UK
- **ACM Sensys 2010** (November 2010) – ETH Zurich, CH
- **WiSIG: Sensing Technology 2010** (September 2010) – Birmingham, UK
- **Wireless Communications to Enable the Internet of Things** (March 2010) – Surrey University, UK

- **Workshop in Speckled Computing (Environment Monitoring)** (December 2009) – University of Edinburgh, UK
 - **Ecif: Introduction to Wireless Sensor Networks** (October 2009) – Coventry University, UK
-

Selected publications follow.

Wireless Sensing For The Built Environment: Enabling Innovation Towards Greener, Healthier Homes

Elena I. Gaura*, James Brusey*, Ross Wilkins*, John Barnham**

* Cogent Computing Applied Research Centre, Faculty of Engineering and Computing,
Coventry University, Coventry, CV1 5FB
{e.gaura,j.brusey,wilkin24}@coventry.ac.uk

** Orbit Heart Of England, 10 Greenhill Street, Stratford-upon-Avon, Warwickshire, CV37 6LG
john.barnham@orbit.org.uk

Abstract

Worldwide carbon reduction targets for the built environment are staggeringly ambitious. If they are to be achieved, orders of magnitude performance increases are required from HVAC systems, construction techniques and insulating materials. Given the limited understanding of many of the newer materials and techniques, objective measurement is fundamental to meeting these targets in time. This paper presents the case for a holistic approach to measurement within the built environment and shows how Wireless Sensor Networks (WSNs) are a prime candidate technology to support such an approach.

WSNs are readily *enablers of understanding* in domains characterised by spatio-temporal, multivariate complexity. Simple, portable and non-intrusive WSN systems, deployed for weeks or years, are powerful tools for empirical environmental and energy performance evaluation of occupied dwellings. Coupled with structured deployment processes and novel empirical evaluation metrics, WSNs enable, for old stock, focused actions towards reduced energy consumption, improved internal environment, lower maintenance costs and maintaining the cost viability of the building asset. They are equally valuable in the context of new builds, for generic and apportioned energy consumption evaluations against the delivered environmental quality and design expectations.

Keywords: WSN, Built Environment, Energy Performance, Occupant Comfort

This item has been removed due to third party copyright. The unabridged version of the thesis can be viewed at the Lanchester Library, Coventry University.

The full text of this article is available free from <https://curve.coventry.ac.uk/open/items/5822a631-7f8e-b83c-ddb4-8e44bd3be4cd/1/>

Inferring Knowledge From Building Monitoring Systems: The Case For Wireless Sensing In Residential Buildings

Elena I. Gaura*, James Brusey*, Ross Wilkins*, John Barnham**

* Cogent Computing Applied Research Centre, Faculty of Engineering and Computing,
Coventry University, Coventry, CV1 5FB
{e.gaura,j.brusey,wilkin24}@coventry.ac.uk

** Orbit Heart Of England, 10 Greenhill Street, Stratford-upon-Avon, Warwickshire, CV37 6LG
john.barnham@orbit.org.uk

Abstract

The built environment offers the Wireless Sensor Networks (WSNs) research and commercial communities, potentially, the best set of applications yet, in terms of market size, revenue and strength of the business cases. The merits of using WSNs, however, to routinely perform empirical evaluations of old and new building stock have not been, as yet, fully appreciated by the domain's specialists (developers, construction contractors, surveyors, stock owners/users and regulatory bodies). It is hypothesised here that, in spite of their technological suitability, evident ability to generate vast amounts of data and commercial readiness, WSNs success (and thus their adoption) as tools for the built environment relies on negotiating the *data to knowledge* gap. The paper proposes a number of empirical metrics for holistic assessment of stock performance in terms of its heating and cooling systems, fabric and estimated occupant comfort. The metrics were developed iteratively in consultation with built environment practitioners.

Keywords: WSN, Built Environment, Energy Performance, Occupant Comfort, Metrics

The full text of this item is available free from
[https://curve.coventry.ac.uk/open/items/
d84e7965-5726-5c83-202c-203a49a4d236/1/](https://curve.coventry.ac.uk/open/items/d84e7965-5726-5c83-202c-203a49a4d236/1/)

This item has been removed due to third party copyright. The unabridged version of the thesis can be viewed at the
Lanchester Library, Coventry University.

Bare necessities—Knowledge-driven WSN design

Elena I. Gaura, James Brusey and Ross Wilkins
Cogent Computing, Coventry University
Priory Lane, Coventry CV1 5FB
{e.gaura,j.brusey,wilkin24}@coventry.ac.uk

Abstract—The viability of wireless sensor applications often hinges on minimising power consumption whilst maximising the informational output. Although many low-level platform-oriented energy saving mechanisms have been developed, considerable savings are possible at application level. This work presents an approach to pushing the calculation of application-level state closer to the information source. The context in which this approach is evaluated is a residential building monitoring application. Combined with the Spanish Inquisition Protocol (SIP), this is shown, based on deployment data, to reduce the average transmission period for temperature data from once every 5 minutes to an average of once every 38 days for an allowed error threshold of 10% on any component of the application-level state. For combined sensing of temperature, relative humidity and CO₂, the average transmission period drops to 13 days. This transmission reduction should considerably extend network life while having minimal effect on the usefulness of the information gathered. Most importantly, the underlying approach generalises to a wide variety of applications.

Index Terms—Wireless Sensor Networks

Please note the full text is available free from
<https://curve.coventry.ac.uk/open/items/788f9526-95d9-4ef4-f63e-b8074c92f6ff/1/>

This item has been removed due to third party copyright. The unabridged version of the thesis can be viewed at the Lanchester Library, Coventry University.

Sustainable future? Building and life-style assessment

Elena I. Gaura, John Halloran, James Brusey, Ross Wilkins, Ramona Rednic
Cogent Computing Applied Research Centre, Coventry University, Coventry, UK
e.gaura@coventry.ac.uk

Abstract—Energy, both in terms of its production and its usage has occupied a prime place in research as well as politics and world economy for the past few years. The majority of nations are aiming to deliver severe carbon cuts in the next few years. However, achieving a carbon-free future needs more than infrastructure investment and novel efficient technologies for buildings, transportation and other large consumer domains. It needs a better understanding of people as consumers, as well as a better understanding of energy waste across the multitude of socio-technical systems around us.

With regards to the built environment and particularly residential buildings, the authors propose that dual, quantitative and qualitative approaches to characterising, assessing and improving occupied buildings are necessary. Such approaches would synchronously cater for understanding i) buildings technical performance (fabric and building heating, cooling and ventilation systems) and ii) occupant's motivation, ability, knowledge and efficacy for adopting low carbon lifestyles. When deployed at scale, the above will enable cost effective, targeted interventions for both building fabric and systems improvement and towards empowering their occupants to live sustainably.

The paper describes such a quantitative and qualitative approach and proposes assessment tools. Further, the authors comment on the potential benefits from monitoring campaigns when deployed at scale.

This item has been removed due to third party copyright. The unabridged version of the thesis can be viewed at the Lanchester Library, Coventry University.

Edge mining the Internet of Things

Elena I. Gaura, James Brusey, Michael Allen, Ross Wilkins, Dan Goldsmith, and Ramona Rednic

Abstract—This paper examines the benefits of *edge mining*—data mining that takes place on the wireless, battery-powered, smart sensing devices that sit at the edge points of the Internet of Things. Through local data reduction and transformation, edge mining can quantifiably reduce the number of packets that must be sent, reducing energy usage and remote storage requirements. Additionally, edge mining has the potential to reduce the risk to personal privacy through embedding of information requirements at the sensing point, limiting inappropriate use. The benefits of edge mining are examined with respect to three specific algorithms: Linear Spanish Inquisition Protocol (L-SIP), ClassAct, and Bare Necessities (BN), which are all instantiations of General SIP (G-SIP). In general, the benefits provided by edge mining are related to the predictability of data streams and availability of precise information requirements; results show that L-SIP typically reduces packet transmission by around 95% (20-fold), BN reduces packet transmission by 99.98% (5000-fold) and ClassAct reduces packet transmission by 99.6% (250-fold). Although energy reduction is not as radical due to other overheads, minimisation of these overheads can lead to up to a 10-fold battery life extension for L-SIP, for example. These results demonstrate the importance of edge mining to the feasibility of many IoT applications.

Please note the full text of this paper is available free from
<https://curve.coventry.ac.uk/open/items/1f766f64-37f9-4abe-a923-55a2f57b5d34/1>

This item has been removed due to third party copyright. The unabridged version of the thesis can be viewed at the Lanchester Library, Coventry University.

Appendix D

Ethical approval

The work in this thesis has been approved through Coventry Universities ethical approval process. The ethical approval, and resident information sheet follows:



Medium - High Risk Research Ethics Approval

Where human participants involved in the research and/or when using primary data - Staff (Academic, Research, Consultancy, Honorary & External), Students (Research & Professional degrees) and Undergraduate or taught Postgraduates directed to complete this category of risk.

Project Title

Generalised approaches to transmissions reduction protocols in fielded Wireless Sensor Networks

Record of Approval

Principal Investigator

I believe that this project does not require research ethics approval. I have completed Sections 1-2 and kept a copy for my own records. I realise I may be asked to provide a copy of this form at any time.	X
I confirm that I will carry out the project in the ways described in this checklist. I will immediately suspend research and request new ethical approval if the project subsequently changes the information I have given in this checklist.	X
I confirm that I, and all members of my research team (if any), have read and agreed to abide by the Code of Research Ethics issued by the relevant national learned society.	X
I confirm that I, and all members of my research team (if any), have read and agreed to abide by the University's Research Ethics, Governance and Integrity Framework.	X

Name: Ross Wilkins

Date: 13/08/2014.....

Student's Supervisor (if applicable)

I have read this checklist and confirm that it covers all the ethical issues raised by this project fully and frankly. I also confirm that these issues have been discussed with the student and will continue to be reviewed in the course of supervision.

Name: Elena Gaura

Date: 19/08/2014.....

Reviewer

Date of approval by anonymous reviewer:

Medium to High Risk Research Ethics Approval Checklist

1 Project Information

Project Ref:	P26330
Full name:	Ross Wilkins
Faculty:	[EC] Faculty of Engineering and Computing
Department:	[UE] Cogent
Module Code:	
Supervisor:	Elena Gaura
Project title:	Generalised approaches to transmissions reduction protocols in fielded Wireless Sensor Networks
Date(s):	01/09/2009 - 12/09/2014
Created:	13/08/2014 11:33

Project Summary

The project investigates the use of transmissions reduction protocols in Wireless Sensor Networks. The work presented makes use of secondary data, and data collected about the building environment to evaluate developed protocols.

Names of Co-investigators (CIs) and their organisational affiliation:	Elena Gaura (Coventry University), James Brusey (Coventry University), John Halloran (Coventry University)
How many additional research staff will be employed on the project?	0
Names and their organisational affiliation (if known):	
Who is funding the project?	EPSRC
Has the funding been confirmed?	Yes
Code of ethical practice and conduct most relevant to your project:	British Computer Society

2. Does this project need ethical approval?

Questions	Yes	No
Does the project involve collecting primary data from, or about, living human beings?		X
Does the project involve analysing primary or unpublished data from, or about, living human beings?		X
Does the project involve collecting or analysing primary or unpublished data about people who have recently died other than data that are already in the public domain?		X
Does the project involve collecting or analysing primary or unpublished data about or from organisations or agencies of any kind other than data that are already in the public domain?		X
Does the project involve research with non-human vertebrates in their natural settings or behavioural work involving invertebrate species not covered by the Animals Scientific Procedures Act (1986)? ¹		X
Does the project place the participants or the researchers in a dangerous environment, risk of physical harm, psychological or emotional distress?		X
Does the nature of the project place the participant or researchers in a situation where they are at risk of investigation by the police or security services?		X
Does the project involve the researcher travelling outside the UK?		X

3 Does the project require Criminal Records Bureau checks?

Questions	Yes	No
Does the project involve direct contact by any member of the research team with children or young people under 18 years of age?		
Does the project involve direct contact by any member of the research team with adults who have learning difficulties?		
Does the project involve direct contact by any member of the research team with adults who are infirm or physically disabled?		
Does the project involve direct contact by any member of the research team with adults who are resident in social care or medical establishments?		
Does the project involve direct contact by any member of the research team with adults in the custody of the criminal justice system?		
Has a Criminal Records Bureau (CRB) check been stipulated as a condition of access to any source of data required for the project?		

If you answered **Yes** to **any** of these questions, please:

- Explain the nature of the contact required and the circumstances in which contact will be made during the project.

¹ The Animals Scientific Procedures Act (1986) was amended in 1993. As a result the common octopus (*Octopus vulgaris*), as an invertebrate species, is now covered by the act.

--

4 Is this project liable to scrutiny by external ethical review arrangements?

Questions	Yes	No
Has a favourable ethical opinion been given for this project by an external research ethics committee (e.g. social care, NHS or another University)?		
Will this project be submitted for ethical approval to an external research ethics committee (e.g. social care, NHS or another University)?		

5 More detail about the project

<i>What are the aims and objectives of the project?</i>
<i>Briefly describe the principal methods, the sources of data or evidence to be used and the number and type of research participants who will be recruited to the project.</i>
<i>What research instrument(s), validated scales or methods will be used to collect data?</i>
<i>If you are using an externally validated research instrument, technique or research method, please specify.</i>
<i>If you are not using an externally validated scale or research method, please attach a copy of the research instrument you will use to collect data. For example, a measurement scale, questionnaire, interview schedule, observation protocol for ethnographic work or, in the case of unstructured data collection, a topic list.</i>

6 Confidentiality, security and retention of research data

Questions	Yes	No
Are there any reasons why you cannot guarantee the full security and confidentiality of any personal or confidential data collected for the project?		
Is there a significant possibility that any of your participants, or people associated with them, could be directly or indirectly identified in the outputs from this project?		
Is there a significant possibility that confidential information could be traced back to a specific organisation or agency as a result of the way you write up the results of the project?		
Will any members of the project team retain any personal or confidential data at the end of the project, other than in fully anonymised form?		
Will you or any member of the team intend to make use of any confidential information, knowledge, trade secrets obtained for any other purpose than this research project?		

If you answered **No** to **all** of these questions:

- Explain how you will ensure the confidentiality and security of your research data, both during and after the project.

--

If you answered **Yes** to **any** of these questions:

- Explain the reasons why it is essential to breach normal research protocol regarding confidentiality, security and retention of research data.

--

7 Informed consent

Questions	Yes	No
Will all participants be fully informed why the project is being conducted and what their participation will involve and will this information be given before the project begins?		
Will every participant be asked to give written consent to participating in the project before it begins?		
Will all participants be fully informed about what data will be collected and what will be done with these data during and after the project?		
Will explicit consent be sought for audio, video or photographic recording of participants?		
Will every participant understand what rights they have not to take part, and/or to withdraw themselves and their data from the project if they do take part?		
Will every participant understand that they do not need to give you reasons for deciding not to take part or to withdraw themselves and their data from the project and that there will be no repercussions as a result?		
If the project involves deceiving or covert observation of participants, will you debrief them at the earliest possible opportunity?		

If you answered **Yes** to **all** these questions:

- Explain briefly how you will implement the informed consent scheme described in your answers.
- Attach copies of your participant information leaflet, informed consent form and participant debriefing leaflet (if required) as evidence of your plans.

--

If you answered **No** to **any** of these questions:

- Explain why it is essential for the project to be conducted in a way that will not allow all participants the opportunity to exercise fully-informed consent.
- Explain how you propose to address the ethical issues arising from the absence of transparency.
- Attach copies of your participant information sheet and consent form as evidence of your plans.

--

8 Risk of harm

Questions	Yes	No
Is there any significant risk that your project may lead to physical harm to participants or researchers?		
Is there any significant risk that your project may lead to psychological or emotional distress to participants or researchers?		
Is there any significant risk that your project may place the participants or the researchers in potentially dangerous situations or environments?		
Is there any significant risk that your project may result in harm to the reputation of participants, researchers, their employers, or other persons or organisations?		

If you answered **Yes** to **any** of these questions:

- Explain the nature of the risks involved and why it is necessary for the participants or researchers to be exposed to such risks.
- Explain how you propose to assess, manage and mitigate any risks to participants or researchers.
- Explain the arrangements by which you will ensure that participants understand and consent to these risks.
- Explain the arrangements you will make to refer participants or researchers to sources of help if they are seriously distressed or harmed as a result of taking part in the project.
- Explain the arrangements for recording and reporting any adverse consequences of the research.

--

9 Risk of disclosure of harm or potential harm

Questions	Yes	No
Is there a significant risk that the project will lead participants to disclose evidence of previous criminal offences or their intention to commit criminal offences?		
Is there a significant risk that the project will lead participants to disclose evidence that children or vulnerable adults have or are being harmed or are at risk of harm?		
Is there a significant risk that the project will lead participants to disclose evidence of serious risk of other types of harm?		

If you answered **Yes** to **any** of these questions:

- Explain why it is necessary to take the risks of potential or actual disclosure.
- Explain what actions you would take if such disclosures were to occur.
- Explain what advice you will take and from whom before taking these actions.
- Explain what information you will give participants about the possible consequences of disclosing information about criminal or serious risk of harm.

--

10 Payment of participants

Questions	Yes	No
Do you intend to offer participants cash payments or any other kind of inducements or compensation for taking part in your project?		
Is there any significant possibility that such inducements will cause participants to consent to risks that they might not otherwise find acceptable?		
Is there any significant possibility that the prospect of payment or other rewards will systematically skew the data provided by participants in any way?		
Will you inform participants that accepting compensation or inducements does not negate their right to withdraw from the project?		

If you answered **Yes** to **any** of these questions:

- Explain the nature of the inducements or the amount of the payments that will be offered.
- Explain the reasons why it is necessary to offer payments.
- Explain why you consider it is ethically and methodologically acceptable to offer payments.

--

11 Capacity to give informed consent

Questions	Yes	No
Do you propose to recruit any participants who are under 18 years of age?		
Do you propose to recruit any participants who have learning difficulties?		
Do you propose to recruit any participants with communication difficulties including difficulties arising from limited facility with the English language?		
Do you propose to recruit any participants who are very elderly or infirm?		
Do you propose to recruit any participants with mental health problems or other medical problems that may impair their cognitive abilities?		
Do you propose to recruit any participants who may not be able to understand fully the nature of the research and the implications for them of participating in it?		

If you answered **Yes** to **any** of the **first four** questions:

- Explain how you will ensure that the interests and wishes of participants are understood and taken in to account.
- Explain how in the case of children the wishes of their parents or guardians are understood and taken into account.

--

12 Is participation genuinely voluntary?

Questions	Yes	No
Are you proposing to recruit participants who are employees or students of Coventry University or of organisation(s) that are formal collaborators in the project?		
Are you proposing to recruit participants who are employees recruited through other business, voluntary or public sector organisations?		
Are you proposing to recruit participants who are pupils or students recruited through educational institutions?		
Are you proposing to recruit participants who are clients recruited through voluntary or public services?		
Are you proposing to recruit participants who are living in residential communities or institutions?		
Are you proposing to recruit participants who are in-patients in a hospital or other medical establishment?		
Are you proposing to recruit participants who are recruited by virtue of their employment in the police or armed services?		
Are you proposing to recruit participants who are being detained or sanctioned in the criminal justice system?		
Are you proposing to recruit participants who may not feel empowered to refuse to participate in the research?		

If you answered **Yes** to **any** of these questions:

- Explain how your participants will be recruited.
- Explain what steps you will take to ensure that participation in this project is genuinely voluntary.

--

13 On-line and Internet Research

Questions	Yes	No
Will any part of your project involve collecting data by means of electronic media such as the Internet or e-mail?		
Is there a significant possibility that the project will encourage children under 18 to access inappropriate websites or correspond with people who pose risk of harm?		
Is there a significant possibility that the project will cause participants to become distressed or harmed in ways that may not be apparent to the researcher(s)?		
Will the project incur risks of breaching participant confidentiality and anonymity that arise specifically from the use of electronic media?		

If you answered **Yes** to **any** of these questions:

- Explain why you propose to use electronic media.
- Explain how you propose to address the risks associated with online/internet research.
- Ensure that your answers to the previous sections address any issues related to online research.

--

14 Other ethical risks

Question	Yes	No
Are there any other ethical issues or risks of harm raised by your project that have not been covered by previous questions?		

If you answered **Yes** to **this** question:

- Explain the nature of these ethical issues and risks.
- Explain why you need to incur these ethical issues and risks.
- Explain how you propose to deal with these ethical issues and risks.

--

15 Research with non-human vertebrates²

Questions	Yes	No
Will any part of your project involve the study of animals in their natural habitat?		
Will your project involve the recording of behaviour of animals in a non-natural setting that is outside the control of the researcher?		
Will your field work involve any direct intervention other than recording the behaviour of the animals available for observation?		
Is the species you plan to research endangered, locally rare or part of a sensitive ecosystem protected by legislation?		
Is there any significant possibility that the welfare of the target species or those sharing the local environment/habitat will be detrimentally affected?		
Is there any significant possibility that the habitat of the animals will be damaged by the project such that their health and survival will be endangered?		
Will project work involve intervention work in a non-natural setting in relation to invertebrate species other than <i>Octopus vulgaris</i> ?		

If you answered **Yes** to **any** of these questions:

- Explain the reasons for conducting the project in the way you propose and the academic benefits that will flow from it.
- Explain the nature of the risks to the animals and their habitat.
- Explain how you propose to assess, manage and mitigate these risks.

--

² The Animals Scientific Procedures Act (1986) was amended in 1993. As a result the common octopus (*Octopus vulgaris*), as an invertebrate species, is now covered by the act.

16 Blood Sampling / Human Tissue Analysis

Questions	Yes	No
Does your project involve blood sampling or human tissue analysis?		X
If your study involves blood samples or body fluids (e.g. urine, saliva) have you clearly stated in your application that appropriate guidelines are to be followed (e.g. The British Association of Sport and Exercise Science Physiological Testing Guidelines (2007) or equivalent) and that they are in line with the level of risk?		
If your study involves human tissue other than blood and saliva have you clearly stated in your application that appropriate guidelines are to be followed? (e.g. The Human Tissues Act, or equivalent) and that they are in line with the level of risk?		

If you answered **No** to **any** of these questions, please provide more information:

--

Note: This checklist is based on an ethics approval form produce by Research Office of the College of Business, Law and Social Sciences at Nottingham Trent University. Copyright is acknowledged.

PARTICIPANT INFORMATION SHEET
COGENT COMPUTING ARC DEPARTMENT
COVENTRY UNIVERSITY

Thank you for considering participating in this research work. This information explains what you will be asked to do. If you have any questions about this please contact Dr James Brusey (j.brusey@coventry.ac.uk; +44 (0) 24 7765 9184).

Please note that you are free to stop taking part in this investigation at any time.

Information about the project/Purpose of the project

The aim of the research is to develop a low cost, robust and long lived energy and environmental monitoring system which allows for automated profiling of resident behaviour, particularly associated with space heating and heating energy miss-use.

It is envisaged that monitoring systems such as the one to be developed, prototyped and trialled here (based on current Cogent design) will become permanent fixtures of new builds to allow for both building performance monitoring and occupiers empowerment towards sustainable living.

It is foreseen that the following parameters will be part of the monitoring solution:

- Electricity at meter level
- Temperature (at several locations)
- Relative humidity (at several locations)
- Air quality (CO₂)
- Gas consumption
- Heating (measured through a heat meter)
- Window opening (measured intermittently with a window sensor)

Why have I been chosen?

You have been approached to take part as you are living in a set of flats that are being renovated with an innovative approach to improve their thermal performance.

Do I have to take part?

You do not have to take part in this research project if you do not want to and you do not need to give any reason if you decide not to take part at any time of your involvement.

What do I have to do?

If you agree to take part, please sign and date at the end of the form. No other action is needed.

What are the risks associated with this project?

There are no risks associated with this study.

What are the benefits of taking part?

You are potentially contributing to help improve the energy efficiency of your home and other similar homes.

Withdrawal options

You are free to stop taking part in this study at any time and you do not have to give any reason for this but you should contact the team (see below) and provide opportunities for removal of equipment.

Data protection & confidentiality

Participant confidentiality will be maintained at all times. For the purpose of confidentiality all data collected will remain anonymous and therefore will be coded. Data will only be accessible to the research team. Data will be stored on a secured computer. Upon completion of the study any electronic file or hard copy containing personal details and anonymity coding will be destroyed.

Who should you talk to if you have questions or you wish to make a complaint

If you have any questions or queries a member of the research team will be happy to answer them. If they cannot help you can speak to any of the key contacts listed below. If you have any questions about your rights as a participant or feel you have been placed at risk please contact James Brusey (details below).

What will happen with the results of the study?

Any data/ results from your participation in the study will be used by Cogent Computing to produce a report for e2Rebuild (an EU project). It may also be published in scientific works. In all cases, your name or identity will not be revealed.

Key contact details

Dr James Brusey, Coventry University

j.brusey@coventry.ac.uk

+44 (0) 24 7765 9184

Mr Ross Wilkins, Coventry University

ab7438@coventry.ac.uk

Please confirm that you have seen the (attached) participant information sheet and that you consent to be part of the study by signing and dating below.

Please note that you may still subsequently withdraw from the study at any time.

Signature _____ Date _____