An OFDMA-Based Data Extraction Protocol for Wireless Sensor Networks

A DISSERTATION SUBMITTED TO

THE INSTITUTE FOR COMMUNICATIONS AND SIGNAL PROCESSING,

DEPARTMENT OF ELECTRONIC AND ELECTRICAL ENGINEERING,

AND THE COMMITTEE FOR POSTGRADUATE STUDIES

OF THE UNIVERSITY OF STRATHCLYDE

IN PARTIAL FULFILMENT OF THE REQUIREMENTS

FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

By

Neil C. MacEwen

September 2009

Declaration

I declare that this thesis embodies my own research work and that it is composed by myself. Where appropriate, I have made acknowledgements to the work of others.

Neil C. MacEwen

Abstract

The field of Wireless Sensor Networks (WSNs) includes a wide range of application areas and presents many different requirements on network and physical layer design and implementation. The Speckled Computing Consortium (SCC) is investigating one form of WSN where each node is severely constrained in size and energy. With this in mind, physical layer architectures have to be carefully designed to be as efficient as possible.

The Orient-2 is a posture tracking system developed within the SCC which uses small, battery-powered nodes to track the movement of a body and relay the information to a central location. Off-the-shelf components limit the manner in which data are extracted, and result in the system being constrained in terms of the number of nodes supported, radio bit rate and update rate. These constraints provide motivation for the development of an application-driven solution which allows these limitations to be relaxed.

A data extraction protocol is proposed which is based on Orthogonal Frequency Division Multiple Access (OFDMA). Although OFDMA allows a receiver to elegantly receive multiple transmissions in parallel, it requires accurate time and frequency synchronisation between users. Standard offset management techniques are overly complex for the simple nodes desired in the Orient-2 system, and as such a system-specific solution is developed. In the proposed protocol, time offsets are managed using receiver-initiated transmissions which result in the offsets being reduced to manageable levels. Subcarrier modulation, frequency offset estimation and

compensation are combined in a novel all-digital transmitter architecture which minimises the effects of the frequency offsets and associated performance loss. Finally, the performance of the frequency estimation, compensation and subcarrier generation architecture is analysed in detail in order to find the minimum complexity node implementation which satisfies the performance requirements of the Orient-2 system.

Acknowledgements

Firstly, I would like to thank my supervisor, Professor Bob Stewart, for giving me the opportunity as a callow undergraduate to taste life in research, and for inviting me into his vibrant group. His support has been invaluable over the past 5 years. I would also like to say thank you to Dr Eugen Pfann, my 2nd supervisor, for being on call to pull me out of numerous mathematical quicksands.

To all those in the lab I owe a massive debt of gratitude; Steve, Graham, Kostas, Tony and Yousif have all contributed to the great working environment I've experienced. Thanks also to my virtual labmates Bobby J and Greg for a little escapism from time to time. To my Specknet buddies in particular, Louise and Faisal, I want to say a huge thank you. Your support has been so important, both technically and otherwise. We got through it! A special mention must also go to the folks down the road for always being available to help, and for latterly providing a distraction from the writing.

I want to say thank you to all those in the Speckled Computing Consortium for providing me with the opportunity to learn about so much more than my own narrow field; in that I've been lucky. In particular, I would like to thank Alex et al. for developing the Orient family of systems and providing me with the inspiration for my own work.

There have been good times and bad times throughout this PhD, and over the past

5 years I've come to realise just how important friends are. You know who you are.

Finally I would like to thank my family. I feel lucky to have grown up with such a great sister and brother, Alison and Iain, and to know I can always count on them, thank you. My parents have been a source of quiet and constant support for all of my life, and I feel privileged to have been given the opportunities they provided. Thank you, Mum and Dad, for helping me reach where I am today.

Acronyms

3GPP	3rd Generation Partnership Project
ADC	Analogue to Digital Converter
BER	Bit Error Rate
BPSK	Binary Phase Shift Keying
BS	Base Station
CDMA	Code Division Multiple Access
CSMA	Carrier Sense Multiple Access
DAB	Digital Audio Broadcasting
DAC	Digital to Analogue Converter
DBPSK	Differential Binary Phase Shift Keying
DDS	Direct Digital Synthesis/Synthesiser
DFT	Discrete Fourier Transform
DPSK	Differential Phase Shift Keying
DQPSK	Differential Quadrature Phase Shift Keying
DS-SS	Direct Sequence Spread Spectrum
FBAR	Thin-film Bulk Acoustic Wave Resonator
FDMA	Frequency Division Multiple Access
FFT	Fast Fourier Transform
ICI	Inter Carrier Interference

IDTF	Inverse Discrete Fourier Transform
IFFT	Inverse Fast Fourier Transform
ISI	Inter Symbol Interference
LTE	Long Term Evolution
LUT	Look Up Table
MAC	Multiple Access Control
NCO	Numerically Controlled Oscillator
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
OOK	On-Off Keying
OTS	Off-The-Shelf
PAN	Personal Area Network
PAPR	Peak to Average Power Ratio
PDF	Probability Distribution Function
PER	Packet Error Rate
PLL	Phase Locked Loop
ppm	parts per million
QPSK	Quadrature Phase Shift Keying
RF	Radio Frequency
ROM	Read Only Memory
SCC	Speckled Computing Consortium
SNR	Signal-to-Noise Ratio
SQNR	Signal-to-Quantisation Noise Ratio
SRO	Super-Regenerative Oscillator
TDMA	Time Division Multiple Access
TSMP	Time Synchronised Mesh Protocol

UMTS	Universal Mobile Telecommunications System
UN	User Network
WLAN	Wireless Local Area Network
WPAN	Wireless Personal Area Network
WSN	Wireless Sensor Network

Mathematical Notation

$*$	Complex conjugate.
c	Speed of light.
Δf	Total frequency offset.
Δf_c	Carrier frequency offset.
$\Delta f_D(t)$	Doppler frequency offset.
Δf_{max}	Maximum frequency offset.
$D_{m,\,i}$	Decision variable for symbol i on subcarrier m.
d_N	Distance between user N and central receiver.
E_b	Energy per bit.
E_{idle}	Energy used by a node in idle state.
E_{rx}	Energy used by a node in receive state.
E_{tx}	Energy used by a node in transmit state.
f_c	Carrier frequency.
f_{clk}	DDS clock frequency.
f_o	DDS output frequency.

$f_{r,\,DDS}$	Residual frequency offset due to subcarrier synthesis error.
f_{res}	DDS frequency resolution.
$f_{r,\,EST}$	Residual frequency offset due to frequency offset estimation error.
f_{rx}	Sampling rate in receiver.
f_s	Sampling rate in node.
f_{sub}	Subcarrier spacing.
I_0	Interference power spectral density.
μ	DDS phase accumulator step size.
N	Number of available subcarriers (size of FFT in receiver).
N_{ch}	Number of available channels.
N_0	Noise power spectral density.
N_u	Number of users.
$\ddot{\Omega}_K$	Normalised Kay frequency estimate.
$\ddot{\Omega}_M$	Normalised Meyr frequency estimate.
P_e	Probability of bit error.
P_I	Interference power.
P_N	Noise power.
P_S	Signal power.
R_b	Bit rate.
R_{up}	Update rate.
$\sigma^2_{\Delta f}$	Frequency offset variance (normalised to subcarrier spacing).

σ^2_{est}	Frequency estimate variance (normalised to sampling rate).
T	OFDM symbol period.
T^N_{ETA}	Time of arrival of a packet from user N at the central receiver.
T^N_{exch}	OFDMA protocol full exchange period for node N.
τ_{max}	Maximum time offset on any user seen by the receiver.
T^N_P	Signal propagation time for user N.
T_{pkt}	Packet duration.
T_s	Sampling period.
T_{tx}	Time taken to receive and process polling signal (perform frequency estimation).
τ_u	Time offset on user u.
T_{up}	Update period.
$V(t)$	Relative velocity between a node and central receiver.
W	Signal bandwidth.
W_I	Interference bandwidth.
$Y^k_{m,\,i}$	Contribution to decision variable on subcarrier m during symbol i due to subcarrier k.

Contents

Chapter 1

Introduction

1.1 Overview

This thesis investigates a novel, receiver-initiated, data extraction protocol that satisfies the need of some Wireless Sensor Networks (WSNs) to extract information efficiently to a central receiver. This application-driven research is targeted at a particular network, the Orient-2 body posture tracking system, which places requirements on information update rate, latency and the number of nodes in the network. Initial implementations of the system use general-purpose, commercially-available technology, and as such have constrained functionality. The OFDMA-based protocol implementation developed for the Orient-2[1] system in this research enables increased capabilities.

1.2 Context of this Work

Wireless Sensor Networks

A WSN can take many different forms and inhabit a wide array of application areas, but at its core such a network provides a link between the physical and digital worlds.

1. At the time of writing, the Orient-2 is the latest in a series of Orient body posture tracking systems.

A network of nodes gathers sensory information from the physical environment, brings it into the digital world and processes it according to the network's application. The wide range of applications that can be supported using a WSN raises a number of interesting design and implementation issues. This research deals with one issue in particular, that of extraction of data from the network to a central receiver.

Speckled Computing and the Orient-2 System

Speckled Computing is a vision of a network of tiny nodes that enables true ubiquitous computing, where computing power is everywhere around us, embedded in the environment. The Speckled Computing Consortium (SCC) is a multi-disciplinary collaboration between five Scottish universities investigating Speckled Computing at all layers of the protocol stack, from physical layer communications right up to applications. Speckled Computing is implemented with a network, called a Specknet, made up of tiny nodes, known as Specks, which are entirely self-sufficient. The long term goal is to develop nodes of small enough dimensions to disseminate entirely into the environment, providing computing power in everyday objects without the user being aware of their existence.

One output of the SCC to date has been the development of the Orient-2 system, which is a wireless body posture tracking network. Nodes possessing movement monitoring sensors (gyroscopes, accelerometers and magnetometers) are placed on a human body, and send sensory information to a central location, which computes a body model and tracks body movement, as shown in Figure 1.1. Although the long term aim is to have perhaps many hundreds of nodes tracking a body with fine resolution, the present system allows posture tracking using a network of fifteen nodes. Initial implementations of the system use commercially available components, and as such, even with this small number of nodes, the current network is limited in terms of update rate by the use of a single, time-shared channel and the maximum data rate of

the radio link. An appropriate design is therefore required to remove these limitations.

Multicarrier Modulation

Multicarrier modulation allows multiple data channels to be transmitted in parallel. In particular, Orthogonal Frequency Division Multiplexing (OFDM) has become very popular over recent years due to its robustness to channel effects and thus increased data rates. Orthogonal Frequency Division Multiple Access (OFDMA) uses the multiple channels for different users, thus allowing a number of users to transmit concurrently. However, time and frequency offset errors result in a loss of orthogonality and thus significant interference between users. The offsets must be minimised to avoid significant performance degradation.

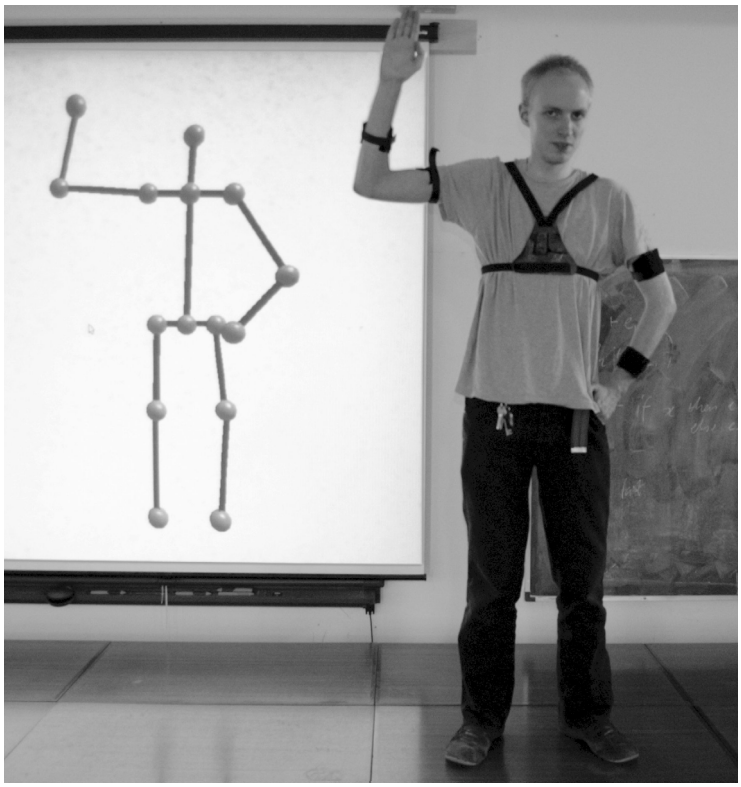

Figure: 1.1: A snapshot of body motion tracking using the Orient-2 System, reproduced from[74], with permission.

An OFDMA-based, Receiver-Initiated Data Extraction Protocol for the Orient-2

By assigning different OFDMA channels to individual users, data can be extracted from multiple nodes in parallel. However, tight constraints on time and frequency synchronisation between users must be met or significant performance degradation will result. Standard synchronisation techniques are unsuitable for the simple nodes desired in the Orient-2 network, and therefore a novel solution is required. This research investigates the implementation of a receiver-initiated protocol which uses OFDMA to enable multiple data extraction channels and relax the constraints placed on the Orient-2 network by the current, off-the-shelf, single-channel implementation. Time and frequency offset problems are managed through the design of the protocol itself, and through the architecture used in each node to transmit to the central receiver.

Orient-2 nodes are battery powered and as such have limited lifetime. In order to maximise the usefulness of the network, energy consumption must be minimised and network design must be as simple as possible. The data extraction protocol and its implementation allow the Orient-2 system to increase its functionality with minimal increased complexity. By designing the protocol and the node architecture specifically for the Orient-2 system, complexity can be reduced wherever possible. In the transmitter a simple, all-digital modulation architecture is used in which the analogue front-end is simply a radio frequency (RF) oscillator and mixer. By performing frequency synchronisation in the transmitter and through the use of differential modulation, the same front-end can be used in the receiver, avoiding the need for complex (and expensive) carrier synchronisation stages. In addition, OFDMA allows the use of an elegant, Fourier Transform based demodulation stage which is easily scaleable with the number of transmitting nodes, as opposed to the bank of filters that would be required for a standard Frequency Division Multiple Access (FDMA) implementation.

1.3 Original Contributions

This thesis investigates the viability of using OFDMA for the extraction of data from a network with certain stringent update rate and latency requirements. The main contributions of this work are considered to be the following.

- A novel, receiver-initiated data extraction protocol using OFDMA to enable multiple communication channels is proposed. It is demonstrated that the protocol is useful in power constrained networks with particular update rate and latency requirements on the extraction of information.

- It is proposed that frequency offsets be both estimated and compensated for in the transmitting node, using a simple, agile digital subcarrier modulation architecture.

- Time and frequency offsets in such a system, and their effects on data extraction performance are considered. The interference caused by frequency offsets is investigated in detail and the results of the analysis used to specify the required accuracy of the frequency offset estimation.

- The implementation of subcarrier modulation and associated frequency offset compensation is investigated. A recommendation is made for a suitable frequency estimator, and the use of a Direct Digital Synthesis (DDS) architecture for subcarrier generation is evaluated in the context of such a network, with a recommendation for a minimum complexity implementation being made.

1.4 Thesis Structure

The main work is presented in Chapters 2-6 before the thesis is concluded and some suggestions for further work are given in Chapter 7. Chapters 2-7 are stuctured as follows.

Chapter 2 places this work in context by introducing WSNs and discusses the many implementation and design issues that arise when considering different applications. The state-of-the-art in WSNs is covered and Speckled Computing is introduced in more detail. Different network scenarios and their requirements for data extraction are considered, along with the associated communication costs. The Orient-2 system is then introduced and the motivation for a data extraction system which will mitigate latency and update rate problems is made clear.

Chapter 3 gives an introduction to multicarrier communications, in particular OFDM and its extension to OFDMA and their applications within Speckled Computing and WSNs. The sensitivity of an OFDMA system to time and frequency offsets is covered in detail before the most common methods of dealing with these problems are discussed.

Chapter 4 introduces the new data extraction protocol, considering time and frequency offsets and deriving limits for network dimensions and on node movement. The application of the protocol with the Orient-2 system is analysed. System implementation and performance issues are then considered, with the Orient-2 system being used as a model network.

Chapter 5 investigates in more detail the effect of frequency offsets in the system, and how they can be managed. A simple expression of the interference power caused by frequency offsets is developed in order to analyse the resulting performance

degradation and obtain a specification for the required accuracy of the frequency estimation. Two commonly used frequency estimation algorithms are then evaluated and a recommendation is made for a suitable estimator.

Chapter 6 analyses the effect on data extraction performance, in terms of bit error rate (BER), of the use of a digital subcarrier generation architecture. A minimum complexity implementation of the node transmitter is found through extensive simulation.

Chapter 7 finishes with a review of the work contained in the thesis, and the original contributions made. Finally, some ideas for future extensions to this work are outlined.

Chapter 2

Wireless Sensor Networks and Speckled Computing

2.1 Introduction

This chapter provides a context for the research discussed in later chapters. The chapter begins by introducing the concept of a Wireless Sensor Network (WSN) and the state-of-the-art in WSNs is described with particular focus on some currently available systems. Speckled Computing and Specknets are then introduced and the particular design issues applicable are explained. The development of different Specknet nodes is covered, before the limitations of the current radio transceiver are described and the motivation for increasing the capabilities of a Specknet node is demonstrated.

WSNs can be targeted at a multitude of application areas; each application area has different requirements regarding the treatment of the data the network gathers. In this chapter, the various issues involved with extracting data from a WSN are explored. Several different network scenarios are investigated with regard to these issues before they are analysed in further detail with reference to the Orient-2 body posture tracking application. This scenario is used as a motivating example to highlight how multiple channels can be exploited to improve data extraction from a WSN.

2.2 Wireless Sensor Networks

In recent years WSNs have become a field of great interest involving research in many different disciplines ranging from sensors to low-power communications to battery technology [1,48,55]. Figure 2.1 shows a generic architecture for a wireless sensor node. It can be seen that sensor network nodes possess sensing/actuation, processing and communications capability; however it is how these nodes are networked and their capabilities utilised that defines a particular network. The wide array of applications in which WSNs can be used ([1,6,22,53]) results in a broad spectrum of different network and node specifications, however in all cases they provide a link between the physical and digital worlds. A selection of WSN application area examples is provided below:

- **Environmental Monitoring** - Indoor and outdoor applications such as temperature sensing, fire/heat detection or wildlife tracking.

- **Agriculture** - Livestock tracking and health monitoring, water delivery systems for irrigation.

- **Surveillance / Battlefield Monitoring** - Intruder tracking, target identification.

- **Asset Tracking** - Cargo tracking, inventory control.

- **Fault Detection** - Real-time monitoring of stress in industrial environments, such as turbines or engines.

- **Human** - Medical and social applications from health monitoring to human interaction tracking.

With such a wide range of applications, WSNs raise a significant number of interesting design and implementation issues, which are explored in the following sections.

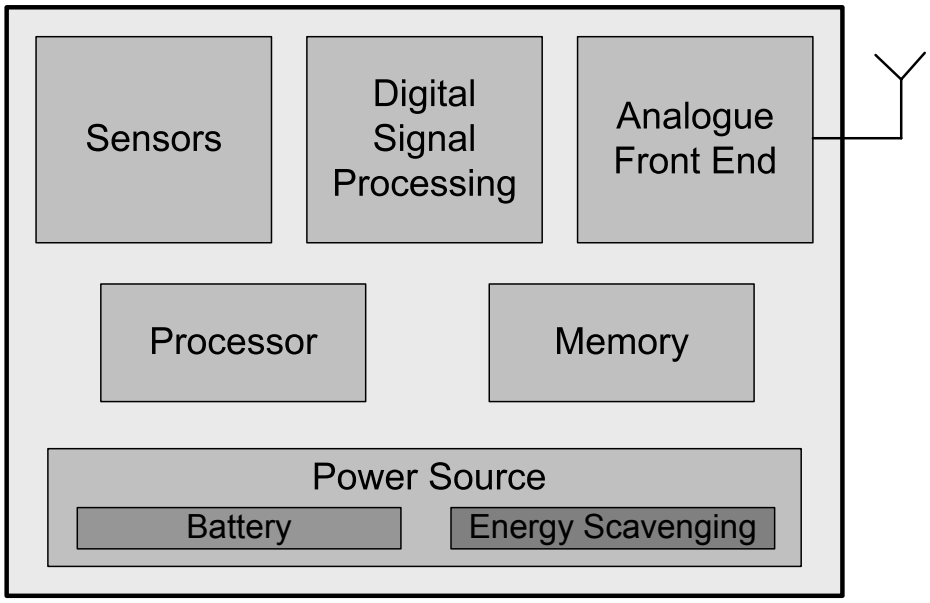

Figure: 2.1: A generic WSN node showing the basic components needed to enable its basic functionality of sensing, processing and communicating.

2.2.1 Power Supply

A minority of sensor networks may have access to a fixed power supply, however for the most part nodes have to exist on a self-contained power source. With nodes generally being very small the available energy storage capacity is limited, and as such the node may have a very limited lifetime. In order to prolong node lifetime as much as possible many networks attempt to scavenge energy from the environment, for example harnessing solar or vibrational energy [1, 49].

2.2.2 Communications

Communication between nodes can be achieved in several different ways [49]. The

most commonly used is radio, which allows for low-power transmission over non-line-of-sight paths. In addition, high frequency carrier frequencies enable the use of very small antennas. Other less commonly used communication mediums include optical and infrared [1]. The choice of communication technique is highly dependent on the network application and the environment in which it will be deployed.

In general, data rate and range requirements for WSNs are low compared to other wireless standards. Figure 2.2 compares these approximate requirements with those of several common digital communication standards.

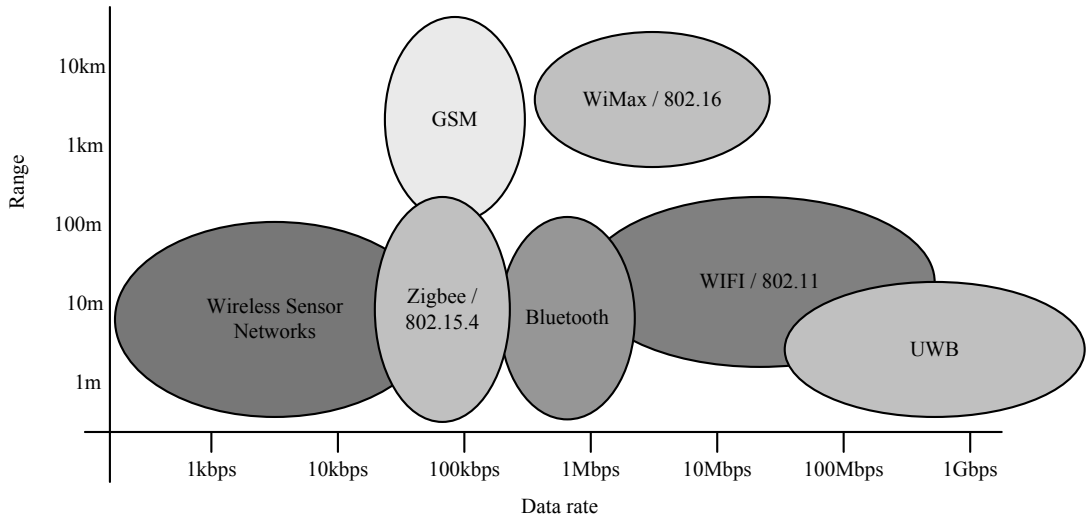

Figure: 2.2: Data rate and range of common digital communications standards.

2.2.3 Physical Design and Network Infrastructure

In many deployed networks environmental concerns will result in specific physical constraints on node design. A network deployed in a demanding, remote environment such as wildlife monitoring in a mountain range, would require robust nodes resistant to challenging weather conditions. In addition, the nodes would require sufficient lifetime to avoid the need for physical access to the network, for example to change batteries or extract data. In this case physical size is not such a demanding issue, however in some applications nodes will be severely size-limited, for example in a

network providing inconspicuous surveillance. Depending on the application a network can be ad-hoc, or can have a centralised structure with a central controlling entity. Both situations place particular needs on other aspects of the network, such as power source and communications. For example a central controlling entity can manage the scheduling of communications, allowing individual nodes to be as simple as possible, and to conserve energy while not communicating. Networks may consist of a handful of nodes, up to several hundreds or thousands [69].

2.2.4 Identification and Localisation

Nodes may need to be individually identifiable and location-aware for the sensory data that they collect to be meaningful. Both of these requirements will significantly increase network control signalling, particularly in an ad-hoc network where nodes need to collaborate to select identities.

2.2.5 Homogeneity

A network can be homogeneous where all nodes possess exactly the same functionality, or heterogeneous where there are specific node types for specific purposes. For example there could be a minority of nodes used for sensing with the remainder used to relay messages and/or process sensed information. A common scenario would be for a network to comprise specialist sensing, sinking and communication relaying nodes.

2.2.6 Data Transfer and Processing

A network can be organised into a source and sink arrangement, where some nodes gather sensory information, and other nodes communicate with the outside, physical world [55]. In a centralised network the controlling entity could be the destination for all data. At the other extreme, the network has no external communication ability or

requirement, in which case the nodes may possess actuators as well as sensors, in order to provide some useful function. Collected sensory data can be processed locally at each node, at some central processing location or in a distributed manner across the whole network. The processing power needed in each node will depend on the data processing model, and the extra cost involved with data transfer in a distributed model will increase the energy consumption in each node. These issues are covered in more detail in Section 2.6.

2.2.7 Sensors

The sensors deployed in a network are dependent on the network application, and many different types exist [49]. Common sensors include temperature and pressure sensors for monitoring applications, or accelerometers and gyroscopes in applications where the movement of a node must be tracked. In many cases the size of the sensor is very important, for example if it has to be embedded into the object it monitors. In all cases, power consumption is of paramount importance to prolong useful node lifetime.

2.2.8 Multiple Access Control and Associated Energy Consumption

It is generally accepted that communication dominates the energy budget of a WSN node [55]. As a result, communication must be minimised as much as possible, while not compromising the needs of the application supported by the network.

A very important function of the network is therefore the Multiple Access Control (MAC) protocol, which controls how and when nodes can access the communications channel. Communication between nodes is the cause of four major sources of energy wastage in a network [73]:

- **Collision** - when two or more nodes try to communicate at the same time, and packets are lost. Nodes are required to retransmit.

- **Overhearing** - a node wastes energy listening to a packet that is destined for another node.

- **Control Packet Overhead** - some MAC protocols require significant control signalling, which causes energy consumption in transmission and reception, and reduces the amount of useful data packets that can be sent.

- **Idle Listening** - when nodes listen for traffic that is never sent.

MAC protocols for WSNs strive to reduce the energy lost to the issues introduced above. A limited set of MAC protocols are considered in further detail in Section 2.5.3, however the interested reader can find an extensive survey of protocols in [28].

An aspect of MAC protocols that targets the reduction of power consumption is duty cycling. Many WSN MAC protocols aim to reduce power by allowing nodes to move into an idle state, in which they are not required to listen for channel activity. Figure 2.3 highlights the energy savings which can be achieved by duty cycling [55].

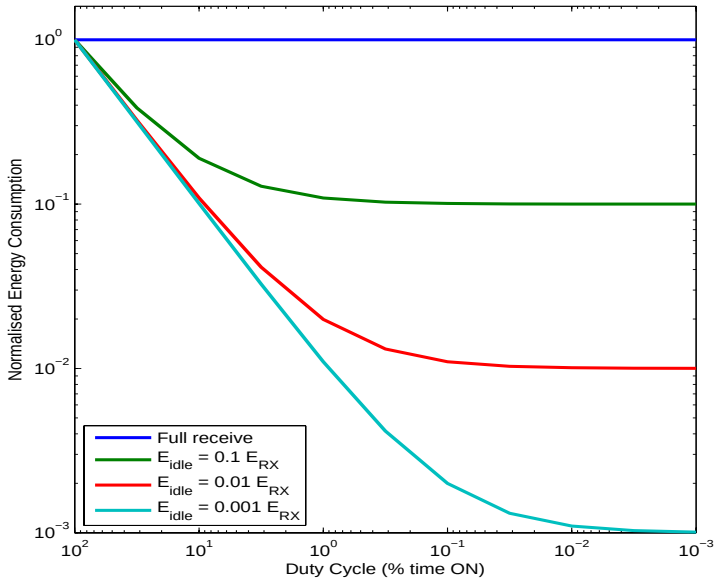

Figure: 2.3: Relationship between node energy consumption and duty cycle period / idle energy requirement.

The energy used for transmission E_{TX} is assumed to be equal to that for reception, E_{RX}, i.e. $E_{TX} = E_{RX} = 1$, and the idle energy consumption, E_{idle}, is shown in the figure as selected fractions of the receive energy consumption. The energy consumption for permanent receive is shown for comparison. The energy consumption can be reduced significantly by decreasing the duty cycle period (the fraction of time the node spends out of idle state, i.e. with receive circuitry turned on) and minimising the energy consumed while in an idle state.

2.3 State of the Art

With WSNs such a popular area of academic study, and deployed networks becoming more and more prevalent, it is interesting to examine currently existing sensor network technology. A selection is presented in the following sections.

2.3.1 IEEE 802.15.4 and ZigBee

ZigBee [95] is a high-level standard that builds on the IEEE 802.15.4 low-power, long battery life, Direct Sequence Spread Spectrum (DS-SS) based communications standard [80]. IEEE 802.15.4 is generally aimed at low data rate, short to medium range applications such as home automation, embedded sensing or industrial control, and as such is well-suited for deployment in WSNs. ZigBee radios are popular in commercially available nodes, including some of the examples detailed below.

2.3.2 Berkeley Wireless Research Centre

The Berkeley Wireless Research Centre, [76], has become one of the foremost centres for WSN research. Research is conducted across all protocol layers, from physical layer communications up to applications. Two of the best-known projects are the Smart Dust and PicoRadio projects.

Smart Dust

The Smart Dust project began in 1997 with the goal of creating cubic millimetre nodes to form a sensing and communication system. Although as yet the goal of such small devices has not been reached, the project resulted in the creation of the TinyOS wireless sensor node operating system [88], and, in 2002, the establishment of Dust Networks [78], which now produces a range of wireless nodes for use in industrial automation. In particular, the company has developed the Time Synchronised Mesh Protocol (TSMP), a robust networking protocol that enables a self-organising and self-healing network. One of their products, the SmartMesh-XT M2030, is shown in Figure 2.4. This node uses the TSMP protocol combined with an IEEE 802.15.4 radio, and can be used in a plug-and-play manner with required sensors or actuators.

Figure: 2.4: Dust Networks SmartMesh-XT M2030 wireless sensor node.

PicoRadio

The PicoRadio project ran from 1999 until 2005, with the goal of investigating low-cost, low-power transceivers for WSNs. Research into many aspects of low-power WSNs was carried out, including synchronisation [3], and RF oscillator design [44].

2.3.3 Crossbow

Crossbow [77] offer a selection of different wireless sensor modules, the latest of

which is the Iris, aimed at low-power WSNs, and is shown in Figure 2.5. The Iris contains an IEEE 802.15.4 compatible radio transceiver, and has expansion connectors for various sensor boards.

Figure: 2.5: Crossbow Iris WSN module [15].

2.3.4 SunSPOTS

Sun Microsystems have developed a node called the SunSPOT [84] which comes in a single package as shown in Figure 2.6. The SunSPOT encompasses an IEEE.802.15.4 compliant radio, and comes with a selection of onboard sensors including an accelerometer and a temperature sensor. The SunSPOT runs a JAVA based operating system.

Figure: 2.6: Sun Microsystems SunSPOT node [60].

Most of the WSN nodes discussed above are in fact relatively similar. Many of the boards are capable of connecting to different sensor boards to provide a wide range of functionality. The Crossbow node, the SunSPOT node and the Dust Networks M2030 node all use the Zigbee radio standard, which as shown in Figure 2.2 transmits over ranges of up to 100m with data rates of up to 100kbps. The Berkeley PicoRadio project aimed to develop even lower power radio technology to transmit over shorter distances and at lower data rates, much like the Speckled Computing project as explained below.

2.4 Speckled Computing

Speckled Computing is a vision of ubiquitous computing, where the digital world is fully embedded into the environment [4,69,83]. Discrete digital devices such as PDAs, mobile phones, computers or memory devices converge and disseminate into the environment, existing in everyday objects. Thousands of tiny nodes, called Specks, form networks called Specknets which provide the building blocks to bridge the divide between the material and the digital worlds.

Speckled Computing is being developed by the SCC [83], a collaboration between five Scottish Universities. The Consortium partners and their respective areas of expertise are detailed below:

- **University of Edinburgh** - Institute for Computing Systems Architecture. Micronet-based asynchronous programmable network architectures [89].

- **University of Strathclyde** - DSP Enabled Communications Group. Wireless digital communications at the physical layer [93].

- **University of Glasgow** - Ultrafast Systems Group. Design, fabrication and on-wafer testing of microwave and millimetre-wave components; RF circuitry design [90].

- **University of St Andrews** - School of Physics and Astronomy [91] and School of Chemistry [92]. Photonics, digital communication by laser and LED. Materials chemistry and solid state electrochemistry (rechargeable energy sources) .

- **Napier University** - Centre for Emergent Computing. Application-level research, self-evolving systems [80].

Figure 2.7 gives a system-level overview of Speckled Computing. Each Speck contains its own source of energy, processor and means of communication, and is small enough to become inconspicuous within an environment. Ultimately, each Speck will be as small as 1 cubic millimetre, resulting in computing power being truly embedded into the environment.

Individually, the Specks have limited processing power, however together they will collaborate and are capable of more powerful computation. The long term goal is of a Specknet as a fully distributed, homogeneous network; each node is exactly the same

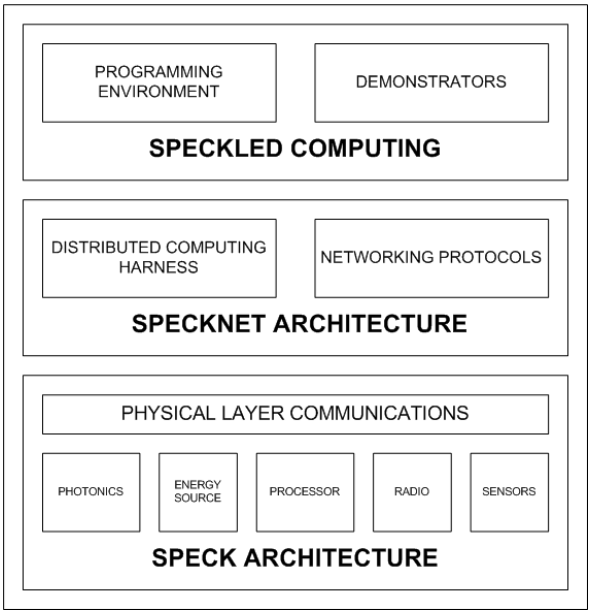

Figure: 2.7: Overview of Speckled Computing, reproduced from [5], with permission.

and as a result there is no central control available to perform processing, schedule communications or assign identities. Specks will communicate wirelessly by radio and optics (laser/infrared) over ranges of the order of tens of centimetres. The severe size restrictions on a Speck result in minimal energy availability, and therefore low-power design is the overriding concern in all levels of the protocol stack. From application development right down to physical layer communications architecture design, energy consumption must be minimised.

The ability to provide computing power embedded discreetly in an environment opens up a whole spectrum of applications. From storing data or embedding intelligence in everyday objects, to unobtrusive condition monitoring of industrial plant, to real-time patient health monitoring, a whole new realm of exciting uses is possible.

2.5 Speck Development

As mentioned above, low-power design is the overriding concern when implementing Specks and Specknets. In this section the progress to date in the low-power implementation of Specks and Specknets is detailed, from power sources to physical layer communications architectures.

2.5.1 Prototype Development

There have been several Speck prototypes produced by the Consortium. Off-the-shelf components (radio, processor) have been used in early prototypes, with custom components intended to be introduced as they are designed. This section details a selection of the prototypes produced to date.

ProSpeckz

The ProSpeckz (Programmable Specks over Zigbee Radio) series of prototypes

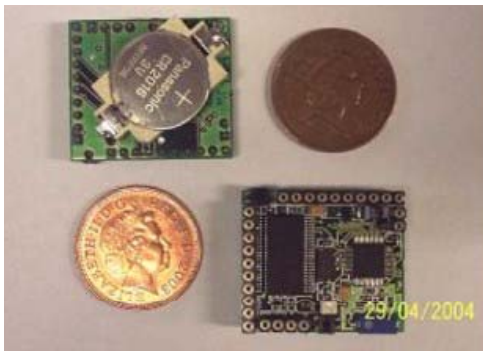

Figure: 2.8: The ProSpeckz-2, an early prototype, reproduced from [5], with permission.

were produced with off-the-shelf components in order to develop and test networking protocols and application concepts. In particular they have been used to examine how Specks can be placed in consumer electronics applications and in verification of the performance of the SpeckMAC protocols [5,70]. The ProSpeckz nodes were also integrated into an 8x8 array, known as the Perspeckz64, for analysis of networking and MAC protocols and visual demonstration purposes [71].

5Cube

Two 5Cube variants have been produced, which are both attempts at creating a 5x5x5mm cubic Speck. The 5CubeOTS is constructed entirely from off-the-shelf components while the 5Cube, based on the 5CubeOTS, replaces various components with custom-designed circuits, such as the early DSP transceiver, which will be covered in Section 2.5.3.

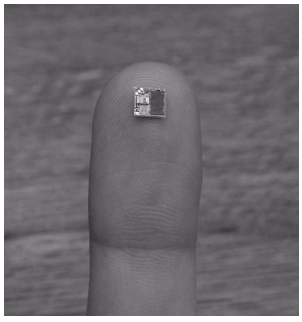

Figure: 2.9: The 5CubeOTS prototype.

Orient

An offshoot of the ProSpeckz prototypes, the Orient series of nodes include accelerometers, gyroscopes and magnetometers to provide a body posture tracking capability. The Orient-2 will be covered in more detail in Section 2.7, while at the time of writing the Orient-3 is in development.

Energy-neutral Platform

A flat form-factor node is currently being investigated as a means to produce an energy-neutral platform. The flat form factor allows a large solar cell to be placed on the surface of the node, and with heavy duty-cycling, the node will be able to function indefinitely with the only energy input coming from solar energy.

2.5.2 Power Source

Limited by the small volume available, the power source used in a Speck has to be of high energy density and rechargeable. As such lithium ion (Li-ion) sources have been investigated. These batteries can be cycled efficiently thousands of times and can possess volumetric energy densities of up to $400mAhl^{-1}$ [61]. For comparison, Nickel Cadmium (Ni-Cd) and Nickel Metal Hydride (Ni-MeH) batteries have energy densities of $150mAhl^{-1}$ and $200mAhl^{-1}$ respectively. The 5Cube prototype uses the smallest commercially available Li-ion cells, the Sanyo ML414 [56], which have a height of 1.4mm and a diameter of 4.8mm. For future implementations other battery systems are being investigated, for example nanotube/wire based batteries and thin-film batteries with potentially double the energy capacity of conventional Li-ion power sources [4].

2.5.3 Multiple Access Control

The MAC layer is extremely important in a Specknet where there is a high density of communicating nodes, and as such significant research into low-power protocols

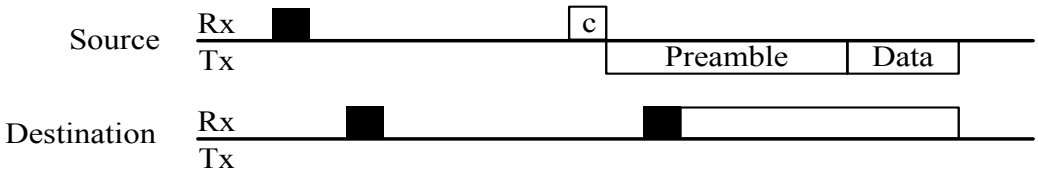

Figure: 2.10: Generic duty-cycled MAC timeline.

has been conducted. Improving on the well known B-MAC protocol [46], a selection of protocols based on Carrier Sense Multiple Access (CSMA) have been developed, [14,70,72]. All protocols are fully distributed, unsynchronised random-access protocols, and are based on nodes periodically waking up and listening to the channel for activity. This allows nodes to cut energy consumption by duty cycling.

With nodes waking at random times, a transmitting node has to transmit over a certain period of time to ensure the destination node will wake and hear the transmission. This concept is shown in Figure 2.10, where the transmitting node samples the channel, and then transmits a preamble signal. The receiving node, on hearing the preamble, will prepare for packet reception. The different protocols use a different preamble format, which instruct the receiver in how to receive the packet. B-MAC transmits a simple preamble of length longer than the duty cycling period. On reception of the preamble the receiver will continuously receive data until the packet is received. The SpeckMAC protocols improve on this concept by allowing the receiving node to spend more time in an idle state, and thus conserve energy. SpeckMAC-B sends a preamble of small wake-up packets, which contain information about when the data packet will be transmitted. The receiving node, on hearing a wake-up packet, will then go back into an idle state and only wake up for the data packet. SpeckMAC-D retransmits the data packet repeatedly, so that when a receiver senses activity, it can listen to the data packet straight away, and then go back into an idle state. SpeckMAC-H is a hybrid combination of SpeckMAC-B and SpeckMAC-D. SpeckMAC-C, SpeckMAC-CB and SpeckMAC-CD, versions of BMAC, SpeckMAC-B and SpeckMAC-D respectively, use Code Division Multiple Access (CDMA) signalling to allow nodes to determine in a shorter time whether they are the intended

destination of the packet, and if not return to an idle state as soon as possible.

Figure 2.11 shows the possible energy savings using the above protocols for an example set of node density and communication density parameters [14]. There is an optimum duty cycle period, $T_{interval}$, for each protocol, however the graph clearly illustrates the energy savings that are possible through careful MAC protocol design.

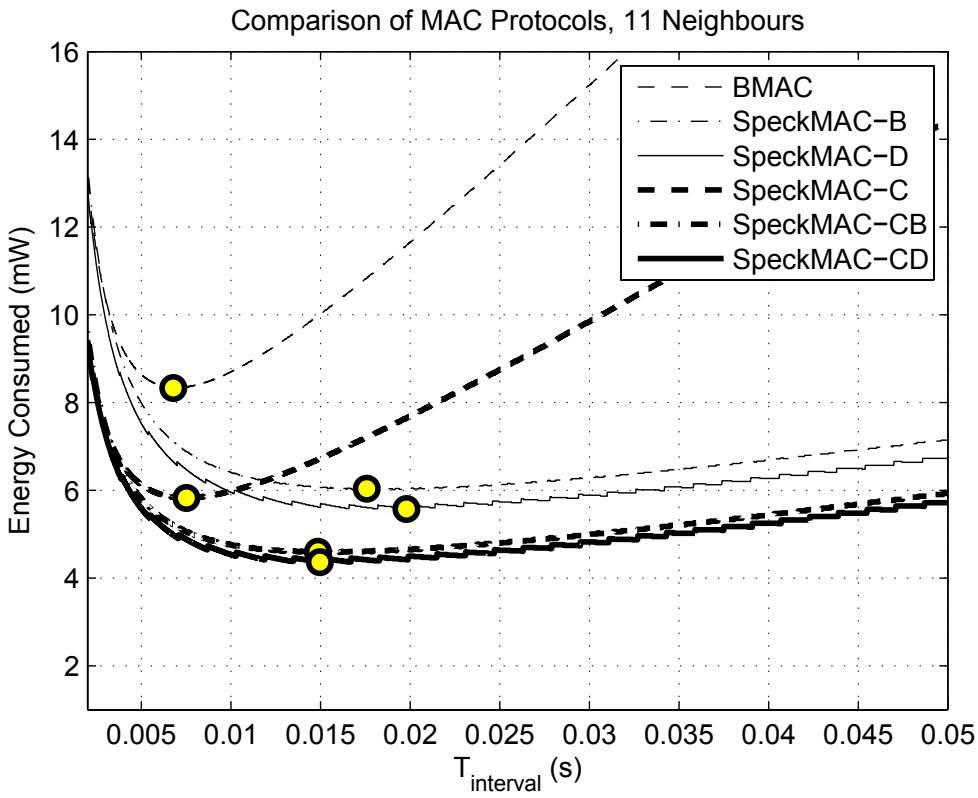

Figure: 2.11: Energy consumption of various MAC protocols. The yellow dots highlight the optimum duty cycle periods. Reproduced from [14], with permission.

2.5.4 Physical Layer Communications

Data rates in a Specknet are in general low (commonly of the order of a few tens of kbps, or at most a few hundreds of kbps) and transmission ranges are short (the dense nature of the network will result in neighbouring Specks being as close as a few centimetres, with a maximum considered range of a few metres). Both laser and radio

communication are being investigated for transmission between Specks. Laser-based communication is also being considered for communication between Specknets, and as such may transmit over a few tens of metres. Physically compact designs are required to comply with the tight size constraints, which of course also result in the need for minimum power consumption due to the small size of the energy source.

Choice of RF Frequency and Antenna Design

The choice of RF carrier frequency is dependent on a few competing parameters. Firstly, the small scale of the Specks requires that the antenna size be as small as possible so as to fit into the fabricated package. With antenna size requirements being inversely proportional to RF frequency this indicates a choice of a high frequency. However, free-space path loss becomes more significant at higher frequencies and thus increased transmit power is required for equivalent performance at lower frequencies. This trade-off is being explored with investigation of radio and antenna designs at a range of frequencies from 2.45 GHz to 24GHz [62].

Laser-based Communication

Optical communication is under investigation as an alternative to radio for communication both between Specks and between Specknets. Laser offers the possibility of reduced power communications due to the collimated nature of its emissions. Laser steering techniques and multiple emitters on each Speck will allow lasers to be directly targeted at a receiver, with minimal energy lost to the environment, and few unintended recipients. This compares favourably to radio transmission, where the (ideally) spherical shape of the emission results in many unintended recipients and inefficient transmission. This is shown in Figure 2.12, where the radio transmission encompasses several unintended recipients and the laser transmission is directly targeted at the desired receiver. This attribute allows laser to communicate over longer distances than radio for the same energy consumption, and therefore lasers are also

being investigated as a means to communicate between separate Specknets. The drawback to this property is of course the need to ensure that the laser is accurately targeted. This can result in increased node complexity, due to the need to either steer the laser beam, or implement multiple transmitters.

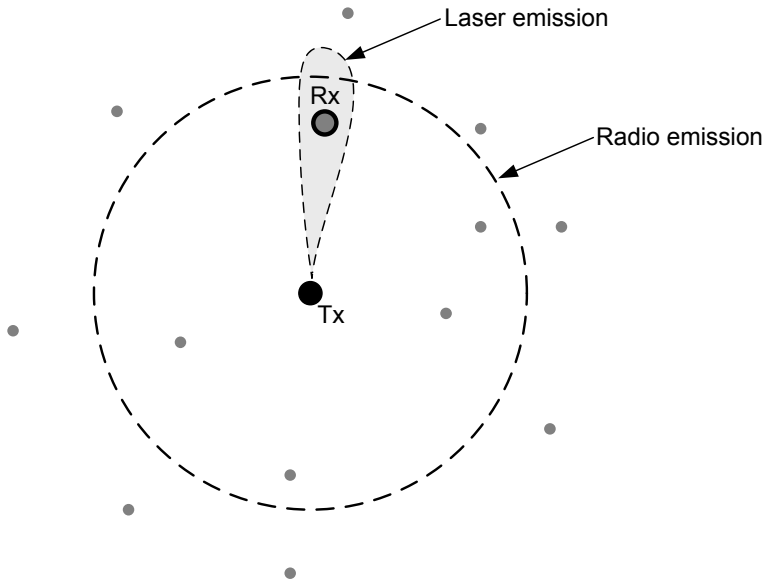

Figure: 2.12: Comparison of the directivity of laser and radio emissions.

Communications Channel Analysis

Specknets will be deployed in many different environments, and will in general transmit over very short distances. Investigation has therefore been undertaken into the characteristics of the communications channel to inform the design of physical layer communications protocols [16]. The resulting channel model is shown in Figure 2.13. Results have shown that over short ranges (tens of centimetres), and at data rates of up to 200kbps, the channel can be considered narrowband, subject to flat fading and with no requirement for equalisation [16]. Path loss was investigated for transmissions made between Specks placed on various surfaces. The likely noise and interference experienced by a Speck-to-Speck transmission using a CSMA MAC protocol was also

analysed, informing design choices in implementing the clear channel assessment stage of the CSMA protocol. Further details can be found in [16].

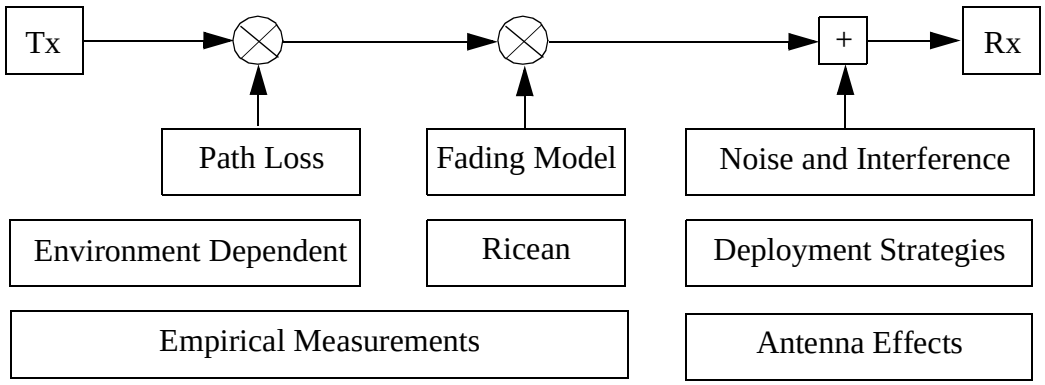

Figure: 2.13: Narrowband channel model for Speck to Speck communication, reproduced from [16], with permission.

Early Radio Design

The first generation radio circuit was designed to minimise complexity and limit power-hungry analogue components such as mixers. An On-Off-Keyed (OOK) system was produced using a simple envelope detector-based receiver, which avoided entirely the need for a digital to analogue converter (DAC) or analogue to digital converter (ADC). The first generation radio circuit is shown in Figure 2.14.

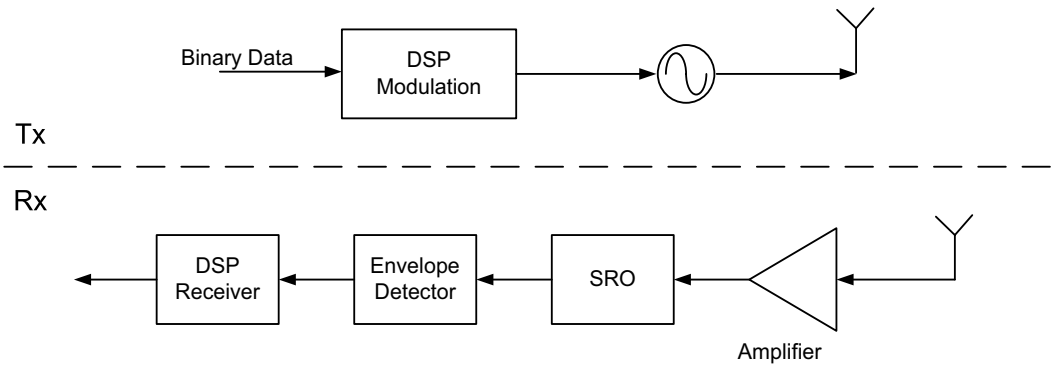

Figure: 2.14: Super-regenerative oscillator (SRO) based OOK transmitter and receiver.

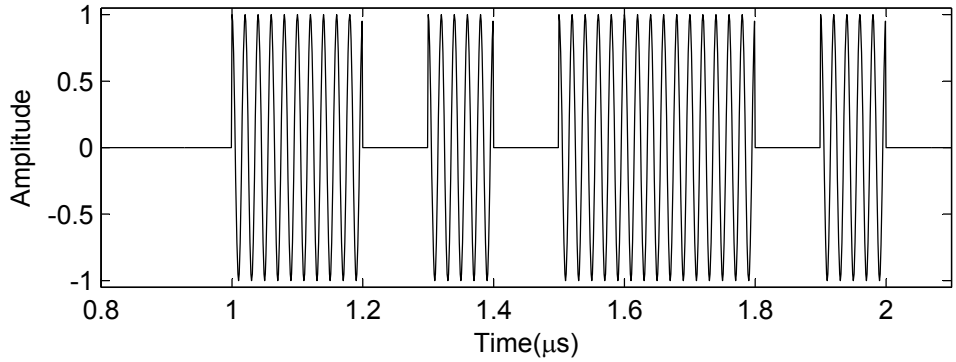

Figure: 2.15: OOK modulated signal.

The OOK transmitter allows for minimal power consumption, with the RF oscillator being switched on and off in synchrony with the baseband binary data, producing an OOK signal as shown in Figure 2.15. In the receiver, the received signal is regenerated using a super-regenerative oscillator (SRO), brought down to baseband non-coherently by envelope detection and sliced to recover the digital signal, which is presented to the DSP receiver [35,36]. Manchester encoding was utilised to minimise synchronisation overhead in the receiver by embedding the clock in the transmitted signal. This results in an increase to the bandwidth of the transmitted signal by a factor of 2, as shown in Figure 2.16, however this technique was deemed suitable for use in a Specknet due to

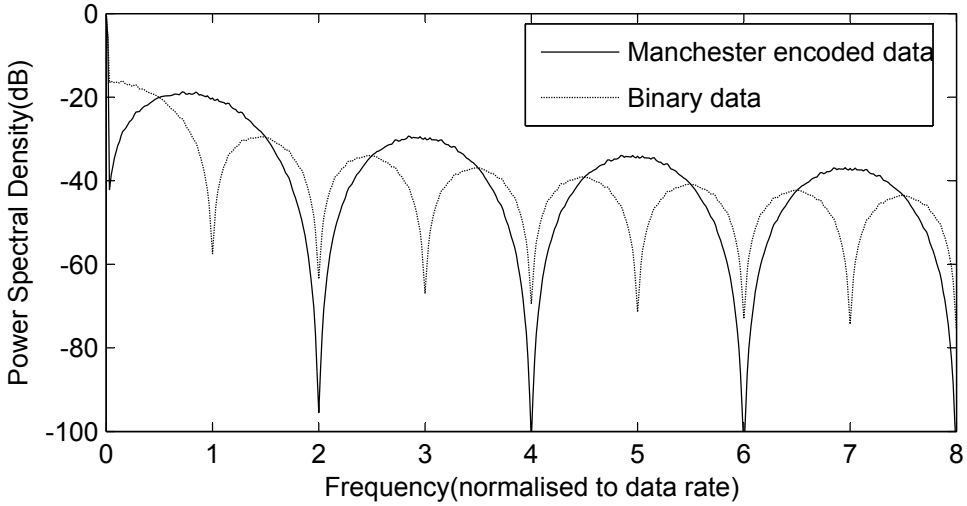

Figure: 2.16: Comparison of bandwidth of binary sequence and Manchester-encoded binary sequence.

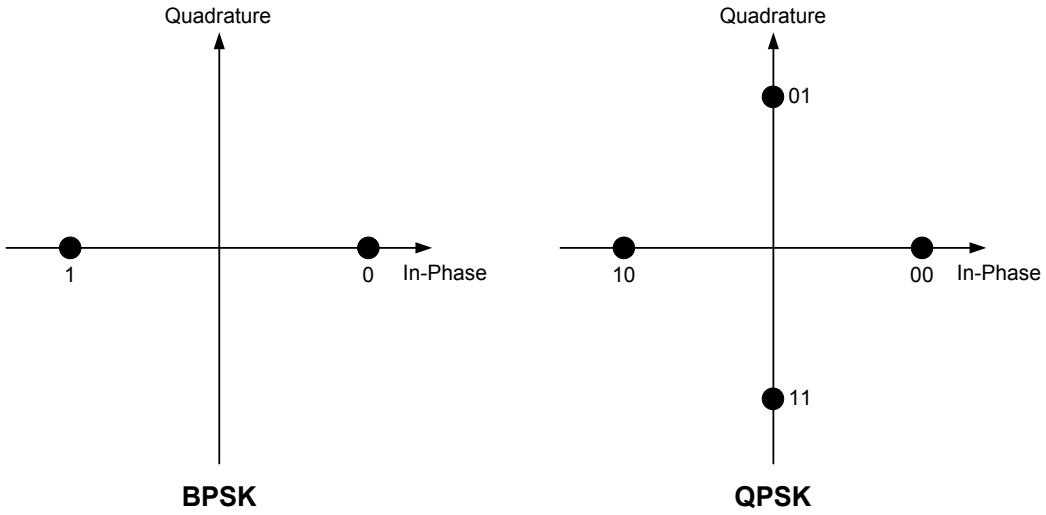

Figure: 2.17: Constellation diagrams for BPSK and QPSK modulation.

the low data rates involved. The implemented Manchester decoder has the added benefit of being tolerant to clock frequency inaccuracies, which enabled the use of a cheap clock source, again minimising complexity and power consumption [34].

Future Radio Design

The current OOK transceiver limits potential to improve physical layer performance using DSP techniques. Different modulation schemes give significantly different performance curves. The theoretical probability of bit error for a non-coherent OOK receiver, as used in the first generation radio design is given by [58]:

$$P_e = Q\left(\sqrt{2\frac{E_b}{N_0}}\right) \tag{2.1}$$

where $Q(x)$ is the complementary error function,

$$Q(x) = \frac{1}{\sqrt{2\pi}}\int_x^\infty \exp\left(-\frac{u^2}{2}\right)du \tag{2.2}$$

and $\frac{E_b}{N_0}$ is the signal to noise ratio normalised to data rate and bandwidth, given by:

$$\frac{E_b}{N_0} = \frac{P_S W}{P_N R_b} \tag{2.3}$$

where P_S is the signal power, P_N is the noise power W is the signal bandwidth, and R_b is the bit rate. Binary Phase Shift Keying (BPSK) or Quadrature Phase Shift Keying (QPSK) are more complicated modulation schemes where the data is encoded into the phase of the carrier wave, as opposed to the amplitude of the carrier as in OOK. The constellation diagrams for BPSK and QPSK are shown in Figure 2.17[1]. The carrier phase is modulated according to the input bits. The probability of bit error for coherent BPSK is equivalent to that for QPSK and is given by:

$$P_e = \frac{1}{2}\exp\left(-\frac{E_b}{2N_0}\right) . \tag{2.4}$$

The bit error probabilities for OOK and BPSK/QPSK are plotted in Figure 2.18. It is clear that significant performance improvement can be achieved by switching to

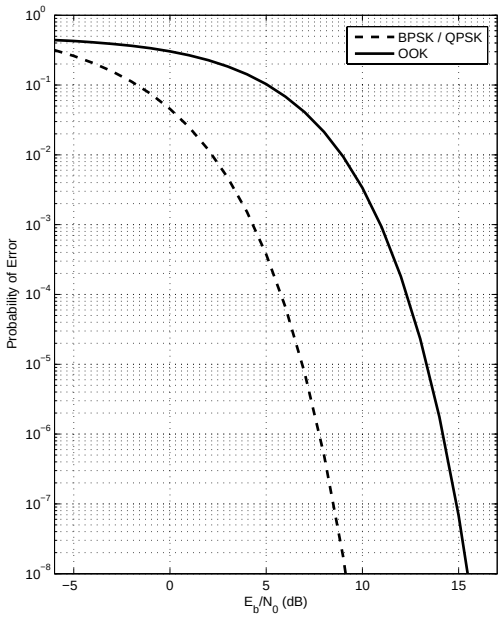

Figure: 2.18: Comparison of BER curves for BPSK/QPSK and OOK.

1. An alternative QPSK constellation can be found by shifting the shown QPSK constellation by $\Pi/4$.

BPSK/QPSK from the non-coherent OOK scheme currently used. The drawback to this switch however is increased complexity in both the transmitter and the receiver. Both the transmitter and the receiver will require a mixer architecture, and there will be significant synchronisation requirements in the receiver. QPSK involves two mixers in the transmitter and receiver to implement the necessary quadrature modulation.

Work is ongoing within the Consortium to design low-power quadrature oscillators and mixers [24] which would support phase modulation and hence enable more complex DSP techniques to be used . CDMA has already been discussed in Section 2.5.3 as a means to improve the energy consumption performance of the SpeckMAC protocol, however it has also been proposed as a method for allowing nodes to transmit in close proximity with lower power, providing some control is exercised over relative positioning of Specks [14].

2.6 Data Management in Wireless Sensor Networks

The goal of every sensor network is to gather information from the environment in which it is deployed. Individual nodes in the network produce data with a frequency dependent on the network's defined application, which also defines how the data are treated. The application's requirements place subsequent demands on the network and its nodes, and define whether the data has to be extracted from the network. An application may be such that the network can make its own decision on the data and act accordingly. On the other hand, the data may need to be extracted from the network and analysed remotely. Extraction can take two principal forms. The first involves the data being extracted via a single or multiple nodes (sinks) reached via a multi-hop route across several links. The second involves the data being transmitted directly from each individual node to a central information gatherer. The following example network scenarios illustrate some of the wide ranging ways with which data may be treated.

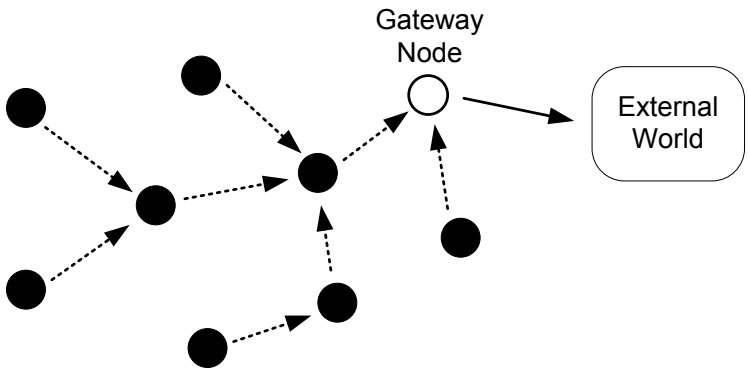

Figure: 2.19: Network with gateway node.

2.6.1 Sensor Network With Gateway Node

This sensor network operates independently by gathering data, which is then relayed via a specific node (or nodes) to an external data gathering centre for analysis. This system setup is shown in Figure 2.19 where all sensing nodes in the network direct their data to a single gateway node which in turn relays the network data to an external entity. This type of network can be seen in [61] where a network set up to monitor volcano activity has individual nodes collecting data, forwarding it to a single gateway node which then relays the data to a remote database.

2.6.2 Sensor Network With Fusion Centres

In this sensor network, as shown in Figure 2.20, data gathered by the nodes is collected at local fusion centres where any processing, extraction or storage of data takes place. Commonly, the fusion centre node will possess a higher level of capability than the other sensor nodes, i.e. a higher processing power and/or communication ability/range. To provide this functionality it may be that the fusion centre requires wired power, limiting the application domain of such a network.

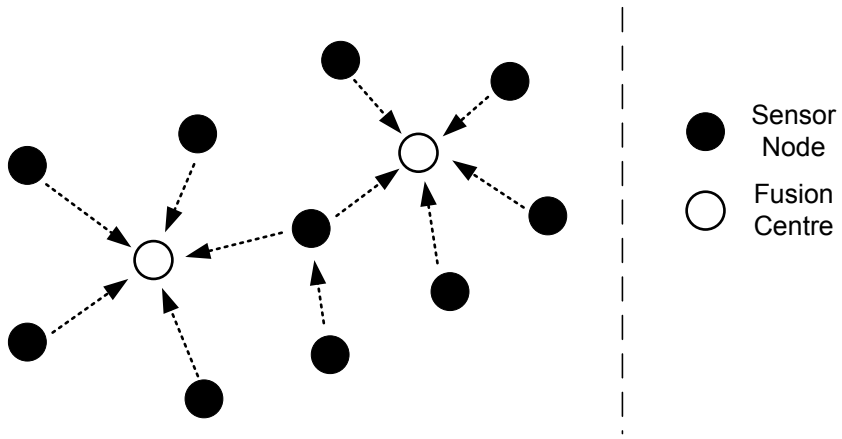

Figure: 2.20: Network with local fusion centres.

2.6.3 Fully Self-Contained Sensor Network

This sensor network also operates independently gathering data, however it is dedicated to one, or at most a few, applications which allow it to be completely self-contained. That is, the network acts on the sensory data it collects without passing the data to any other entity. A simple example of this type of network could be a fire monitoring network in which each node monitors its local area, and triggers some alarm on fire detection, with nodes operating cooperatively to reduce false alarm probability. Figure 2.21 shows an example of such a network, where nodes communicate as required to share information and make joint decisions.

2.6.4 Remote Sensor Network With Periodic Data Extraction

A set of possible sensor network applications including, for example, environmental monitoring, battlefield surveillance or planetary exploration result in a remote network where it is impossible, or unnecessary, to continuously extract information from the network. In this case the nodes in the network will store data until an external source is ready to extract it. This external source can take the shape of a

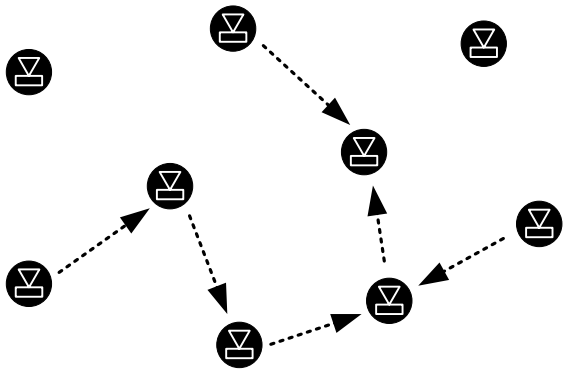

Figure: 2.21: A fully self-contained network with no link to an outside world; each individual node has an output device, such as an alarm.

stationary or mobile (robotic or human transported) sink. This application area is seen in the study of maximising data extraction [54], as over a period of time some nodes in the network may expire, and the data can be lost.

2.6.5 Leaf Network

As shown in Figure 2.22 the nodes in this sensor network report directly to a central

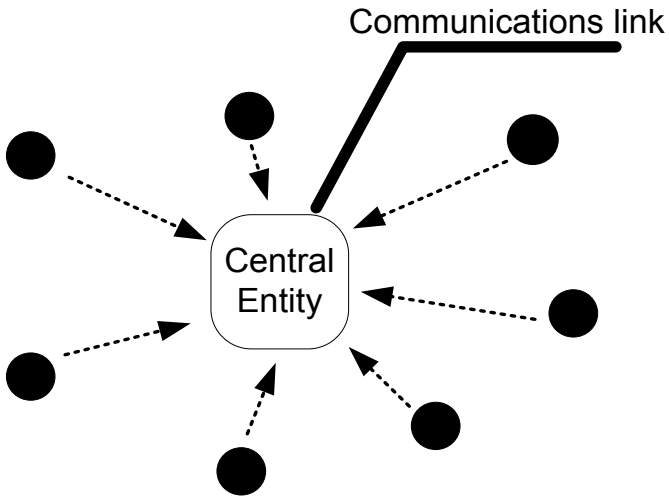

Figure: 2.22: Leaf network.

entity without using any form of hopping communication. The central entity is commonly an energy-rich node with communication links to the external world, or in some cases is even external to the network itself, and the data flow can be frequent, as in the gateway network scenario, or infrequent, as in the remote network scenario.

2.6.6 Latency and Update Rate

Some sensor network applications will place requirements on the latency and update rate of data being transmitted to the required location. For clarity, we refer to latency as the time in seconds between sensor data being generated and it being received, and update rate as the frequency in Hertz at which the system generates a set of data. Some systems will require real-time or near real-time performance, where data is generated and then processed as quickly as possible, giving low latency requirements. Update rates will vary depending on the application, for example a temperature monitoring network may only require data to be generated perhaps every minute, while a motion tracking application could require several updates per second, depending on the speed of the motion.

2.6.7 Effect on Communication Costs

Communication between nodes has been shown to be the most energy expensive requirement in a WSN [55]. Therefore the communication burden on a node must be minimised to prolong lifetime. In the worst case, raw sensor data has to be gathered at a central point via a multi-hop link, which requires nodes to act as relays, increasing the communication cost on the network as a whole. In [48] it is shown for one scenario that 3 million instructions could be executed by a single node for the same energy cost as it would take to transmit 1000 bits over 100m. This result shows that the processing power of the nodes can be used to reduce the amount of communication that must take place. That is, the processing of the data can be undertaken locally at nodes, either individually or with a limited amount of near neighbours, minimising the amount of

data that must be transmitted. In addition, it can be assumed in some sensor network applications that data generated by the many sensors will be highly correlated, and can thus be reduced using data aggregation techniques [22,30]. Traffic can also be reduced by careful design of data dissemination protocols [23,25]. In some applications, however, in particular where energy constraints on the nodes are very strict, expensive multi-hop communication can be avoided by extracting data (raw or processed/ compressed) directly from nodes to a central location. This system shifts the energy requirements and complexity to the data gatherer, which may not be energy limited, and allows individual network nodes to be as simple as possible [59].

2.6.8 Managing Data in a Specknet

As seen in Section 2.5, nodes in a Specknet are heavily energy constrained, and minimising inter-Speck communications as much as possible is an advantage. In suitable applications, a dense network of Specks will communicate amongst themselves as little as possible, and transmit data to a central location as and when it is required. This allows individual nodes to expend minimal energy on communications, and thus prolong their lifetime. One application of a Specknet in which data are extracted to a central location for analysis is a body posture tracking system. Such a system provides motivation for creating an efficient method of data extraction, as it possesses the usual energy limitations on the nodes, as well as requirements on the latency and update rate of transferring data to a central location. The Orient system possesses all these characteristics, and has been the stimulus for the application-driven research presented in this thesis. The Orient-2 system is analysed for data transfer requirements in more detail below.

2.7 The Orient-2 Body Posture Tracking System

The Orient-2 system [74,75] is a real-time posture tracking system that has been

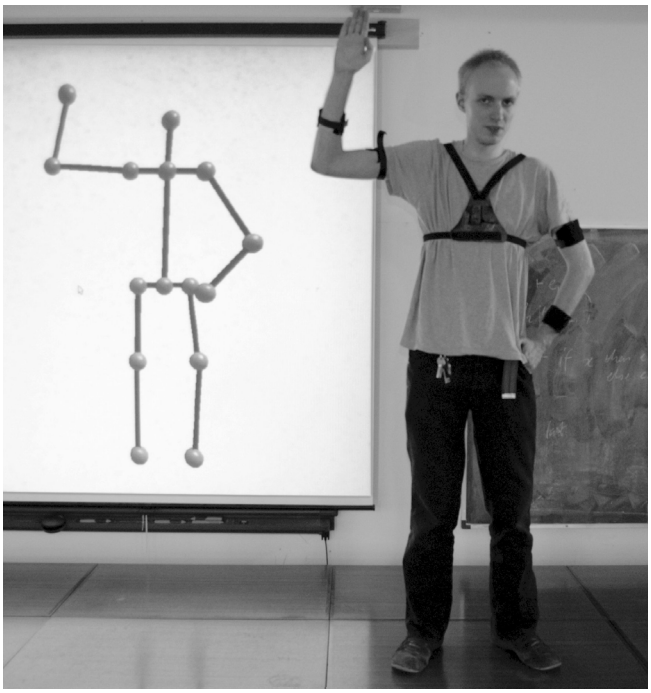

Figure: 2.23: A snapshot of body motion tracking using the Orient-2 System, reproduced from[74], with permission.

developed within the Speckled Computing Consortium. The development of the device has provided valuable insights into chip manufacture, network and physical layer issues, application development, and many other factors involved in device conception. In particular, the system has interesting requirements on the treatment of the data produced by the orientation sensors.

Although the eventual aim is to have a body covered in many tiny specks, the current system allows for full body motion tracking using a network of 15 devices, with motion sensing achieved using data from an accelerometer, two magnetometers and three gyroscopes. Data are gathered at a central location, and mapped onto a rigid body model, as shown in Figure 2.23. A single device is shown in Figure 2.24. The device comprises a 16-bit microcontroller, a 32kHz crystal oscillator, a lithium polymer battery, a 250kbps transceiver and the sensors. Further details can be found in [74].

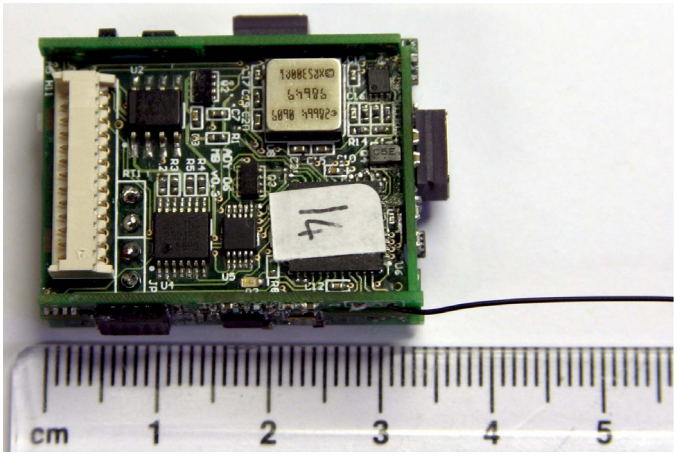

Figure: 2.24: Orient-2 device, reproduced from [74], with permission.

2.7.1 Data Management

The devices are arranged in a leaf network and each node sends its sensory data to a central device (in this case, a laptop) which computes the body model. A CSMA scheme is used for system setup, before the network implements a Time Division Multiple Access (TDMA) scheme during motion capture. In the implemented system, nodes compute an orientation estimate locally and transmit this to the central host, reducing the required information transfer rate significantly [74].

In a fully centralised system, where nodes transmit raw sensor data to the central location, the raw data rate is around 19kbps at an update rate of 180Hz (nine 12-bit numbers are sent, 180Hz is commonly used in commercially available systems [26]). However, a rate of 64Hz is deemed sufficient for animation purposes, and by performing some in-situ processing to transmit only the orientation estimate, the data rate is reduced to around 4kbps per device (four 16-bit numbers are transmitted), only 21% of the raw rate. This is known as a semi-distributed system.

The remaining option, a fully distributed system, where the body model would be computed in-network, is undesirable because of the significant requirement for communication within the network, as nodes must now receive orientation estimates from other nodes, as well as transmitting their own estimate. Assuming packet

reception energy cost to be approximately the same as packet transmission cost (as is the case for the radio currently used in the system), this results in an approximate doubling of overall communication energy consumption [75]. It is clear therefore that the semi-distributed system is preferable.

2.7.2 Current System Limitations

The current implementation is limited in terms of update rate and latency by the time-shared use of a single radio channel by multiple devices. These problems can be mitigated to some extent by increasing the data rate, however this approach has obvious implications on power consumption, which will rise with communication rate. Another approach is to implement multiple, parallel radio channels, for example by using Frequency Division Multiple Access (FDMA). This would enable the individual devices to remain relatively simple and low-power, while pushing any additional complexity onto the data gatherer.

The system is also currently limited in tracking multiple bodies. Sharing the single radio channel not only between individual sensor devices but also between multiple users would result in large latencies and reduced achievable update rates. Again, multiple radio channels can be used to mitigate these limitations. The multiple channels can be used in two different manners. Firstly between individual nodes on one body, termed the Personal Area Network (PAN). In this case latency can be reduced when forming a snapshot of the posture of a single body, i.e. all nodes can report to the receiver in parallel and higher update rates can be achieved. In the second case the multiple channels can be used between multiple body networks reporting to a central location. A network of different bodies is termed a User Network (UN).

The two network scenarios are summarised in Table 1. Using different combinations of TDMA and FDMA gives different latency performances for each scenario.

PAN	UN	Resulting network
TDMA	TDMA	No guarantee on latency
TDMA	FDMA	Low latency across users
FDMA	TDMA	Low latency for individuals

Table 1: Summary of network scenarios and resulting latency performance

2.7.3 Using Multiple Frequency Channels for Data Extraction

As highlighted in Table 1, FDMA can be used to improve the latency, and associated update rate, performance of the Orient-2 system. Indeed, this approach can be used in any case where it is useful to extract data from multiple nodes concurrently. Figure 2.25 illustrates the case where N_u nodes desire to transmit to a central location with an update rate of $R_{up} = \frac{1}{T_{up}}$. Here T_{pkt} is the packet duration for each node, which is dependent on packet length and bit rate, and there are N_{ch} channels available. The maximum possible update rate is:

$$R_{up} = \frac{N_{ch}}{N_u T_{pkt}}. \tag{2.5}$$

The limitations of having only a single channel are illustrated in Figure 2.26. In the figure, (2.5) is plotted for a selection of radio bit rates, a packet length of 256 bits and a single channel, i.e. $N_{ch} = 1$. For a certain desired update rate, there is a point beyond which a single channel does not provide sufficient capacity, and multiple channels must be introduced. For example, the Orient-2 update rate is limited to 64Hz by the radio bit rate, as shown on the graph by the dotted line. For this given packet length, and fifteen nodes, it would be impossible to implement the TDMA scheme currently used with anything less than a 250kbps radio. In systems where radio bit rates are limited, multiple channels become a necessity where there is a requirement for a certain update rate. In addition, assigning a channel to each node enables the minimum

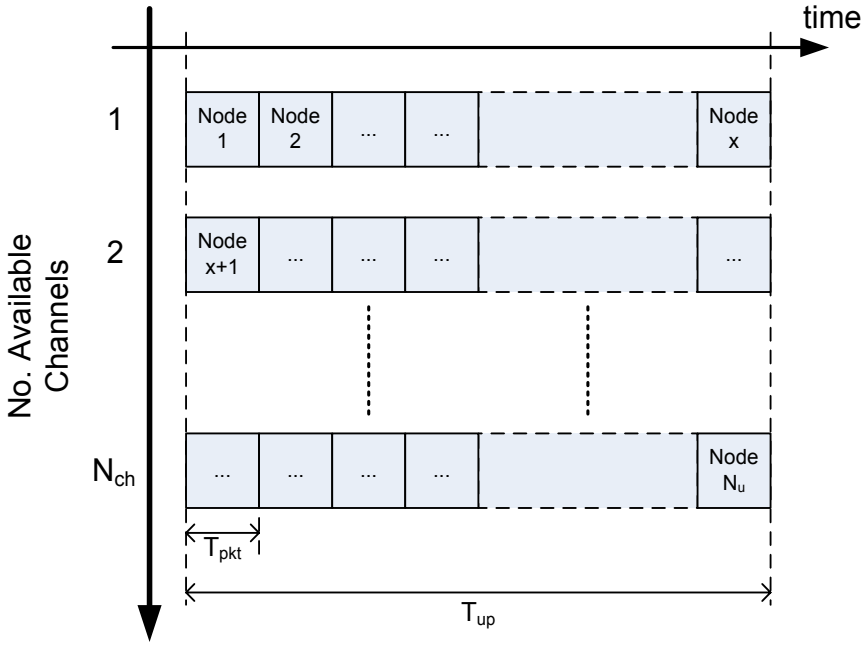

Figure: 2.25: N_u users transmitting to a central location over multiple channels.

possible extraction time, by allowing nodes to transmit in parallel and minimising the network-wide data extraction time period to the minimum value of T_{pkt}. This allows the maximum update rate to be achieved, which is limited only by the amount of data that have to be extracted, rather than by any radio or modulation scheme performance characteristics.

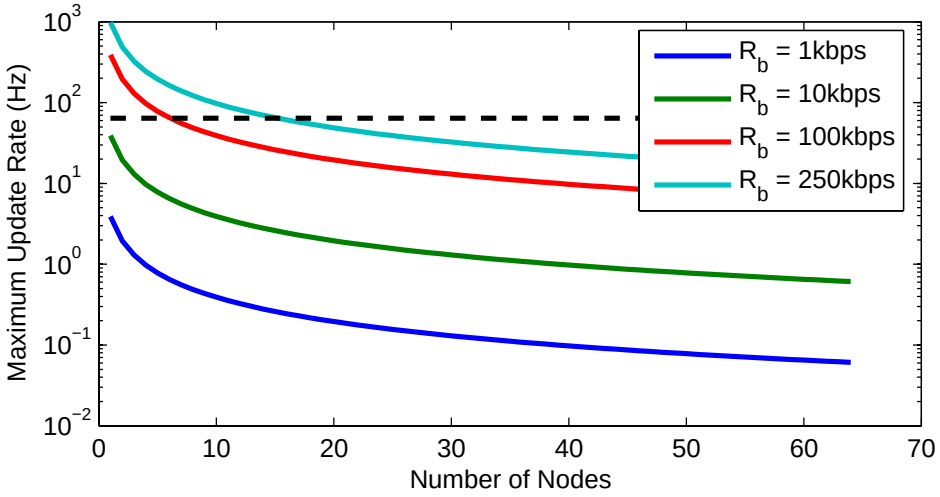

Figure: 2.26: Maximum achievable update rate over a single channel for a packet size of 256 bits. The dotted line represents the Orient-2 update rate of 64Hz.

2.8 Concluding Remarks

In this chapter, WSNs have been introduced, and in particular Speckled Computing, which is the context of this work. WSNs place constraints on all aspects of network design, from node power supply to communications protocols; in particular it was shown that the small size of the nodes severely limits the possible energy consumption of each node. This low power constraint influences the design process in all layers, from application development down to physical layer communications architecture.

The first generation Speck radio prototype was introduced and its limitations clearly outlined. It has been demonstrated that the current OOK-based transceiver limits potential to improve physical layer performance using DSP techniques. Progressing from the current simple envelope-detector based radio circuit to an oscillator-mixer architecture would enable future DSP designs to include CDMA or OFDMA.

The issues involved with the processing, communication and extraction of data from a WSN were then examined. It was explained that communication costs dominate the power budget in energy-constrained sensor nodes, and therefore communication must be minimised. With reference to a particular application of Speckled Computing, the Orient-2 body posture tracking system, the use of multiple frequency channels to extract data in a parallel manner to a central location has been reviewed. It has been shown that this concept is particularly advantageous in systems with strict latency and update rate requirements, and where nodes are energy constrained.

Chapter 3

Multicarrier Modulation

3.1 Introduction

This chapter provides an introduction to multicarrier modulation techniques, and in particular Orthogonal Frequency Division Multiple Access (OFDMA), which provides the basis for the work explored in this thesis. Firstly, a special case of single-user multicarrier modulation, known as Orthogonal Frequency Division Multiplexing (OFDM) is introduced, and it is shown how this techniques is extended to multiuser applications. The sensitivity of OFDMA and OFDM to time and frequency offsets is then analysed, and it is shown how these errors in synchronisation cause a loss of performance. Common techniques for dealing with the offsets are introduced, before the use of OFDMA in Speckled Computing is considered.

3.2 Single-User Multicarrier Modulation

Multicarrier modulation involves the transmission of multiple, parallel frequency channels instead of a single, wider band channel. Multicarrier modulation schemes are used in many current systems such as Digital Audio Broadcasting (DAB) [94], and the IEEE 802.11 Wireless Local Area Network (WLAN) standard [79].

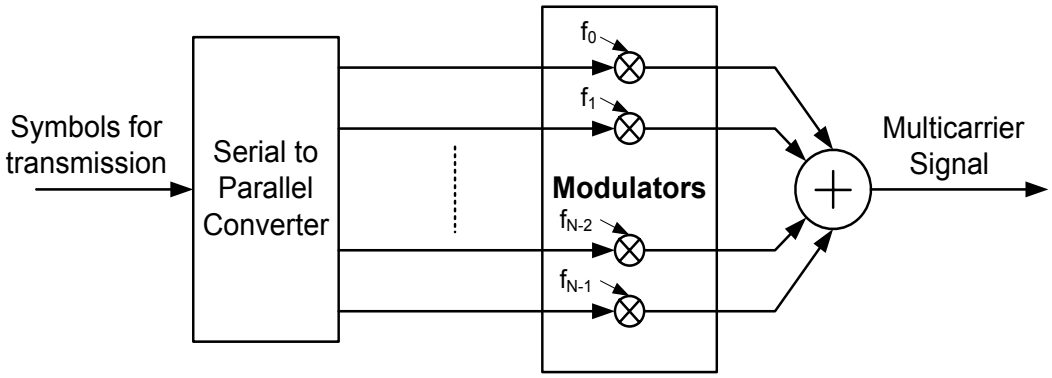

Figure: 3.1: Basic multicarrier transmitter block diagram.

A basic multicarrier transmitter is shown in Figure 3.1. Data are split into several parallel streams and modulated onto different frequency carriers, known as subcarriers. Figure 3.1 suggests any set of frequencies are used for channel modulation, and when the idea was first introduced in the late 1950s and was known as Frequency Division Multiplexing (FDM) [42], this was the case. Originally implemented using banks of analogue oscillators, in 1971 Weinstein proposed the use of the Discrete Fourier Transform (DFT) to digitally modulate data onto multiple orthogonal frequency carriers [67]. Bingham later realised that processing technology had progressed to such a degree that digitally implemented multicarrier techniques using the Fast Fourier Transform (FFT) were realisable [10]; the FFT is a fast and efficient method of implementing the DFT [13]. Using DFT techniques results in a special case of FDM, where all the subcarriers are orthogonal. This modulation scheme has become known as Orthogonal Frequency Division Multiplexing (OFDM).

3.2.1 Orthogonal Frequency Division Multiplexing

A block diagram of a basic OFDM system is shown in Figure 3.2. The transmitter performs parallel-to-serial conversion on complex symbols, and modulates them directly onto subcarriers using an inverse DFT (IDFT, commonly implemented using the inverse FFT (IFFT)) before converting the IDFT outputs to a serial transmitted

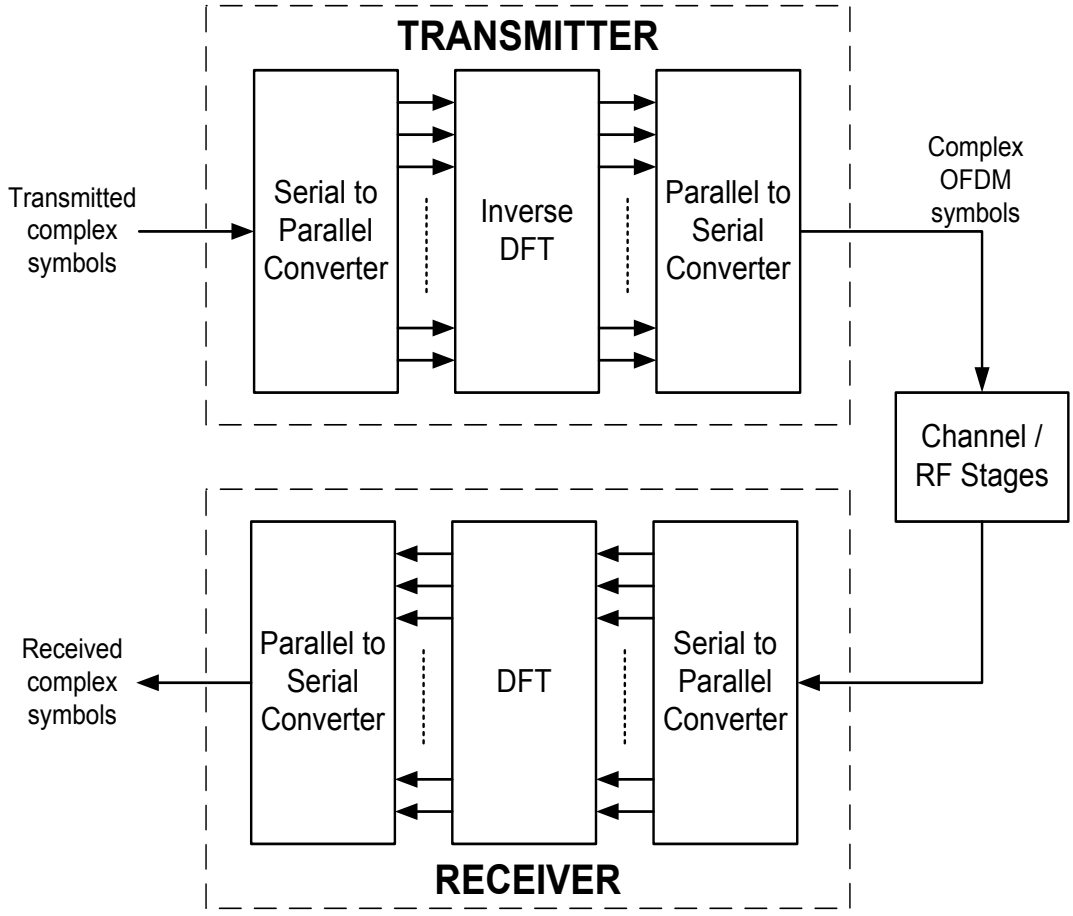

Figure: 3.2: Block diagram of an OFDM transmitter and receiver.

waveform. The receiver takes blocks of samples and processes them with the DFT, effectively sampling at regular frequencies using the DFT, before parallel-to-serial converting to recover the transmitted complex symbols.

A transmitted baseband-equivalent OFDM symbol is defined as

$$s[n] = \sum_{k=0}^{N-1} c_k e^{j2\pi kn/N} \tag{3.1}$$

where N is the number of subcarriers, f_s is the sampling rate, $T = N/f_s$ is the OFDM symbol period and $0 \leq n < N$. The IDFT is given by

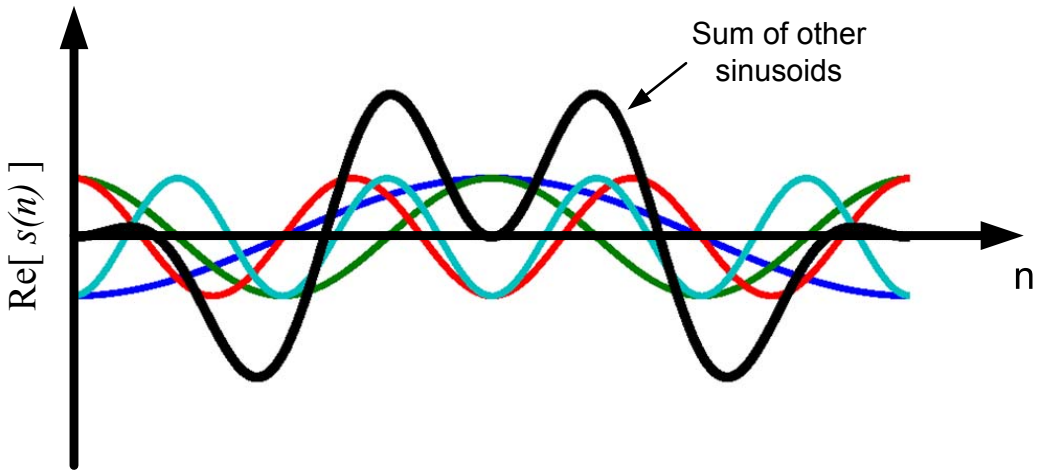

Figure: 3.3: Real part of the output of an OFDM transmitter IFFT block showing the sum of sinusoids.

$$x[n] = \frac{1}{N}\sum_{k=0}^{N-1} X_k e^{j2\pi kn/N}. \tag{3.2}$$

Comparing (3.1) with (3.2), it is clear that an OFDM transmitter can be implemented using the IDFT.

Looking at (3.1) it can be seen that the OFDM symbol is a sum of N modulated sinusoids, as is shown in Figure 3.3 for the first four (non-zero frequency) subcarriers of an OFDM symbol. It can be seen that a single OFDM symbol consists of an integer number of periods of each sinusoid; the generated frequencies are integer multiples of the subcarrier spacing, f_{sub}, which is the inverse of the OFDM symbol period:

$$f_{sub} = \frac{1}{T} \tag{3.3}$$

The figure also illustrates one of the major drawbacks of OFDM, namely that it can possess a large Peak to Average Power Ratio (PAPR). As an OFDM symbol can be made up of a large number of sinusoidal signals, when they add constructively large peak amplitude values are seen. A large PAPR causes a power amplifier to operate in an inefficient manner.

The spectrum of each modulated subcarrier forms a sinc shape due to the rectangular

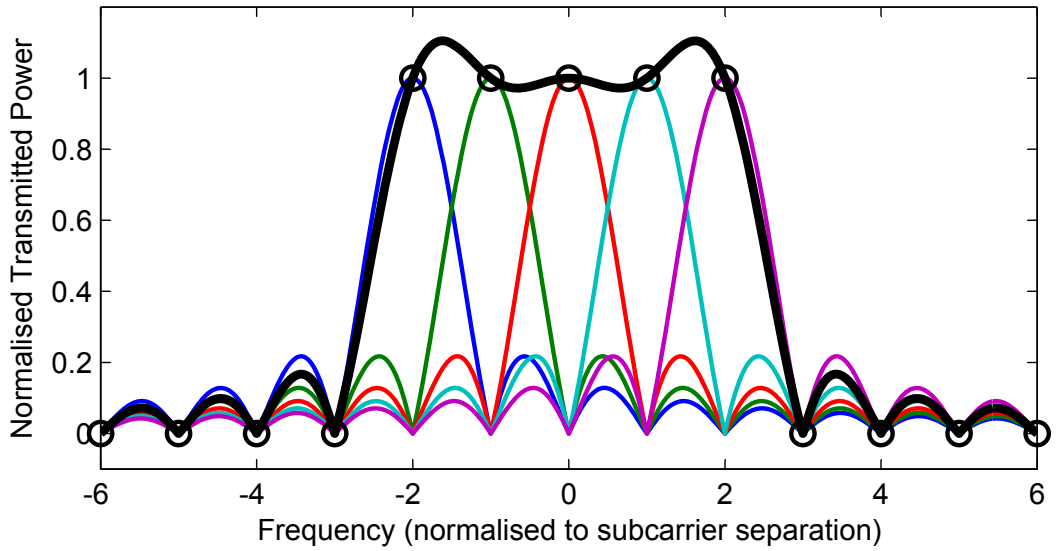

Figure: 3.4: Spectrum of an OFDM signal, also showing individual subcarriers.

windowing inherent in the IDFT operation, as shown in Figure 3.4, which shows the spectrum of five modulated subcarriers and their sum. Each subcarrier therefore has a peak at one frequency, and nulls evenly spaced at integer multiples of the subcarrier spacing f_{sub}, as shown by the circles in the figure. The nulls in the spectrum of one subcarrier therefore coincide with the peaks of the other subcarriers. The orthogonal property of the subcarriers allows them to be placed, in the frequency domain, as close to each other as possible without causing mutual interference. There will be no Inter-Carrier Interference (ICI) when symbol samples are taken by the receiver using the DFT. The receiver will sample and demodulate subcarrier m of the sampled transmitted signal $s[n]$, as given in (3.1), by performing a DFT as follows:

$$\begin{aligned} Y_m &= \frac{1}{N}\sum_{n=0}^{N-1} s[n]e^{-j2\pi mn/N} \\ &= \frac{1}{N}\sum_{n=0}^{N-1}\sum_{k=0}^{N-1} c_k e^{j2\pi(k-m)n/N} \end{aligned} \tag{3.4}$$

where Y_m is the decision variable for subcarrier m. In ideal conditions, $Y_m = c_m$, as

$$\frac{1}{N}\sum_{n=0}^{N-1} e^{j2\pi(k-m)n/N} = \begin{cases} 1 & k = m \\ 0 & k \neq m \end{cases}. \qquad (3.5)$$

Equation (3.5) illustrates the orthogonal nature of the subcarriers.

3.2.2 Benefits of Multicarrier Modulation

The splitting of a single carrier signal into multiple subcarriers is of great benefit in multipath channels. A multipath channel results in a frequency-selective response, as shown in Figure 3.5. With different levels of attenuation occurring at different frequencies, in a single carrier system a complex equaliser can be necessary to mitigate channel effects. By using a multicarrier signal, however, the received signal can be

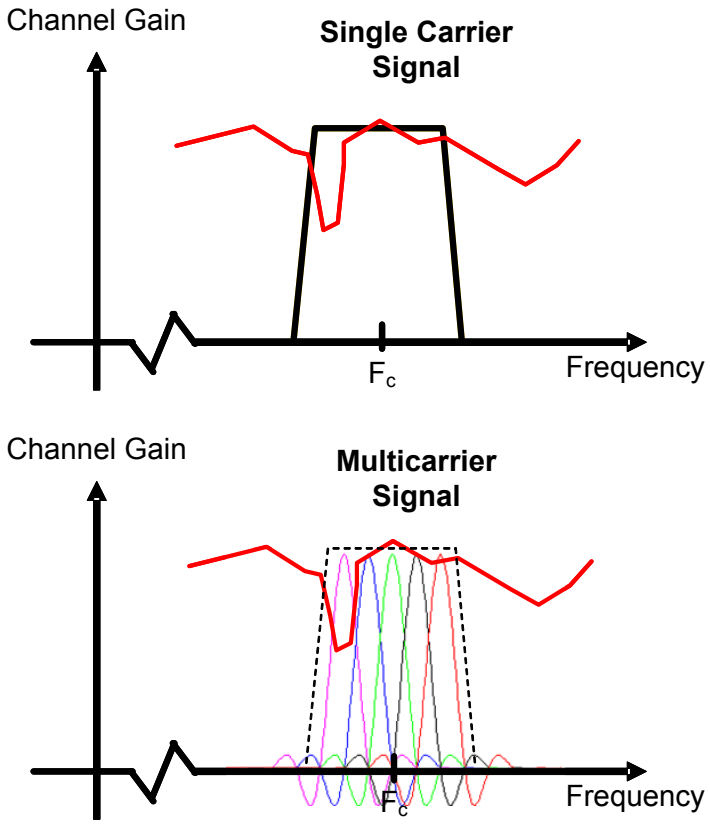

Figure: 3.5: Comparison of single and multicarrier signals in a frequency-selective channel.

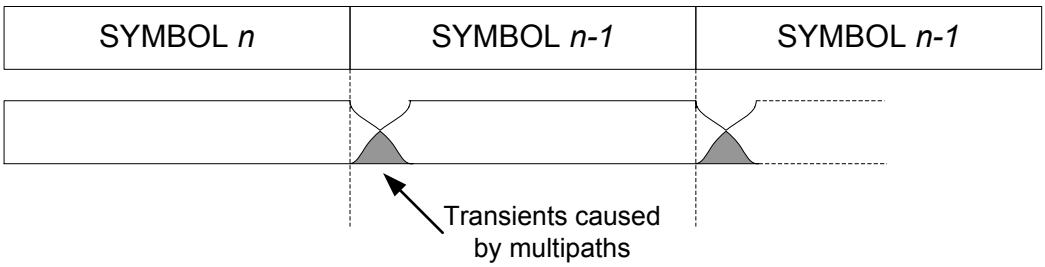

Figure: 3.6: Multipaths cause symbols to merge into one another.

treated as a collection of individual, flat-fading channels, which are easily equalised.

OFDM is also extremely resistant to Inter-Symbol Interference (ISI) caused by multipath delayed arrivals. Multipath arrivals cause transients at the beginning and end of each symbol, changing the sinusoidal shape of the symbols and causing one symbol to merge into the next, as shown in Figure 3.6, destroying the orthogonality between subcarriers. For the same system bandwidth, OFDM has a much longer symbol length than a single carrier system and therefore the effect of this symbol merging is less important. For example, an OFDM system with N subcarriers has a symbol length of N times the equivalent BPSK or QPSK system. For further protection against ISI caused by multipath channels, a guard period is inserted between OFDM symbols [10]. Usually, a section of the end of the symbol is copied and used as a prefix to the symbol, and is known as a cyclic prefix. This process is shown in Figure 3.7. The cyclic prefix is removed before demodulation and, as long as the cyclic prefix is longer than the channel response, orthogonality is maintained. The cyclic prefix is also frequently used in symbol synchronisation [8], and provides protection against time offsets [31]. In an OFDM receiver it is very important that accurate time and frequency synchronisation is maintained. If it is not, both ISI and ICI can be problematic [21,47]. The issues surrounding time and frequency offsets will be covered in Section 3.3, when the use of OFDM as a multiple access technique is considered.

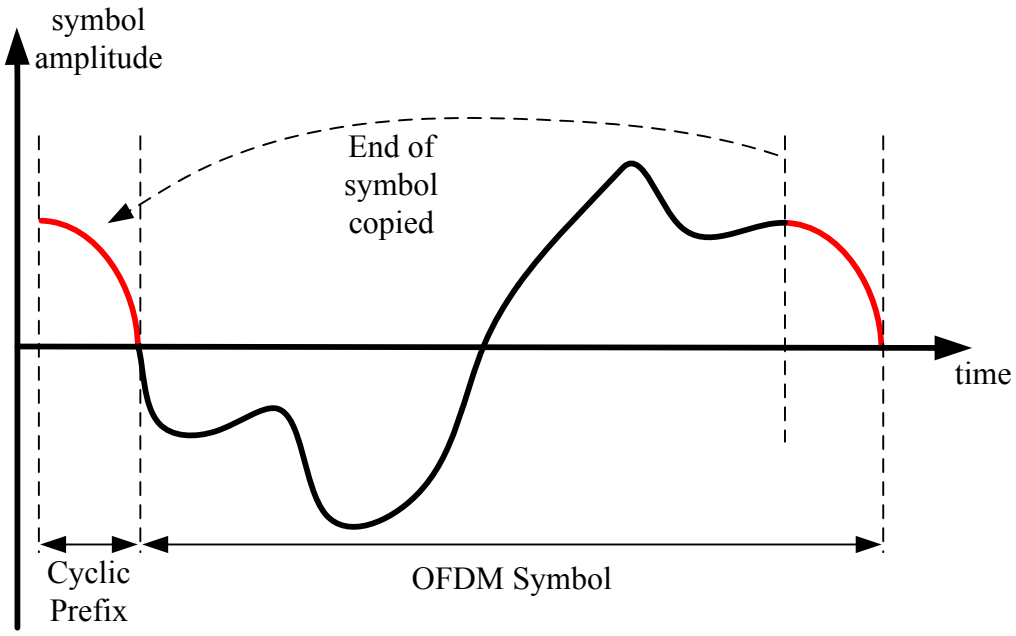

Figure: 3.7: A cyclic prefix is formed by taking a portion from the end of a symbol and affixing it to the start.

3.2.3 Use of OFDM in Speckled Computing

The main benefit of OFDM, as explained previously, is its robustness to the effects caused by transmitting over multipath channels. In single carrier systems, the maximum achievable data rate and transmission range are limited by the problems outlined above, namely frequency-selective fading and ISI. Using OFDM and exploiting its inherent mitigation of these problems allows higher data rates and longer transmission ranges. It is for these reasons that OFDM is being considered for development of high data rate standards such as the 3rd Generation Partnership Project (3GPP) Long Term Evolution (LTE) project for evolution of the Universal Mobile Telecommunications System (UMTS) [87].

Speckled Computing has very different requirements on a physical layer protocol. In general, data rates are low and ranges are short, as seen in Chapter 2, and as such the added complexity of OFDM implementation is not worthwhile. OFDM can, however,

be used as a multiple access technology, for example by assigning different subcarriers to different users. This will provide the main subject of investigation in this thesis as OFDM techniques are applied to the problem of extracting data efficiently and at the lowest possible cost from the Orient-2 network. Subsequent sections therefore investigate the use of OFDM techniques in multiuser scenarios.

3.3 Orthogonal Frequency Division Multiple Access

OFDM can be extended in various ways to cater for multiple users [52]. The simplest method is to assign discrete sets of subcarriers to individual users, thus implementing a form of FDMA. This technique provides the basis for this work and is thus covered in some detail in the succeeding sections.

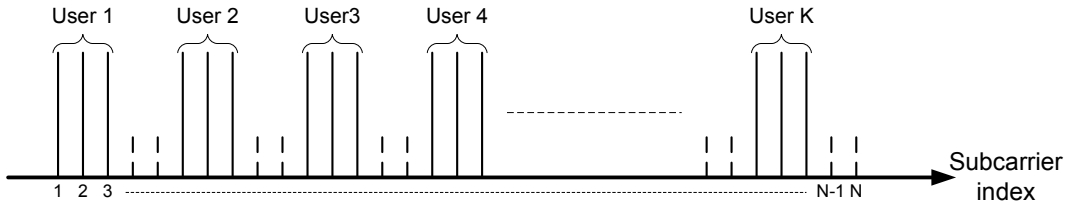

Figure: 3.8: OFDMA block-based subcarrier assignment, with 3 subcarriers per user and a guard subcarrier on each side.

3.3.1 Multiuser Orthogonal Frequency Division Multiplexing

OFDM can be applied in a multiuser context by assigning different subcarriers to different users, as in Figure 3.8. This is known as Orthogonal Frequency Division Multiple Access (OFDMA). In the example in the figure, the subcarriers have been assigned to users in blocks of 3 subcarriers, with a single subcarrier used as a guard on either side. The number of subcarriers allocated to each user for data and guards is not necessarily fixed, and indeed significant investigation has been carried out into subcarrier allocation strategies, see [32] and references therein.

In an ideal case, a receiver can receive the superposition of multiple users' signals

and demodulate them together using a DFT, as if the signal was a single OFDM transmission. In practical situations, however, Doppler shifts and oscillator instabilities result in different frequency offsets between users and a receiver, and different transmission times and propagation delays result in time offsets between users. As mentioned briefly above, an OFDM-based signal is very sensitive to time and frequency offsets, and as such users in an OFDMA system must be carefully synchronised. The sensitivity of an OFDMA system to time and frequency offsets between users is analysed in the next section.

3.3.2 Sensitivity to Synchronisation Errors

The effects of time and frequency synchronisation errors have been extensively covered in the literature for both OFDM and OFDMA, [21,31,41,47,66]. Here, the effects caused by time and frequency synchronisation errors on the demodulation of a single subcarrier are examined. The analysis of a single subcarrier is carried out in the continuous time domain as it provides an intuitive, and visual, understanding of the effects of time and frequency offsets. The results shown here can be readily applied to both full OFDM and OFDMA systems, as each subcarrier can be analysed individually and then the results superimposed as both signals are defined as sums of subcarriers. In an OFDMA system, time and frequency offsets will be constant across discrete sets of subcarriers, while in an OFDM system, time and frequency offsets will be constant across all subcarriers.

The system used to analyse the effects of time and frequency offsets is shown in Figure 3.9. A stream of complex symbols, $c_{k_u,\,i}$, (each of period T) is modulated by a subcarrier at frequency k_u/T Hz. The signal is then subject to a time offset, τ_u, and a frequency offset, Δf_u, before being demodulated by Fourier Transform to produce a decision variable, Y_m. Note that for generality τ_u and Δf_u are both defined relative to the symbol period T, and as such

$$0 \le [\tau_u, \Delta f_u] < 1 . \tag{3.6}$$

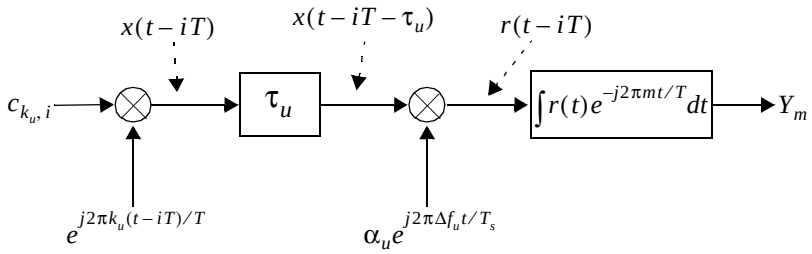

Figure: 3.9: System used to analyse the effects of time and frequency offsets on the decision variable of a single subcarrier, showing the addition of the time and frequency offsets.

In a practical system the decision variable would then be input to a decision device to recover the transmitted symbols.

A transmission by a single user on a single subcarrier takes the following form:

$$x(t-iT) = c_{k_u,i}e^{j2\pi k_u(t-iT)/T} \tag{3.7}$$

where $c_{k_u,i}$ is the complex data symbol transmitted during the i^{th} OFDM symbol on the $k_u{}^{th}$ subcarrier. The received signal (under otherwise ideal channel conditions, and including time and frequency offsets) is:

$$r(t-iT) = c_{k_u,i}e^{j2\pi k_u(t-iT-\tau_u)/T}e^{j2\pi\Delta f_u t/T} \tag{3.8}$$

and the decision variable at the output of the Fourier Transform for the i^{th} symbol on the m^{th} subcarrier is:

$$Y_{m,i}^{k_u} = \frac{1}{T}\int_{iT}^{(i+1)T} r(t-iT)e^{-j2\pi mt/T}dt\,. \tag{3.9}$$

Due to the time offset, the transform window is effectively offset by $\tau_u \times T$, and some of the previously transmitted symbol is included in the window, as shown in Figure 3.10. Without loss of generality, we analyse the decision variable for the $(i=0)^{th}$ symbol on subcarrier *m*.

$$Y_{m,0} = Y_{m,0}^{P} + Y_{m,0}^{C} \tag{3.10}$$

where $Y_{m,0}^{P}$ is the contribution from the previous symbol and $Y_{m,0}^{C}$ is the contribution

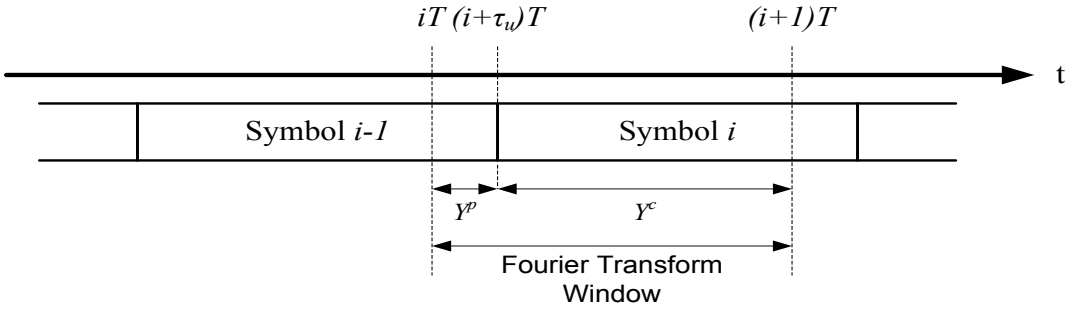

Figure: 3.10: Receiver timeline showing the time offset of the received signal relative to the Fourier Transform window.

from the current, desired symbol, as shown in Figure 3.10.

$$\begin{aligned} Y^{P}_{m,0} &= \frac{1}{T}\int_{0}^{\tau_u} c_{k_u,-1}\, e^{j2\pi k_u(t+T-\tau_u)/T} e^{j2\pi\Delta f_u t/T} e^{-j2\pi mt/T}\, dt \\ &= \frac{1}{T}\int_{0}^{\tau_u} c_{k_u,-1}\, e^{j2\pi[(k_u-m+\Delta f_u)t + k_u(T-\tau_u)]/T}\, dt \end{aligned} \tag{3.11}$$

and

$$\begin{aligned} Y^{C}_{m,0} &= \frac{1}{T}\int_{\tau_u}^{T} c_{k_u,0}\, e^{j2\pi k_u(t-\tau_u)/T} e^{j2\pi\Delta f_u t/T} e^{-j2\pi mt/T}\, dt \\ &= \frac{1}{T}\int_{\tau_u}^{T} c_{k_u,0}\, e^{j2\pi[(k_u-m+\Delta f_u)t - k_u\tau_u]/T}\, dt \end{aligned} \quad . \tag{3.12}$$

(3.11) and (3.12) reduce to:

$$Y^{P}_{m,0} = c_{k_u,-1}\frac{\tau_u}{T}\mathrm{sinc}\left(\frac{(k_u-m+\Delta f_u)\tau_u}{T}\right) e^{j\pi(2k_u(T-\tau_u)+(k_u-m+\Delta f_u)\tau_u)/T} \tag{3.13}$$

and

$$Y_{m,0}^{C} = c_{k_u,0}\left(\frac{T-\tau_u}{T}\right)\mathrm{sinc}\left((k_u - m + \Delta f_u)\left(\frac{T-\tau_u}{T}\right)\right) e^{j\pi(-2k_u\tau_u + (k_u - m + \Delta f_u)(\tau_u + T))/T} \quad (3.14)$$

where

$$\mathrm{sinc}(x) = \frac{\sin(\pi x)}{\pi x}. \quad (3.15)$$

In an ideal, fully-synchronised system, $\tau_u = \Delta f_u = 0$, and

$$Y_{m,0} = \begin{cases} c_{k_u,0} & m = k_u \\ 0 & m \neq k_u \end{cases}, \quad (3.16)$$

again showing the orthogonal nature of the subcarriers in ideal conditions, as the signal transmitted on the single subcarrier k_u has no influence on any other subcarriers. Equations (3.13) and (3.14), essentially expressing the frequency domain representation of an OFDM or OFDMA signal, also confirm its sinc-shaped spectrum.

If $[\tau_u, \Delta f_u] \neq 0$, the decision variable on what is termed the *desired subcarrier*, i.e. $Y_{k_u,0}$, where $m = k_u$, is subject to a phase rotation and an attenuation, which are dependent on the offsets and the subcarrier frequency k_u/T, but which are constant in time. It is also subject to ISI from the previous symbol. In addition, the output on the other subcarriers, $m \neq k_u$, is non-zero, i.e. there is ICI. The magnitude of the ICI on neighbouring subcarriers and the loss of energy on the desired subcarrier caused by time and frequency offsets is depicted in the following figures, which show the magnitude of the decision variable output, $|Y_m|$, on several subcarriers for different time and frequency offsets, assuming a single transmitted subcarrier.

3.3.3 Time Offsets

Figure 3.11 shows the magnitude of the decision variable output when considering only a single symbol, i.e. the previous symbol takes a zero value and there is no ISI. The time offset is varied relative to the OFDM symbol period, *T*, and zero frequency

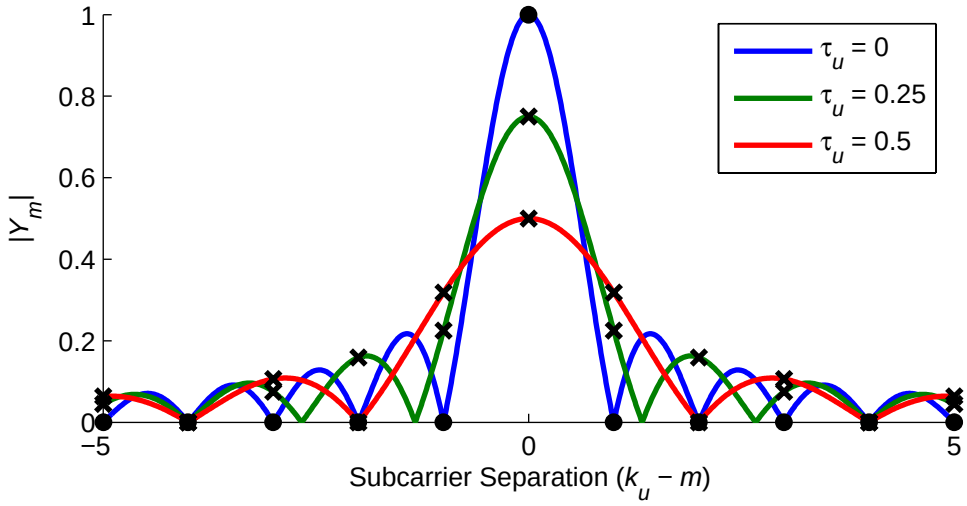

Figure: 3.11: Magnitude of the decision variable, Y_m, for varying time offsets, showing interference on neighbouring subcarriers; previous symbol is zero.

offset is assumed. Considering a single symbol in isolation allows a clear visual understanding of what occurs as time synchronisation is lost. For zero offset, shown in the figure by the black dots, the entire received energy is concentrated on the desired subcarrier (where $k_u = m$), and no interference is seen by the other subcarriers, as expected. As the time offset increases, it can be seen that there is an amplitude loss on the desired subcarrier, and an increase in ICI on neighbouring subcarriers, as shown by the crosses. The worst interference is experienced by those subcarriers adjacent to the subcarrier with the offset, and the interference decreases with spectral distance. The worst case offset of half the symbol period results in a 3dB loss of energy on the desired subcarrier.

The actual behaviour of the interference caused by time offsets is also dependent on the relationship between the current and previous symbols. Indeed, if the previous and current symbols are the same, the only effect of the time offset will be a phase shift on the decision variable, and there will be no ICI. Using a cyclic prefix has the same effect; employing a cyclic prefix of length longer than the maximum expected time offset results in no loss of orthogonality, and only a phase shift. The decision variable can then be phase rotated to correct the error, or a differential modulation scheme can be used. Figure 3.12 shows the decision variable output when previous symbols with

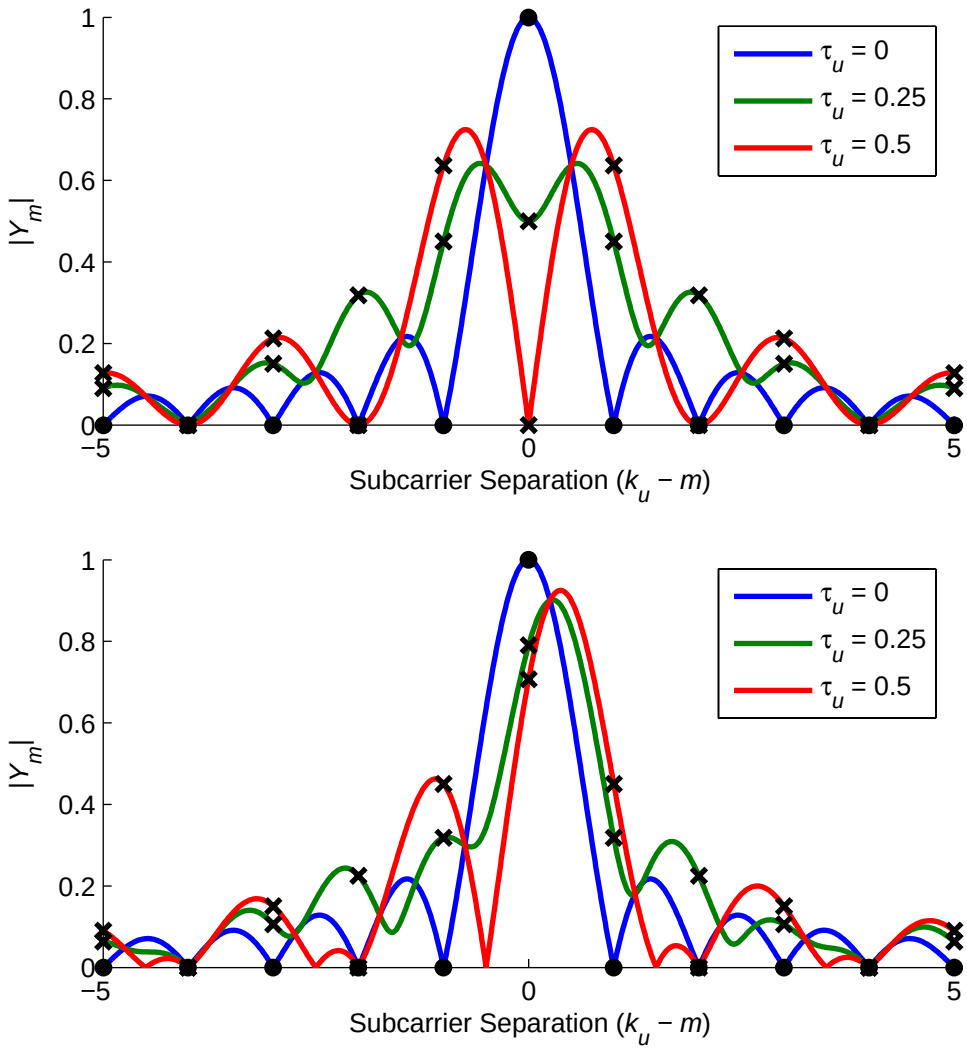

Figure: 3.12: Magnitude of the decision variable, Y_m, for varying time offsets, using previous symbols with a phase shift of π (top) and $\pi/2$ (bottom) from the current symbol.

a phase shift of π and $\pi/2$ from the current symbol are transmitted, with the dots and crosses having the same relevance as in Figure 3.11. A phase shift of π between the symbols produces the worst result, with a significant amount of interference being seen on neighbouring subcarriers, and a large amplitude loss on the desired subcarrier.

3.3.4 Frequency Offsets

Assuming no time offset means that there is no contribution to the decision variable

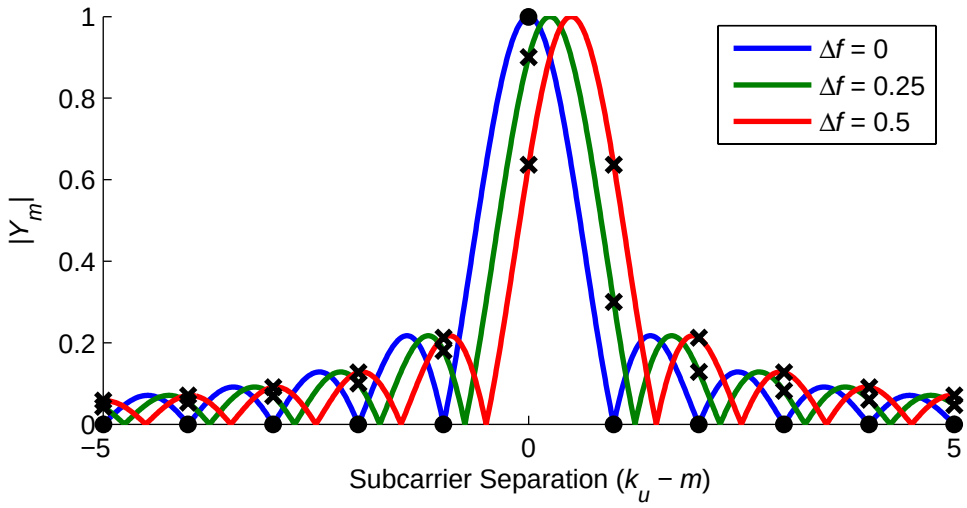

Figure: 3.13: Magnitude of the decision variable, Y_m, for varying frequency offsets, showing interference on neighbouring subcarriers.

from a previous symbol. The decision variable therefore reduces to:

$$Y_{m,0} = Y^{C}_{m,0}\Big|_{\tau_u = 0} = c_{k_u,0}\,\mathrm{sinc}(k_u - m + \Delta f_u)e^{j\pi(k_u - m + \Delta f_u)}. \qquad (3.17)$$

Equation (3.17) shows that, when there is a frequency offset (which may be due to oscillator inaccuracies or Doppler effects), that the decision variable on the desired subcarrier is again subject to an error in phase and amplitude, and is non-zero when $k_u \neq m$. Figure 3.13 shows the magnitude of the decision variable output for various subcarriers for a selection of frequency offsets. The dots give the ideal sample points and the crosses give the sample points under the influence of frequency offsets. The result in this case is an amplitude loss on the desired subcarrier, and ICI on neighbouring subcarriers.

Both time and frequency offsets on a subcarrier cause performance degradation on the subcarrier in question and neighbouring subcarriers. ICI and ISI increase the noise seen on a decision variable and thus reduce the effective Signal to Noise Ratio (SNR), degrading performance. Considering these effects when multiple subcarriers are in use, it can easily be inferred that the noise introduced will rapidly become significant if offsets are large.

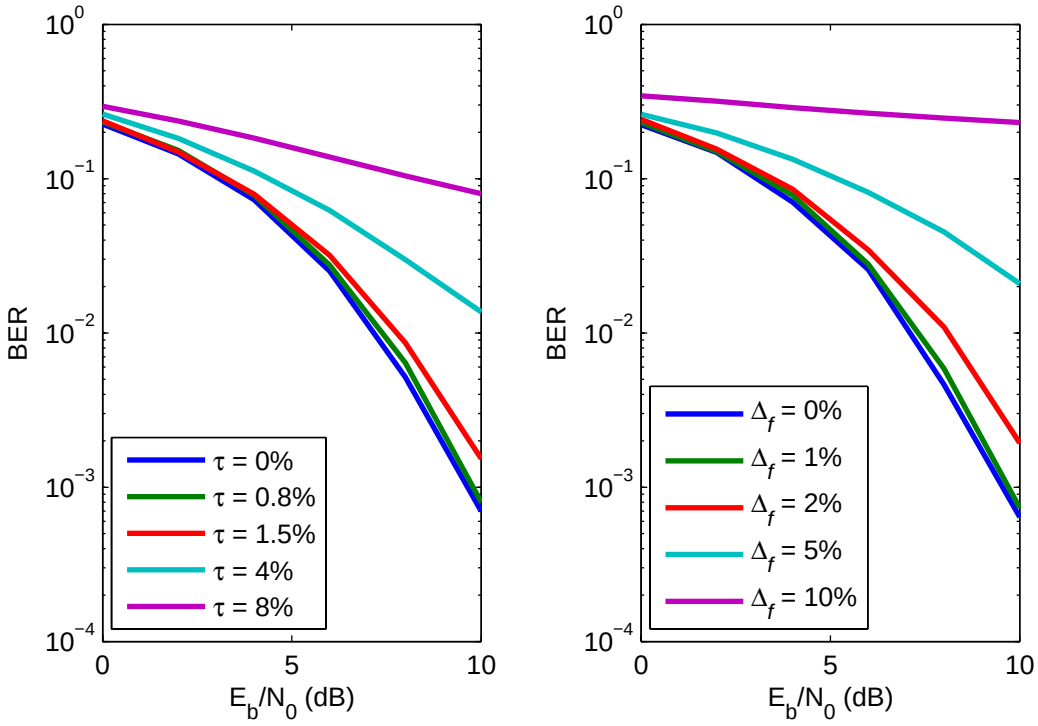

Figure: 3.14: Performance degradation on a single user among fifteen users, due to time (left) and frequency (right) offsets. Offsets given are relative to symbol period and subcarrier spacing respectively.

3.3.5 Performance impairments

Figure 3.14 gives examples of the performance degradation caused by a lack of synchronisation in an OFDMA system where each user transmits a single subcarrier. The BER performance for Differential Quadrature Phase Shift Keying (DQPSK) is plotted against E_b/N_0 at the input to the receiver FFT. The curves are calculated by Monte Carlo simulation for a single user on the 8th subcarrier among 15. The worst case result is shown by giving all users the same time offset for the time offset curves, and by arranging the frequency offsets as in Figure 3.15. Each user has the same magnitude of frequency offset while the direction (positive or negative) of the offset for each user is chosen to maximise interference on user 8.

Both graphs show the significant performance degradation which is experienced as

offsets increase. For this reason time and frequency offsets in an OFDMA system must be minimised, and synchronisation across users is crucial. Note, however, that the performance loss is minimal for small offsets, indeed both graphs show a very small degradation for offsets of less than approximately 1%.

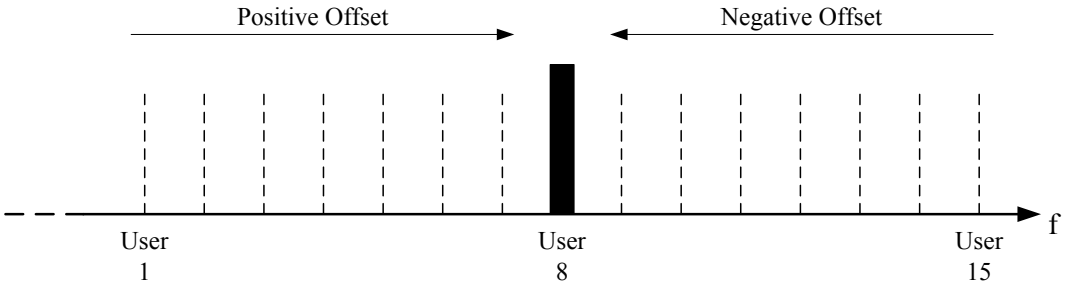

Figure: 3.15: Interfering users are given frequency offsets to cause the worst case performance on the middle user, user 8.

3.4 User Synchronisation in an OFDMA System

As seen in Section 3.3.5 synchronisation across users in an OFDMA system is very important to avoid performance loss. The synchronisation requirements for an OFDMA system depend on the scenario in which the network is deployed. There are essentially three cases where such a system could be used to provide multiple access, as shown in Figure 3.16; the downlink of a centralised system, the uplink of a centralised system, or in an asynchronous system where nodes can transmit freely amongst themselves. Synchronisation in OFDMA systems has been extensively surveyed in [41], and will be reviewed here.

3.4.1 Centralised Downlink

In a cellular downlink scenario a base station transmits to other, possibly mobile, nodes in the network. Different subcarriers are allocated to different users, and the base station transmits on all subcarriers at once, transmitting to all users in parallel. The

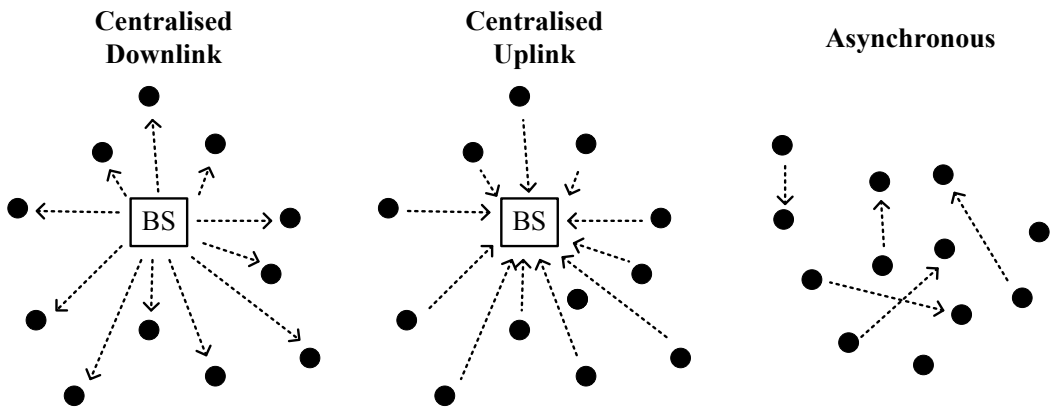

Figure: 3.16: The three main network scenarios in which OFDMA can be used as a multiple access technique.

signal arriving at each node can be treated as an OFDM single user transmission, as each subcarrier is subject to the same time and frequency offset. Single user OFDM synchronisation techniques can therefore be used in this case, and are extremely well covered in the literature. Synchronisation techniques commonly use the redundancy present in the cyclic prefix [8] and/or a preamble of training symbols [38,39,57] to estimate the offsets. The cyclic prefix method negates the need to send additional training symbols, however is limited to detection of frequency offsets up to half of the subcarrier spacing. Often short training symbols are sent to increase the frequency range for initial coarse estimation [39]. The frequency offset can be compensated by phase rotation of the DFT outputs, and the time offset can be minimised by shifting the DFT window.

3.4.2 Fully Asynchronous Network

In ad-hoc networks, for example some WSNs, there is no central entity controlling communications, and nodes can transmit to any other node, at any time. All OFDMA transmissions must be synchronised in time and frequency, in order to avoid the interference issues explored previously. The PACWOMAN project is investigating ad-hoc Wireless Personal Area Networks (WPAN) [82]. OFDMA is investigated in [27],

where a combination of TDMA and OFDM is implemented. In the system a single node is nominated as a time and frequency reference, and any new node coming into the network must synchronise to this node. Through this procedure, which involves a significant amount of control overhead, network-wide synchronisation is achieved.

3.4.3 Centralised Uplink

Synchronisation for the uplink in a centralised network is complex, as users arrive at a receiver with different time and frequency offsets. The downlink approach cannot be used in the uplink, as correcting one user's offset will misalign any other user that may have been synchronised. The simplest solution to this problem was introduced in [9], where blocks of subcarriers belonging to discrete users are separated by a bank of bandpass filters. Time and frequency offsets for each user are then estimated using standard OFDM techniques, and returned to users in a downlink control channel. Users then adjust their oscillator frequencies and transmit times to compensate for the offsets.

Offsets cannot be estimated in this way in systems where users are assigned non-block-based sets of subcarriers, however techniques exist for estimation of offsets in this case, [11,40]. The process of returning offset estimates to users introduces an obvious overhead to network communications, therefore, more recently, techniques have been investigated which estimate and compensate for offsets directly in the base station, see examples in [41]. Another approach is for users to use their offset estimates from the downlink communications as references for uplink transmission. In [31] the authors describe a system where users time synchronise to the downlink transmission, and time offsets reduce to the order of the propagation delay between users and the base station.

Although the Orient-2 system takes this form, the complexity involved in implementing many of these solutions, such as the control overhead involved in signalling time and frequency offsets to each user, makes them unsuitable. A novel solution to these issues is therefore investigated in this thesis.

3.4.4 Using the Cyclic Prefix to Mitigate the Effect of Time Offsets

The analysis in Section 3.3.2 assumed a continuous-time Fourier Transform in the receiver, in order to clearly show the ICI and ISI caused by time and frequency offsets. In any practical system, the receiver will be implemented using a DFT (or, most likely, an FFT), in which case the time offset seen by the receiver becomes a function not only of the physical time offset, but also of the number of subcarriers. This becomes important in systems with a small number of subcarriers, as even a small delay can result in a significant time offset relative to the OFDM symbol period *T*, if no cyclic prefix is used. This is illustrated in Figure 3.17, where a simple example of only four subcarriers is shown (thus the receiver samples at four samples per symbol). Any delay of up to one sample period results in a sample being taken from the previous symbol, and thus an effective delay, as experienced by the FFT, of one whole quarter of the symbol period. If a cyclic prefix is used that is longer than the maximum possible time offset, the first sample will be taken from the cyclic prefix, rather than a previous symbol, and the only effect of the offset is a phase rotation of the decision variable. In systems where users synchronise to the downlink transmission, ensuring the cyclic prefix is longer than twice the maximum propagation delay can render timing

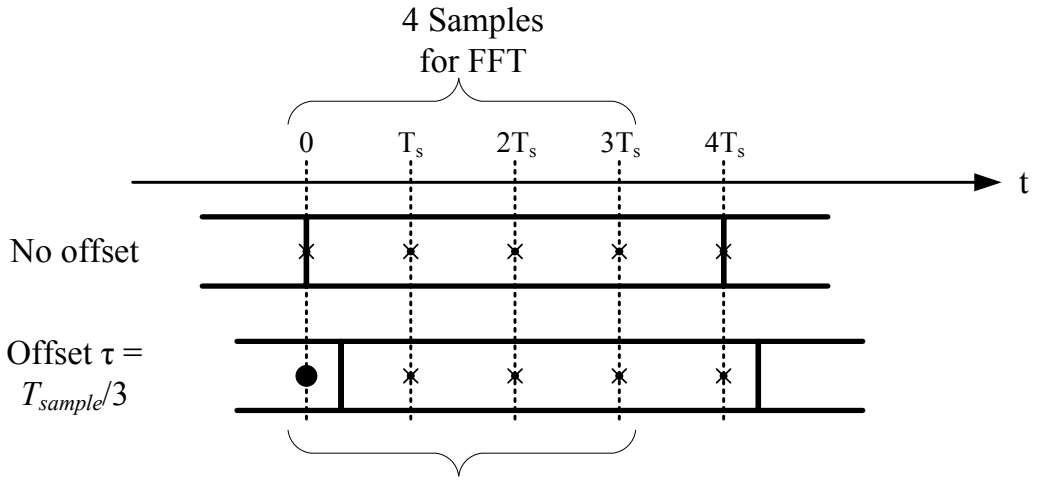

Figure: 3.17: The discrete nature of the receiver results in time offsets measured in multiples of the sampling period. In this example, with only four samples per symbol, the FFT experiences an effective offset of one quarter of the symbol period (one sample) for any physical offset of up to one sample period.

synchronisation in the receiver unnecessary; this type of system is often referred to as a *quasi-synchronous* system [7,41]. The phase shift can be corrected in the receiver, or differential modulation can be used, in which case a constant phase shift does not pose a problem, as explained below.

3.4.5 Carrier Phase Offsets

Time offsets result in users arriving at the receiver with random phase offsets. It is impossible for the receiver to phase synchronise to all users, as synchronising to one user will not resolve the phase problem with all other users. A coherent receiver is therefore not possible. In a quadrature system, the result of a carrier phase offset is a rotation of the constellation, as shown in Figure 3.18 for a QPSK modulated signal afflicted with a carrier phase offset of θ. This issue can be resolved by transmitting differentially phase modulated signals, (known as Differential Phase Shift Keying (DPSK)) where each symbol is encoded with reference to the previous symbol, rather than mapped to a fixed constellation with reference to some fixed phase. For example, for Differential Binary Phase Shift Keying (DBPSK), to transmit a binary '1' the carrier phase is shifted +180° from the phase of the previous symbol. To transmit a binary '0', the phase is shifted by 0° (i.e. the phase is unchanged). To demodulate the

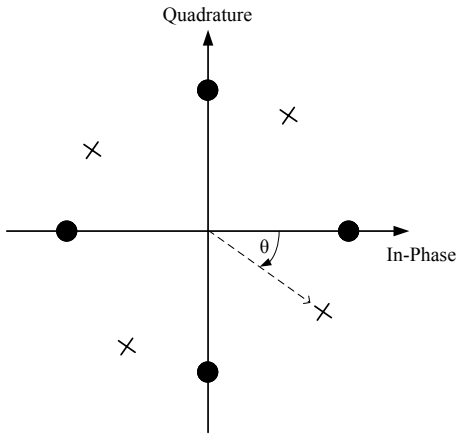

Figure: 3.18: QPSK constellation showing ideal signal points (circles) and points rotated as a result of a carrier phase offset of θ (crosses).

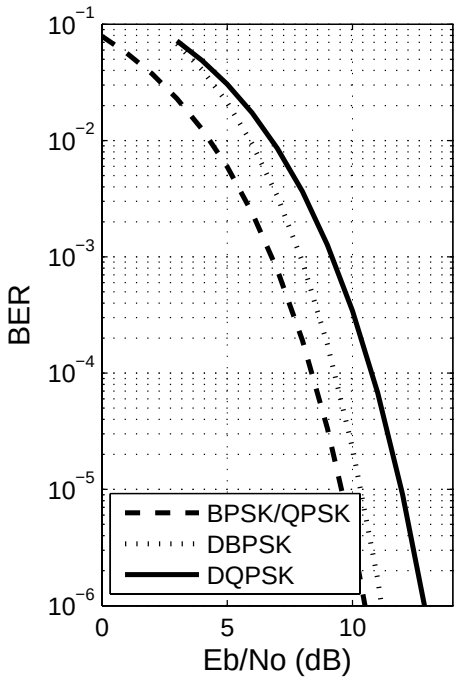

Figure: 3.19: Comparison of probability of bit error curves for BPSK/QPSK, DBPSK and Gray-coded DQPSK.

signal the receiver compares the phase of two symbols instead of comparing the phase of a single symbol to some global reference phase, and therefore the receiver oscillator does not have to be in phase synchronisation with the transmitted carrier.

The penalty of using differential modulation is a loss in BER performance as shown in Figure 3.19. For example the loss is approximately 2dB at a BER of 10^{-4}. This drop in performance is due to the increased noise involved in the decision process, the reference phase is taken from a noisy previous symbol, rather than a clean locally generated zero phase [50]. The theoretical probability of bit error in white Gaussian noise for DBPSK is

$$P_e = \frac{1}{2} e^{-E_b/N_0} \tag{3.18}$$

and for Gray coded DQPSK is

$$P_e = Q_1(a, b) - \frac{1}{2} I_0(ab) e^{-\frac{1}{2}(a^2 + b^2)} \tag{3.19}$$

where $Q_1(a, b)$ is the Marcum Q function and $I_0(x)$ is the modified Bessel function of order zero. For further information see [50], where the parameters *a* and *b* are also defined as:

$$\begin{aligned} a &= \sqrt{2\frac{E_b}{N_0}\left(1 - \sqrt{\frac{1}{2}}\right)} \\ b &= \sqrt{2\frac{E_b}{N_0}\left(1 + \sqrt{\frac{1}{2}}\right)} \end{aligned}. \tag{3.20}$$

3.4.6 Use of OFDMA in Speckled Computing

A Specknet is by definition a low-data rate, short range, asynchronous network. OFDMA could be used in a similar manner to the PACWOMAN WPAN network introduced in Section 3.4.4, however the control overhead required is excessive for use in energy-limited nodes such as Specks. In some Specknet applications, in particular the Orient-2 body posture tracking system which was explored in Chapter 2, it is useful to be able to extract data from the network in a parallel and efficient manner, in which case OFDMA provides suitable functionality, provided time and frequency offsets are dealt with. It was shown in Section 3.4.3 that there are many methods for managing the effects of time and frequency offsets in a centralised uplink system such as the Orient-2. Most, however, involve too much complexity for nodes which need to be kept as simple as possible. As a result, a novel method of dealing with the time and frequency offsets must be proposed. Such a solution is introduced in the next chapter.

3.5 Concluding Remarks

This chapter has provided an introduction to multicarrier modulation techniques, in particular OFDM and its multiuser counterpart OFDMA. OFDM modulation was reviewed, and its suitability for high data rate networks was made clear. The extension of OFDM to OFDMA was then investigated, with particular emphasis on the

importance of time and frequency synchronisation between users. The detrimental effects of the offsets and the subsequent performance degradation were investigated, and some synchronisation techniques for different network types were reviewed.

It was shown how the cyclic prefix can be used to form a quasi-synchronous system, where the effect of time offsets is to introduce a phase rotation to the received data symbols. This can be overcome by using differential modulation techniques. It was also briefly suggested that OFDMA can be used to implement multiple channels for data extraction from a WSN, and particularly the Orient-2 system, provided the frequency offsets can also be managed. Standard synchronisation techniques are, however, overly complicated for use in simple nodes, and therefore a novel solution must be developed. In the next chapter, a receiver-initiated data extraction protocol is proposed which addresses these issues.

Chapter 4

An OFDMA-Based Data Extraction Protocol

4.1 Introduction

It was shown in Chapter 2 that parallel frequency channels can be useful in networks such as the Orient-2 with update rate, latency and energy limitations. Chapter 3 introduced OFDMA and suggested how it could be used to create an efficient orthogonal FDMA system. In this chapter, a receiver-initiated OFDMA data extraction protocol is introduced, which allows data to be extracted from multiple nodes in parallel using a simple digital receiver. As shown in Chapter 3 OFDMA has very tight time and frequency synchronisation requirements, thus the protocol is designed to elegantly manage these issues. In particular, a novel method of combining frequency offset compensation with subcarrier generation is proposed. Various implementation issues for the data extraction system are then investigated, with close reference to the Orient-2 network.

4.2 An OFDMA-Based Data Extraction System

As shown in Chapter 3, OFDMA is a spectrally efficient method of implementing an FDMA system. In addition, it can be received efficiently using Fourier Transform based processing, commonly implemented with an FFT. As suggested previously for

the Orient-2 system, it can be used to extract data from multiple nodes in a sensor network in a parallel manner. By allocating a single subcarrier to each user, the sensor network nodes can remain relatively simple, needing only an extra digital subcarrier modulation stage in addition to their normal radio front end. As also shown in Chapter 3, however, transmissions from multiple users must be synchronised in time and frequency in order to avoid multiple access interference. In a data extraction context these limitations can be overcome if the extraction takes place in a receiver-initiated process. That is, the receiving entity polls a selection of nodes who reply concurrently with the required data. Assuming certain limits on propagation distances and required data rates, the time offsets between user signals arriving at the receiver can be sufficiently small. In addition, by structuring the polling signal in such a way that a node can make an estimate of the frequency offset between its oscillator and that of the receiver, the transmitting node can adjust its transmission frequency to compensate for the offset. This also results in a simple receiving node, which has no requirement for complex synchronisation circuitry and demodulates data from multiple sources using efficient FFT-based hardware.

The timeline for this process is shown in Figure 4.1. In the diagram, the process is initiated by the receiver sending a polling signal. On reception of the polling signal a node estimates the carrier frequency offset, compensates for it and then sends its data. In the Figure, T_N is the time between the beginning of the transmission of the polling signal and reception of the data from node *N*. This includes twice the propagation delay

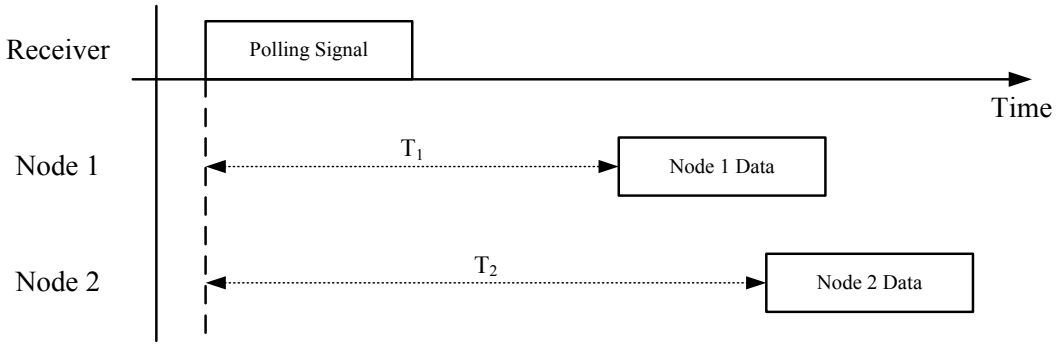

Figure: 4.1: Receiver timeline of the polling data extraction protocol.

between node *N* and the receiver (propagation time of the polling signal and then the data packet), the time taken by the node to estimate the frequency offset, and the time taken to then prepare and send a packet. It is assumed that the estimation, preparation and transmission time is equal for all nodes. As a result, time offsets between nodes are directly proportional to the difference in propagation delay (and thus distance) between each node and the receiver. The issues relating to time offsets, as well as phase and frequency offsets, are examined in more detail in the following sections.

The block diagram of the transmitter structure for each individual node is shown in Figure 4.2. The additional complexity needed to implement the OFDMA data extraction system is a digital mixing stage, where the baseband data is modulated onto a subcarrier whose frequency includes compensation for the frequency offset. The binary data is first baseband modulated with the desired baseband modulation before being digitally mixed onto the desired subcarrier frequency f_{sub}, where f_{sub} includes the frequency offset compensation, see Section 4.2.2 for more details on this operation. The subcarrier generator is therefore sensitive to the frequency offset estimate calculated on reception of the polling signal. The modulated subcarrier signal is then digital to analogue converted before being mixed up to the carrier frequency f_c.

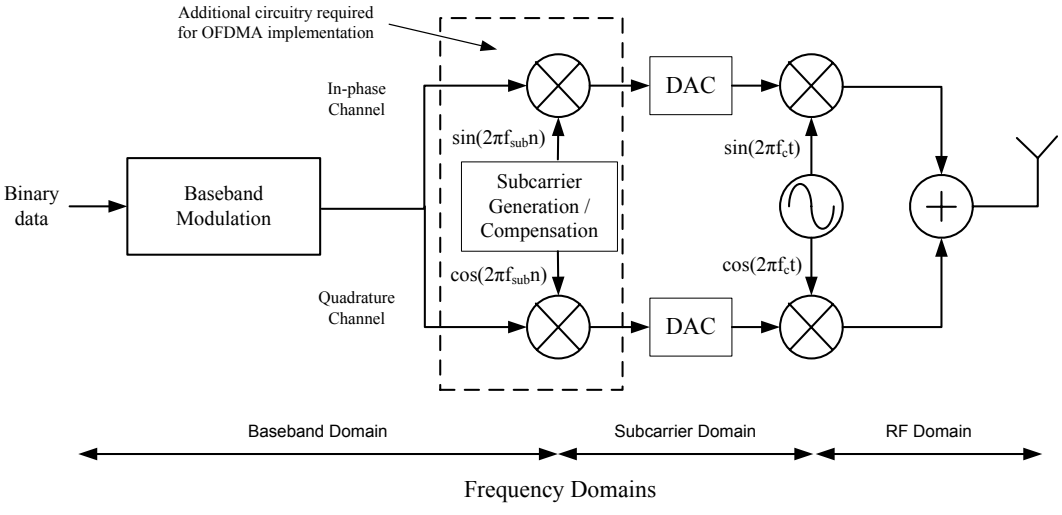

Figure: 4.2: Block diagram of individual node transmitter.

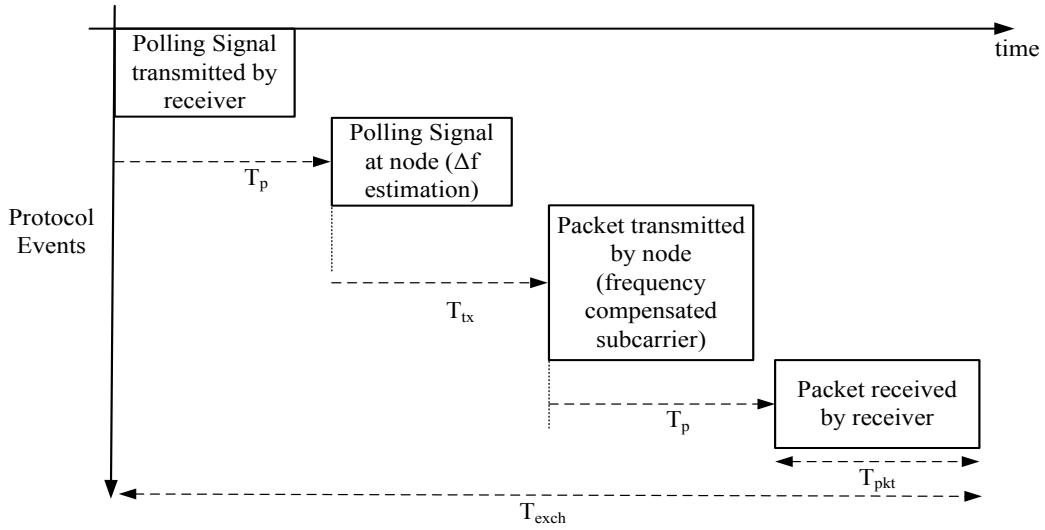

Figure: 4.3: Timeline of the receiver-initiated communication exchange.

4.2.1 Dealing With Time Offsets

The time offsets between users are directly proportional to the difference in distance between each user and the receiver. The communication exchange protocol is shown in Figure 4.3 and includes, as explained previously, transmission of the polling signal, frequency offset estimation and packet preparation and finally packet transmission. The period of the whole exchange for node *N*, T_{exch}^{N}, is:

$$T_{exch}^{N} = 2T_{p}^{N} + T_{tx} + T_{pkt} \tag{4.1}$$

and a packet from node *N* will arrive at the receiver after a period T_{ETA}^{N}, given by:

$$T_{ETA}^{N} = 2T_{p}^{N} + T_{tx} \tag{4.2}$$

where T_{p}^{N} is the propagation time between user *N* and the receiver, T_{tx} is the time taken to receive and process the polling signal (i.e. estimate and compensate for the frequency offset) and T_{pkt} is the length of a packet.

With reference to Figure 4.1, Figure 4.3 and (4.2), the time offset between the 2 users, τ_{12}, is:

$$\begin{aligned}\tau_{12} &= T_2 - T_1 \\ &= T_{ETA}^2 - T_{ETA}^1 \\ &= (2T_p^2 + T_{tx}) - (2T_p^1 + T_{tx}) \\ &= 2(T_p^2 - T_p^1) \\ &= 2\frac{(d_2 - d_1)}{c}\end{aligned} \tag{4.3}$$

where d_N is the distance between user *N* and the receiver and c is the speed of light. In many systems this difference in distance will be small given the close proximity of nodes in a sensor network. In addition, as shown in Chapter 3, the extent of the time offset problem is closely related to the size of the offset relative to the symbol length, or required data rate. Thus the time offset issue is less of a problem in systems where data rates and distances are low, such as many forms of WSNs. In these cases, time offset problems can be managed through the use of a very short cyclic prefix, as shown in Section 3.4.

An Example of Time Offsets in a Realistic System

The worst case offset occurs between the nodes with the biggest difference in distance from the receiver. Figure 4.4 uses (4.2) and (4.3) to show the maximum possible difference in distance between two nodes for given desired bit rates and different maximum offset limits, which can be used as an approximate guide to maximum network size. BPSK modulation is assumed, giving one bit per OFDM symbol. As seen in Chapter 3, a cyclic prefix of length longer than the maximum predicted time offset can be used to implement a quasi-synchronous system. There is a trade-off between the redundancy needed to implement a cyclic prefix of sufficient length, and the network dimensions for a given bit rate. One way a network designer can use the results in the figure is as a guide to maximum network dimensions. By selecting a cyclic prefix length that is reasonable for the design specification, the maximum difference in distance between two nodes and the receiver can be found for

the desired bit rate. Two examples are marked with crosses on the figure. Firstly, for a maximum acceptable cyclic prefix length of 1% of the OFDM symbol period, and a desired bit rate of 150kbps, the maximum difference in distance between two nodes is approximately 10m. A reduced cyclic prefix length of 0.1% of the OFDM symbol period can be used to transmit at a bit rate of 10kbps over a similar distance. Increasing the network dimensions requires an increase in cyclic prefix length or a decrease in achievable bit rate.

4.2.2 Correcting Frequency Offsets

A frequency offset between an individual transmitter and the receiver causes a loss in performance, as was discussed in Chapter 3. This spectral effect is reviewed in Figure 4.5. Assuming an FFT-based receiver, the spectrum is essentially sampled at regular frequencies, where each frequency belongs to an individual user. In the figure,

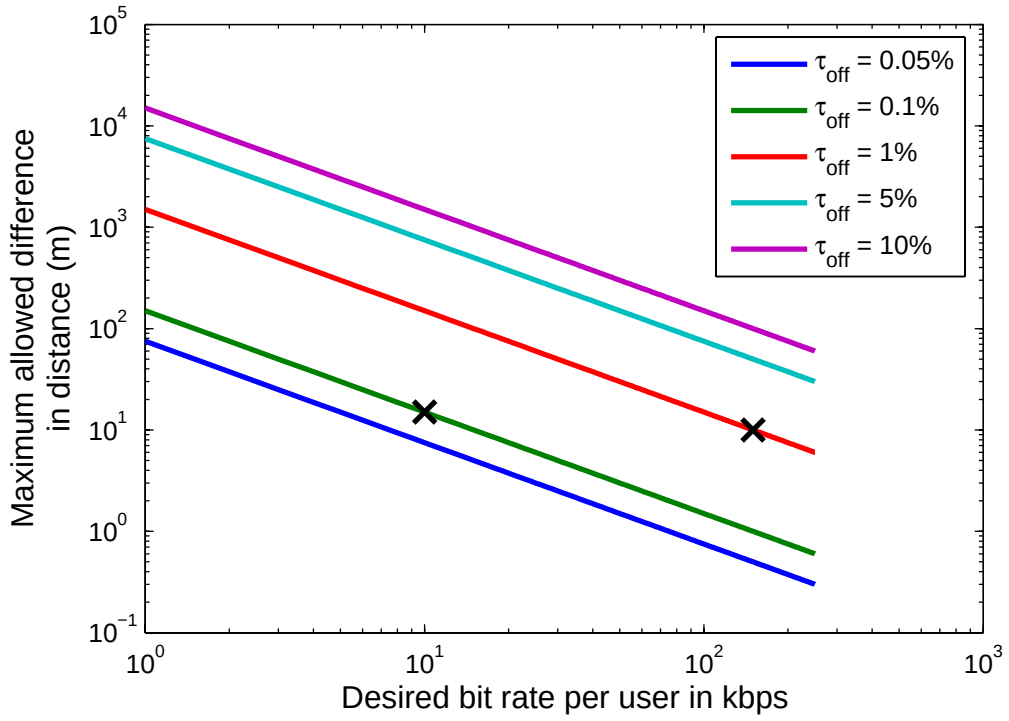

Figure: 4.4: Maximum allowed difference in distance between two nodes communicating with a central receiver. τ_{off} gives the maximum permitted offset (equal to the maximum desired cyclic prefix length), as a fraction of OFDM symbol length.

the central user (shown with a dotted spectrum) is suffering from a frequency offset, which causes a loss of orthogonality with the other users. This results in a loss of energy on the desired subcarrier (given by the central circle), and multiple access interference on adjacent users, as shown in one example by the cross marker on the plot. Unless oscillators can be assumed to be perfectly matched and nodes stationary (thus producing no Doppler shift), this is an unavoidable occurrence, and must be dealt with in either the receiver or in the transmitter, as reviewed in Section 3.4. Correction of the frequency offset can however be dealt with in the transmitter by compensating for the offset when digitally modulating onto the desired subcarrier.

Figure 4.2 presented the proposed system as a two-stage carrier modulation. Firstly the user's modulated symbols are digitally mixed onto the allocated subcarrier frequency, before the signal is digital to analogue converted and mixed onto an analogue RF carrier. The flexibility of the digital mixing stage allows frequency offset compensation to be elegantly combined with the subcarrier frequency generation.

Let f_{sub} be the subcarrier frequency allocated to a single user, and f_c be the network-wide carrier frequency, as transmitted by the receiver. The transmitting user has an offset (including oscillator error and Doppler shift) with respect to the receiver of Δf and therefore transmits:

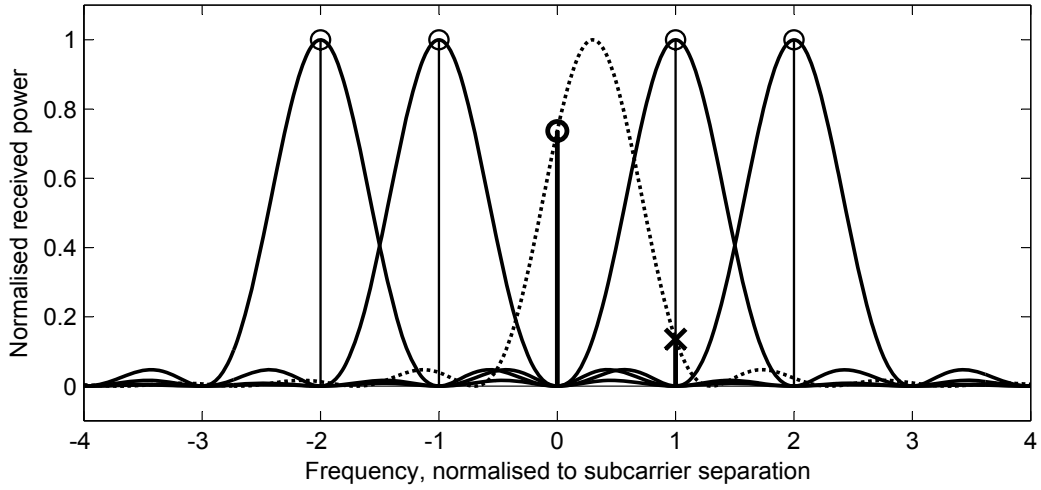

Figure: 4.5: Frequency spectrum showing a single user (dotted line) with an offset of 30% of the subcarrier spacing.

$$x(t) = Re\{s(t)e^{j2\pi(f_c + \Delta f)t}\} \tag{4.4}$$

where $s(t)$ is the baseband signal,

$$s(t) = c(t)e^{j2\pi f_{sub}t} \tag{4.5}$$

and $c(t)$ is a sequence of complex data symbols.

At the receiver the oscillator runs at target frequency f_c, giving a downmixed baseband signal (after filtering) for input to the FFT of:

$$\begin{aligned} r(t) &= x(t)e^{j2\pi f_c t} \\ &= s(t)e^{j2\pi \Delta f t} \end{aligned} \tag{4.6}$$

The correct transmitted baseband signal $s(t)$ is therefore shifted in the frequency domain by the offset, causing the aforementioned interference and associated performance loss. By making an estimate Δf of the frequency offset Δf on reception of the polling signal, the transmitter can adjust its subcarrier frequency to compensate as shown in Figure 4.6.

Thus the received baseband signal becomes:

$$x'(t) = Re\{s'(t)e^{j2\pi(f_c + \Delta f)t}\} \tag{4.7}$$

where the transmitted baseband signal was

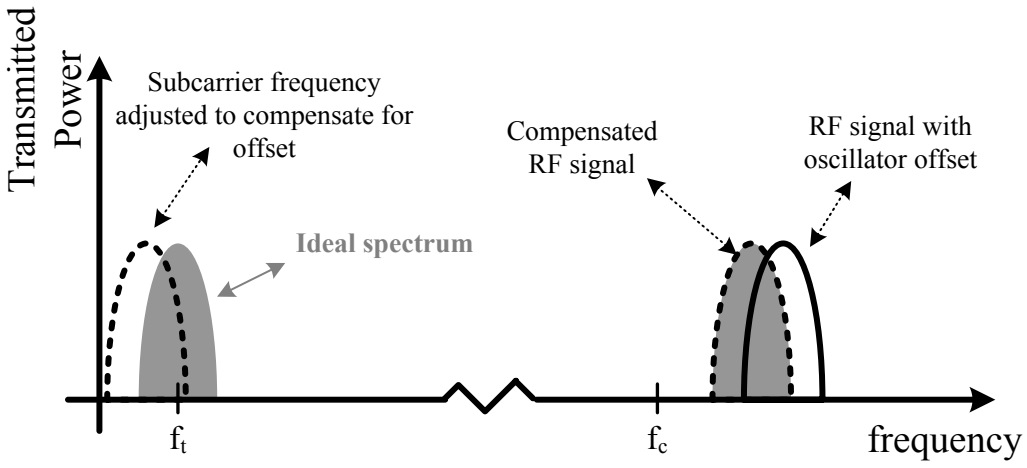

Figure: 4.6: Spectral representation of complex offset correction.

$$s'(t) = c(t)e^{j2\pi(f_{sub} - \hat{\Delta}f)t} \tag{4.8}$$

and $\hat{\Delta}f$ is the estimated compensation frequency, ideally

$$\hat{\Delta}f = \Delta f \tag{4.9}$$

giving

$$\begin{aligned} \hat{r}(t) &= x'(t)e^{j2\pi f_c t} \\ &= c(t)e^{j2\pi(f_{sub} - \Delta f + \Delta f)t} \\ &= s(t) \end{aligned} \tag{4.10}$$

The transmitted baseband signal can, in the ideal case where $\hat{\Delta}f = \Delta f$, be perfectly recovered in the presence of frequency offsets by compensating for the offset in the transmitter.

4.3 Source of the Frequency Offset

The total frequency offset between a transmitting node and the receiver is made up of two components, and is defined as:

$$\Delta f = \Delta f_c + \Delta f_D(t) \tag{4.11}$$

where Δf_c is the fixed carrier frequency offset component, and $\Delta f_D(t)$ is the Doppler offset component, which may change with time, and is defined as [58]:

$$\Delta f_D(t) = \frac{f_c}{c} V(t) \tag{4.12}$$

where $V(t)$ is the relative velocity between a transmitting node and the receiver (positive when the distance between them is reducing).

The carrier frequency offset results from the difference between the transmitter and receiver oscillator frequencies, and is thus directly related to errors in carrier frequency generation. Traditionally, carrier signals are generated using a frequency synthesiser

circuit [19], as shown in Figure 4.7. A phase locked loop (PLL) structure allows the generation of frequencies at integer multiples of a reference frequency, commonly generated by a crystal oscillator. In essence, the circuit works by attempting to lock the divided output frequency to the frequency reference, thus producing an output frequency at a multiple of the reference. As the carrier is generated directly from the crystal, the tolerance of the carrier is the same as the crystal tolerance. Crystal tolerances are generally expressed in parts per million (ppm).

The main disadvantages of the use of a frequency synthesis circuit in carrier generation are high phase noise, high power and a long start-up time. A simple, low-power solution is to use a free-running LC oscillator without a frequency reference, however the low centre frequency accuracy and high phase noise of such a solution is prohibitive for use in a system that requires a relatively accurate carrier signal. Recent advances in Thin-Film Bulk Acoustic Wave Resonators (FBARs) are producing a new category of oscillator that promises a low-power and accurate RF source [12,44]. Fabrication tolerances are currently around 500ppm, however, as shown in [44] the centre frequency can be accurately calibrated using small tuneable capacitors.

The Doppler Component

The Doppler component is included in the frequency offset estimate made by the node on reception of the polling signal. However, if the velocity changes over the period of the communication exchange, a residual frequency offset due to the Doppler

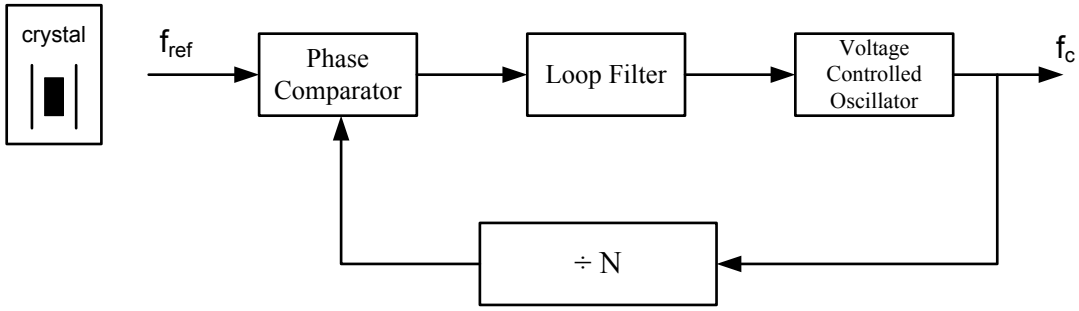

Figure: 4.7: Frequency Synthesis Circuit.

component is unavoidable. It is desirable to minimise the influence of the unavoidable, residual Doppler frequency offset, thus:

$$\frac{\mathrm{d}}{\mathrm{d}t}\Delta f_D T_{exch} \ll f_{sub} \tag{4.13}$$

i.e. the change in Doppler offset over the period of the communication exchange, $\frac{\mathrm{d}}{\mathrm{d}t}\Delta f_D T_{exch}$, must be small compared to the subcarrier spacing. This condition can also be expressed as:

$$\frac{\mathrm{d}V}{\mathrm{d}t} \ll \frac{f_{sub}c}{f_c T_{exch}} \tag{4.14}$$

giving a limit on the possible rate of change of relative velocity (acceleration) of the node to the receiver. This result is analysed with respect to the Orient-2 system in Section 4.4.3.

The frequency estimation technique must be chosen carefully in order to satisfy the trade-off between the limitations of the low-power, low-complexity nodes, and the required frequency estimate accuracy for satisfactory performance. Frequency estimate accuracy requirements, along with the choice of frequency estimation algorithm and associated design issues, are investigated in Chapter 5.

4.4 Application of the Protocol to the Orient-2 System

In this section the protocol is considered explicitly with regard to its use within the Orient-2 system. The magnitude of the time and frequency offsets that are experienced in such a network are used to place limits on the size and mobility of the network. System parameters such as desired data and error rates are defined and used to specify parameters such as the required subcarrier spacing.

4.4.1 Performance Requirements and Subcarrier Spacing

The subcarrier spacing determines the available bandwidth per user and thus the

maximum possible bit rate, depending on the modulation format. Each user transmits one data symbol per OFDM symbol, where the period of the OFDM symbol T is the inverse of the subcarrier spacing f_{sub}. The bit rate per user can therefore be found from the following relationship:

$$R_b = Mf_{sub} \tag{4.15}$$

where M is the modulation order ($M = 1$ for BPSK or DBPSK, $M = 2$ for QPSK or DQPSK). The Orient-2 system currently has a required data rate of approximately 4kbps, and therefore henceforth in this work a subcarrier spacing of 10kHz is assumed, giving a maximum bit rate of 10kbps for $M = 1$ and 20kbps for $M = 2$. Both maximum achievable bit rates give a reasonable overhead for any desired increase in data rate.

As explained in Chapter 3 a quasi-synchronous system is achievable when the cyclic prefix is longer than the maximum possible time offset present at the receiver, which is easily achieved in many sensor networks where transmission ranges are often small. In such cases there is a rotation of the constellation, which can either be corrected in the receiver for each user, or differential modulation can be used at the expense of a performance loss (see Section 3.4.5). For the Orient-2 system it is considered more important to minimise the complexity in the receiving node thus henceforth in this thesis a quasi-synchronous network with the use of DQPSK is assumed, giving a maximum bit rate of 20kbps.

A target BER of 10^{-4} and a packet length of 256 bits are assumed for the following analysis in this thesis. The BER can be related to Packet Error Rate (PER) by the following relationship:

$$PER = 1 - (1 - BER)^{N_b} \tag{4.16}$$

where N_b is the number of bits in the packet. A PER of approximately 2.5% is achieved with these parameters. This figure can be improved if necessary through the use of error coding, or by decreasing the packet length. Given the bit rate of 20kbps, the 256 bit packet allows for an update rate per user of around 78Hz (assuming a single

update per packet).

4.4.2 Network Size Limitations Due to Time Offsets

Figure 4.8 shows a simple 2D network topology for an Orient-2 system which allows worst-case time offsets to be defined. The maximum possible distance of a node to the body centre can be imagined as a sphere surrounding the body, represented by the dotted circle. The worst case situation (in terms of time offsets between nodes) is when two nodes are situated diametrically opposite each other in a straight line path from the receiver. In this case the difference in propagation distance between the two nodes is the diameter of the sphere, *d*. The receiver is situated at an arbitrary distance from the Orient-2 network, and the maximum offset seen by the receiver corresponds to twice the diameter of the circle:

$$\tau_{max} = \frac{2d}{c}. \tag{4.17}$$

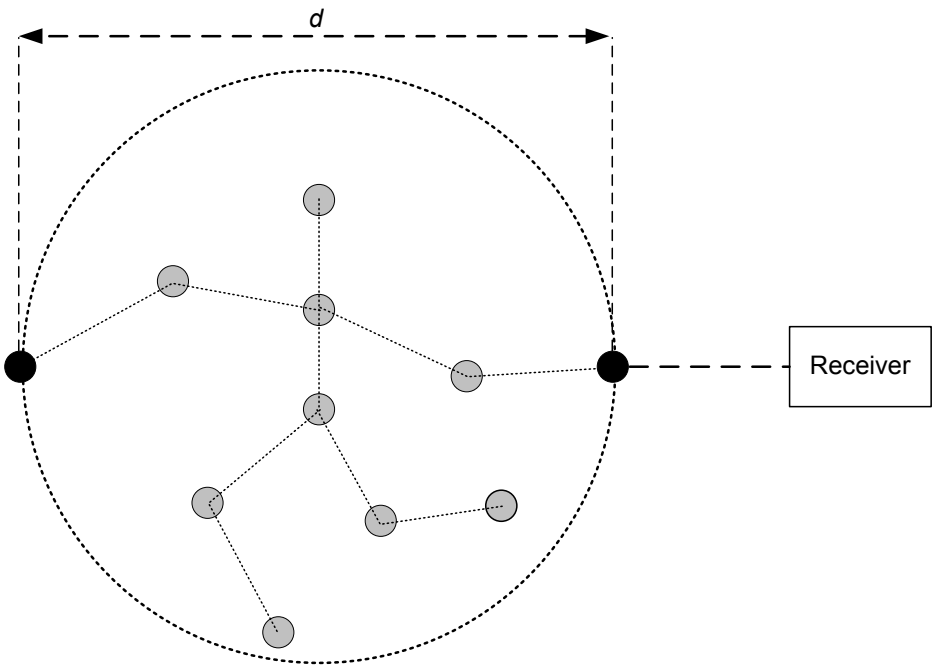

Figure: 4.8: Example Orient-2 network showing maximum node separation *d*.

Assuming the maximum required data rate to be 20kbps and DQPSK modulation, the required symbol period is 100μs. In order to keep the offset below 0.1% of the symbol period, giving a reasonable size of cyclic prefix of minimum size 0.1% of the symbol period, the maximum offset must be:

$$\tau_{max} = \frac{100\mu}{1000} = 100ns \tag{4.18}$$

giving a maximum diameter, corresponding to a maximum node separation, of d = 15m. This is more than sufficient for a single body network, however it places limitations on the area that can be covered by a multiple body network for this data rate.

4.4.3 Network Mobility Limitations Due to Frequency Offsets

The radio transceiver used in an Orient-2 node, the Texas Instruments (formerly Chipcon) CC1100, generates its RF carrier signal by frequency synthesis, and as such has a tolerance of 40ppm [85]. At the maximum possible RF frequency of around 900MHz, in the current system the maximum carrier frequency offset, Δf_c, between two users is approximately 70kHz. In this research a significant overhead is added to this figure, and the maximum offset is assumed to be 200kHz, allowing for a 40ppm oscillator at 2.4GHz, or perhaps a cheap, inaccurate oscillator at lower frequency.

In order to keep the residual Doppler component of the frequency offset as small as possible, as shown previously, the following inequality must be satisfied,

$$\frac{\mathrm{d}V}{\mathrm{d}t} \ll \frac{f_{sub}c}{f_c T_{exch}}. \tag{4.19}$$

Assuming a node to receiver distance of 10m, a 32 byte packet at 20kbps, a processing time of twice the packet length and a maximum offset of 1% of the subcarrier spacing (see Chapter 5 for further investigation of practical offset values), the maximum allowed acceleration is more than 250ms^{-2}. This translates to a maximum change in velocity of more than 25ms-1 over the period of a communication exchange, which is more than sufficient. For example, studies on karate experts have shown peak hand

accelerations of about $35ms^{-2}$ [64].

The residual Doppler component can therefore be considered negligible, and the OFDMA data extraction system is possible for the Orient-2 network assuming accurate correction of frequency offsets of the order of 100kHz, using the current radio specifications. Any implementation of the Orient-2 system using an OFDMA data extraction protocol would, however, need to use a different radio front-end, which may have different characteristics, therefore a maximum frequency offset of 200kHz is considered.

4.4.4 Wake-up Technique and Subcarrier Allocation

Consideration must be given to the practicalities involved in such a protocol such as subcarrier allocation. Each user must be assigned a distinct subcarrier from the set of possible subcarrier frequencies. Perhaps the subcarriers could be allocated intelligently in order to minimise interference between users, however in a dynamic network such as the Orient-2 this could involve a significant overhead. In this research it is assumed that the allocation has already been achieved and is static, perhaps applied during some start up phase initiated by the central receiver.

In order for nodes to conserve energy, it is desirable that they can move into a sleep mode when not transmitting. There are two options for releasing a node from this sleep state. Firstly it may have an ultra low-power secondary radio, which will detect when the central node is initiating a data extraction and wake the node up, or secondly it may operate some duty-cycle, the benefits of which were covered in Chapter 2, waking periodically to listen for a polling signal from the central receiver. If the second technique is used, the central receiver must be aware of the duty cycle and transmit over a sufficient length of time to wake all nodes up. Obviously, this approach has the drawback of a significant overhead.

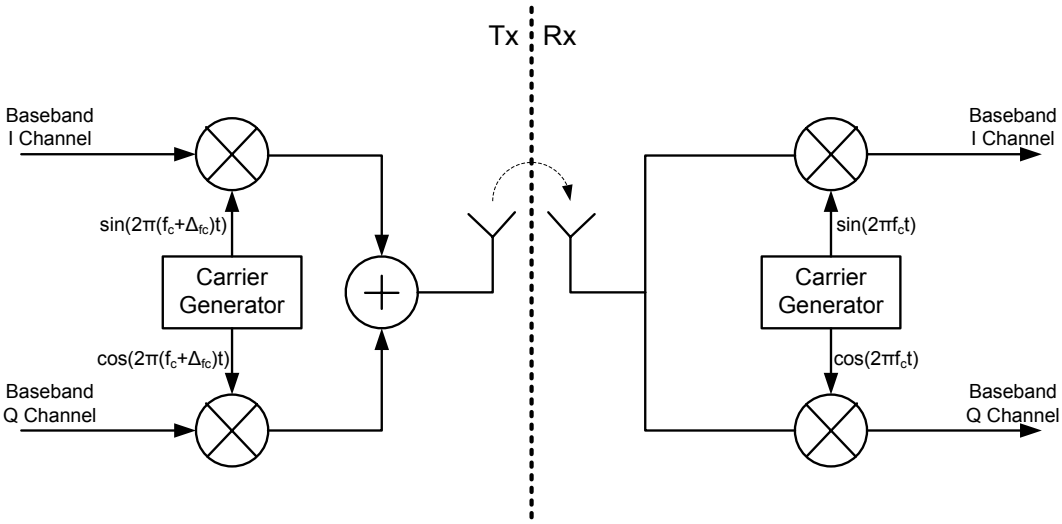

Figure: 4.9: RF section model.

4.4.5 RF Transceiver and Baseband Equivalent Channel Model

A simple block diagram of the RF section of a transmitter/receiver pair is assumed, as shown in Figure 4.9. This model allows the whole link between transmitter and receiver to be viewed as a baseband model, as it is assumed that the RF section introduces only a frequency offset. The I and Q baseband signals are mixed onto an RF carrier generated by a carrier generating circuit. Assuming the receiver carrier frequency as reference, the transmitter circuit generates a carrier at the target frequency f_c with some degree of error Δ_{f_c}. In the receiver the received signal is mixed down using a locally generated carrier signal at frequency f_c. It is this difference between transmitter and receiver carrier frequencies which produces the carrier frequency offset component of the considered frequency offset seen at the receiver.

Using the above model, the channel between the transmitter baseband output of user *u* and the central receiver baseband input can be modelled as in Figure 4.10. The transmitted baseband signal suffers from time and frequency offsets and path loss effects, before the received signal is formed by superposition with other users and the addition of noise. It is assumed that data rates and transmission distances are low

enough to render multipath negligible [17].

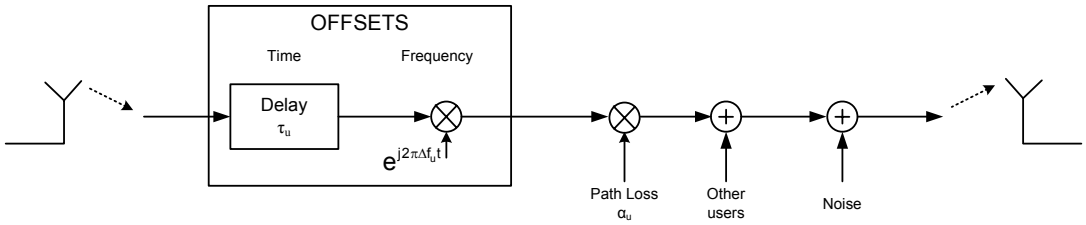

Figure: 4.10: Complex baseband equivalent channel model.

Path Loss

In the data extraction scenario discussed, users experience different path loss attenuations according to the distance between the individual node and the receiver. However, assuming the distance between the central receiver and the network is large compared to the node separation distances, that a line of sight path exists between each node and the central receiver, and that antennas are ideal, the differences in path loss attenuations will be small [59]. In addition, assuming that frequency offsets are adequately managed, multiple user interference will be minimal, and path loss effects will affect only the performance of the considered user. Without these assumptions some form of power control would be required, which could be transmitted from the receiver to the transmitting nodes within the polling signal.

4.5 Concluding Remarks

A receiver-initiated OFDMA data extraction protocol has been proposed, and its limitations in terms of time and frequency offsets determined. Its application in the Orient-2 system was shown, and an elegant, fully digital solution to frequency offset correction was introduced.

The issues involved in the implementation of the protocol were then explored, with

particular reference to the Orient-2 network. The performance requirements were outlined, in the context of a target BER and PER, and a typical subcarrier spacing for the Orient-2 system was given. A baseband channel model was introduced, and it was shown how DQPSK modulation can be used to solve the problem of carrier phase offsets. Finally, the use of the cyclic prefix to mitigate the effect of time offsets between users was explained, and guidelines were given on network size and mobility limitations.

Chapter 5

Frequency Offset Analysis and Estimation

5.1 Introduction

The receiver-initiated OFDMA protocol for data extraction was introduced in Chapter 4, and a method of correcting for a frequency offset in the transmitter was proposed. In this chapter the performance loss in the Orient-2 system due to inaccuracies in frequency estimation and compensation is considered, highlighting the crucial nature of accurate estimation of the frequency offset between a transmitting node and the central receiver. Furthermore, two simple frequency estimation algorithms are considered, with variance being used as a performance metric to make a recommendation for a suitable estimator implementation.

5.2 Frequency Estimation and Synthesis Accuracy Requirements

The frequency offset was shown in Chapter 4 to originate from oscillator inaccuracies and Doppler shifts due to node movement. Assuming that node movement is slow enough to render residual Doppler effects negligible, as was investigated in Section 4.3, it is assumed there is no residual frequency offset due to Doppler effects, and the total residual offset becomes dependent solely on inaccuracies in the frequency offset estimate and in subcarrier synthesis in the transmitter.

$$\Delta f = f_{r,\,DDS} + f_{r,\,EST} \tag{5.1}$$

where $f_{r,\,DDS}$ is the residual error in subcarrier frequency synthesis and $f_{r,\,EST}$ is the residual error in frequency offset estimation.

Assuming the effect of time offsets to be negligible due to the use of a cyclic prefix and differential modulation, the residual frequency offsets in the received signal produce a performance loss for each user. Indeed, there is a trade-off between the expense of accurate frequency estimation and synthesis and the performance loss, with associated increased transmission power requirement. As shown in the Monte Carlo simulations of Figure 3.14, frequency offsets experienced by users cause a loss in BER performance. The cumulative interference power from other users (as a result of their individual frequency offsets), and a decrease in power on the desired subcarrier due to the frequency offset on that subcarrier, result in an increased transmit power being required to achieve the same desired BER performance. The impact of the multiple user interference can be highlighted by considering the ratio $\frac{\overline{E_b}(\Delta f)}{N_0 + I_0(\Delta f)}$, where $\overline{E_b}$ is the reduced energy per bit (reduced due to a frequency offset on the desired subcarrier), and I_0 is the interference power spectral density, given by,

$$I_0(\Delta f) = \frac{P_I(\Delta f)}{W_I} \tag{5.2}$$

where P_I is the interference power and W_I is the bandwidth of the interference (in this case the subcarrier spacing, $1/T$). $\overline{E_b}$ and I_0 are both dependent on the frequency offsets present in the system, and for a zero frequency offset $I_0(0) = 0$ and $\overline{E_b}(0) = E_b$. Therefore in the ideal case,

$$\frac{\overline{E_b}(\Delta f)}{N_0 + I_0(\Delta f)} = \frac{E_b}{N_0} \tag{5.3}$$

and we achieve the theoretical BER curve, as given by (3.19) for DQPSK.

The Monte Carlo simulation assumed a worst case frequency offset scenario. However, as explained above, in reality the frequency offset for a single user is a function of the accuracy of frequency estimation and synthesis. As shall be seen in

Section 5.3 a frequency estimator is assumed to produce a zero-mean estimate with a given variance. The frequency offset of a user using this estimator will therefore possess the same variance and have mean $f_{r,\,DDS}$.

The BER performance for a user can be evaluated using the theoretical formula of (3.19), provided the total interference from other users is Gaussian. The statistical characteristics of the contribution from interfering users are investigated in the following section.

5.2.1 Investigation of Interference from Multiple Users

Assuming the receiver samples the received signal at a critical sampling rate (i.e. $f_s = Nf_{sub}$) where N is the FFT length and $N_u \le N$) the contribution to the decision variable for symbol *i* on subcarrier *m* from subcarrier *k* will be:

$$Y_{m,\,i}^{k} = \frac{1}{N} c_{k,\,i} \sum_{n=0}^{N-1} e^{j2\pi(k-m+\Delta f_k)n/N} \tag{5.4}$$

where Δf_k is the frequency offset on subcarrier *k*. For a single subcarrier *m* the decision variable is a combination of the desired component, $D_{m,\,i}^{S}$, plus interference from other subcarriers, $D_{m,\,i}^{I}$, i.e.

$$D_{m,\,i} = D_{m,\,i}^{S} + D_{m,\,i}^{I} \tag{5.5}$$

where

$$D_{m,\,i}^{S} = Y_{m,\,i}^{m} \tag{5.6}$$

and

$$D_{m,\,i}^{I} = \sum_{\substack{k=1 \\ k \neq m}}^{N_u} Y_{m,\,i}^{k} \tag{5.7}$$

where N_u is the total number of users in the system, and the D.C. subcarrier is assumed

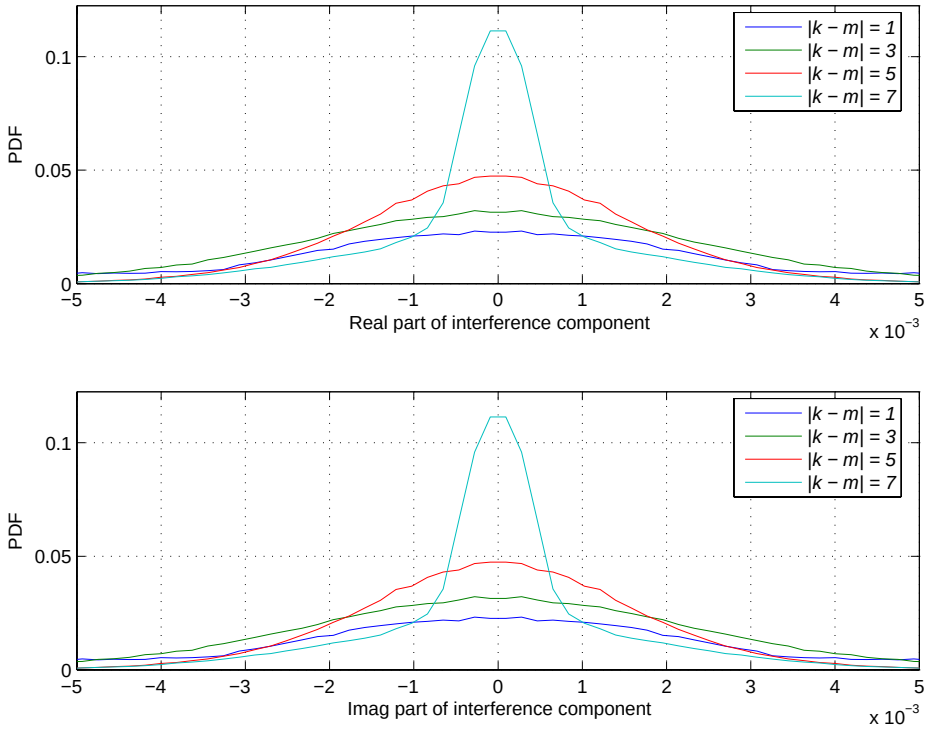

Figure: 5.1: PDFs of the real and imaginary parts of the ICI caused by frequency offsets on other users at different subcarrier separations, Y_m^k.

to be unused.

The Probability Density Functions (PDFs) of the real and imaginary parts of $Y_{m,i}^k$ for $N_u = 15$ and QPSK modulated symbols are shown in Figure 5.1 for different subcarrier separations $|k-m|$, and offsets taken from a zero-mean Gaussian distribution with variance $\sigma_{\Delta f}^2 = 0.0001$ (corresponding to a standard deviation $\sigma_{\Delta f}$ of 100Hz for a subcarrier separation of 10kHz, ideal subcarrier synthesis is assumed). The PDFs of the interference are obviously not Gaussian, however we are interested in the sum of the contributions $D_{m,i}^I$, which the Central Limit Theorem says will tend to a Gaussian distribution as $N_u \to \infty$. The PDFs of the real and imaginary parts of the total interference is the convolution of the individual contributions and is shown in Figure 5.2 along with a Gaussian distribution of equal power. It can be seen that the total interference is approximately Gaussian.

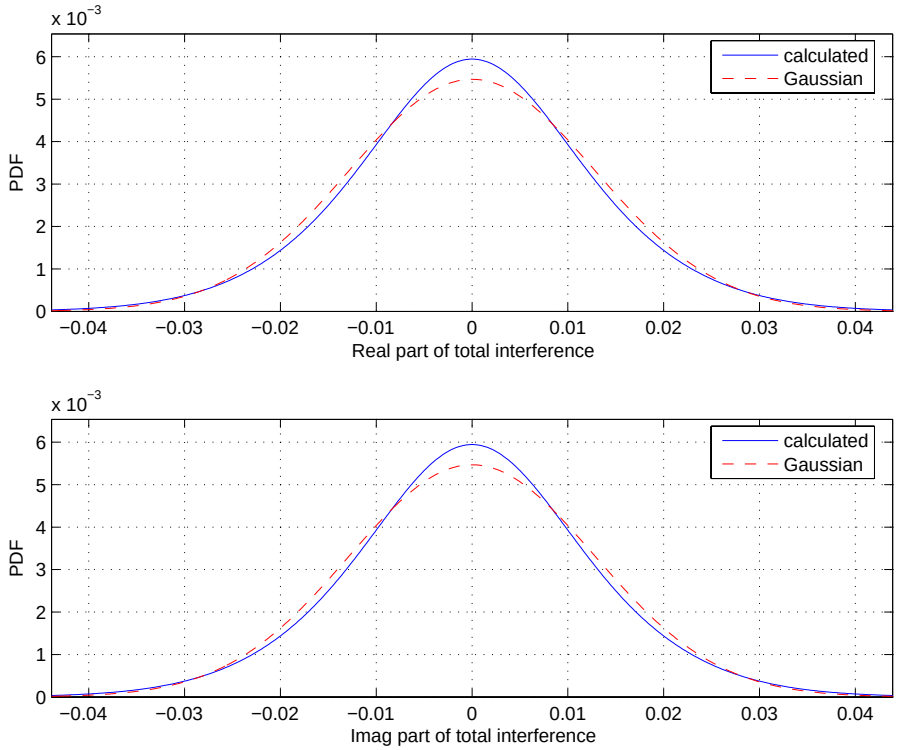

Figure: 5.2: PDFs of the real and imaginary parts of the total interference on a single subcarrier due to Gaussian frequency offsets on other subcarriers, $D^{I}_{m,i}$. The PDFs are approximately Gaussian.

The expected total interference power from multiple users due to a given variance of frequency offset can be calculated by Monte Carlo simulation. The interference power on a single subcarrier among 15 users is shown in Figure 5.3 for varying offset variances. Finding an accurate curve fit allows the interference power to be calculated without performing a Monte Carlo simulation for every offset variance. The expected interference power can be approximated by the following simple, linear relationship[1].

$$P_I(\Delta f) \approx 3.173\sigma^2_{\Delta f} + 0.0002815 \tag{5.8}$$

where $\sigma^2_{\Delta f}$ is the variance of the frequency offset. This relationship is depicted in Figure 5.3. This approximation is evaluated for the very small variance values shown

1. Calculated using The Mathworks MATLAB curve fitting toolbox [86]

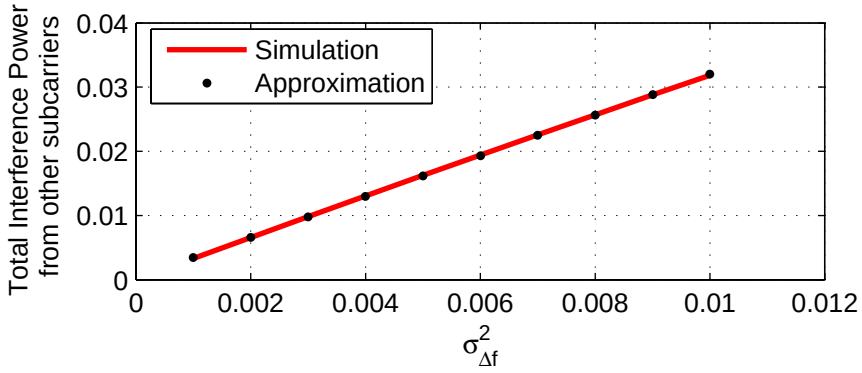

Figure: 5.3: Expected interference power, calculated by both Monte Carlo simulation and by approximation, on a single user for different offset variances.

as the estimator will need to operate at similar values for adequate BER performance, as will be shown in subsequent sections. As seen in Chapter 3, interference is most significant from users which are spectrally close to the user of interest, and becomes less significant as the interferer becomes more distant. As a result, the total interference power does not change significantly as the number of users increases, and in fact the relationship given in (5.8) can be used for any number of users above approximately 16, as can be seen in Figure 5.4 for a selection of offset variances. Figure 5.4 shows the total interference power calculated by Monte Carle simulation on a single user in systems containing different numbers of users. Each interferer has an offset taken from a Gaussian distribution with the shown variance. The figure shows that the interference

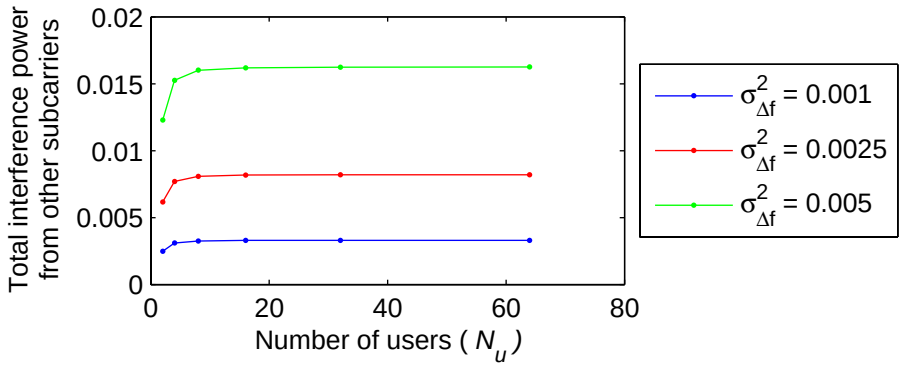

Figure: 5.4: Total interference power from other users does not increase significantly with the number of users for N_u greater than approximately 16.

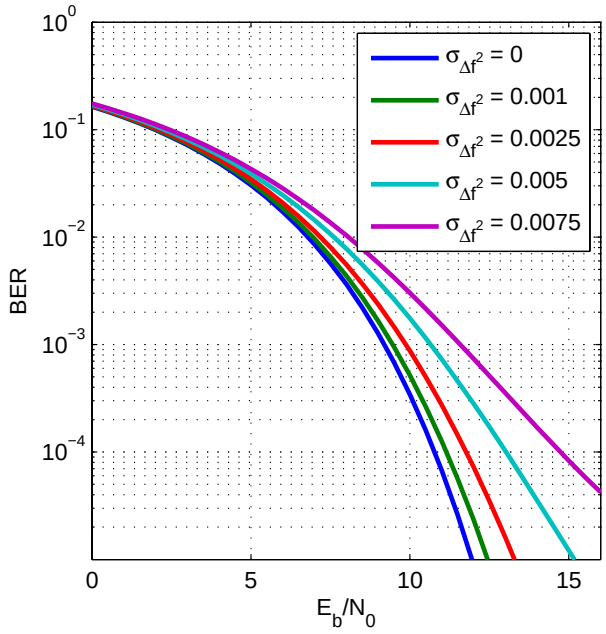

Figure: 5.5: Expected BER curves for a single user among 16 with offsets from a distribution with variance $\sigma^2_{\Delta f}$.

power doesn't increase as the numbers of users rises above 16. This is an important result as it means that the analysis carried out in this thesis can be readily applied to future Orient systems. The long term goal of the Speckled Computing project is to involve a large number of nodes, and future implementations of the Orient-2 system will involve a much greater number of sensors. Increasing the number of users in the system above the current figure of 15 will not degrade the performance seen by a single user.

Figure 5.5 shows the expected BER performance on a single user among 15 for a set of offset variances. As expected, the larger the variance of the offset, the worse the BER performance. The degradation in E_b/N_0 in the presence of frequency offsets can be seen by analysing more closely the BER curves, as shown in Figure 5.6. Assuming a desired BER of 10^{-4}, which has a required E_b/N_0 of 10.8dB, frequency offsets with variance 0.001 result in a drop in E_b/N_0 of 0.37dB, while the drop is 2.24dB at offsets with variance 0.005, as summarised in Table 2.

Residual frequency offsets must therefore be minimised to avoid the need to

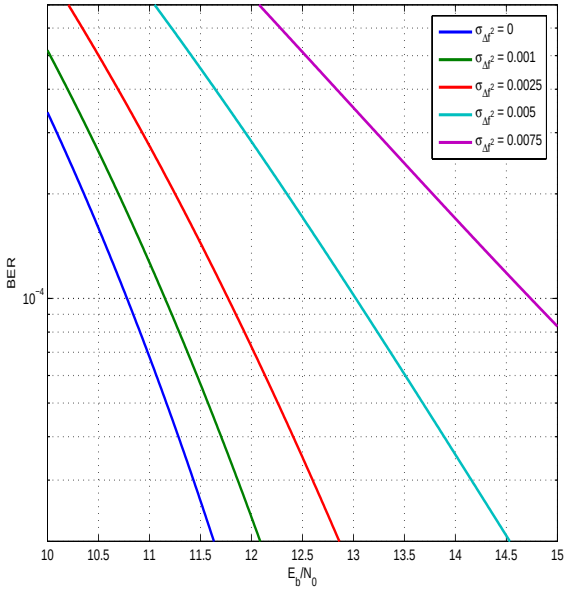

Figure: 5.6: BER performance around 10^{-4} for varying offsets from a distribution with variance $\sigma^2_{\Delta f}$.

increase transmit power to achieve the same BER. Frequency offsets are minimised by synthesising more accurate subcarrier frequencies, and estimating the carrier and Doppler frequency offsets accurately in each node. The performance curves with respect to frequency offsets therefore give a target variance for frequency offset estimation.

Offset Variance $\sigma^2_{\Delta f}$	E_b/N_0 **Degradation (dB)**
0.001	0.37
0.0025	0.99
0.005	2.24
0.0075	3.95

Table 2: Summary of E_b/N_0 degradation for a target BER of 10^{-4} given frequency offset variance $\sigma^2_{\Delta f}$.

5.2.2 Target Frequency Estimate Variance

The offset variance must be translated into a variance that can be applied to the frequency estimate. In the central node the offset variance is relative to the subcarrier spacing f_{sub}, while in each node the estimate variance is a function of the sampling rate used to generate the estimate, f_s. The two variances can be related as follows;

$$\sigma_{est}^2 = \left(\frac{f_{sub}}{f_s}\right)^2 \sigma_{\Delta f}^2. \tag{5.9}$$

5.3 Frequency Estimation

In the previous section it was shown that the interference on a single user from a number of other users is Gaussian, and as such expected BER curves for given offset variances can be calculated from theoretical models. For a given desired BER of 10^{-4}, the E_b/N_0 degradations resulting from a range of estimate variances were calculated, the importance of accurate frequency estimation therefore being made clear. A choice of suitable frequency estimation algorithm will now be investigated.

In this research only the estimation of the frequency offset is considered. In certain cases, full packets would need to be received by the nodes, and thus phase and frame synchronisation would also be required. For example to communicate control information such as subcarrier allocations. In this work, however, it is the uplink communication for data extraction that is of interest, so estimation of the frequency offset is considered sufficient. On reception of the polling signal from the central receiver to the network, each node must estimate its frequency offset in order to minimise interference between users on the uplink transmission. In order to achieve minimum complexity in each node the central node is assumed to transmit a burst of unmodulated carrier wave, which results in a problem of frequency estimation of an unmodulated sinusoid. Many frequency estimators generate an estimate through analysis of sample correlations, see [45] and references therein. Two efficient and low-

complexity frequency estimation algorithms are the Meyr [37] and Kay [29] estimators, which are based on estimating the phase difference between pairs of samples. In [2] both estimators are compared for performance and power consumption for a variety of different estimator parameters and the Kay estimator is found to be the most suitable for the particular requirements defined for that research, i.e. with the lowest energy cost. In this work an estimator must be selected with the goal of achieving a sufficiently accurate estimate to minimise interference in the OFDMA data extraction system.

Assuming, as explained previously, that the central receiver sends a burst of unmodulated quadrature carrier wave, the sampled, received baseband signal is

$$r_n = e^{j\Delta\omega n T_s} + w_n \tag{5.10}$$

where $\Delta\omega = 2\pi\Delta f$ is the carrier frequency offset, T_s is the sampling period and w_n is complex, white Gaussian noise. To illustrate the function of the estimators consider a simple example involving only two samples of the received signal, where the noise is assumed to be 0, $w_n = 0$.

$$\begin{aligned} r_{n-1} &= e^{j\Delta\omega(n-1)T_s} \\ r_n &= e^{j\Delta\omega n T_s} \end{aligned} \tag{5.11}$$

Both estimators operate on the phase difference between a pair of samples, which is proportional to the frequency of the sinusoid.

$$\begin{aligned} \Phi &= \angle r_n - \angle r_{n-1} = \Delta\omega n T_s - \Delta\omega(n-1)T_s \\ &= 2\pi\Delta f T_s \end{aligned} \tag{5.12}$$

The frequency offset, normalised to the sampling period T_s, Φ, can then be used for whatever function needed, or an absolute frequency offset value can be found by dividing by $2\pi T_s$.

To provide robustness in the presence of noise, the phase difference between pairs of samples is averaged over *L* pairs, giving an estimator "length" of *L*. The estimators for the normalised frequency offset $\hat{\Omega} = \Delta\omega T_s$ are:

1. Kay Estimator

$$\hat{\Omega}_K = \sum_{n=1}^{L-1} b_n \arg\{r_n r^*_{n-1}\} \ . \tag{5.13}$$

2. Meyr Estimator

$$\hat{\Omega}_M = \arg\left\{\sum_{n=1}^{L-1} b_n r_n r^*_{n-1}\right\} . \tag{5.14}$$

where

$$b_n = \frac{6n(L-n)}{L(L^2-1)} . \tag{5.15}$$

Both estimators have a limited estimation range

$$|\hat{\Omega}| < \pi \tag{5.16}$$

i.e.

$$|\Delta f| < \frac{f_s}{2} . \tag{5.17}$$

The weighting function as given in (5.15) can be changed to an unweighted average, i.e. $b_n = \frac{1}{L-1}$, giving a reduction in required hardware at the expense of some performance . Additionally, the variance of the estimators can be improved by a factor D at the expense of an equivalent loss of range by examining the phase difference between samples with a lag of D samples. I.e. the product $r_n r^*_{n-D}$ is used instead of $r_n r^*_{n-1}$, and the range becomes $|\hat{\Omega}| < \frac{\pi}{D}$.

5.4 Choosing an Estimator

Parameters are chosen to implement estimators for the Orient-2 network, as

summarised in Table 3. As the central receiver is assumed to possess unlimited energy, it is not limited in terms of transmit power, and thus the SNR at each node can be assumed to be flexible. For completeness two significantly different SNRs are analysed to cover a wide range, 12dB and 20dB. The maximum frequency offset to be estimated is assumed to be 200kHz, which allows for a 40ppm oscillator accuracy at a carrier frequency of 2.4GHz. Using a lower carrier frequency (as is the case in the current Orient-2 radio) would allow for less accurate oscillators to be used with the same estimator. The value of the maximum frequency allows $D = \{1, 2\}$ to be used. A target offset variance is chosen to give an E_b/N_0 degradation of approximately 1dB.

Parameter	Symbol	Unit	Value
Maximum frequency offset	Δf_{max}	kHz	200
Desired bit error rate	BER	-	10^{-4}
Signal to noise ratio at node	SNR	dB	12,20
Sampling frequency	f_s	MHz	1
Subcarrier spacing	f_{sub}	kHz	10
Number of users/nodes	N_u	-	15
Target offset variance	$\sigma^2_{\Delta f}$	-	2.5×10^{-3}
E_b/N_0 degradation	-	dB	0.99
Target estimate variance	σ^2_{est}	-	2.5×10^{-7}

Table 3: Orient-2 data extraction system parameters

Using the above parameters with (5.9) gives a target offset standard deviation of 500Hz and a target estimate variance of

$$\sigma^2_{est} = \left(\frac{f_{sub}}{f_s}\right)^2 \sigma^2_{\Delta f} = 2.5 \times 10^{-7}. \tag{5.18}$$

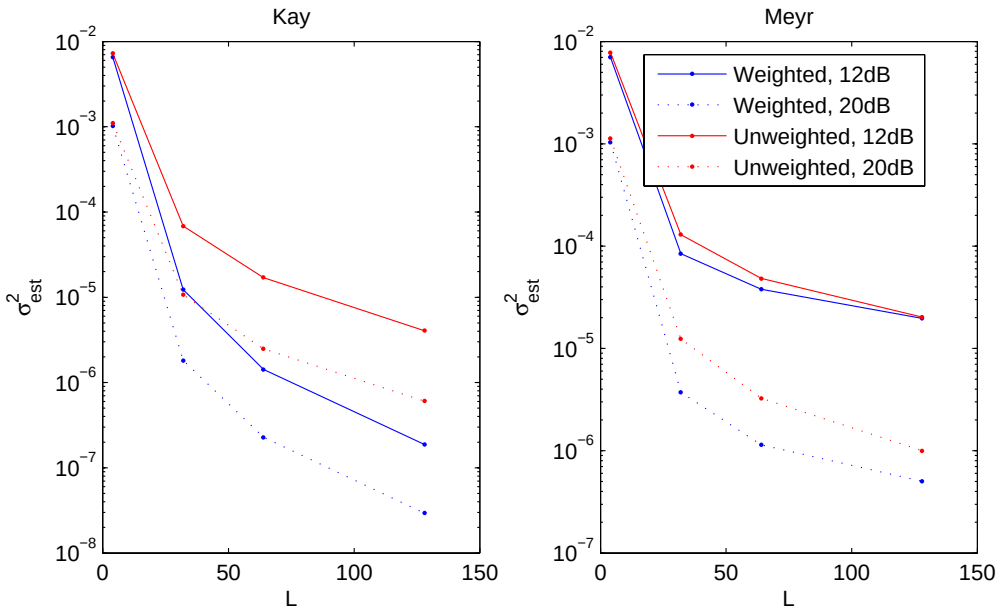

Figure: 5.7: Kay and Meyr weighted and unweighted estimator performances at 12dB and 20dB SNR.

The performances of the various estimators for the given parameters are shown in Figure 5.7. With an SNR of 12dB at the node, the only suitable estimator which can achieve the required variance with the parameters set out in Table 3 with reasonable length is the Kay weighted estimator. If the central receiver increases its transmit power until the SNR at the node is 20dB, different combinations of estimator algorithm, *L* and *D* can be used. All feasible estimators are summarised in Table 4, along with their respective variances.

The fact that the Kay algorithm is, of the estimators analysed, the only suitable estimator in a 12dB SNR environment, in addition to the findings in [2] (where the Kay estimator is found to be the most energy efficient), makes it a good selection as a suitable estimator for all round performance. If desired, and if the SNR at the node can be assumed to be at least 20dB, then the estimator can be reduced from a minimum length of 118 to only 64 samples.

SNR(dB)	Estimator	L	D	Variance
12	Kay weighted	118	1	2.45×10^{-7}
20	Kay weighted	64	1	2.45×10^{-7}
20	Kay unweighted	203	1	2.41×10^{-7}
20	Kay unweighted	146	2	2.42×10^{-7}
20	Meyr weighted	92	2	2.46×10^{-7}
20	Meyr unweighted	172	2	2.46×10^{-7}

Table 4: Possible estimator implementations

5.5 Concluding Remarks

In this chapter the frequency offset estimation requirements for the data extraction protocol for the Orient-2 system have been investigated. The size of the frequency offset on a single user signal arriving at the receiver (neglecting Doppler offset) depends on the accuracy of frequency estimation and subcarrier synthesis in the transmitting node. Considering, in isolation, the frequency offset due to inaccurate frequency offset estimation, it has been shown that the interference on a single user due to offsets on other users is Gaussian, and that it can be modelled using a simple equation dependent only on the variance of the frequency offset affecting each interfering user. In addition, it has been shown that the interference power does not greatly increase if the number of users in the Orient-2 system increases above the fifteen currently operating. This allows the results in this thesis to be applied to future Orient systems with greater numbers of nodes.

The interference results were then used to select a suitable frequency estimation algorithm. A Kay estimator of length 118 samples can be recommended as it provides sufficient accuracy to meet the desired BER degradation limit over a wide range of

SNRs. If the central receiver increases its transmit power the Kay estimator length can be significantly reduced, or a Meyr estimator can be used. The Meyr estimator, however, requires in all cases a much greater number of samples to achieve the same accuracy.

Chapter 6

Transmitter Implementation and Performance

6.1 Introduction

The receiver-initiated OFDMA data extraction protocol for the Orient-2 system was introduced in Chapter 4, and node implementation was briefly considered. The implementation of the frequency offset estimation was investigated in Chapter 5. In this chapter, the generation and transmission of the uplink signal is considered and the minimum complexity transmitter is found which satisfies the performance requirements of the Orient-2 system.

Firstly, a transmitter architecture is introduced in further detail, and the generation and modulation of subcarrier signals is considered. The subcarrier signal is synthesised using Direct Digital Synthesis, which introduces various unwanted artefacts into the transmitted signal, degrading performance. There is a trade-off between the costs of the subcarrier synthesis architecture and the errors introduced by achieving the generation as simply as possible. The performance degradation as a result of minimised hardware cost is therefore investigated in order to make a recommendation for a minimum complexity node transmitter implementation.

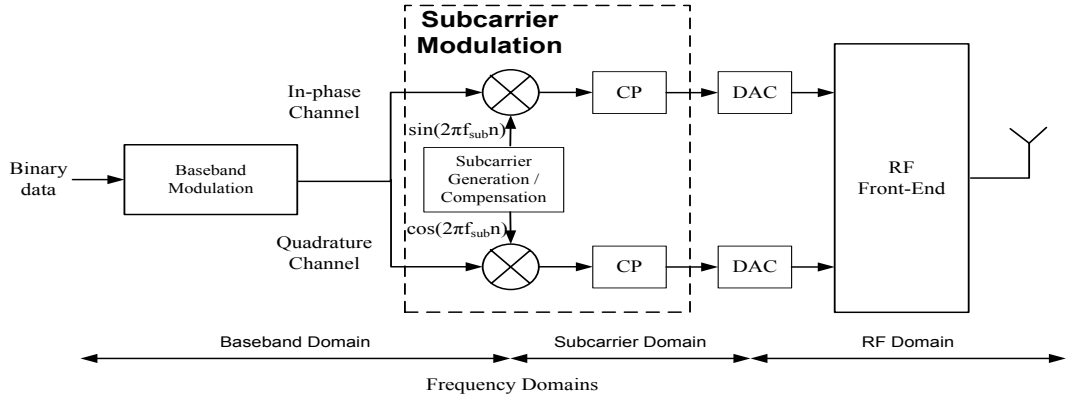

Figure: 6.1: Block diagram of individual node transmitter.

6.2 Transmitter Architecture

Figure 6.1 presents the proposed node transmitter architecture. Binary data is baseband modulated and presented to the subcarrier modulation section, inside the dotted box. Within this section the modulated symbols are mixed onto the desired subcarrier frequency, a cyclic prefix is attached, and the signal is digital-to-analogue converted, mixed onto the RF carrier and transmitted via the antenna. This chapter investigates the implementation of the subcarrier modulation section and its effect on system performance.

The subcarrier modulation section essentially performs two main functions, subcarrier synthesis and then subsequent modulation of the generated sinusoidal signal with the baseband data symbols. The effect of the DACs is not considered, i.e. it is assumed that they posses suitable Nyquist filters and introduce no significant degradation to the transmitted signal other than that caused by amplitude quantisation.

6.3 Subcarrier Generation

A subcarrier signal (real or complex depending on the baseband modulation) must

be produced at the node's desired subcarrier frequency, which is assumed to be frequency offset compensated. The signal must be generated as efficiently as possible while satisfying resolution and range requirements. The resolution of the generated sinusoids must be fine enough to select the desired subcarrier frequency with sufficient accuracy. An error in subcarrier frequency will be seen by the receiver as a frequency offset, with associated performance loss, as seen in Chapter 3. In addition, the frequency range of the generation hardware must be sufficient to cover all possible subcarrier requirements that the Orient system may produce.

There are several ways of digitally generating sinusoidal signals [51], the most popular being a structure consisting of a phase accumulator and phase-to-amplitude mapper, as shown in Figure 6.2. This structure is known as a Numerically Controlled Oscillator (NCO) or Direct Digital Synthesiser (DDS).

The phase to sine amplitude mapping block can be implemented in several different ways, [65], with the most common being a sine look-up table (LUT) addressed by the output of the phase accumulator [43,63].

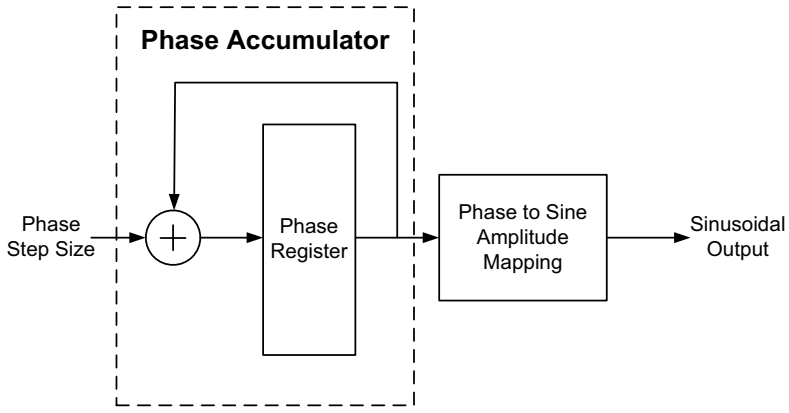

Figure: 6.2: Basic DDS architecture.

6.3.1 Direct Digital Synthesis by Phase Accumulator and Sine LUT

A more detailed model of the sine LUT-based DDS structure is shown in Figure 6.3. The phase accumulator uses overflow accumulation to model the periodic nature of the

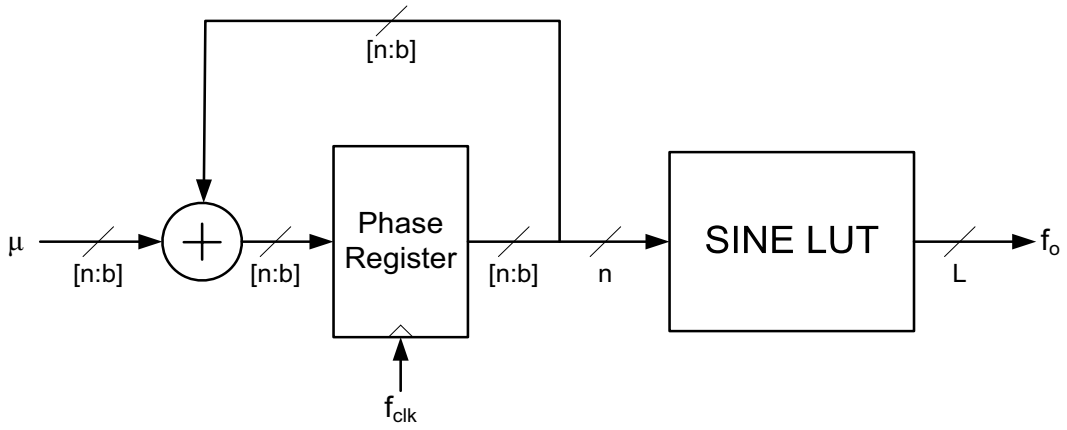

Figure: 6.3: Sine LUT-based DDS showing wordlengths.

sinusoid. The output frequency, f_o, of the DDS is governed by the phase accumulator step size, μ, which is of fixed point format with n integer bits and b fractional bits, and the size of the LUT, which is 2^n. The output frequency is given by the following relationship:

$$f_o = \frac{f_{clk}}{2^n}\mu \tag{6.1}$$

where f_{clk} is the clock frequency, or sampling rate, of the phase accumulator. The frequency resolution of the DDS is given by the minimum synthesisable frequency, which is produced when the step size is minimum, i.e.

$$f_{res} = \mu_{min}\frac{f_{clk}}{2^n} = \frac{f_{clk}}{2^n}2^{-b} = \frac{f_{clk}}{2^{n+b}}. \tag{6.2}$$

The output frequency range must be restricted, to avoid aliasing,

$$f_o < \frac{f_{clk}}{2} \tag{6.3}$$

and the sine LUT amplitude values are quantised to L bits.

There are two places in the architecture which introduce distortion to the sinusoidal output. Firstly the phase accumulator output of n integer and b fractional bits is truncated to n bits to form the LUT address. Using fractional bits in the accumulator allows for a finer frequency resolution, at the expense of inaccuracies in the output

waveform. The second source of error in the output is from the quantisation of the sine LUT entries. These errors will introduce a performance degradation on the OFDMA system, and so there is a trade-off with hardware complexity. In the following sections these issues are investigated to find the minimum complexity transmitter implementation while satisfying the Orient-2 system requirements.

6.4 Operation of the DDS

The operation of the DDS can be described using the phase wheel in Figure 6.4. The phase wheel shown is for a simple case, $n = 4$, corresponding to a 16-entry sine LUT, and as such there are 16 possible phase points on the wheel. Each point corresponds to a certain amplitude of the sine wave, as can be seen on the amplitude axis. Starting at the zero point, the phase moves round the wheel in increments of the step size, μ, and at each increment the phase is truncated to the previous phase wheel point, which in the example shown corresponds to the integer part of μ. One complete circuit of the circle corresponds to one period of the sine wave, and thus it can be intuitively seen that the frequency of the output sine wave may be varied by altering the step size. A

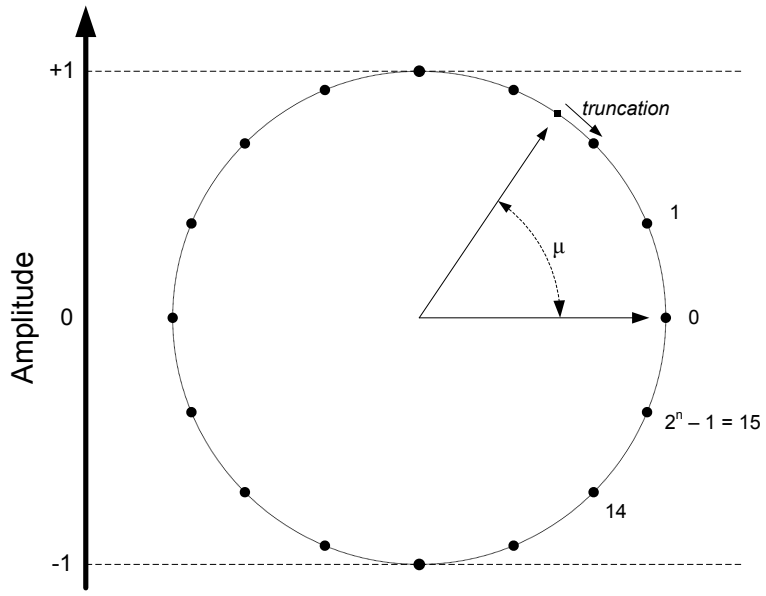

Figure: 6.4: Phase accumulator phase wheel for a 16 entry LUT.

larger step size results in a full circle being completed more quickly, and thus a shorter sine wave period. A cosine wave can be synthesised by rotating the phase wheel by $\pi/2$. At each truncation there is an error introduced into the phase of the wave, and it is this error which causes distortion in the output waveform, in addition to the quantisation of the sine LUT values.

6.4.1 Spectral Distortion Due to Phase Truncation

In order to achieve fine resolution, a large accumulator wordlength may be required. Using this whole word to address the LUT could result in a large amount of required storage. For example a 16 bit word would require a 2^{16} entry LUT, which may not be desirable, particularly in an energy constrained node. The accumulator value is therefore truncated to form the LUT address. This allows for a smaller LUT, at the expense of an error in the phase which results in distortion of the output waveform. An in-depth analysis of these errors is beyond the scope of this thesis, however it is treated in some detail in [43]. The phase error introduced by each truncation is equal to the fractional part of the phase accumulation word, i.e. *b* bits. The cyclical nature of the accumulation means that after some time the error signal will start to repeat.

The periodic nature of the errors results in the appearance of line spectra (spurs) in the spectrum of the output. The magnitude and position of these spurs is dependent on the accumulator wordlength, $n + b$, the LUT address size, n, and the step size, μ, as it is these parameters which will determine the cyclical nature of the error. The involvement of the step size in the definition of the spurs means that for every desired frequency a different set of spurs exists. Figure 6.5 shows the spectrums for two different synthesised sine waves at 20kHz and 40kHz. A DDS with $n = 5$, $b = 7$ and a 100kHz clock source was used to generate the waveforms. In order for amplitude quantisation effects to be considered negligible the sine wave values in the LUT were quantised with 32 bits resolution. The figure clearly shows the spurs created by periodic nature of the phase truncation errors.

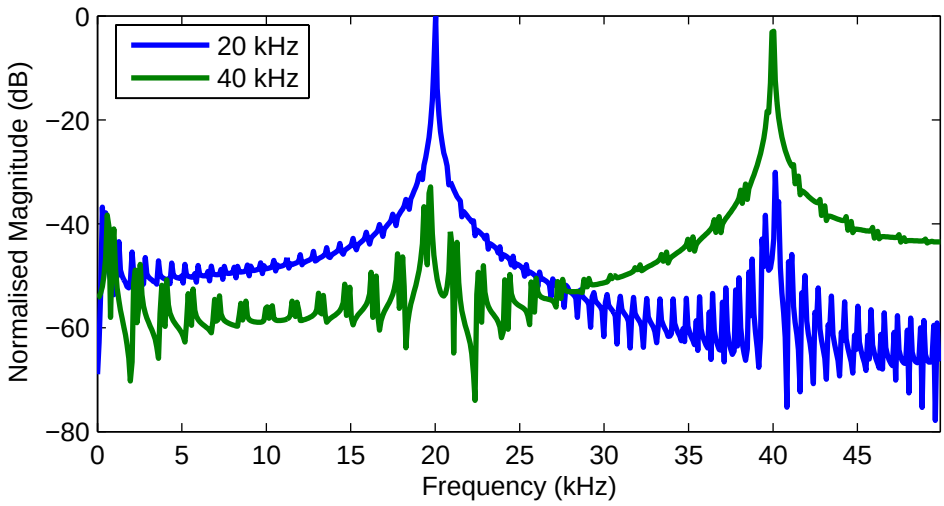

Figure: 6.5: Spectral spurs as a result of phase truncation in a DDS.

6.4.2 Spectral Distortion Due to Amplitude Quantisation

The values of the sinusoidal LUT entries must be quantised for storage. Assuming a simple 3-bit midriser quantiser, [58], as shown in Figure 6.6, the quantisation process introduces an error into the output waveform as shown in Figure 6.7. The error can also be seen on the quantiser characteristic figure, where the error for each sample is the difference between the ideal, dotted line quantiser, and the non-ideal, staircase quantiser. Much like the phase truncation error, the quantisation error of a sinusoid is

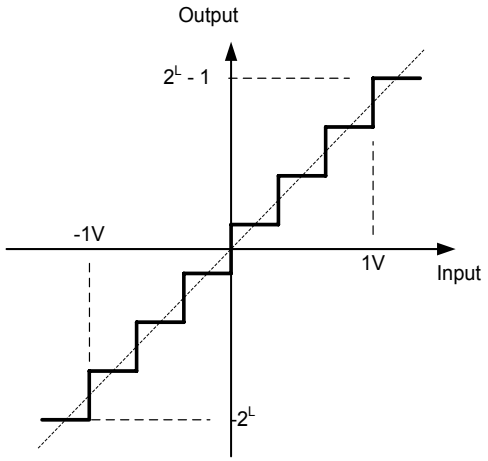

Figure: 6.6: Midriser quantisation characteristic.

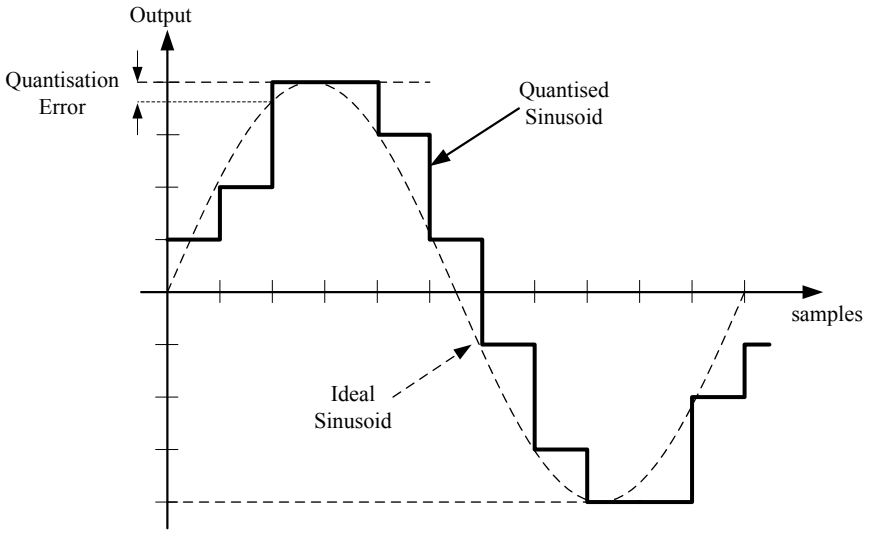

Figure: 6.7: Sine wave quantised with 3 bits resolution.

periodic, and results in spurs appearing in the quantised signal spectrum, as shown in Figure 6.8 for two very different quantisers. The figure shows a 24kHz sinusoid sampled at 100kHz and amplitude quantised with 3 and 16 bits resolution. The improvement obtained by increasing the number of bits used for quantisation is clearly noticeable. The Signal to Quantisation Noise Ratio (SQNR) for a sinusoid quantised over the full range of a midtread quantiser can be determined directly by the following equation [58]:

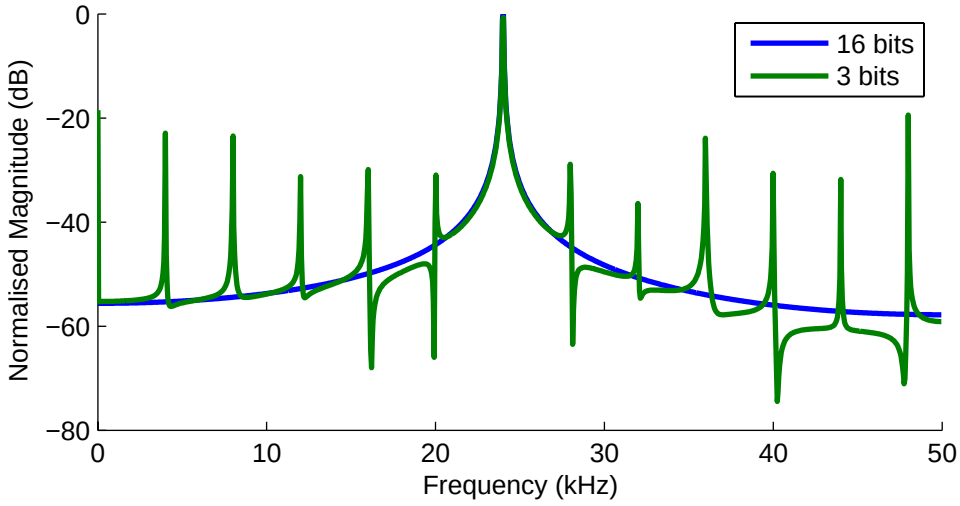

Figure: 6.8: Spectrum of an amplitude quantised sine wave showing spurs.

$$SQNR = 1.76 + 6.02dB. \quad (6.4)$$

Thus the SQNR of a quantised sinusoid is improved by around 6dB for every bit of resolution in the quantiser, and gives a good indicator of the quality of the synthesised signal. However, in the OFDMA system, the location of the spurs relative to other users is of most interest when analysing multiuser interference.

6.4.3 Implications for the OFDMA System

The effect of the DDS-generated subcarrier signals on the performance of the OFDMA data extraction system has two components. Firstly, the frequency resolution limitation of the DDS will result in a fixed offset on each user, which must be minimised in order to avoid significant performance degradation. Secondly, the spurs caused by amplitude and phase quantisation causes interference among users (as the data symbols are modulated by each spur, in addition to the desired subcarrier frequency). Both effects reduce performance, and as a result there is a trade-off between performance and implementation cost.

6.4.4 Reducing the Magnitude of the Spurs

As shown above, the spurs result from the periodic nature of the amplitude and phase quantisation errors. To reduce their impact, dithering can be used to break up the periodic nature of the quantisation errors and reduce the amplitude of the spurs, at the expense of a raised noise floor [20]. Figure 6.9 shows the same system as in Figure 6.5, but with dithering applied before phase truncation. The dithering applied is $\pm 2^{-b}$. It is noticeable that the magnitude of the highest spurs are significantly reduced and that the noise floor is higher.

6.4.5 Generating In-Phase and Quadrature Components

In order to synthesise a quadrature subcarrier signal, both a cosine and a sine wave

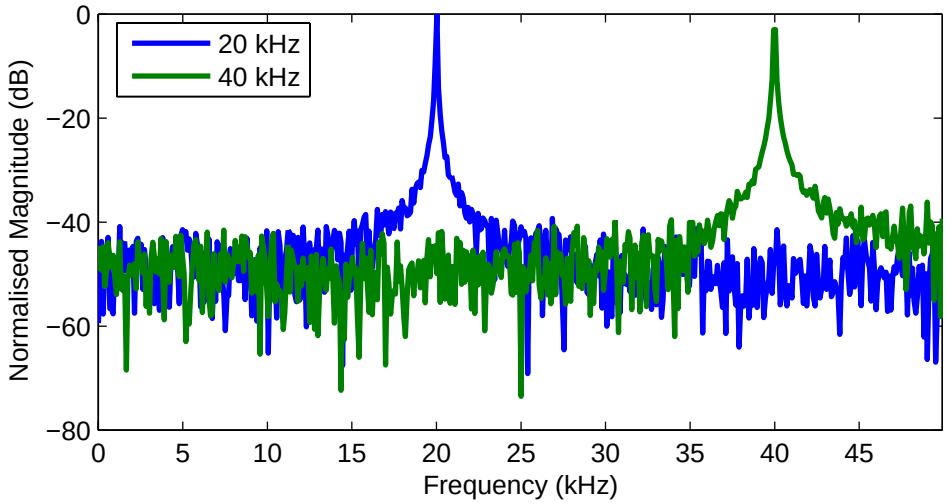

Figure: 6.9: Spectral spurs can be used by dithering before phase truncation at the expense of a raised noise floor.

must be generated. This results in the need for a double LUT architecture, with one containing sine entries, and the other cosine entries. The same accumulator address is input to both LUTs.

6.5 Subcarrier Frequency

As stated in Section 4.4.4, it is assumed that each user has been assigned a distinct subcarrier frequency. The transmit frequency must then be adjusted to include the frequency offset, as calculated during the offset estimation stage. Any resulting negative frequency can be implemented by negating the quadrature channel, or by running the accumulator with a negative step size if the chosen implementation allows.

6.6 Subcarrier Modulation Architecture

The data symbols belonging to each user must differentially phase modulate the synthesised subcarrier. Data bits are Gray coded onto a QPSK constellation as shown in Figure 6.10, the symbols are differentially encoded, and then mapped onto a

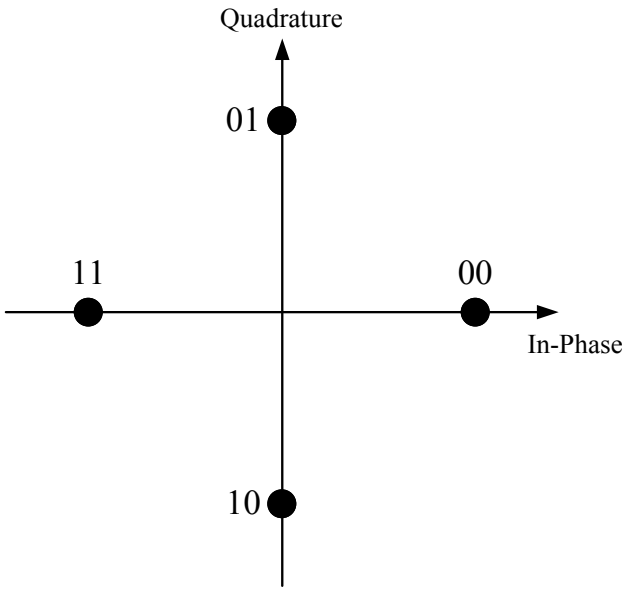

Figure: 6.10: Gray coded constellation for DQPSK symbols.

subcarrier. Table 5 shows the combinations of cosine and sine components that are required to generate the in-phase and quadrature channels of a subcarrier for a given symbol.

Using the combinations shown in Table 5, an efficient multiplexing architecture can be utilised as shown in Figure 6.11. The outputs of the sine and cosine LUTs can be negated as necessary and directly fed into the multiplexers.

Symbol	Complex Representation	In-Phase	Quadrature
0	1	cos	sin
1	i	-sin	cos
2	-1	-cos	-sin
3	-i	sin	-cos

Table 5: Combinations of sine and cosine components to produce a modulated subcarrier

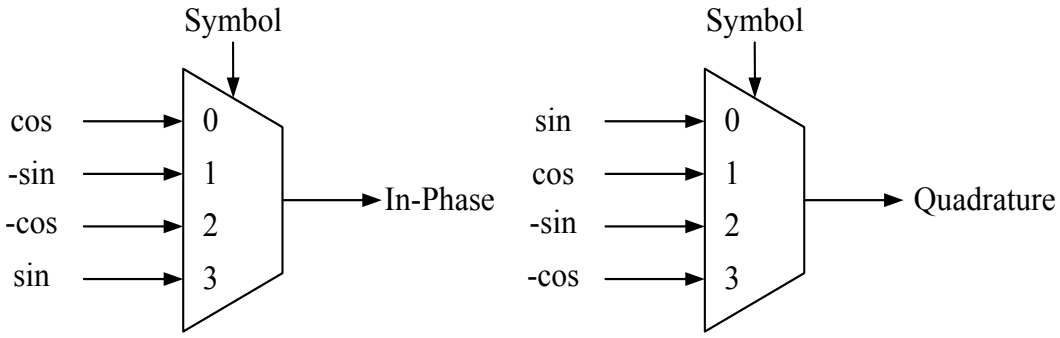

Figure: 6.11: Simple multiplexing architecture for modulation of the subcarrier.

6.7 Cyclic Prefix

To insert a cyclic prefix, a number of samples are copied from the end of a symbol and attached as a prefix to the same symbol. As seen in Chapter 4, the cyclic prefix should be chosen to provide a length of longer than the maximum possible time offset between users, which is a function of the network size. The subcarrier generation sampling frequency should be chosen as an integer multiple of the subcarrier spacing to facilitate this process. This generates an integer number of samples per OFDM symbol, and makes the extraction and addition of the cyclic prefix straighforward.

Parameter	Symbol	Unit	Value
Maximum frequency offset	Δf_{max}	kHz	200
Desired bit error rate	BER	-	10^{-4}
Sampling frequency /DDS clock frequency	f_s / f_{clk}	MHz	1
Subcarrier spacing	f_{sub}	kHz	10
Number of users/nodes	N_u	-	15

Table 6: Orient-2 data extraction system parameters

6.8 DDS Implementation

As seen above, there are many factors involved in the implementation of the subcarrier synthesis which will affect the performance of the OFDMA data extraction system. Phase accumulator wordlength, LUT size, amplitude quantisation and dithering will all result in different BER performances. Allowing for the almost 1dB degradation due to the frequency estimation inaccuracy, the DDS must be implemented as cheaply as possible while targeting minimal further performance loss. In this section DDS implementation for a node transmitter in the Orient-2 network is investigated with the goal of minimising implementation complexity as much as possible with tolerable performance degradation. The parameters given in Chapter 5 are therefore still valid, however the salient parameters are given here in Table 6 for convenience. Clock sources are assumed to be ideal.

Receiver Assumptions

In the following an ideal receiver is assumed, which allows performance to be analysed solely with respect to the transmitter implementation. It is therefore assumed that the receiver operates an ideal $(N_u + 1)$-point DFT at the critical OFDM sampling rate of $f_{rx} = (N_u + 1)f_{sub}$. In addition it is assumed that the receiver has already achieved frame and (OFDM) symbol synchronisation with the received signal, and that there are no time offsets between users (which is reasonable given the use of DQPSK in a quasi-synchronous system).

6.8.1 Frequency Range

The required range of the DDS is dependent on the maximum required subcarrier frequency it must produce, and the maximum offset it must compensate for. The maximum required frequency output of the DDS is therefore

$$N_u f_{sub} + \Delta f_{max} = 350 kHz \tag{6.5}$$

meaning that a clock speed of the same order as used for frequency estimation can be used for frequency synthesis, giving a maximum output frequency of

$$f_{o,\,max} = \frac{f_{clk}}{2} \approx 500kHz \tag{6.6}$$

which is more than adequate.

6.8.2 Frequency Resolution

In order to define the frequency resolution required in the DDS the BER degradation due to subcarrier inaccuracies is investigated. As analysed in Section 5.6, the frequency offset of a user arriving at the receiver is $\Delta f = f_{r,\,DDS} + f_{r,\,EST}$, where Δf is Gaussian with variance $\sigma^2_{\Delta f}$ and mean $f_{r,\,DDS}$. The expected interference power is again calculated by Monte Carlo simulation for $\sigma^2_{\Delta f} = 2.5 \times 10^{-3}$ and $f_{r,\,DDS}$ taken from a uniform distribution across the range $\left\{\frac{-f_{res}}{2}, \frac{f_{res}}{2}\right\}$. The total interference is seen to be zero-mean Gaussian, as the expected value of $f_{r,\,DDS}$ is zero. The frequency resolution is calculated from (6.2) as

$$f_{res} = \frac{f_{clk}}{2^{n+b}} \tag{6.7}$$

where $f_{clk} = 1MHz$. Figure 6.12 shows the total interference power on a single user among fifteen for various accumulator wordlengths and $f_{sub} = 10kHz$. It can be seen

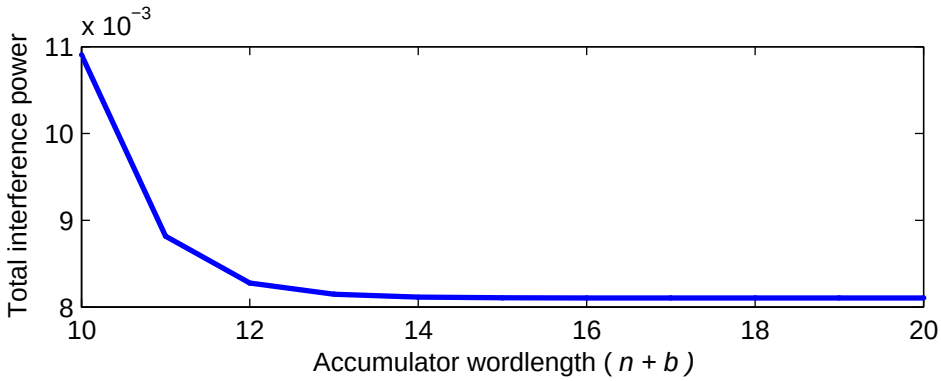

Figure: 6.12: Total interference power on a single user among fifteen with $\Delta f = f_{r,\,DDS} + f_{r,\,EST}$.

that there is no significant benefit to increasing the accumulator wordlength above 13 bits ($f_{res} = 122Hz$), as the interference power approaches the ideal value (corresponding to $f_{res} = 0Hz$). Indeed, the difference between the ideal value and the interference power for a wordlength of 13 bits is of the order of 10^{-5}.

In order to confirm the findings above the interference power values found by simulation are used to produce BER curves with the theoretical formula given in (3.19), which is valid as the total interference is again zero-mean Gaussian distributed. The performance curves around the target BER of 10^{-4} for a selection of different frequency resolutions are shown in Figure 6.13. It can clearly be seen that the BER performance is minimally degraded for accumulator wordlengths greater than around 13 bits, confirming the results in Figure 6.12. The performance degradation for 11 bits is approximately 0.1dB, for 12 bits approximately 0.03dB, and for the longer wordlengths, as can be seen in the figure, the degradation is almost negligible. A

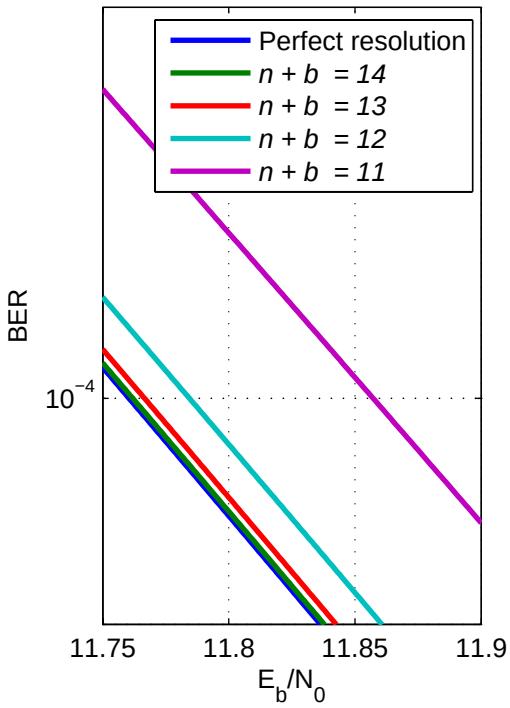

Figure: 6.13: BER degradation due to offsets produced in subcarrier synthesis.

reasonable choice of phase accumulator wordlength would, therefore, be 13 bits, giving a frequency resolution of $f_{res} = 122Hz$.

6.8.3 Look-up Table Size

The accumulator wordlength was investigated above and it was observed that a wordlength of 13 bits was a good choice. The next stage in the DDS chain takes the phase value from the accumulator and maps it to a corresponding sine and cosine amplitude value. Using a full period table, representing every possible phase that the accumulator is capable of generating, would require $2^{13} = 8192$ entries. The size of the LUT can be reduced to a quarter of the full size by using a small phase-to-quadrant mapping circuit, however this has no effect on BER performance. As explained in Section 6.4.1, the phase accumulator output can be truncated in order to reduce the size of the LUT at the expense of some loss in quality of the synthesised sinusoid. In this work, in order to minimise complexity in a transmitting node, it is desirable to reduce the LUT size as much as possible without sacrificing performance. Reducing the LUT address wordlength (*n*) by one bit corresponds to a 50% reduction in the size of the LUT.

Figure 6.14 shows the expected total interference power on a single user among fifteen for varying LUT address sizes. As each synthesised frequency possesses a

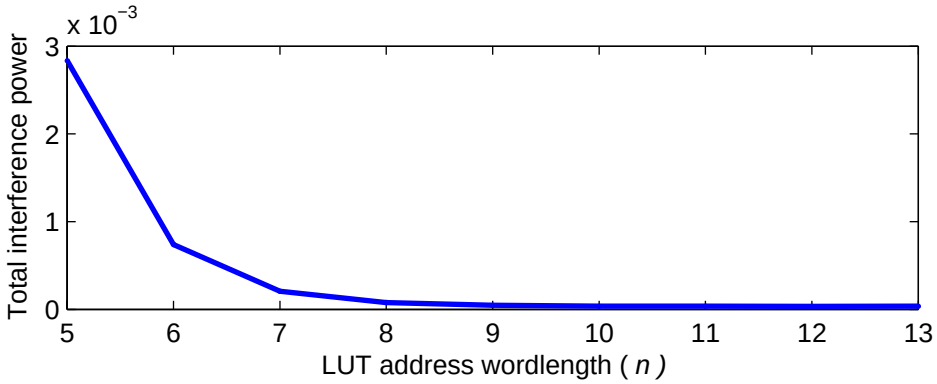

Figure: 6.14: Expected total interference power on a single user among fifteen for varying LUT sizes.

different quality of spectrum, the interference powers are calculated by Monte Carlo simulation with random frequency offsets taken from a uniform distribution with range [-200,200] kHz applied to each user, and other parameters as given in Table 6. For each simulation run, therefore, a user generates a frequency made up of the assigned subcarrier frequency and two frequency offset components; the overall frequency offset, Δf, and the residual offset after offset estimation, $f_{r,\,EST}$.

LUT entries were quantised with 32 bits resolution, which results in a resolution fine enough to assume amplitude quantisation effects negligible. From Figure 6.14 it may be concluded that there is minimal benefit to increasing the LUT address wordlength above 7 or 8 bits (corresponding to a LUT sizes of 128 or 256 entries respectively).

The above findings are again confirmed in Figure 6.15 where theoretical BER curves are plotted for the various LUT sizes. The leftmost plot shows that there is not a significant degradation with decreasing LUT size until a 64 entry LUT is reached ($n = 6$). Examining the rightmost plot, where the E_b/N_0 range is concentrated around the theoretical value of 10.8dB for the target BER of 10^{-4}, reveals that although there is a drop in performance if the LUT is reduced to 64 entries, there is minimal

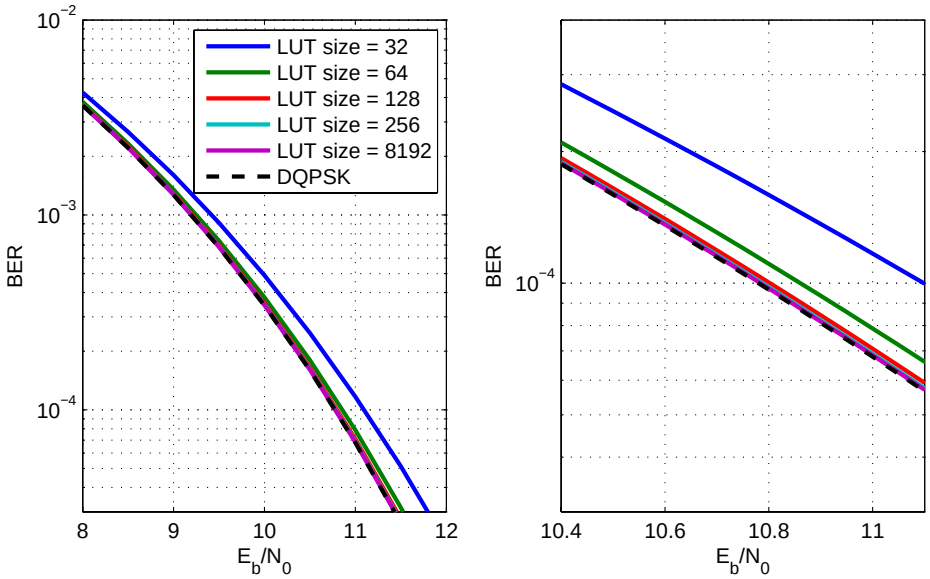

Figure: 6.15: BER performance for varying LUT sizes.

benefit to increasing the LUT size above 128 entries. A reasonable choice of LUT size would therefore be 128 entries, giving a fixed point format accumulator of $n = 7$ integer and $b = 6$ fractional bits.

6.8.4 Sine amplitude resolution.

It was seen above that the LUT can be reduced from the maximum 8192 entries to only 128 entries with minimal performance loss. As explained previously, the investigation was carried out with the LUT entries quantised to 32 bits, in order to examine only the effect of the LUT size on BER performance. It is desirable, however, to reduce this wordlength, L, to the minimum possible length so as to reduce DAC cost and power consumption, and the size of the LUT. Using the same methods as above, the expected interference power on a single user among fifteen is shown in Figure 6.16 for a selection of quantisation wordlengths. It can be seen that there is minimal benefit to increasing the quantisation wordlength to above 6 or 7 bits.

The above results are confirmed through theoretical BER analysis in Figure 6.17. The rightmost plot again shows the BER curves around the desired BER of 10^{-4}. It can be seen that the degradation only becomes significant once the wordlength is reduced to 5 bits. At 6 bits and above the degradation is minimal, and thus 6 bits is a reasonable recommendation for the amplitude quantisation of the LUT entries.

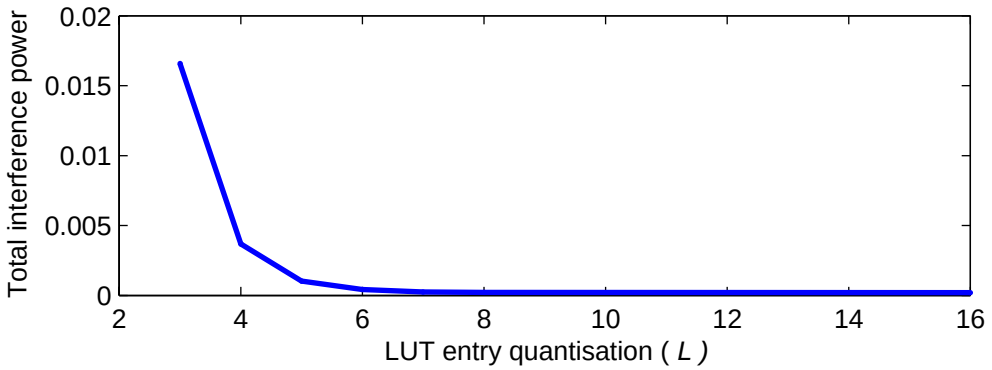

Figure: 6.16: Expected total interference power on a single user among fifteen for varying LUT entry quantisation wordlengths (L).

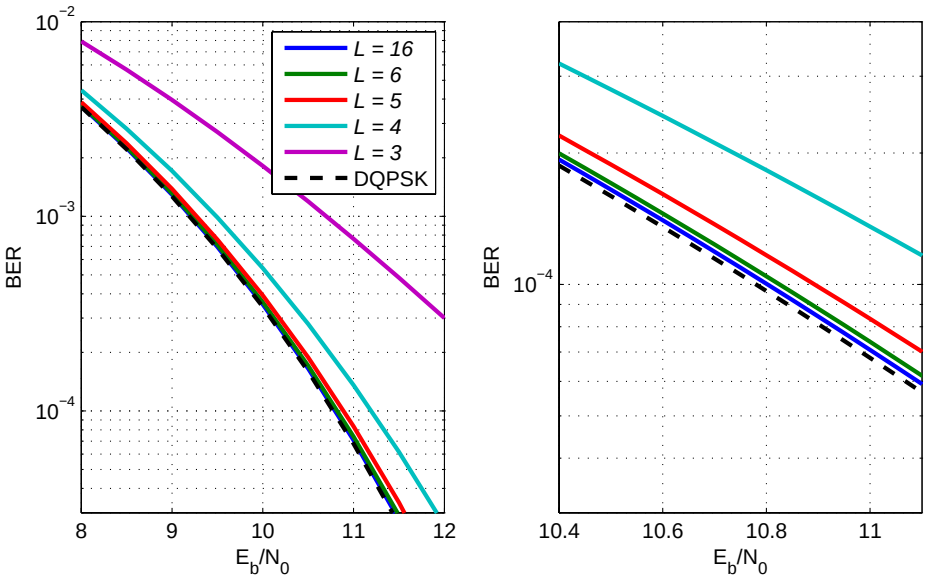

Figure: 6.17: BER performance for varying LUT entry quantisation.

6.8.5 Dithering

In this section, dithering is applied to the phase accumulator output before truncation, thus reducing the magnitude of spectral spurs at the expense of a raised noise floor. The spectrum of a selection of dithered and non-dithered transmitted signals is shown in Figure 6.18, where the user signals are created using the minimum complexity DDS parameters found above. It is obvious that the spurs have been eliminated, at the expense of a small rise in the noise floor of a few dB. It is also interesting to link this to expected BER performance. Figure 6.19 compares the expected interference power on a single user among fifteen for dithered and non-dithered signals. Perhaps unexpectedly, slightly more interference is experienced when each user uses dithering. This is due to the random nature of the spurs on the output of each user. The output spectrum of each user is dependent on the frequency that they synthesise, which is in turn entirely dependent on the frequency offset. This results in spur locations being essentially random. For a spur to cause significant interference to a certain user, it must be very close to the frequency of the user subcarrier. The

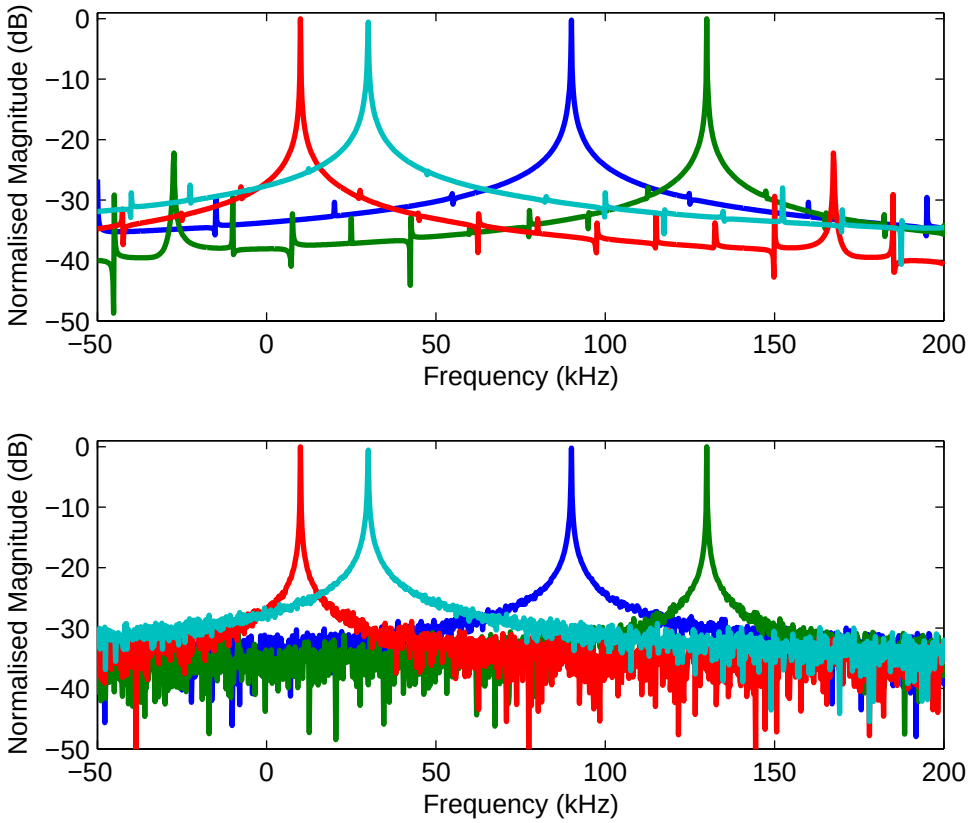

Figure: 6.18: Specta of a selection of non-dithered (top) and dithered (bottom) user signals.

probability of this occurring is small. In the case of the dithered signal, however, a raised noise floor is present at all frequencies, and will increase the expected interference power irrespective of which subcarrier frequencies are transmitted.

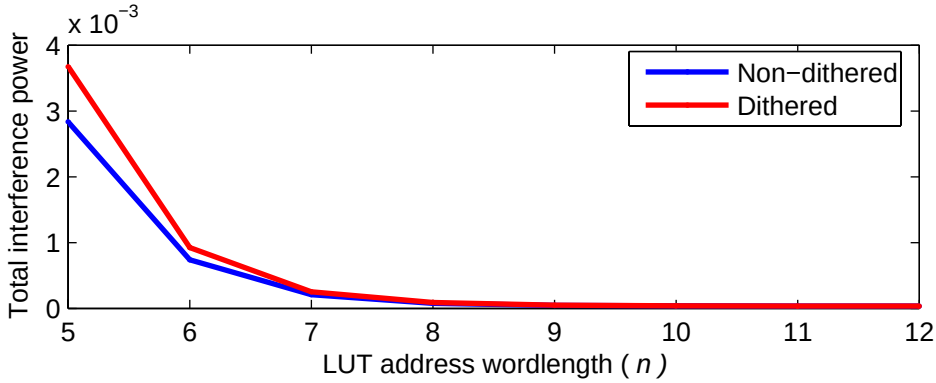

Figure: 6.19: Expected interference power on a single user among fifteen for varying LUT address wordlengths (n), dithered and non-dithered.

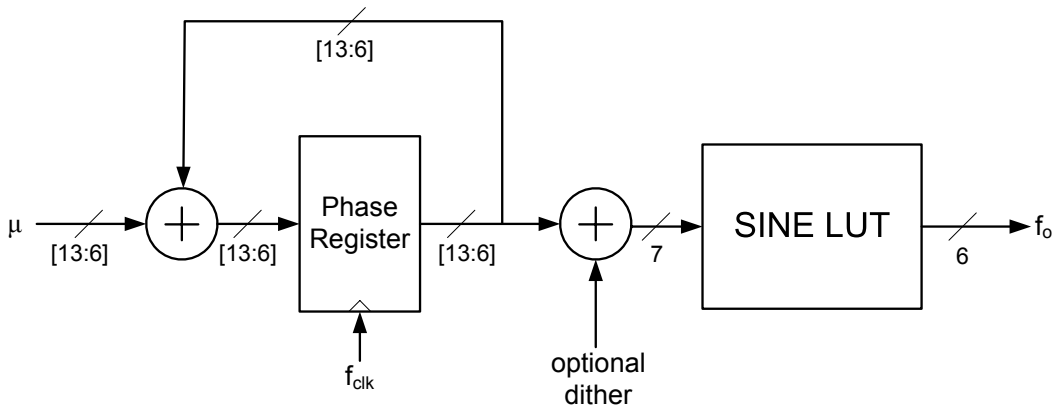

Figure: 6.20: Transmitter DDS architecture showing minimum complexity wordlengths.

The use of dithering, however, is still worthy of consideration. Although the probability of a spur and subcarrier frequency clash is small, and on average the expected interference is low, there will be occasions when a large spur interferes with a user subcarrier and causes catastrophic performance degradation. For the cost of a small loss in expected performance (the increase in interference shown in Figure 6.19 for $n = 7$ results in an E_b/N_0 degradation of less than 0.01dB at a BER of 10^{-4}), dithering ensures that this never occurs.

6.9 Design Review

In this chapter the minimum complexity implementation of a node transmitter for the Orient-2 system has been investigated. It has been shown that by designing specifically for the Orient-2 system a very simple transmitter DDS architecture can be implemented with minimal performance degradation, as shown in Figure 6.20. In summary, a DDS with a 13 bit phase accumulator and a 128 entry sine LUT with entries quantised to only 6 bits can be implemented. Dither can optionally be added before phase truncation, but in the average case provides limited benefit.

6.10 Concluding Remarks

In this chapter the implementation of a node transmitter for the Orient-2 system has been investigated. The node transmitter comprises subcarrier generation and modulation stages, which can be achieved using a sine LUT based DDS and a simple multiplexing architecture, respectively. The implementation of the DDS has been analysed in detail and and by designing specifically for the Orient-2 system, a low-complexity implementation has been found which results in minimal performance degradation compared to the ideal case. It has been shown that a DDS consisting of a 13 bit phase accumulator and a 128 entry LUT with entries quantised to 6 bits of resolution is sufficient, and causes only a small amount of BER degradation. In addition, it has been seen that, although in the average case dithering provides limited benefit, it can be used to ensure users never suffer the small possibility of catastrophic link failure due to spectral spurs.

Chapter 7

Conclusions

7.1 Résumé

In this section, the work which has been presented in this thesis is reviewed.

Wireless sensor networks were introduced with particular attention paid to **Speckled Computing**, a vision of tiny nodes enabling computing power to be embedded into the environment.

Data extraction for WSNs was discussed. The Orient-2 body posture tracking system was introduced, and the requirement for a solution to the problem of nodes communicating data to a central location under latency and update rate constraints was discussed.

Multicarrier modulation was reviewed, and its applicability within a Specknet was discussed. In particular OFDM and its multiuser extension, OFDMA, were described, and the need for tight time and frequency synchronisation between users was stressed. Standard time and frequency synchronisation techniques were investigated, and the motivation for developing a novel, Orient-2 specific solution to the offset problems was developed.

An OFDMA-based data extraction protocol was proposed, where nodes transmit on individual subcarrier frequencies to a central FFT-enabled receiver. The limitations of the protocol were investigated as applied to the Orient-2 system, and it was explained how it addresses the problem of time and frequency synchronisation between users in a novel manner.

An elegant transmitter architecture was then investigated, which combines subcarrier modulation and frequency offset estimation into a single procedure. The implementation of the transmitter was considered in some detail, with a recommendation being made for a minimum complexity architecture which results in minimal performance degradation.

7.2 Conclusions

In this section, the main results of this work are described.

- A novel data extraction protocol in which users transmit on individual OFDMA subcarriers has been introduced. The protocol allows multiple nodes to communicate with a central receiver concurrently, thus resolving previously existing limitations on update rate and latency.

- The viability of the protocol for use with the Orient-2 body posture tracking system has been investigated in detail and the main issues resolved. With the use of a receiver-initiated communications exchange, as detailed in Figure 4.3, time offsets between received signals are mitigated through the use of a cyclic prefix and differential modulation. Frequency offsets are estimated in the downlink and compensated for in each node’s transmission.

- A novel method of combining subcarrier modulation with frequency offset compensation has been investigated. A single, elegant, entirely digital architecture has been introduced which ensures frequency synchronisation between users.

- It has been shown that the interference on a single user due to offsets on other users is Gaussian, and that it can be modelled using a simple equation dependent only on the variance of the frequency offset affecting each user. A frequency offset variance, $\sigma_{\Delta f}^2$, of 2.5×10^{-3} (normalised to the subcarrier spacing of 10kHz) results in around 1dB of E_b/N_0 degradation at the target BER of 10^{-4}.

- It has been shown that the interference power on a single user does not significantly increase if the number of users increases above the fifteen used in the current implementation. This is important as it allows the results found in this thesis to be readily applied to future systems containing greater numbers of nodes.

- The implementation of the protocol for the Orient-2 system has been investigated, with recommendations made for implementation of the frequency estimation and frequency compensation/subcarrier modulation stages of the communications exchange.

- The interference on a user due to frequency offsets has been analysed in detail in order to place requirements on frequency offset estimation accuracy. The defined frequency offset variance requirements have been translated into a specification of an estimator that possesses a variance of lower than 2.5×10^{-7} (normalised to the assumed sampling rate of 1MHz). For the defined requirements, it has been shown that the Kay estimator provides good all round

performance. Assuming an SNR of 12dB at the node, an estimator of length 118 samples is required. If desired, the SNR can be increased, and the estimator length reduced.

- The implementation of the transmitter was considered in some detail, with a recommendation being made for a low-complexity architecture which limits performance degradation when compared to the ideal case. Subcarrier synthesis is achieved using a DDS with the minimum complexity wordlengths as shown in Figure 6.20. A DDS with a 13 bit phase accumulator and a 128 entry sine LUT with entries quantised to only 6 bits can be implemented with minimal performance degradation.

- Dithering can be used to eliminate the spectral spurs produced by the DDS at the expense of a raised noise floor. It can therefore be used to ensure that users never experience the small possibility of catastrophic link failure due to spectral spurs exactly coinciding with another user's subcarrier frequency.

7.3 Future Work

This section considers the limitations of the protocol as it has currently been investigated, and suggests other ways in which it could be applied.

- There are some assumptions made within this work which when removed provide interesting problems to be resolved. Firstly, the difference in path loss between different users and the receiver is considered negligible. Removing this assumption would require power control to be considered. In this case an efficient solution would be needed that avoided significant control signalling overhead. Secondly, subcarrier allocation is considered fait accompli, perhaps achieved during some start-up phase. It is interesting to consider how this could

be achieved, and indeed if allocation of subcarriers to users could be intelligently managed to minimise interference, and reduce frequency synchronisation constraints.

Both of the above problems would be interesting to consider, and could provide a starting point for further investigation of the use of the polling signal to communicate control signalling.

- A comparison with other methods of generating parallel channels would be worthwhile. For example how would the protocol introduced in this thesis compare to a similar scheme using CDMA, or indeed standard FDMA? A study of this kind would require investigation of the receiver in addition to the node transmitters.

- The protocol could be investigated for use within a WSN, as a receiver-initiated MAC protocol for multi-hop communications. There is potential for the investigation of minimal-complexity FFTs within energy-constrained nodes for the implementation of a similar protocol.

- Finally, a physical implementation of the protocol for use with the Orient-2 system would be an interesting challenge. The Orient-2 system is already proving to be an effective method of tracking body posture, and with further improvements, and an implementation of the protocol described in this thesis, future generations of the system should find application in many domains.

References

[1] Akyildiz, I.F.; Weilian Su; Sankarasubramaniam, Y.; Cayirci, E., "A Survey on Sensor Networks," *IEEE Communications Magazine*, vol.40, no.8, pp. 102-114, Aug 2002.

[2] Ammer, M.J.; Rabaey, J., "Frequency offset estimation with improved convergence time and energy consumption," *IEEE Eighth International Symposium on Spread Spectrum Techniques and Applications*, pp. 596-600, 30 Aug.-2 Sept. 2004.

[3] Ammer, J.; Rabaey, J., "Low Power Synchronization for Wireless Sensor Network Modems," *IEEE Wireless Communications and Networking Conference*, pp. 670-675, Vol. 2, 13-17 March 2005.

[4] Arvind, D. K.; Elgaid, K.; Krauss, T.; Paterson, A.; Stewart, R.; Thayne, I., "Towards An Integrated Design Approach to Specknets," *IEEE International Conference on Communications*, 2007, pp.3319-3324.

[5] Arvind, D.K.; Wong, K.J., "Speckled Computing: Disruptive Technology for Networked Information Appliances," *IEEE International Symposium on Consumer Electronics*, 2004, pp. 219-223.

[6] Arampatzis, Th.; Lygeros, J.; Manesis, S., "A Survey of Applications of Wireless Sensors and Wireless Sensor Networks," *Proceedings of the 2005 IEEE International Symposium on Intelligent Control, 2005, Mediterranean Conference on Control and Automation*, pp.719-724, 27-29 June 2005.

[7] Barbarossa, S.; Pompili, M.; Giannakis, G.B., "*C*hannel-independent Synchronization of Orthogonal Frequency Division Multiple Access Systems,"*IEEE Journal on Selected Areas in Communications*,vol.20, no.2, pp.474-486, Feb 2002.

[8] van de Beek, J.J.; Sandell, M.; Borjesson, P.O., "Ml Estimation of Time and Frequency Offset in OFDM Systems," *IEEE Transactions on Signal Processing*, vol.45, no.7, pp.1800-1805, Jul 1997.

[9] van de Beek, J.-J.; Borjesson, P.O.; Boucheret, M.-L.; Landstrom, D.; Arenas, J.M.; Odling, P.; Ostberg, C.; Wahlqvist, M.; Wilson, S.K., "A Time and Frequency Synchronization Scheme for Multiuser OFDM," *IEEE Journal on Selected Areas in Communications*, vol.17, no.11, pp.1900-1914, Nov 1999.

[10] Bingham, J.A.C., "Multicarrier Modulation for Data Transmission: an idea whose time has come," *IEEE Communications Magazine*, vol.28, no.5, pp.5-14, May 1990.

[11] Zhongren Cao; Tureli, U.; Yu-Dong Yao; Honan, P., "Frequency Synchronization for Generalized OFDMA Uplink," *IEEE Global Telecommunications Conference*, vol.2, pp. 1071-1075, 29 Nov.-3 Dec. 2004.

[12] Chee, Y.H.; Niknejad, A.M.; Rabaey, J., "A Sub-100 μW, 1.9-GHz CMOS Oscillator using FBAR Resonator," *Radio Frequency integrated Circuits (RFIC) Symposium*, 12-14 June 2005.

[13] Cooley, J, W., and J. W. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series," *Math. Comput.* 19: 297–301, 1965.

[14] L. H. Crockett - “On Code Division Multiple Access Applied to Specknets.” PhD thesis, University of Strathclyde, 2008.

[15] Crossbow Technology Iris Wireless Sensor Network Module, data sheet available: http:// www.xbow.com/Products/Product_pdf_files/Wireless_pdf/IRIS_Datasheet.pdf.

[16] Darbari, F.; Glover, I. A.; Stewart, R. W., "Channel and Interference Analysis for Wireless Sensor Networks," *IEEE International Conference on Communications*, 2007, pp.3289-3294, 24-28 June 2007.

[17] Darbari, F - “Wireless Channel Modelling for Specknet.” PhD thesis, University of Strathclyde, 2008.

[18] Dust Networks SmartMesh-XT M2030 wireless sensor node, data sheet available: http:// www.dustnetworks.com/docs/M2030.pdf.

[19] W. F. Egan, *Frequency Synthesis by Phase Lock*, John Wiley & Sons, 1981.

[20] Flanagan, M.J.; Zimmerman, G.A., "Spur-reduced digital sinusoid synthesis," *IEEE Transactions on Communications*, vol.43, no.7, pp.2254-2262, Jul 1995.

[21] Gudmundson, M.; Anderson, P.-O., "Adjacent Channel Interference in an OFDM System," *46th IEEE Vehicular Technology Conference*, vol.2, pp.918-922 , 28 Apr-1 May 1996.

[22] J. Heidemann, F. Silva, C. Intanagonwiwat, R. Govindan, D. Estrin, D. Ganesan, “Building Efficient Wireless Sensor Networks with Low-Level Naming,” *Proceedings of the eighteenth ACM symposium on Operating Systems Principles*, Banff, Canada, 2001.

[23] J. Heidemann, F. Silva, D. Estrin, “Matching Data Dissemination Algorithms to Application Requirements,” *Proceedings. 1st International Conference on Embedded Networked Sensor Systems*, pp. 218 - 229, LA, 2003.

[24] C. J. Hwang, “24GHz Low Power UWB Compact TX/RX MMIC Design”, *Presentation at the 6th Workshop in Speckled Computing*, 2007. Available http://www.specknet.org/publications/ 6thWorkshop2007/Hwang.pdf.

[25] Intanagonwiwat, C.; Govindan, R.; Estrin, D.; Heidemann, J.; Silva, F., "Directed Diffusion for Wireless Sensor Networking," *Networking, IEEE/ACM Transactions on* , vol.11, no.1, pp. 2-16, Feb 2003.

[26] Intersense Wireless InertiaCube3, Datasheet. Available: http://www.intersense.com.

[27] Jiménez, V.; Armada, A.; “Multi-user Synchronisation in Ad Hoc OFDM-based Wireless Personal Area Networks”, *Wireless Personal Communications*, Volume 40, Number 3, February 2007.

[28] R. Jurdak, C. Lopes, and P. Baldi. “A Survey, Classification and Comparative Analysis of Medium Access Control Protocols for Ad Hoc Networks”, *IEEE Communications Surveys and Tutorials*, vol. 6, no. 1, 2004.

[29] Kay, S., "A Fast and Accurate Single Frequency Estimator," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol.37, no.12, pp.1987-1990, Dec 1989.

[30] Krishnamachari, L. Estrin, D. Wicker, S. “The Impact of Data Aggregation in Wireless Sensor Networks”, *Proceedings. 22nd International Conference on Distributed Computing Systems*, Vienna, Austria, 2002.

[31] S. Kaiser, W. Krzymien, “Performance Effects of the Uplink Asynchronism in a Spread Spectrum Multi-Carrier Multiple Access System”, *European Trans. on Telecommunications*, July-August 1999, pp. 399-406.

[32] Kulkarni, G.; Adlakha, S.; Srivastava, M., "Subcarrier Allocation and Bit Loading Algorithms for OFDMA-Based Wireless Networks," *IEEE Transactions on Mobile Computing*, vol.4, no.6, pp. 652-662, Nov.-Dec. 2005.

[33] M. Kuorilehto, M. Hannikainen and T. D. Hamalainen, "A Survey of Application Distribution in Wireless Sensor Networks," *EURASIP Journal on Wireless Communications and Networking*,Issue 5, 2005, pp. 774 - 788.

[34] MacEwen N C , Crockett L H, Pfann E and Stewart R W, "Symbol Synchronisation Implementation for Low-Power RF Communication in Wireless Sensor Networks" *39th IEEE Asilomar Conference on Signals, Systems and Computers*, Asilomar, CA, USA, October – November 2005.

[35] McGregor, I, Wasige, E and I Thayne, "Sub Milli-Watt, 2.4 GHz, Super-Regenerative Transceiver with Ultra Low Duty Cycle", *Asia Pacific Microwave Conference*, Yokohama, Japan, December 2006.

[36] McGregor, I, Whyte, G, Elgaid, K, Wasige, E and Thayne, I, "A 400uW Tx/380uW Rx 2.4GHz Super - Regenerative GaAs Transceiver," *Proc European Microwave Confernence*, pp 1523-1525, Manchester, UK, Sept 2006.

[37] H. Meyr, M. Moeneclaey, and S. A. Fechtel, *Digital Communication Receivers: Synchronization, Channel Estimation and Signal Processing*, Wiley Press, 1998.

[38] Hlaing M.; Bhargava, V.K.; Letaief, K.B., "A Robust Timing and Frequency Synchronization for Ofdm Systems," *IEEE Transactions on Wireless Communications*, vol.2, no.4, pp. 822-839, July 2003.

[39] Moose, P.H., "A Technique for Orthogonal Frequency Division Multiplexing Frequency Offset Correction," *IEEE Transactions on Communications*, vol.42, no.10, pp.2908-2914, Oct 1994.

[40] Morelli, M.; "Timing and Frequency Synchronization for the Uplink of an OFDMA System," *IEEE Transactions on Communications*, vol.52, no.1, pp. 166-166, Jan. 2004.

[41] Morelli, M.; Kuo, C.-C.J.; Pun, M.-O., "Synchronization Techniques for Orthogonal Frequency Division Multiple Access (OFDMA): A Tutorial Review," *Proceedings of the IEEE* , vol.95, no.7, pp.1394-1427, July 2007.

[42] R. R. Mosier and R. G. Clabaugh, "Kineplex, a Bandwidth-Efficient Binary Transmission System", *AIEE Transactions*, Vol. 76, January 1958, pp. 723 - 728.

[43] Nicholas, H.T.; Samueli, H., "An Analysis of the Output Spectrum of Direct Digital Frequency Synthesizers in the Presence of Phase-Accumulator Truncation," *41st Annual Symposium on Frequency Control*, pp. 495-502, 1987.

[44] B. P. Otis; J. M. Rabaey. "A 300µW 1.9GHz CMOS Oscillator Using Micromachined Resonators," *IEEE Journal of Solid-State Circuits*, vol.38, no.7, pp. 1271-1274, July 2003.

[45] I. Perisa and J. Lindner, "Frequency Offset Estimation Based on Phase Offsets Between Sample Correlations," *Proc. of The 13th European Signal Processing Conference (EUSIPCO 2005)*, Antalya, Turkey, September 2005.

[46] Polastre, J., Hill, J., and Culler, D. 2004. "Versatile Low Power Media Access for Wireless Sensor Networks". *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*, Baltimore, MD, USA, November 03 - 05, 2004.

[47] Pollet, T.; Van Bladel, M.; Moeneclaey, M., "BER Sensitivity of OFDM Systems to Carrier Frequency Offset and Wiener phase noise," *IEEE Transactions on Communications*, vol.43, no.234, pp.191-193, Feb/Mar/Apr 1995.

[48] G. J. Pottie and W. J. Kaiser. "Embedding The Internet: Wireless Integrated Network Sensors," *Communications of the ACM*, Vol. 43, no. 5, pp. 51–58, May 2000.

[49] G. Pottie and W. Kaiser, *Principles of Embedded Network Systems Design*, Cambridge University Press, 2005.

[50] J. G. Proakis, *Digital Communications*, 4th Ed., McGraw-Hill, 2001.

[51] Reinhardt, V.S., "Spur Reduction Techniques in Direct Digital Synthesizers," *Proceedings of the 1993 International Frequency Control Symposium*, pp.230-241, 2-4 Jun 1993.

[52] Rohling, H.; Gruneid, R., "Performance Comparison of Different Multiple Access Schemes for the Downlink of an OFDM Communication System," *IEEE 47th Vehicular Technology Conference*, vol.3, no., pp.1365-1369, 4-7 May 1997.

[53] Romer, K.; Mattern, F., "The Design Space of Wireless Sensor Networks," *IEEE Wireless Communications*, vol.11, no.6, pp. 54-61, Dec. 2004.

[54] Sadagopan, N; Krishnamachari, B, "Maximizing Data Extraction in Energy-limited Sensor Networks," *INFOCOM, The Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societie*s, Hong Kong, 2004.

[55] B. M. Sadler, "Fundamentals of Energy-Constrained Sensor Network Systems," *IEEE Aerospace and Electronic Systems Magazine*, Vol. 20, Iss. 8, part 2, pp. 17-35, Aug. 2005.

[56] Sanyo ML414, specification sheet: http://www.sanyo.com/batteries/specs.cfm.

[57] Schmidl, T.M.; Cox, D.C., "Robust Frequency and Timing Synchronization for Ofdm," *IEEE Transactions on Communications*, vol.45, no.12, pp.1613-1621, Dec 1997.

[58] B. Sklar, *Digital Communications: Fundamentals and Applications*, 2nd Ed., Prentice-Hall, 2001.

[59] Sivanadyan, T.; Sayeed, A., "Active Wireless Sensing For Rapid Information Retrieval In Sensor Networks," *The Fifth International Conference on Information Processing in Sensor Networks*, April 2006.

[60] Sun Microsystems SunSPOT node, documentation available: http://www.sunspotworld.com/docs/

[61] J.M Tarascon and M. Armand, "Issues and Challenges Facing Rechargeable Lithium Batteries," *Nature*, **414**, 359-361, 2001.

[62] Thayne, I.; Elgaid, K.; Holland, M.; McLelland, H.; Moran, D.; Thoms, S.; Stanley, C., "50nm GaAs mHEMTs and MMICs for Ultra-Low Power Distributed Sensor Network Applications," *Proceedings, 2006 International Conference on Indium Phosphide and Related Materials*, pp. 181-184, 7-11 May 2006.

[63] Tierney, J.; Rader, C.; Gold, B., "A Digital Frequency Synthesizer," *IEEE Transactions on Audio and Electroacoustics*, vol.19, no.1, pp. 48-57, Mar 1971.

[64] Walker, J. D., "Karate Strikes," *American Journal of Physics*, October 1975, Volume 43, Issue 10, pp. 845-849.

[65] Vankka, J., "Methods Of Mapping From Phase To Sine Amplitude In Direct Digital Synthesis," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol.44, no.2, pp.526-534, March 1997.

[66] Wei, L.; Schlegel, C., "Synchronization Requirements For Multi-user OFDM on Satellite Mobile and Two-path Rayleigh Fading Channels," *IEEE Transactions on Communications*, vol.43, no.234, pp.887-895, Feb/Mar/Apr 1995.

[67] Weinstein, S.; Ebert, P., "Data Transmission by Frequency-Division Multiplexing Using the Discrete Fourier Transform," *IEEE Transactions on Communication*s, vol.19, no.5, pp. 628-634, Oct 1971.

[68] Werner-Allen, G; Lorincz, K; Ruiz, M; Marcillo, O; Johnson, J; Lees, J; Welsh, M, "Deploying A Wireless Sensor Network on an Active Volcano," IEEE Internet Computing. Vol. 10, no. 2, pp. 18-

25. Mar.-Apr. 2006.

[69] Wong, K.J.; Arvind, D.K., "Specknets: New Challenges for Wireless Communication Protocols," *Third International Conference on Information Technology and Applications*, 2005, vol.2, pp. 728-733, 4-7 July 2005.

[70] Wong, K. and Arvind, D. K. SpeckMAC: Low-power Decentralised MAC Protocols for Low Data Rate Transmissions in Specknets. *In Proceedings of the 2nd international Workshop on Multi-Hop Ad Hoc Networks: From Theory To Reality,* Florence, Italy, May 26, 2006.

[71] Wong, K. and Arvind, D. K. "Perspeckz-64: Physical Test-bed for Performance Evaluation of MAC and Networking Algorithms for Specknets". *Proceedings of the 2nd international Workshop on Multi-Hop Ad Hoc Networks: From Theory To Reality,* Florence, Italy, May 26, 2006.

[72] Wong, Kai-Juan; Arvind, D.K, "A Hybrid Wakeup Signalling Mechanism for Periodic-Listening MAC Algorithms," *15th IEEE International Conference on Networks*, 2007, pp.467-472, 19-21 Nov. 2007.

[73] Wei Ye; Heidemann, J.; Estrin, D., "An Energy-efficient MAC Protocol for Wireless Sensor Networks," *Proceedings of the Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies*, vol.3, pp. 1567-1576, 2002.

[74] A. Young, M. Ling, D.K. Arvind, "Orient-2: A Realtime Wireless Posture Tracking System Using Local Orientation Estimation," *Proceedings, 4th Workshop on Embedded Networked Sensors*, Cork, Ireland, 2007.

[75] A. Young, "Wireless Realtime Motion Tracking System using Localised Orientation Estimation," PhD thesis, University of Edinburgh, 2008.

WEBSITES

[76] Berkeley Wireless Research Centre, http://bwrc.eecs.berkeley.edu/

[77] Crossbow Technology Inc, http://www.xbow.com

[78] Dust Networks Inc, http://www.dustnetworks.com

[79] IEEE 802.11 Working Group, http://www.ieee802.org/11/

[80] IEEE 802.15 WPAN™ Task Group 4, http://www.ieee802.org/15/pub/TG4.html.

[81] Napier University, Centre for Emergent Computing, http://www.cec.soc.napier.ac.uk/

[82] PACWOMAN: Power Aware Communications for Wireless OptiMised personal Area Networks, http://www.imec.be/pacwoman/

[83] Speckled Computing, http://www.specknet.org/

[84] SunSPOT Project, Sun Microsystems, http://www.sunspotworld.com/

[85] Texas Instruments CC1100, http://www.ti.com/lprf

[86] The Mathworks MATLAB Curve Fitting Toolbox, http://www.mathworks.com/products/curvefitting/description2.html

[87] Third Generation Partnership Project (3GPP), http://www.3gpp.org

[88] TinyOS Alliance Community Forum, http://www.tinyos.net/

[89] University of Edinburgh, Institute for Computing Systems Architecture, http://www.icsa.informatics.ed.ac.uk/

[90] University of Glasgow, Ultrafast Systems Group, http://www.gla.ac.uk/departments/electronicsandelectricalengineering/research/micronanotechnology/ultrafastsystems/

[91] University of St Andrews, School of Chemistry, http://ch-www.st-andrews.ac.uk/

[92] University of St Andrews, School of Physics and Astronomy, http://www.st-andrews.ac.uk/physics/

[93] University of Strathclyde, DSP Enabled Communications Group, http://www.dspec.org/

[94] World Digital Media Broadcasting Online, http://www.worlddab.org/

[95] ZigBee Alliance, http://www.ZigBee.org