Wireless Sensor Networks and their Industrial Applications

Alex Mason

A thesis submitted in partial fulfilment of the requirements
of Liverpool John Moores University for the degree of
Doctor of Philosophy

Radio Frequency and Microwaves Group
General Engineering Research Institute (GERI)

June 2008

Abstract

Wireless Sensor Networks (WSN) represent a relatively modern concept which has captured the interest of many in the research community. Coupled with appropriate hardware, they offer great flexibility in terms of their applicability to solving real world problems. This can be seen with applications ranging from environmental issues to healthcare and even artificial intelligence.

Much of the work relating to WSN has been predominantly in the research domain, and so it is the purpose of this study to investigate ways in which they can be applied to solve industrial issues. This study particularly considers inventory management in the airline and packaged gas industries where there are many common fundamental requirements. A prototype system is presented which includes a database to record and obtain relevant tracking data in order to facilitate asset identification. Information of how this system may be applied within each industry is also included, in addition to how a WSN can be utilised to fulfil the specific needs of individual industries through the use of custom built hardware and sensors. Initial experimental results of this system are also given along with experimental results pertaining to the suitability of WSN devices in industry.

Despite WSN devices still being relatively new many advances have been made in order to make them more powerful and also smaller. However, as the size of the devices has decreased very little has been done with regards to critical components such as the antenna. As a result this work looks at the production of an industrially suitable antenna in terms of its design, construction and testing.

Finally, wireless sensing in the automotive industry is briefly discussed. The application of WSN in the automotive industry aims to improve recent spot weld monitoring techniques which determine the quality and integrity of a spot weld in real-time. Since such systems are wired, it is thought that WSN technology may finally make them feasible for retrofitting to existing spot welding machinery.

Acknowledgements

There are a great number of people to whom I would like to offer my thanks, beginning with my supervisors Professor A. Al-Shamma'a and Dr. A. Shaw who have been close at hand when I required the benefit of their knowledge and expertise, whilst also not stifling the project such that it could not grow.

In everyday work there are those in the RF & Microwave Group who have assisted with helping me to form ideas, or by providing technical assistance based around their particular area of expertise. In particular I would like to say a huge thanks to Dr. Jeff Cullen for his general help in many aspects of this project. Then there are Dr. Steve Wylie and Dr. Bob Stuart, both of whom have provided tremendous amounts of support in the area of microwave and RF theory which has been extremely beneficial in attempting to understand antennas and how to put them into practise.

In terms of academic assistance within Liverpool John Moores University, I would like thank Dr. D. Harvey for kindly letting me utilise his RF chamber for some of my measurements.

Of course, then there are my friends and family – without them really I would not have survived the past three years. Eduardo, thanks for being a great friend through what has been a trying time – also your help in the last stages of my work were much appreciated; I wish you good luck in the final stages of your studies. To my parents, thanks for keeping a roof over my head for my entire life which has lead to this point; I can't thank you enough for that, nor the huge amounts of support and tolerance you have shown. And my ickle sister Chantelle, just when you thought I had forgotten – thanks for making me smile when I needed to, and simply being the best little sister a big brother could have! Last but most definitely not least I would like to thank my partner, Laura-Jayne, for her eternal patience, love and understanding in this past 18 months; hopefully the near future will prove that it has all been worth it!!

Table of Contents

List of Figures

List of Tables

List of Abbreviations

AC	Alternating Current
ADC	Analogue to Digital Channel
AODV	Ad-hoc On Demand Distance Vector
API	Application Programming Interface
ASCII	American Standard Code for Information Interchange
CAD	Computer Aided Design
CAN	Controller Area Network
COTS	Cheap Off The Shelf
CPW	Co-Planar Waveguide
CRC	Cyclic Redundancy Check
DARPA	Defense Advanced Research Projects Agency
DC	Direct Current
DCV	Destination Coded Vehicle
DSDV	Destination Sequenced Distance Vector
EAN	European Article Numbering
EEPROM	Electrically Erasable Programmable Read Only Memory
EM	Electro-Magnetic
FCPW	Folded Co-Planar Waveguide
FEM	Finite Element Model
FSK	Frequency Shift Keying
GIS	Geographical Information System
GPRS	General Packet Radio Service
GPS	Global Positioning System
GUI	Graphical User Interface
HF	High Frequency
HFSS	High Frequency Structure Simulator
HHU	Hand Held Unit
HKIA	Hong Kong International Airport
HPBW	Half Power Beam Width
IC	Integrated Circuit

ID	IDentification
ISO	International Organisation of Standardisation
JIT	Just In Time
LCD	Liquid Crystal Display
LED	Light Emitting Diode
LF	Low Frequency
LOS	Line Of Sight
MAC	Medium Access Control
MAS-Net	Mobile Actuator Sensor Network
MEMS	Micro-Electro-Mechanical System
MIB	Mote Interface Board
MMCX	Micro-Miniature CoaXial
NEC	Numerical Electromagnetics Code
NESC	Network Embedded Systems C
OFCOM	Office of COMmunications
OS	Operating System
PCB	Printed Circuit Board
PDA	Personal Digital Assistant
PEC	Perfect Electrical Conductor
PIC	Programmable Interface Controller
PID	Product IDentification
PVC	Poly-Vinyl Chloride
RAM	Random Access Memory
RF	Radio Frequency
RFID	Radio Frequency IDentification
RISC	Reduced Instruction Set Computer
ROM	Read Only Memory
RSSI	Received Signal Strength Indicator
SMA	Sub-Miniature, version A
SoC	System on Chip
SWR	Standing Wave Ratio
TTL	Transistor-Transistor Logic
UART	Universal Asynchronous Receiver Transmitter

UHF	Ultra High Frequency
UK	United Kingdom
UPC	Universal Product Code
UPS	United Parcel Service
US	United States
USB	Universal Serial Bus
UWB	Ultra Wide Band
VNA	Vector Network Analyser
VSWR	Voltage Standing Wave Ratio
WINS	Wireless Integrated Network Sensor
WLAN	Wireless Local Area Network
WSN	Wireless Sensor Network
XOR	eXclusive OR

1. Introduction

1.1. The Initial Aim

This project was initially started with the sole aim of investigating a way in which to improve the effectiveness of luggage handling in airports. Luggage is essential to the vast majority of passengers, and so its effective handling is a primary concern of airlines, particularly since recent reports [1-2] suggest that lost luggage alone costs the airline industry in excess of $1 billion annually. This is a figure which is set to rise year upon year [3] if some action is not taken to rectify the way in which luggage is tracked.

The ability to quantify a cost to the airline industry is useful in bringing attention to the luggage issue, but it does not highlight other implications of poor luggage handling. Lost luggage cannot simply be defined as a cost to the airline, it must also be considered as an inconvenience to the passengers affected; arriving at a holiday or meeting destination without ones luggage could have devastating consequences.

Aside from lost luggage, the majority of commercial airlines have the ability to carry large numbers of passengers; a Boeing 747-400 [4] can carry more than 500, and the smaller Boeing 737 [5] nearly 200. Even if each passenger has only one item of luggage stored in the cargo hold of each aircraft, there is a great deal of luggage, little of which will be particularly distinctive. If there is a problem with a passenger (i.e. – they fail to board the aircraft, or they have to leave before take-off) then their luggage must be removed from the cargo hold. At present such an action involves a member of staff physically searching through the cargo hold looking for that passenger's luggage based upon its appearance. This is an extremely labour intensive chore, one which must be done as quickly as possible often resulting in damage to luggage in the process. Of course, speed is essential since an aircraft being held at an airport longer than necessary costs the airline money, whilst also delaying current passengers and potentially those who are to board subsequent flights. As with the inconvenience of lost luggage, these potential delays, indirect costs and further inconveniences are

difficult to quantify: they are however incredibly important, and add weight to the argument that airlines need more effective ways of handling and tracking their luggage.

Therefore, the initial aim of this project was to look at creating a system which could be suitable for use in airports which would eliminate, or at least minimise, lost luggage whilst also considering those factors which cannot be easily quantified such as the passenger experience of the luggage handling process and the ability to quickly find luggage in the aircraft cargo hold.

1.2. Further Goals

Once the investigation began into the idea of luggage tracking, it became apparent that the airline industry was not alone in suffering due to the implementation of ineffective tracking technology. The packaged gas industry has for many years relied solely on manual methods of asset tracking – some larger organisations have begun to implement similar technology to that of airlines but are still finding it to be extremely error prone and labour intensive. As a result, the packaged gas organisations have a problem with gas cylinders being removed from storage compounds without authorisation (i.e. – stolen). Furthermore, there is typically no record of when the theft occurs so there can be a large time gap before the gas distributor becomes aware of it.

The issues in both industries could be solved through the implementation of more advanced inventory management technology which allows real-time monitoring. Real time monitoring will allow greater visibility of assets (i.e. – luggage or gas cylinders) as they move within an organisation. This means that, for example, the act of an item going astray can be detected in real time which could allow corrective action.

A major similarity between assets in the airline and packaged gas industries is the environment in which they are kept. Conveyor systems which automatically move luggage around airport buildings are highly metallic, as well as the aircraft luggage hold. Gas cylinders themselves are made from metal, and are often stored in metal

cages. Since many recent automated inventory management technologies make use of radio waves for wire free identification of individual assets, metallic environments can cause significant problems – this will be shown at various points in this work. This is an issue which needs to be overcome to allow the effective use of automated identification in these environments and is the reason that this project is strongly focused on the concept and implementation of Wireless Sensor Networks (WSN) as it is hoped this technology can overcome communication problems brought about through difficult communication environments.

The packaged gas industry, whilst having some key needs in common with the airline industry, could benefit from a more effective and possibly feature rich asset management system than it currently has. One such area is the idea of automation; having the ability to automate asset identification to the point where personnel are no longer required in great numbers. In addition, gas cylinders could be equipped with sensors in order to provide information upon the condition of a cylinder. As an example, a gas cylinder could be equipped with a pressure sensor which could be used to determine if a cylinder is full or empty, or even to determine if it is leaking over time; this later feature could be more than merely useful since it could, in the case of harmful gases, avert a crisis situation.

A key part of any radio communication system is the antenna; its design determines how effectively a device can communicate wirelessly. However, there are many other factors to consider, such as the intended application. For example, mobile phones have very small antennas in order ensure that the device can be kept compact – the same idea applies with WSN devices when they are considered in an asset management environment, since the more compact the antenna, the less obtrusive the device is as a whole. Both the airline and packaged gas industries require a robustness, and so this project discusses the design, implementation and testing of a compact antenna suitable for communication in WSN; one which cannot easily be snagged or otherwise be damaged throughout the life of an asset.

Over time, as the project developed around the idea of utilising WSN, a further seemingly unrelated application was explored. The automotive industry relies heavily

on spot welding (sometimes also referred to as resistance welding) in order to build its cars. Indeed, it is estimated that a single car can have in excess of 5000 such spot welds! There has been work conducted over a number years looking at ways in which to improve the spot welding process in order to both make it more reliable [6] and also more energy efficient [7]. However, the systems developed for these purposes are difficult to deploy in industry since they are typically quite large and relay sensor data for analysis via wired means. For the highly mobile robotic spot welders used in industry, this is not sufficient and so this project also considers the implementation of WSN in the automotive industry for spot welding purposes. The aim here is not to develop sensors for spot welding, but to show that WSN could indeed be suited to the metallic environment, afforded by the spot welding machinery, and provide a tool for existing sensor systems to become useable on industrial spot welding machines.

1.3. Defined Project Targets

Whilst the text thus far has served to introduce this project it is important that there is a clear understanding of the project targets. Therefore, these targets are clearly defined below, in bullet form.

- Consider the existing technologies for inventory management and asset identification, in addition to the current situation within the airline and packaged gas industry.
- Investigate the use of WSN, their key features and their suitability to the proposed industrial problems.
- Implement an inventory management system using WSN, which makes use of the advantages afforded by the technology.
- Consider industrial concerns regarding antenna design and obtain knowledge pertaining to the fundamentals of antenna theory, design and simulation.
- Create an industrially suitable antenna which can be implemented alongside WSN hardware.

- Consider the performance of the WSN hardware from an industrial point of view, as well as demonstrating the use of the aforementioned inventory management system.
- Design and demonstrate a proof of principle system which utilises WSN for monitoring the spot welding process.

1.4. Chapter Overview

Chapter 2 is concerned with the topics of inventory management and asset identification methods. Much effort is concentrated here on discussing the problems faced during inventory management, as well as introducing the technologies by which it may be currently achieved. The situation within the airline and packaged gas industries is also discussed with regards to how assets are currently managed and what potential issues this brings about.

Chapter 3 takes an in depth look at the WSN concept and seeks to outline its advantages as well as possible hurdles which may be considered as part of this work. The technology used to implement WSN is also discussed, and it concludes with a comparison of existing technologies (as discussed in Chapter 2) and WSN devices.

After much discussion on the topic of inventory management and WSN, Chapter 4 is focused on the implementation of a prototype system. This system, which is made up of a number of components, is presented in detail in order to attempt to demonstrate the features of the WSN concept which have been utilised.

Chapter 5 provides an insight into antenna theory, beginning with the reasoning behind the creation of a new antenna. Much of this information provides the basis for further work which is presented in Chapter 6, which explains the design steps taken in creating an industrially suitable antenna.

Chapter 7 attempts to consolidate the work, providing various test results and discussing the feasibility of the work in industry. In addition, the final part of this

chapter provides information about the topic of implementing WSN in the spot welding industry and demonstrates the proof of principle system created within the scope of this work.

Chapter 8 seeks to conclude the work by describing what it has achieved, as well as seeking to convey how this work could feasibly fit into real world applications. In addition, it looks at future topics which could be useful for future exploration in this area of research.

2. Asset Tracking

2.1. What is Asset Tracking?

An asset is defined by the Oxford English Dictionary [8] as a "useful or valuable person or thing". In the context of this work, asset tracking refers to the knowledge an organisation possesses with regards to the location of their property as a result of its movement. One should not confuse an asset with the term inventory, since although the two are commonly interchangeable; inventory is formally defined as "a list of items held in stock" [9]. Therefore in the airline industry an asset may be an item of luggage whilst the inventory would be a list of all of the items of luggage currently in an airline's possession. Similarly in the packaged gas industry a gas cylinder is an asset, whilst a listing of all of the gas cylinders would be the inventory.

Such a distinction is important, since asset tracking is not a standalone procedure but more of a component of a larger inventory management process. Inventory management is crucial to organisations for a number of reasons, and is ultimately the catalyst for this project. An inventory management system allows an organisation to know what assets it has available. However, this knowledge alone is often not particularly effective because there is the possibility of assets being stolen, misplaced or lost, all of which add cost – a feature which was alluded to for both the airline and packaged gas industries in Chapter 1. Therefore the addition of extra tools, such as asset tracking, into the inventory management process allows organisations to reduce such costs through maintaining greater asset visibility. Furthermore, concepts such as just-in-time (JIT) operation [10] can be implemented more effectively, thus allowing organisations, where applicable, to reduce their total inventory.

2.2. The Need for Automation in Asset Tracking by Example

There is the need for automation in asset tracking for a number of reasons. Rather than blandly detail these reasons, a simple example is explored of a warehouse which stores boxes. The contents of the boxes are unimportant – it is sufficient to say that

the contents of the boxes are manufactured at some location, packaged, and then delivered to the warehouse. When they are required by a customer, they are dispatched from the warehouse to the customer. This simple example is illustrated in Figure 2.1.

2.2.1. No Asset Tracking

In this situation, no inventory management system is employed, and hence there is no asset tracking. Therefore each time a box requires dispatching, it will have to be searched for by a member of the warehouse staff to see if it is actually available. If the box is not available, a back order will have to be placed further up the supply chain. Of course, this causes a further issue, because if there is no inventory management how does the warehouse know when, or even if, it has received this back order?

This simple scenario highlights the need for inventory management in two ways. Firstly, the process of searching for items as they are required is particularly inefficient, and wastes man hours which may not otherwise be required, or could be applied more effectively to other tasks. Secondly, the process is inefficient from a customer's point of view since there is no control in place to expedite the dispatch of back orders. In effect such an order could be delivered to the warehouse and then lay untouched for some period of time before it was queried once again by the customer.

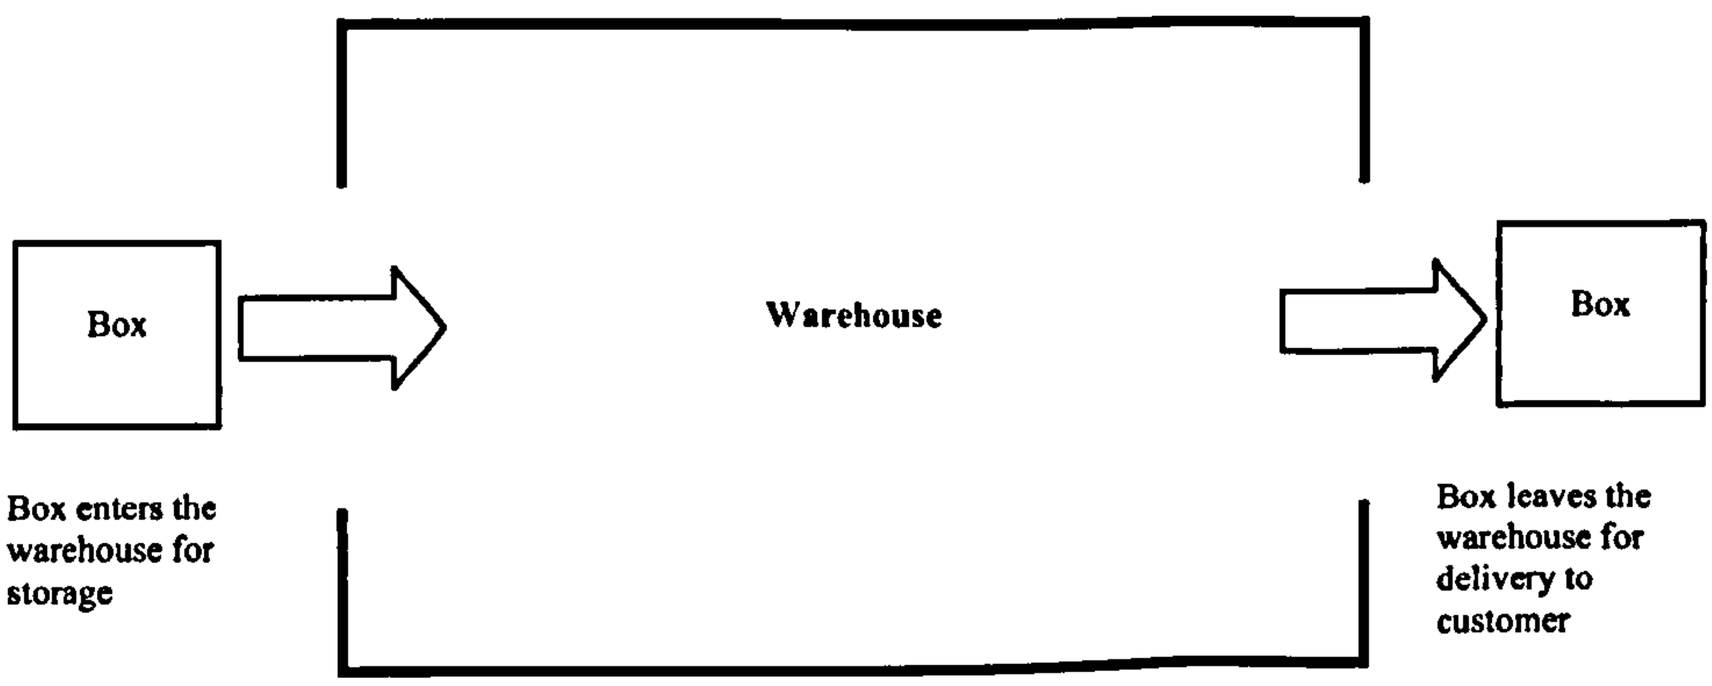

Figure 2.1 – Simple overview of a box warehouse

2.2.2. Manual Asset Tracking

Now that reasoning has been established for the need of an inventory management system, imagine that the same warehouse has installed a system which allows personnel to manually record all incoming and outgoing boxes. All of these records are then computerised in an automated way to keep the scenario relatively simple. As a result of this change, there is now an inventory available at all times. In addition, the boxes can be tracked at two points – entering and leaving the warehouse – so there is a crude asset tracking system in place. This means that the availability of boxes is always known with some degree of certainty, without having to search the warehouse each time a request is received. In addition, back ordered items can be recorded as coming into the warehouse; having the computer system alert personnel to the back order being available is then a trivial matter.

However, there are still a number of issues as there is a high reliance on personnel to reliably record each box that enters and leaves the warehouse. In addition, once a box enters the warehouse, there is little way of knowing where it is stored, particularly if the warehouse is very large. This leads to two potential issues; human error in recording the boxes movement, and misplacement of boxes. One could tackle the later problem by splitting the warehouse into sections and tracking the movement of boxes within these sections but, since all recording is done manually, this would be expensive in terms of labour as it would consume significant amounts of time. In addition, there is no mechanism in place to deter theft since boxes could be removed from the warehouse without a record being made which would lead to an inaccurate inventory. There would be no knowledge of this theft until there was an inventory check, which may take place only a few times a year.

2.2.3. Automated Asset Tracking

As a final iteration of this scenario, consider the case of the warehouse upgrading now to an automated system for recording the entry and departure of boxes from the warehouse. This too is linked to a computerised system which stores the inventory. In this case, boxes can be delivered and dispatched without the need for human

recording of these events. Time is saved since the automated system is much quicker than a human carrying out the same task by hand on a box per box basis – the automated system allows for however many boxes being delivered to be scanned automatically at once. In addition, there is no reason why this automation could not be extended in order to achieve higher resolution visibility of box movement inside the warehouse in order to safeguard against boxes being misplaced. Theft could be deterred too, since the system would know when an item left the warehouse. Analysis of the computerised inventory system in association with closed circuit television camera recordings, for example, could lead to efficient recovery of stolen goods.

It is important to note that this scenario is limited in its scope, but serves to give the reader a taste for the rationale behind asset tracking, and why it could be important. Of course, large organisations do not generally consist of a single warehouse. In their supply chain they may have factories where different components are manufactured, warehouses for storing components and manufactured goods, as well as outlets to allow consumers to purchase their goods. Therefore, asset tracking as part of an inventory management system allows visibility of any asset within their supply chain. Automated asset tracking can enhance the resolution of this visibility, as well as allowing the inventory to be updated efficiently and reliably in real-time. Further benefits can also be gained, as highlighted by the example, in terms of a reduction in the misplacement and theft of assets.

With regards to both the airline and packaged gas industries, all of these outlined benefits associated with implementing automated asset tracking are relevant. Further advantages may be derived from the system in terms of JIT operations for the packaged gas industry, but the key aim at this stage is to focus on those benefits which could be utilised generically by both industries. Before looking at the current situation in these industries however, is it important to have knowledge of those technologies which are presently available to support the implementation of asset tracking.

2.3. Asset Tracking Technologies

2.3.1. Bar Codes

A bar code is a machine readable representation of information [11] which utilises an optical bar code reader in order to be read. Bar codes are used extensively in asset tracking as well as in point of sale systems in shops and supermarkets; nearly all packaged products use bar codes which can be utilised throughout a supply chain. The idea was first patented in 1952 but did not achieve commercial success until the 1980's; in 1970 it was estimated that just 1% of retail stores had bar code scanners, whilst by 1981 this had grown to 10% [12]. Since then, bar codes have become by far the most common technology used in inventory management and asset tracking, and it is quite rare to find an organisation which does not utilise bar codes in some form.

Two of the most widespread bar code symbologies [11,13-14] are Universal Product Code (UPC) and European Article Numbering (EAN). Each symbology has different versions, each of which use a different number of digits for identification purposes. Both symbologies are compatible since there has been significant cooperation between parties involved in the standardisation of both UPC and EAN – these parties have now merged to become known as GS1[1] [15]. This means that EAN bar code scanners can also read the UPC symbology, and after a mandate in 2005 [11], the opposite should also be largely true.

An example of a bar code is shown in Figure 2.2. Bar codes use the width and gaps of printed parallel lines to encode data. When light from an optical bar code reader shines onto a bar code it is reflected back to the reader in differing amounts due to the white spaces and black bars. A photo detector inside the reader then outputs an analogue signal which represents the bar code pattern. A signal processor converts this signal to a digital form; if the output voltage of the photo detector is above a decision threshold then a white gap is detected, if it is below then a black bar is

[1]It is important to note that GS1 is not an acronym but stands for "One global standard, one global system and one global organisation."

detected. As a result, the digital pulse widths correspond to the scanned bars and gaps, which can be decoded using the appropriate symbology.

Figure 2.2 depicts a UPC Version A type bar code which comprises ten identification digits, and two overhead digits. There are longer guard bars in both the left, right and centre positions of the label, acting as dividers, which allow for it to be read in any orientation. Either side of the bar code there are quiet zones, where no printing is allowed so as not to confuse the bar code reader.

As can be seen from Figure 2.2, the bar code number is made up of four individual numbers, each with their own meaning. The number system character denotes the type of product, with zero being given to most nationally branded items. Exceptions include random weight items, such as meat and cheese as well as drug and health related items. The manufacturer code is a five digit number assigned to a manufacturer by GS1-US[2] who oversees the UPC standard. The item code is a five digit code assigned by the manufacturer in order to identify its product. The final digit is the modulo-10 [13] checksum, which enables the bar code scanner to verify the accuracy of its read.

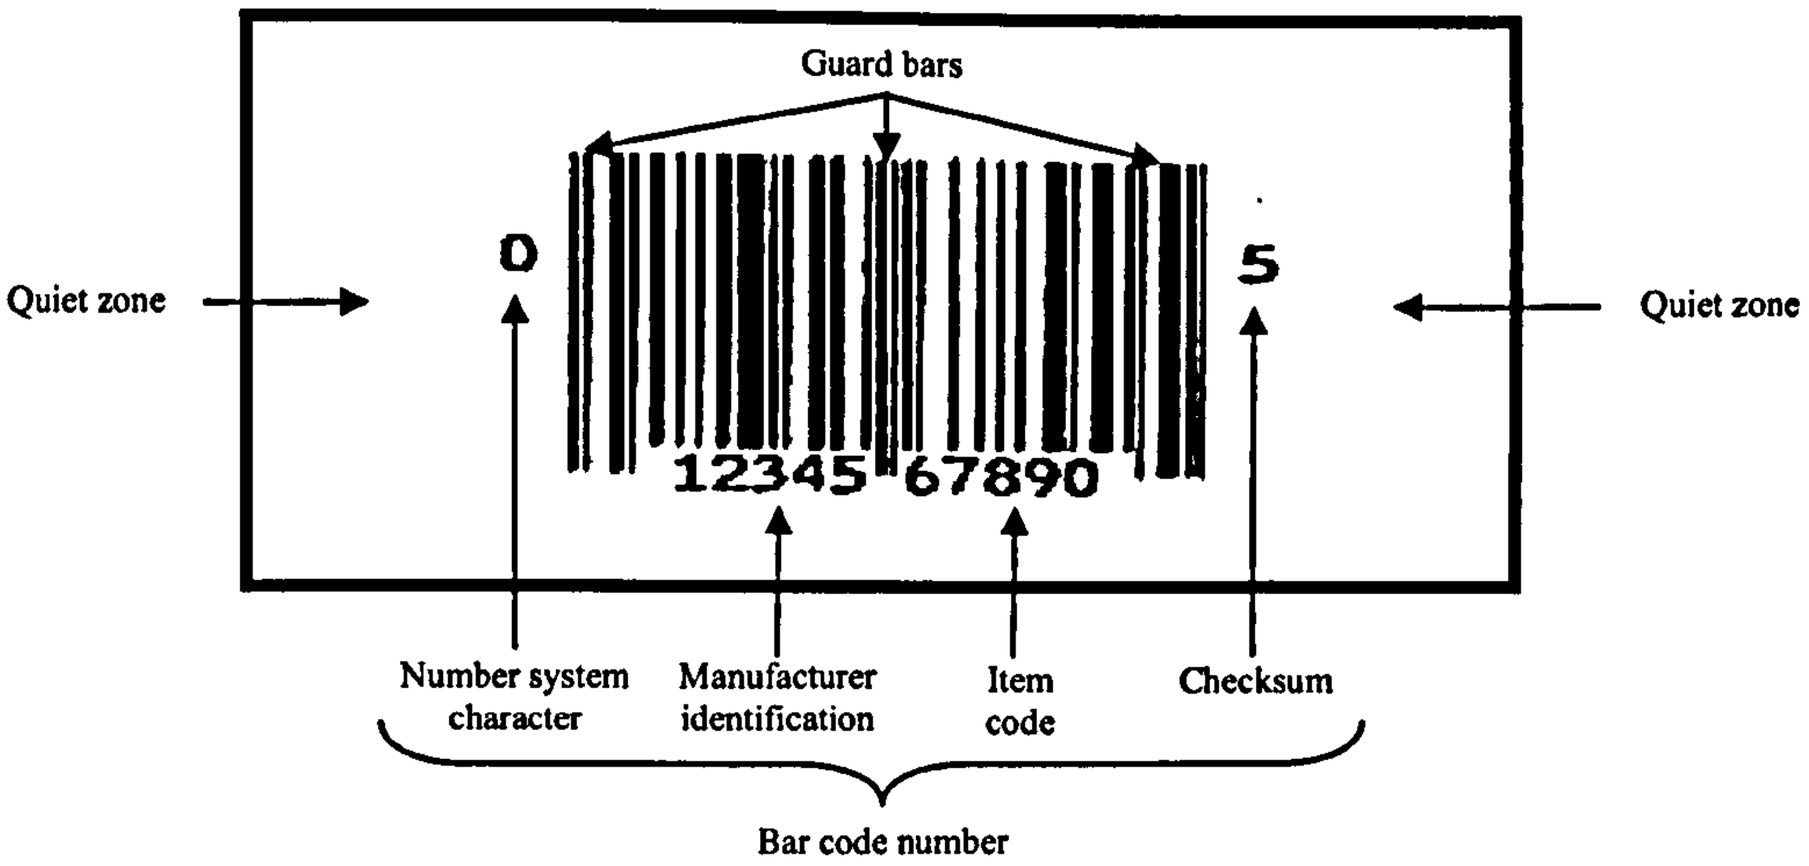

Figure 2.2 – UPC Version A

[2] The United States arm of GS1.

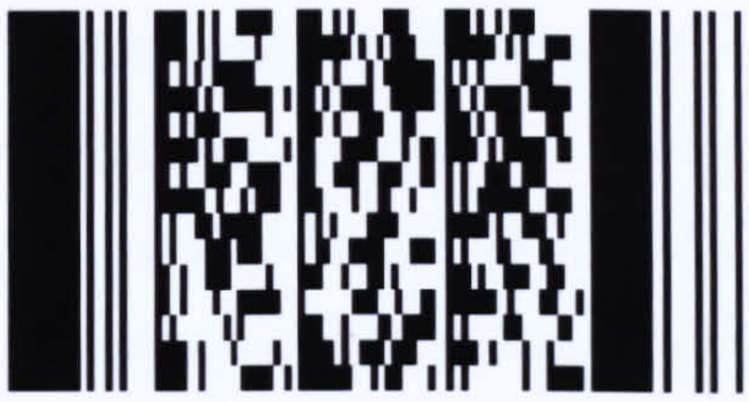

Figure 2.3 – A 2-D bar code

Figure 2.4 – Industrial (left) and handheld (right) bar code readers

The UPC bar code shown in Figure 2.2 is an example of a single dimension (1D) bar code, since data is encoded across it horizontally only. However, this is not the only way in which to encode data onto a bar code, since both the vertical and horizontal dimensions can be utilised on a 2D bar code, an example of which is shown in Figure 2.3. It is unlikely that the 2D bar code will completely replace its 1D counterpart however, as 1D bar codes spread data over the whole height of the bar code leading to a high degree of redundancy. This means the bar code can be read even with considerable degradation, so if an application requires only a few characters then a 1D bar code is probably the best solution.

In terms of read range, this varies depending upon the type of scanner being used. Handheld readers have a range of up to approximately 50cm, whilst larger industrial units claim read ranges of up to 3m. Examples of such units are shown in Figure 2.4, with the industrial unit being a Microscan MS-890 [16] and the handheld unit being a Microscan Q-Basic [17].

Whilst read range is one of the issues with bar codes when trying to completely automate asset tracking, there are also a number of other issues to contemplate. To begin with, the technology requires line-of-sight (LOS) in order to operate, since it relies upon visible light. Therefore if an asset's bar code is not aligned with the reader then it is unlikely to be identified or tracked successfully. In addition, bar codes are easily damaged because they are typically printed on paper labels. Whilst 1D bar codes offer a high degree of redundancy, a completely torn or scuffed bar code is unlikely to be readable.

Both of these issues lead to the need for human intervention, which is why many distributors offer handheld readers such as that shown in Figure 2.4. These devices still require LOS, and so may require the human operator to move the asset in question in order to identify it. Despite this, there are examples of bar codes being used for automated asset tracking in electronics [18] and automotive [19] manufacturing, as well as in asset shipments [20]. However, many such systems tend to be semi-automated since they often require human oversight and participation, much in the same way as the example given in Section 2.2.2.

2.3.2. Radio Frequency Identification

In a similar way to bar codes, Radio Frequency Identification (RFID) technology is made up of a tag and reader device. RFID actually has roots which lay as far back as World War II [21], where the Germans, Japanese, Americans and British were all making use of radar as an early warning system for incoming enemy aircraft. A discovery made by the Germans was that if pilots rolled their planes as they were returning to base the radio signal reflected back from them changed. This became the way in which German aircraft identified themselves to their airbases, and was essentially the first example of a passive RFID system.

RFID, in the form used today, was first patented in 1973 and there have been many applications developed with this technology as the focal point, such as animal tagging [22] and car security [23]. In more recent times the improvement of technology and the development of standards has helped to popularise it. This has been the case more

so since retail giant Wal-Mart demanded [24] their top suppliers use RFID in order to assist their attempts of achieving automated asset tracking. Whilst it was predicted that such a mandate alone would create the need for around 1 billion RFID tags each year [25], RFID does not have anywhere near the same level of market penetration as bar codes – reasons for this will become clear as the technology is discussed further.

When an RFID tag detects a radio frequency signal from a reader it wakes up and transmits a reply. If a tag uses the energy from the reader as its only power source, then it can be considered as passive. Some tags have a power source but still utilise the energy from the reader for transmission and are regarded as semi-passive, as they only use the power source to drive an integrated circuit (IC). Those tags which utilise a power source for transmission are called active tags. Each different type of tag comes in a variety of shapes and sizes to suit a wide variety of applications, a selection of which are shown in Figure 2.5. Irrespective of shape, RFID tags have a number of features in common; an IC on which to store data (e.g. – an identification number similar to that found as part of a bar code), a coupling device or antenna, and some sort of housing which may vary from a glass tube to a thin sheet of paper or plastic (see Figure 2.6).

Figure 2.5 – A typical RFID reader (left) and a selection of RFID tags (right)

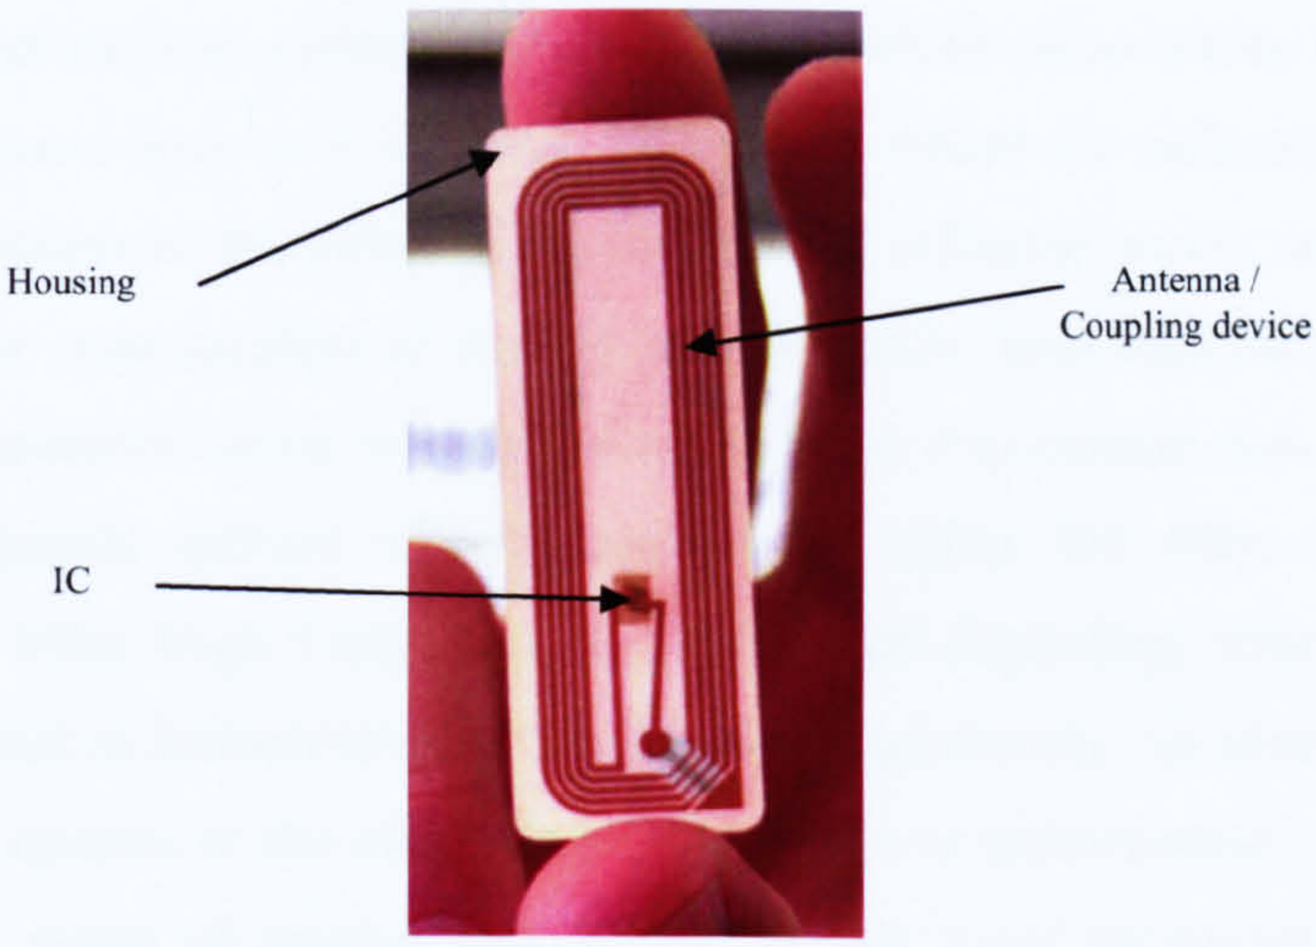

Figure 2.6 – The basic structure of an RFID tag

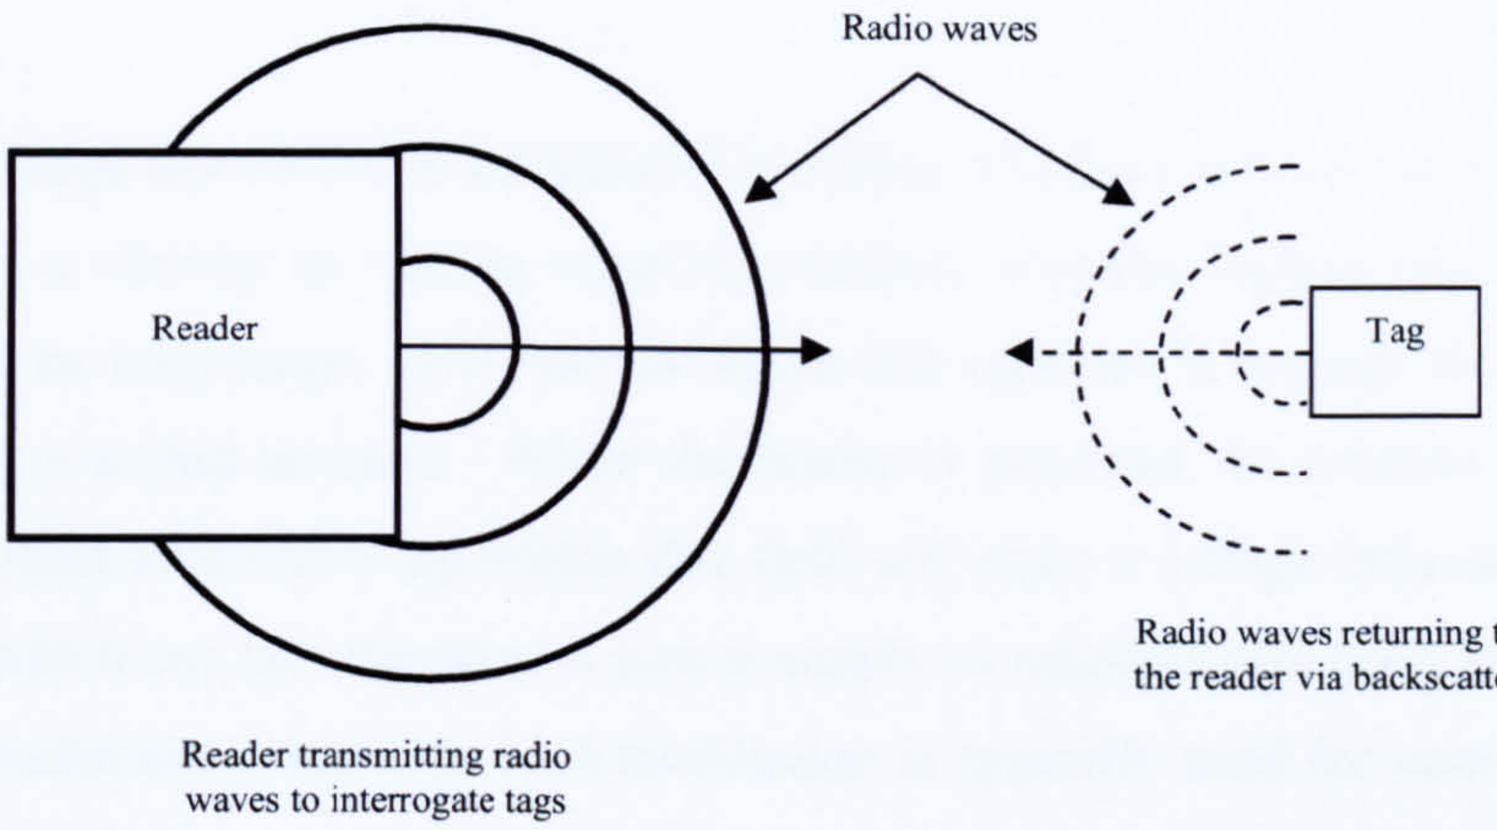

Figure 2.7 – Backscatter coupling

In order for a tag to communicate with a reader, the two must be coupled. The coupling mechanism determines the way in which the circuits on the tag and reader influence each other in order to send and receive information as well as power. There are a number of coupling mechanisms which can be used – backscatter [26] and inductive [27] coupling are discussed here as they are supposed to offer read ranges which can be considered long-range (significantly more than 1m) and remote (1cm to 1m) respectively. Anything which has a read range less than 1cm can be considered as using a vicinity coupling method, where the reader and tag have to be almost touching for communication to occur. Vicinity coupling is not discussed futher in this work since its range is insufficient for asset tracking in the context of this project.

Backscatter coupling is so called due to the way in which radio waves from a reader device are scattered back to it by a tag, that is, the waves are reflected back to the source. This process is illustrated in Figure 2.7. The reflected waves are of the same frequency as the ones incident to the tag and, typically, load modulation is used to alter some characteristic of the returning wave to allow it to contain data. This type of coupling is typically utilised with frequencies of either 868 MHz or 915 MHz (referred to as Ultra High Frequency, or UHF, tags) depending upon whether the technology is used in Europe or the United States, respectively. In addition, there are tags which can operate at the microwave frequencies of approximately 2.5 GHz and 5.8 GHz. The range of passive tags which operate using backscatter coupling is approximately 3m, although with battery assistance ranges of up to 15m can be achieved [28].

A shorter range alternative to backscatter coupling is inductive coupling which can be defined as a vicinity to remote coupling method, whereas backscatter is generally deemed to be long range. The way in which this operates is through the reader and tag having a coiled antenna. When the reader is powered, its antenna generates a magnetic field such that a tag within this field will have a voltage induced in its own antenna, which can be utilised as a power supply to transmit data back to the reader. As with backscatter coupling, load modulation is typically used for communication, although the frequency of the reader and tag are not necessarily the same. Typical frequencies used for inductive coupling are 125 kHz (Low Frequency or LF) and 13.56 MHz (High Frequency or HF) [28].

Unlike bar codes, which have well defined and recognised standards, standardisation is an area of RFID which can lead to a great deal of confusion – so much so that it has slowed the uptake of the technology somewhat [29-30]. There are a handful of different standards which define RFID tags for different purposes. For example, ISO 14223 [31] defines the air interface[3] used in tags for animals, ISO 15693 [32] defines tags for use in vicinity cards (e.g. – identification cards), whilst ISO 18000 [33]

[3] The description of the way in which a tag communicates with a reader, including its power source, coupling method and operating frequency amongst other characteristics.

defines the air interface for a variety of tags operating over a range of different frequencies.

EPCGlobal Inc. [34], a collaboration between GS1 and industrial partners (such as Sony, Proctor & Gamble and Hewlett Packard), have sought to define a method of classifying tags, as shown in Table 2.1. These classifications however do not all have a set of standards associated with them, since the latest set of standards [35] cover only Class 0 and Class 1 UHF tags. The reason for this is that the classes are meant as a route forward; as new technology becomes available new classes of tag can be utilised whilst still providing some level of compatibility with previous ones.

Looking positively at RFID, it does hold the potential to overcome a major shortcoming of bar codes in that it does not require LOS. Radio waves can penetrate materials – some better than others – and so it is possible to see how this could allow RFID to become part of an automated asset tracking system. Furthermore, in harsh environments RFID tags can be encapsulated in protective material or embedded within an asset to prevent them becoming damaged in a way which would make a bar code unreadable. Further potential is realised in the fact that the IC present on RFID tags can be utilised as a storage space, so not only can RFID tags store more data than bar codes, they can also be rewritten or updated. Effectively, if one thinks back to the scenario outlined in Section 2.2.3, RFID has the potential to realise automated asset tracking. Many boxes could be taken through the entrance and exit points simultaneously without LOS being required. There is no need for human interaction or knowledge of the process.

Inevitably however, RFID is not without its faults. One of the biggest of those is with regards to a lack of universal standards. Not only has this lead to the slowing of RFID uptake, it has also lead to systems being deployed which are not interoperable, whether it is because they use different frequencies or differing air interfaces.

Class	Description
Class 0	Passive, read only
Class 1	Passive, write once
Class 2	Passive, write once with encryption
Class 3	Semi-passive, rewritable, integrated sensors
Class 4	Active, rewritable, integrated sensors, ad hoc networking

Table 2.1 – EPCGlobal Inc. classifications for RFID tags

Standards are not the only issue however as there has been concerns about privacy [36-37] with RFID. One worry is that RFID could be used to track not only assets but also consumers who have purchased these assets. Take for example clothing chains who wish to use RFID tags on clothes [38] – if a consumer buys an item of clothing on a payment card, their payment details could be linked to the RFID tag on the clothes. As a result, the consumer could, in theory, have their movement and spending habits tracked resulting in an invasion of their privacy. Parts of the EPCGlobal Inc. standards deal with this issue by requiring RFID tags which conform to their standards to have a remote destruction feature which can be activated with an appropriate password. However, one might wonder what the incentive is for organisations to deactivate tags via this method, and with so many standards, who is to say that all organisations will choose to use those put forth by EPCGlobal? Some also identify security as an issue with RFID, although this is not something which tends to apply to asset tracking since the tags store only an identification number. Application of RFID in biometric passports has come under fire however, since the tags are intended to carry data specific to the passport owner which could be read by a third party [39].

RFID, as previously explained, makes use of radio waves which is a type of Electromagnetic (EM) radiation. It is therefore constrained by characteristics of this radiation which will be discussed in detail in Chapter 5. At this time however, one is aware of the fact that metallic environments are said to pose a significant problem for radio wave communication as was mentioned in Chapter 1. This fact applies to RFID tags, which suffer significantly from signal attenuation due to metallic materials.

Table 2.2 [40] shows that commonly used UHF tags suffer from poor range when placed on a metal surface. This means that whilst RFID tags may have a place in some industries, there are significant challenges to overcome for it to be applicable universally. It could be concluded that normal passive RFID tags would not be well suited to the highly metallic environments found in the airline and packaged gas industries.

In an attempt to overcome these issues there has been research in recent years to create metal mountable tags; the range of such a tag is also shown in Table 2.2. Whilst this tag offers a considerable performance boost in terms of read range in almost all environments it is still limited, under the best conditions, to approximately 3m. In addition, these tags cost £3.80 [41], are rigid in their structure, and are 90mm in length. This makes mounting them on smaller assets (e.g. – luggage) difficult. By comparison, a normal tag may cost only £0.03 [42] if bought in large quantities.

	Plastic	**Metal (direct contact)**	**Metal (0.06" stand off)**	**Glass**	**Card-board**	**Plywood**	**Air**
Normal plastic tag	1.37	0.15	0.15	2.44	1.52	1.83	1.07
Encapsulated metal mount tag	3.81	3.05	2.44	3.96	3.96	3.66	3.96

Table 2.2 – UHF RFID tag read ranges in metres

Whilst it was hoped that RFID may be the successor to bar codes one has to wonder if this will actually ever happen – some have been so bold as predict that it will not [43]. Although it offers some benefit in comparison to bar codes, it costs significantly more to implement and it is hard to establish in some cases if the benefits justify these costs. Many have also realised that the implementation is actually quite challenging, as has been seen with Wal-Mart's suppliers who have been slow to react to its demands [44]. Whilst Wal-Mart is still going ahead with its RFID implementation, albeit much slower than first anticipated, large companies such as United Parcel Service (UPS) have ruled out the use of RFID altogether for the moment for individual parcel tracking [45]. This, coupled with the fact that there are alternative niche technologies (discussed in the next section) offering specific benefits to some organisations, means that the universal use of RFID as a bar code replacement is probably quite a long way off.

2.3.3. Other Tracking Technologies

Whilst bar codes, and to a lesser extent, RFID are the most widespread technologies available for inventory management and asset tracking, there are a few which offer promise to selected applications.

The first type of technology to be discussed here is that which utilises existing wireless local area network (WLAN) infrastructures[4] to track assets as they move within the boundaries of an organisation complex (e.g – an office building or warehouse). Possibly the earliest example of this was Microsoft's RADAR [46], which demonstrated the ability to track assets using the strength of the signal between a computer WLAN card, and a base station. Through correlation of data from three base stations, this work claims to be able to calculate the position of a wireless device to within an average of 2-3 metres. The scope of this work was to add value to networks within organisations in order to facilitate activities such as printing to the nearest printer and building navigation. Commercial products based around this idea

[4] Therefore operating at approximately 2.4GHz.

[47-49] have come about which allow more general purpose tracking with a company named Aeroscout [47] potentially being the largest supplier.

Aeroscout have taken the step of developing tags which can be attached to assets such as large shipping containers and lorry trailers [50] – the tags are quite large however (see Figure 2.8) measuring 62×40×17mm and cost approximately £50 each. This size is most likely attributed to the battery which is claimed to be able to last up to 3 years at duty cycles less than 1%. It is possible to add further accessories to tracking systems which make use of Aeroscout tags such as base stations which allow for more accurate positioning through measuring the time of flight of a signal. Obviously this adds extra cost and moves away from the idea of utilising existing infrastructure. Aeroscout claim that the tag is accurate to approximately 8-10m using signal strength, and 3-5m using time based measurements.

Moving away from WLAN technologies, there have been a couple of systems [51-52] which have demonstrated the use of Ultra Wideband (UWB) communications for tracking and positioning, with the Ubisense Location System [51] claiming accuracy better than 30cm. The system uses four base stations which enclose the area of interest for tracking with each base station having four receivers which allow for the measurement of both the angle and time of flight of a signal. The tags themselves are of a similar size to those from Aeroscout, but offer only a one year lifespan.

Figure 2.8 – Aeroscout Tag

Traditional radio transmissions require high frequencies in order to be able transmit large amounts of data because their bandwidth is relatively narrow – typically in the range of kHz to MHz. However, UWB transmissions use large amounts of bandwidth to send data (500MHz, or 25% of the channel centre frequency [53]). This is thought to be the reason for the enhanced accuracy of UWB positioning systems, since the use of many different frequencies is thought to overcome the problem of multipath-fading; a characteristic of EM wave propagation discussed in detail in Chapter 5. However UWB is not considered further in this work since its use in the UK has only recently been allowed by the Office of Communications (OFCOM) [54].

The final technology to be discussed here is one named RuBee [55], which surfaced within the last year. The RuBee tag (see Figure 2.9) operates at 131 kHz, which is similar to the frequency utilised by LF RFID tags, and has a data rate of 1200bps (bits per second). This is incredibly small compared to the previous technologies discussed here which operate at typical WLAN data rates of up to 54Mbps. As a result of its low data rate, the system is limited to reading only 6 tags per second which appears to drop dramatically in metallic environments as a report [56] by the producers of the tag suggests that it takes over 120s to discover just 13 tags. RuBee has been demonstrated with relatively low volume tracking applications such as livestock tracking [57], and it is unlikely that it would be well suited to high volume applications such as tracking luggage or gas cylinders.

Figure 2.9 – RuBee tracking tags

Figure 2.10 – RuBee and a forklift truck

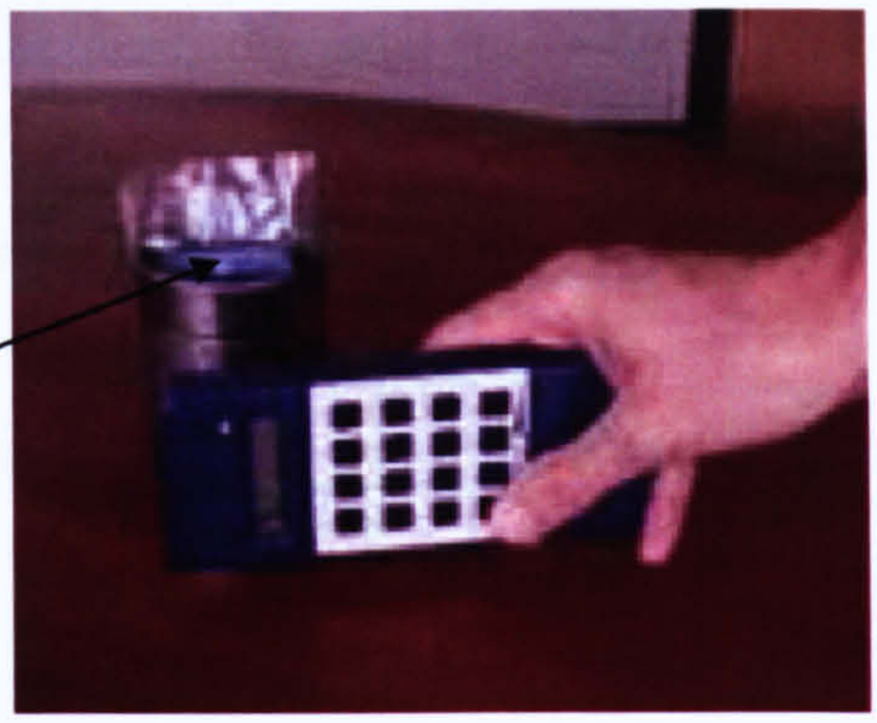

Figure 2.11 – RuBee in metal and water

Much work has gone into making this technology as rugged as possible, as can be seen in Figures 2.10 and 2.11. The former shows the tag can survive impact by a 10-ton fork lift truck, whilst the later appears to show that it can transmit whilst mostly surrounded by metal, a feature which is claimed to be a result of the low frequency communications. The tags themselves are estimated to cost around £5[5] each and, for the larger tags, have a battery life of 5 years.

As can be seen, there are a number of technologies available to assist with implementing asset tracking as part of an inventory management system. The next step is to look at the current situation for both the airline and packaged gas industries in order to ascertain the requirements for asset tracking in those scenarios.

2.4. Luggage Handling in the Airline Industry

2.4.1. The Luggage Handling Process

The process of moving luggage around an airport is complex, and often goes largely unnoticed when one is a passenger. However, a luggage handling system has to perform three vital functions for a passenger's luggage:

[5] Pete Abell, Visibility Assets, Business Development, 195 Bunker Hill Ave., Stratham, NH 03885.

- Transport it from the check in area to the departure gate, ready for loading onto an aircraft.
- Move it from the arrival gate to the carousel, or claim area, for collection by a passenger.
- In the event of a transfer, move it from the arrival gate to the subsequent departure gate.

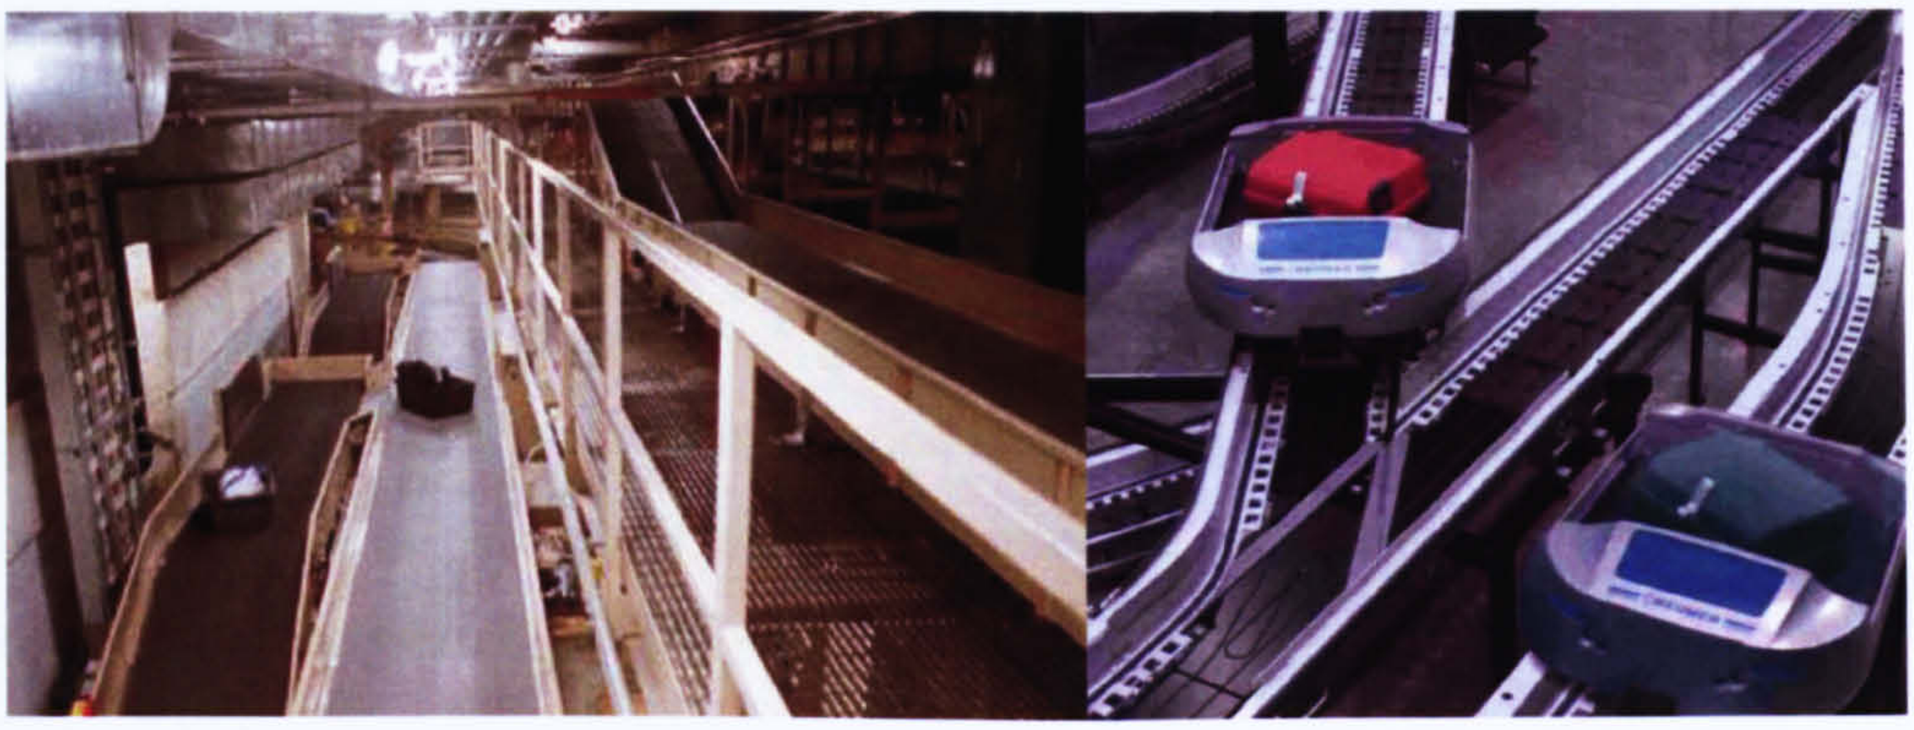

Figure 2.12– Conveyor (left) and DCV (right) networks for transporting baggage

Figure 2.13 – Typical luggage loading and unloading setup

Transportation of luggage around the airport is carried out through a series of conveyor and DCV (Destination Coded Vehicle) networks, shown in Figure 2.12. These serve to take luggage to its intended destination automatically, be it a departure gate or the luggage carousel for claiming by the owner. Once luggage reaches a departure gate, it must be transferred to the aircraft cargo hold onboard a luggage truck. It will then be loaded onto the aircraft manually by luggage handlers using a conveyor. Figure 2.13 shows a typical setup for loading and unloading an aircraft of luggage. Before takeoff, if a passenger is not aboard the aircraft their luggage must be removed and this will be done manually by a luggage handler searching visually for the item based upon its appearance and identification tags.

Once the aircraft has landed at the destination airport, luggage will be manually unloaded and taken to the arrival gate were it once again enters the network of conveyors and DCV's, often being destined for a luggage carousel where a passenger may collect it. The passage of a typical item of luggage through a luggage handling system is illustrated in Figure 2.14.

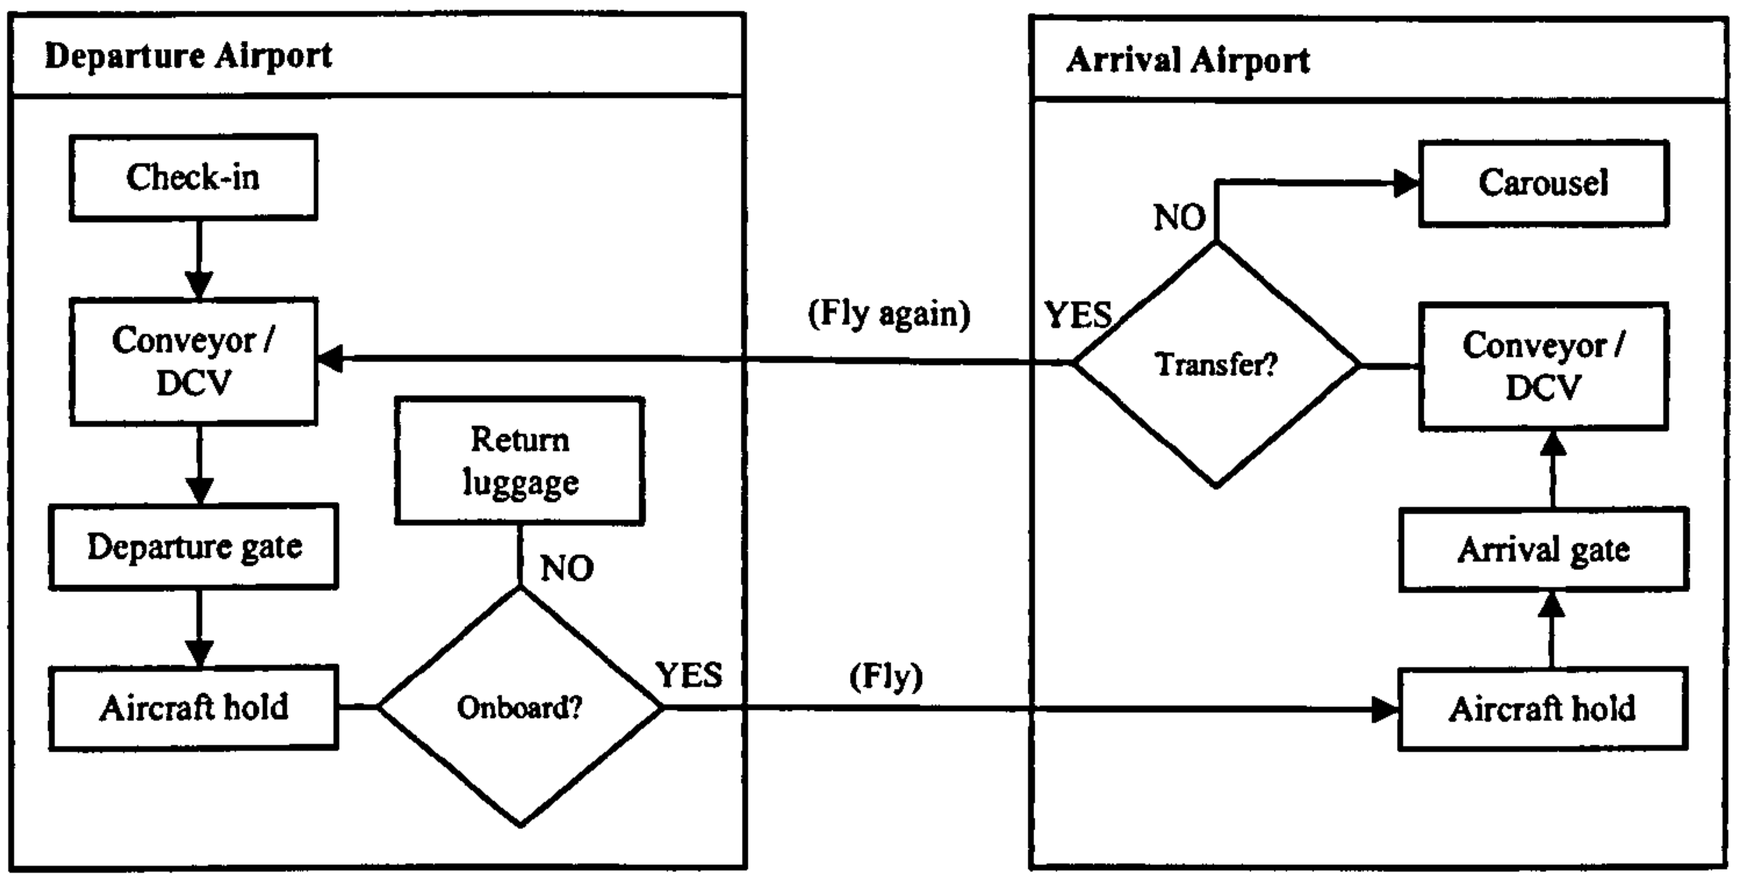

Figure 2.14 – Luggage flow in airport handling systems

2.4.2. Current Situation

In order to track luggage in an automated fashion as it moves around the airport conveyor and DCV system, the vast majority of airports utilise bar code technology. Bar codes are attached to luggage at check-in and these are then scanned by large arrays of scanners, shown in Figure 2.15, as they move. Such tracking allows computer systems to control the conveyor and DCV systems in order to route luggage to the correct departure gates and carousels as appropriate.

Within this automated process there is the possibility of the bar code tag becoming damaged or the luggage obscuring the bar code through rolling or tipping over. Both of these events have a chance to occur particularly when luggage moves from conveyors to DCVs and back again, as well as when travelling along conveyors which have an incline. Such occurrences can lead to the system being unable to read the bar code labels, and therefore not being able to transfer luggage to its appropriate destination in time.

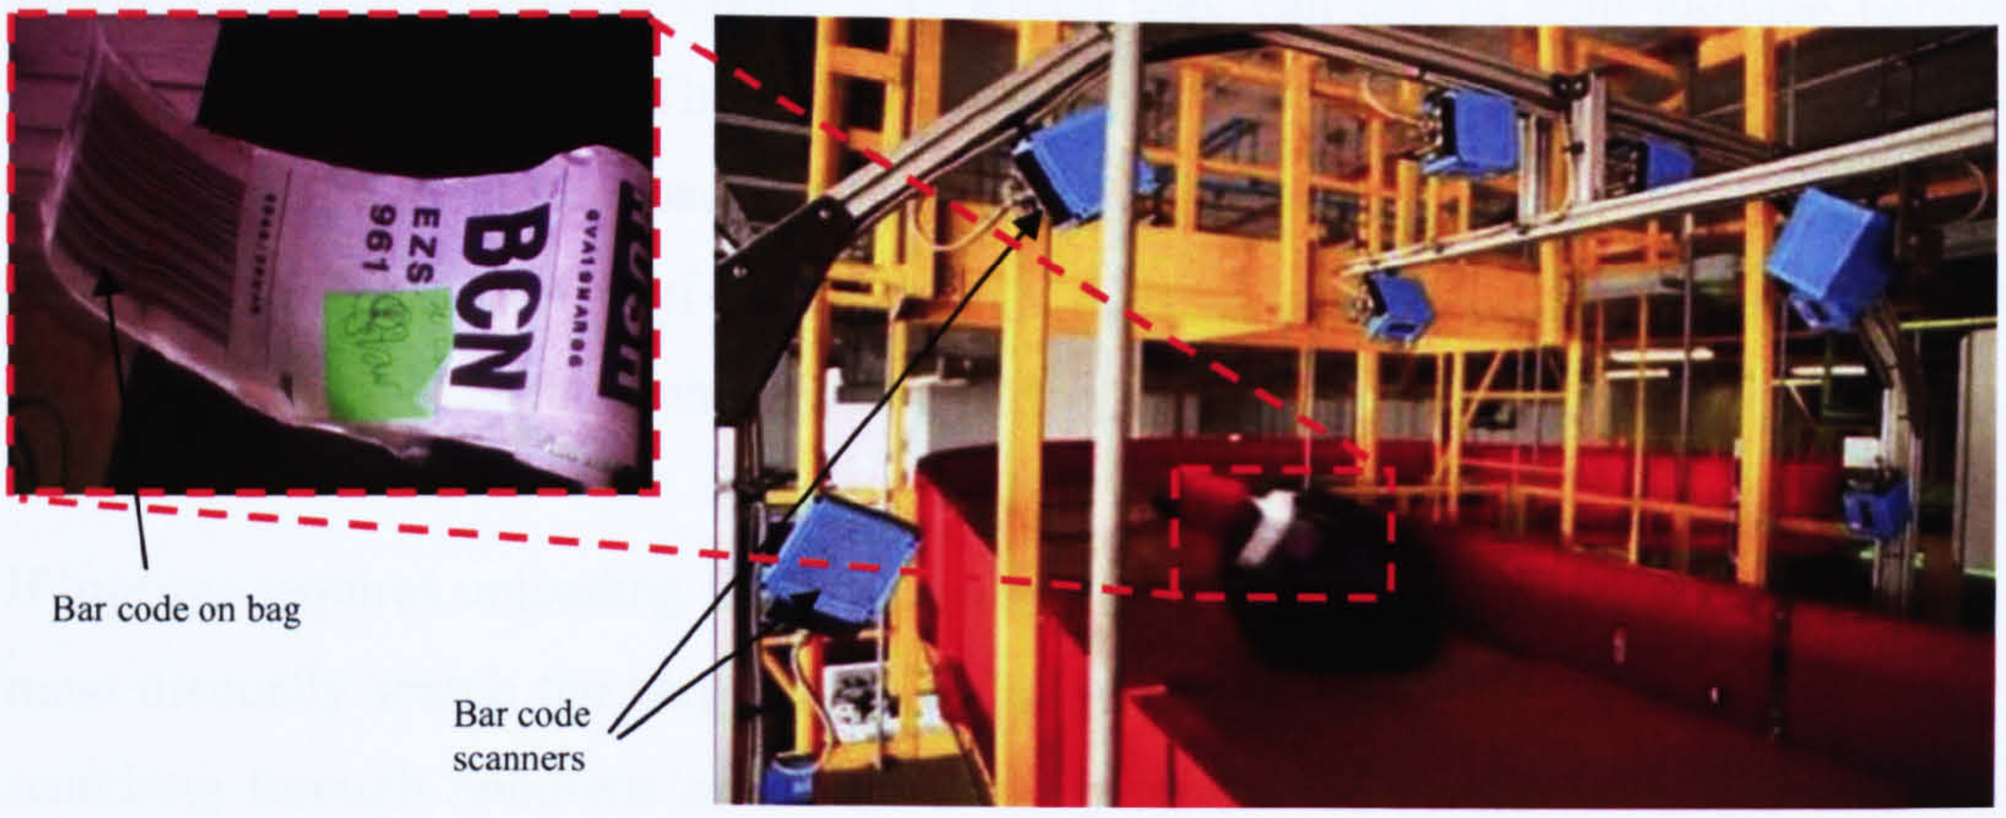

Figure 2.15 – Luggage on a conveyor being tracked via its bar code

Once at the departure gate, there is further potential for luggage being lost. Luggage handlers have the responsibility [58] to check luggage present at the gate against the luggage inventory which is compiled from check-in information. This only checks that the luggage present should be on the current flight, and not if luggage is missing – in many cases there simply is not time to go looking for luggage which is missing as there may only be an hour or two from a passenger checking in to the aircraft taking off.

Luggage which is verified for a flight is then taken to an aircraft on a small luggage truck. Often these trucks have no sides, as this makes it easier for staff to load and unload luggage, so there is the potential for an item of luggage to fall off unnoticed. Furthermore, once at the aircraft, luggage is not typically verified again – it is simply loaded into the cargo hold. This leaves the possibility for rogue[6] luggage to be placed onto the luggage truck whilst in transit and is hence a security risk.

Airlines in New Zealand have had to find a way [59] around this problem since it is law there that they must authenticate luggage for travel at the point it enters the cargo hold. As a result, airline employees have been issued with hand-held bar code scanners (similar to that in Figure 2.4) which they can use to scan luggage before loading it onto an aircraft. The hand-held scanner verifies whether or not a particular item of luggage should be loaded - this decision is made by the scanner itself and not by the personnel in charge of it. Whilst this does go some way to solve the issue of rogue luggage, it is not something which is deployed world-wide.

If luggage requires unloading from the aircraft prior to take off, then luggage handlers must manually search the cargo hold for the item of luggage which can often entail searching through hundreds of similar looking items. Since such searches must be made quickly in order to prevent delays, luggage can be trampled, thrown and generally mishandled in an attempt to find the correct luggage. Obviously this is detrimental to the passengers whose luggage is being mistreated, but it also poses a

[6] Generally undesirable items, although with recent terror threats explosive devices come to mind.

health risk for the luggage handlers in terms of possible injury related to the extra manual labour [60].

Unloading luggage from an aircraft upon landing leads to issues similar to those found when loading. Luggage can be lost in transit, either on its way from the aircraft to the arrival gate, or once back in the airport conveyor system. There is also the possibility that luggage could be left in the cargo hold mistakenly. Since the luggage is not tracked as it leaves the hold, only when it arrives at the airport building, there is no system in place to make luggage handlers aware of this situation.

When luggage arrives at the carousel for claiming by its owners, there is often a large crowd of people waiting. This can be very frustrating as many people try to collect their bags at once, pushing and shoving to keep themselves first in line. There may be instances were the wrong bags are claimed, since a lot of luggage looks similar, and there may even be malicious attempts to steal luggage. In terms of the luggage carousel, there is no system in place to make the process of claiming luggage easy and relatively stress free.

A report from Delta Airlines [61] states that it only achieves 80-85% accuracy when tracking luggage using bar code systems. Despite this low level of accuracy, the majority of airports have stayed with bar code technology. Delta Airlines has been trialling RFID technology and are impressed by the results. They have reported that in comparison with their existing bar code system, the use of RFID during trials on 40,000 bags yielded a successful scan rate of over 99% - much better than their results with bar codes.

McCarren International Airport in Las Vegas [62] and Hong Kong International Airport (HKIA) [63] are the only two airports to have implemented RFID tagging on a post trial basis. It is unclear what features and benefits are provided by the technology, although it does not appear that there is much benefit to passengers, particularly since airports which support connecting flights are still using bar codes. Therefore the tags issued at these airports include both bar code and RFID technology, as shown in Figure 2.16.

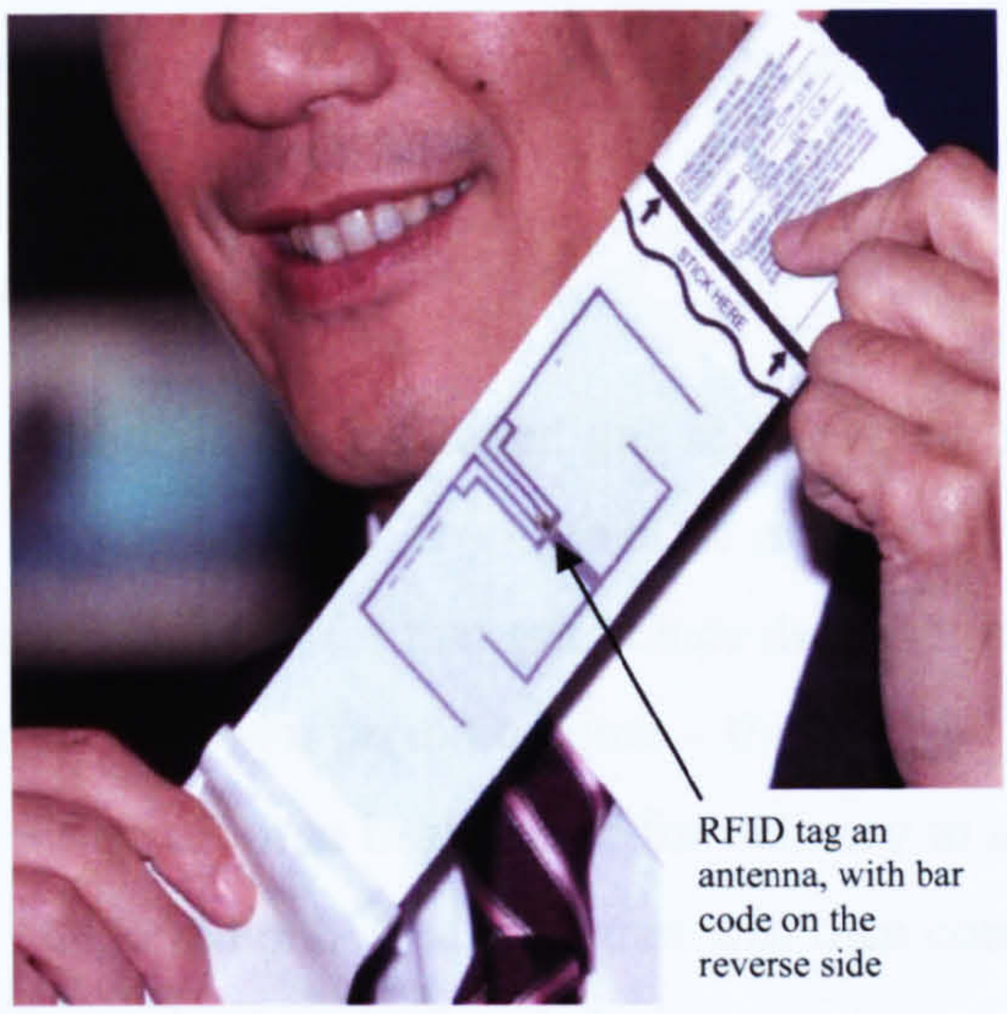

Figure 2.16 – Luggage tags used in Hong Kong International Airport (HKIA)

HKIA suggests that successful reads of bags have risen from an average of 80% using bar codes to 97% using RFID.

In the UK, Heathrow [64] airport recently began trialling RFID; however reports [65] suggest that the largest issue in replacing bar codes is the cost of implementing an alternative – something which the airlines are reluctant to do with the onset of increased competitiveness, as well as rising fuel prices. In addition it appears that the different available standards and frequencies have been an issue, since the tags must be legally useable worldwide. To some extent this has been overcome through the production of tags which operate at multiple frequencies; however they are costly (approximately 14 pence each in volumes of 10 million).

In its current state, the tracking of luggage is a hybrid of tracking methodologies. On one hand, luggage is tracked automatically through the airport but on the other it is not tracked at all once it leaves the airport building – the assumption being that luggage which has reached the departure gate will be loaded onto the aircraft, and will all be unloaded at the destination airport. This simply is not the case, leading to the possibility of luggage being lost in transportation, maybe even stolen. To add to this, the method of luggage tracking used largely (i.e. – bar codes) leads to such a low level of accuracy that it is surprising that the airline industry only loses the small percentage of the luggage that it does.

2.5. Asset Tracking in the Packaged Gas Industry

2.5.1. Gas Cylinder Life Cycle

Looking now at gas cylinders in a packaged gas scenario, an empty gas bottle must first be filled with an appropriate gas. It will then be marked for delivery to a regional storage centre based upon demand for that gas. Once delivered to a storage centre, the gas will be stored for an unknown period of time – this may be days, months or even years in some cases. When a bottle is allocated for delivery to a customer it is loaded onto a delivery truck and delivered. The customer will then consume the gas for their purposes – once it is completely consumed they may contact the gas company and request the empty cylinder be picked up. This empty cylinder will then be returned for refilling, but first it is inspected for any damage. If it is not serviceable, then it will either be repaired or disposed of. This entire process is illustrated in Figure 2.17.

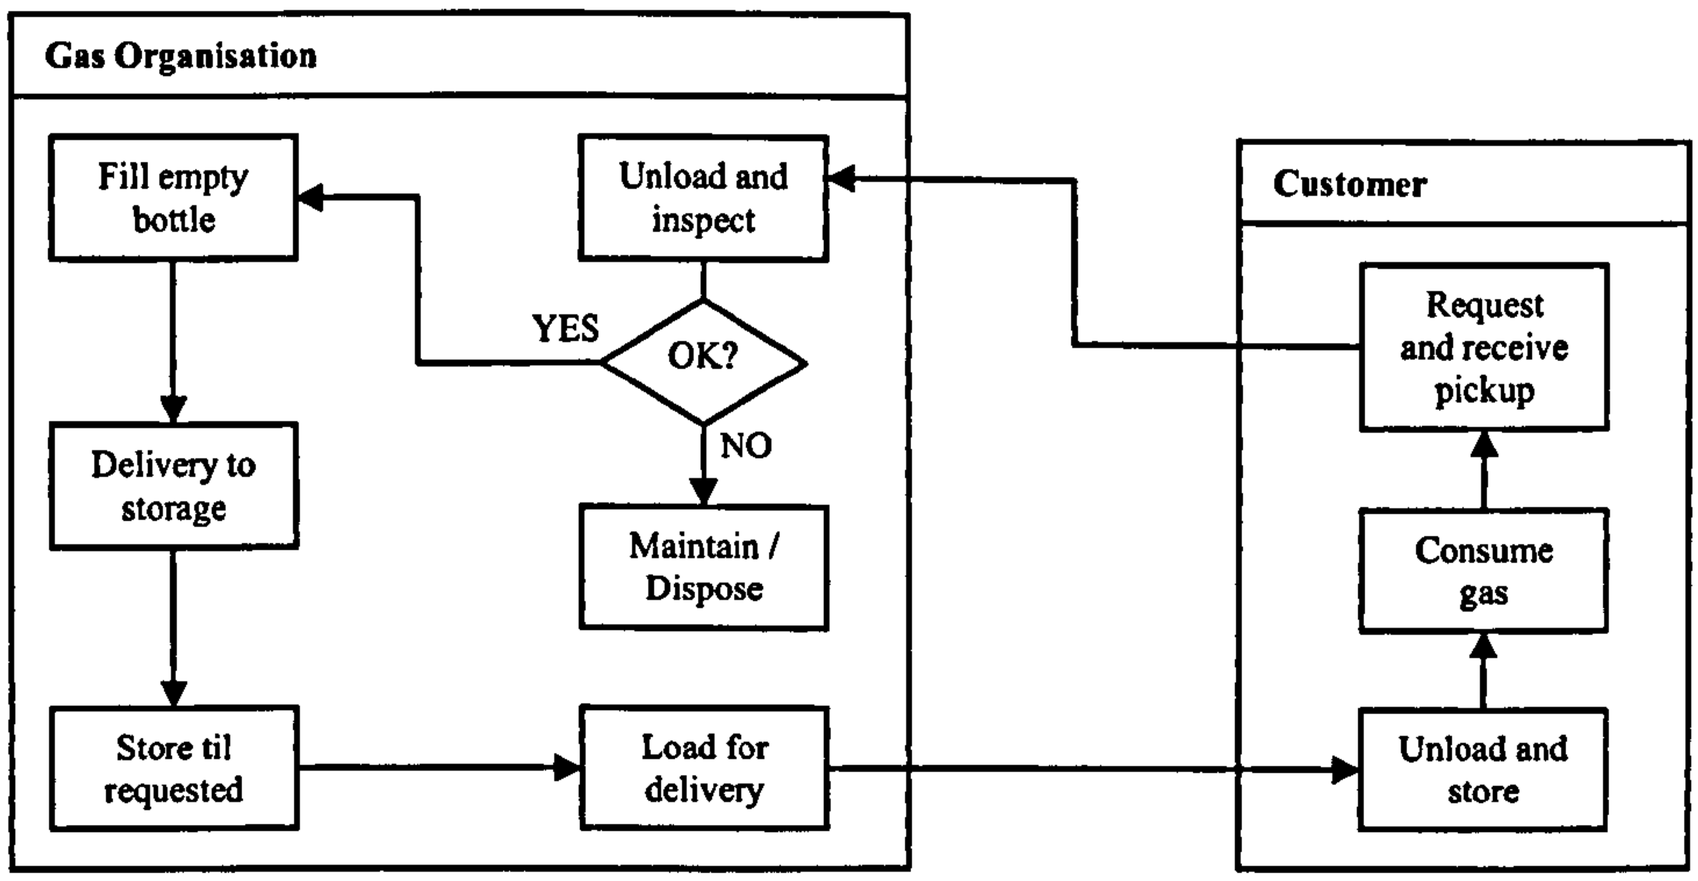

Figure 2.17 – Movement of a gas cylinder

2.5.2. Current Situation

Until quite recently, many within the packaged gas industry did not apply any method of tracking to their gas cylinders. This led to it being impossible to trace individual cylinders, meaning that an organisation would not know whether a particular cylinder was in stock, or with a customer. Storing information about individual cylinders was impossible, which is unacceptable when one considers the importance of cylinder maintenance. Cylinders have to be regularly tested to see if they are corroded and if they are able to withstand their rated pressure[7]. These tests are vital in order to ensure that cylinders are safe for use; if a cylinder cannot be identified, how does one prove that such tests were carried out?

In terms of inventory management, many resorted to a best effort method by simply recording stock levels. This afforded the ability to estimate whether or not a particular type of gas cylinder might be in stock. However, this system was open to errors in both the manual counting of stock levels and their communication. In addition, malpractice by customers was difficult to trace and so theft of gas was common.

Figure 2.18 – Gas cylinders in storage racks

[7] Via a method referred to as hydrostatic testing.

Figure 2.19 – Gas cylinders on a delivery truck

As a result larger suppliers of gas cylinders, such as Air Products [66], have started using bar code systems. In addition, numerous third party companies, Pollution Control Industries [67] for example, have similar technologies which they will assist organisations in deploying. However these systems are largely manually operated, with personnel being equipped with hand held scanners (similar to that in Figure 2.4) in order to maintain inventory. Therefore, whilst it is now possible to track a cylinder as it moves through the organisation, is sent to customers, and is maintained, there is still the potential for error and theft. As a comparison, the example given in Section 2.2.2 probably best fits the situation at the moment in large packaged gas organisations, whilst many smaller organisations have more in common with the scenario described in Section 2.2.1.

The problems faced as a result of poor tracking have been recognised by the packaged gas industry, but there appears to be little motivation to move to RFID technology at the moment, owing to reliability issues. Table 2.2 presented an RFID tag which had a read range, when attached to metal, of 3m. However, this is when in LOS conditions which is rare with gas cylinders, as demonstrated in Figures 2.18 and 2.19. As a result, organisations such as Air Products who have trialled passive RFID systems are failing to see the potential in its application since it simply is not proving reliable for

them[8]. This is corroborated further by the fact that organisations offering RFID technologies are specifically marketing them for short range applications [68]. Due to this, they are really offering nothing more than bar code technology whilst charging more for the privilege. A report [69] suggests that these systems have a maximum range of 10cm, which is inadequate for anything other than manual asset tracking considering that the majority of gas bottles have a diameter greater than this.

[8] Information obtained after speaking to R. Wiktorowicz, Air Products PLC, Crocked Lane, Chineham, Basingstoke, Hampshire, RG24 8FE, UK

3. Wireless Sensor Networks

After considering the technologies currently available to assist industries in tracking assets, an alternative concept was sought as an improvement. The specific aim was to find an idea or technology which overcame the issues relating to the environment of assets in the airline and packaged gas industries, whilst maybe providing further benefits. As a result, this work found a potential candidate in the area of Wireless Sensor Networks (WSNs), an exciting and ever expanding topic predominantly limited to the research domain at the moment. However, it is felt that this technology has huge potential for many applications, with asset tracking being just one of those; spot welding in the automotive industry is another application which is discussed later in Chapter 7.

3.1. What is a Wireless Sensor Network?

A Wireless Sensor Network consists of many spatially distributed devices which act autonomously and communicate using radio waves. They utilise sensors to monitor one or more aspects of the environment in which they are placed. The network created between the devices can typically be described as ad hoc or spontaneous, since there is no formal infrastructure; inter-device connections are often temporary, lasting long enough only to exchange data as it becomes available.

Unlike dumb networks, which operate on an end-to-end principle, WSN can be considered as smart. Each device can be said to possess intelligence, since they are all able to perform tasks in response to events; such events may occur as the result of sensor stimuli or a data packet received from a neighbouring device. This has meant that the concept of WSN can be used in a broad range of applications, a number of which are discussed briefly in this chapter.

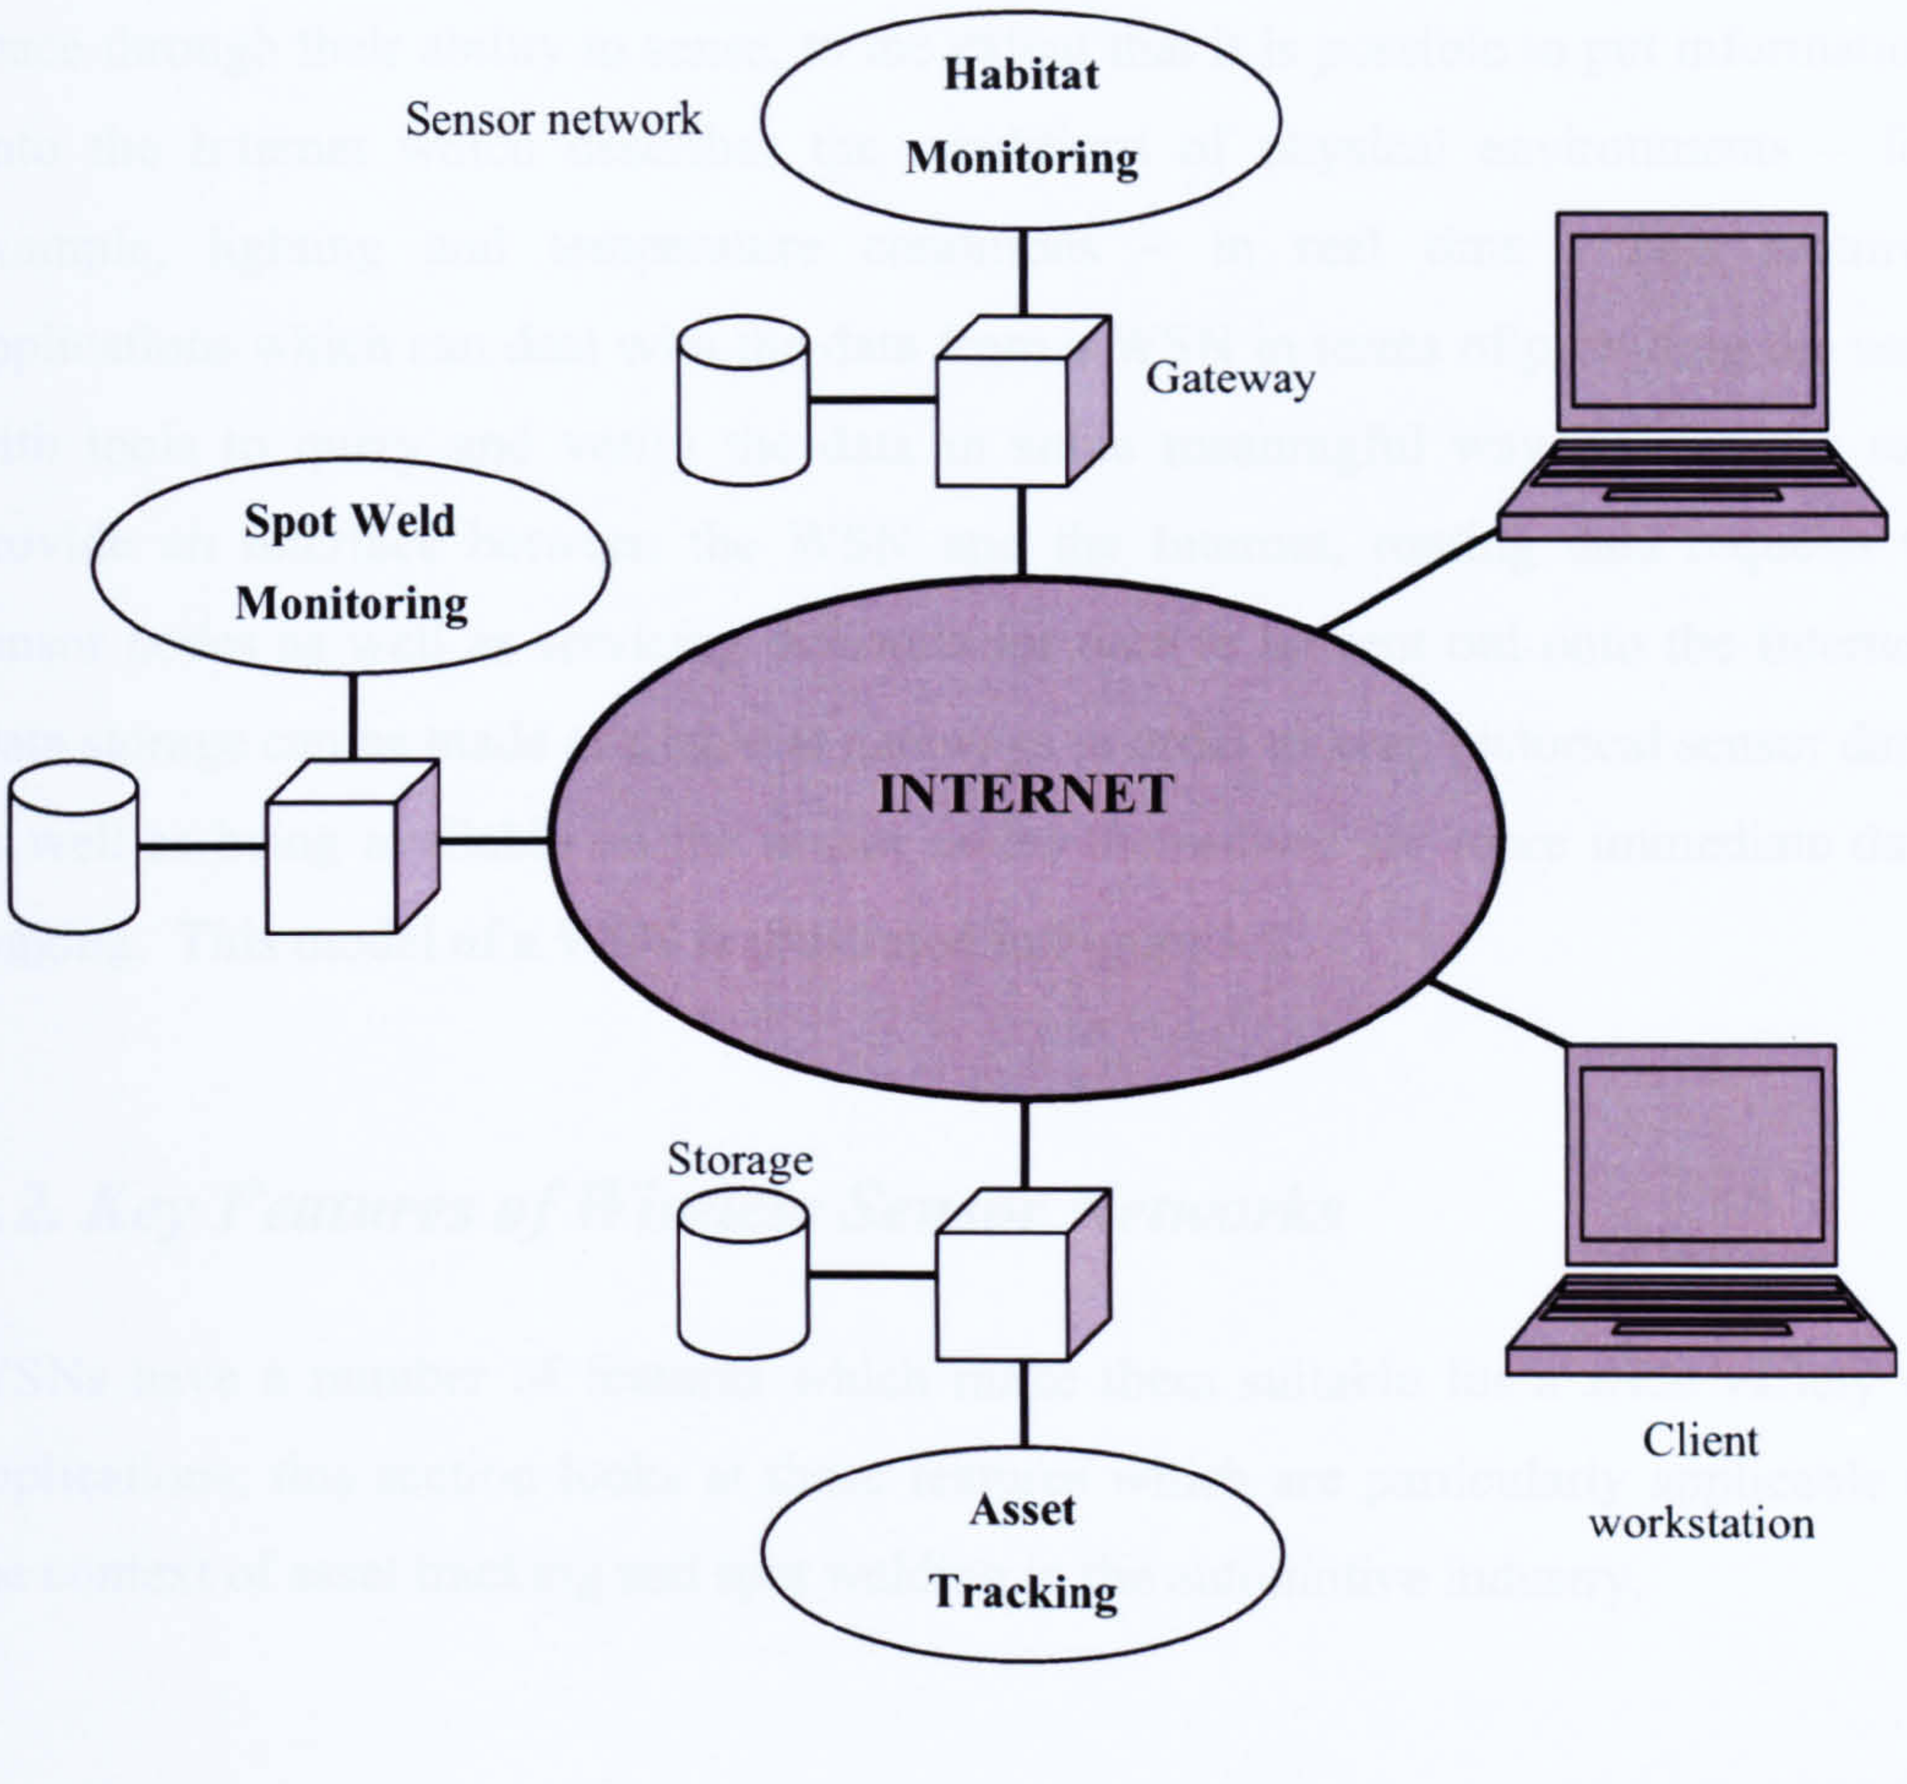

Figure 3.1 – WSN expanding the Internet through invading physical spaces

The sensor devices, or nodes, are typically equipped with a battery power source, a small microcontroller, a small quantity of memory as well as the obvious radio transceiver for wireless communication and an optional sensor (or maybe even an array of sensors).

They are often designed to be inexpensive and as small as possible, whilst being able to operate autonomously for long periods of time. As a result there are inevitable constraints such as limits to their processing capabilities; these are discussed later in this chapter.

WSNs offer a way of extending the current networks and the Internet into physical space through their ability to sense, to the extent that it is possible to put information onto the Internet which describes the conditions of physical environments – for example, lighting and temperature conditions – in real time. This requires applications which can deal with the data from a WSN in terms of providing the user with tools to query and verify the data in some meaningful way. Gateways can provide an interface between the WSN and the Internet, routing data requests to sensor nodes as well as servicing demands for data to be sent out onto the Internet. Data storage can be made available at gateways in order to keep historical sensor data, as well as being available on the sensor nodes themselves for more immediate data logging. This model of a WSN is illustrated in Figure 3.1.

3.2. Key Features of Wireless Sensor Networks

WSNs have a number of features which make them suitable for a wide variety of applications; this section looks at those features which are particularly applicable in the context of asset tracking and spot welding in the automotive industry.

3.2.1. Mesh Networking

The ability of sensor nodes to communicate and form networks is an essential part of their function within a WSN. They require this ability in order to communicate with their neighbours on the network, as well as with the base station which serves as a gateway between a WSN and a workstation or the Internet, as illustrated in Figure 3.1. The nodes within a WSN typically form a mesh topology [70]. This allows the WSN to route data between nodes on the network by relaying, or hopping, data from node to node until an appropriate destination is reached. This is often referred to as multihop communication [71] or multihop routing [72], and is illustrated in Figure 3.2.

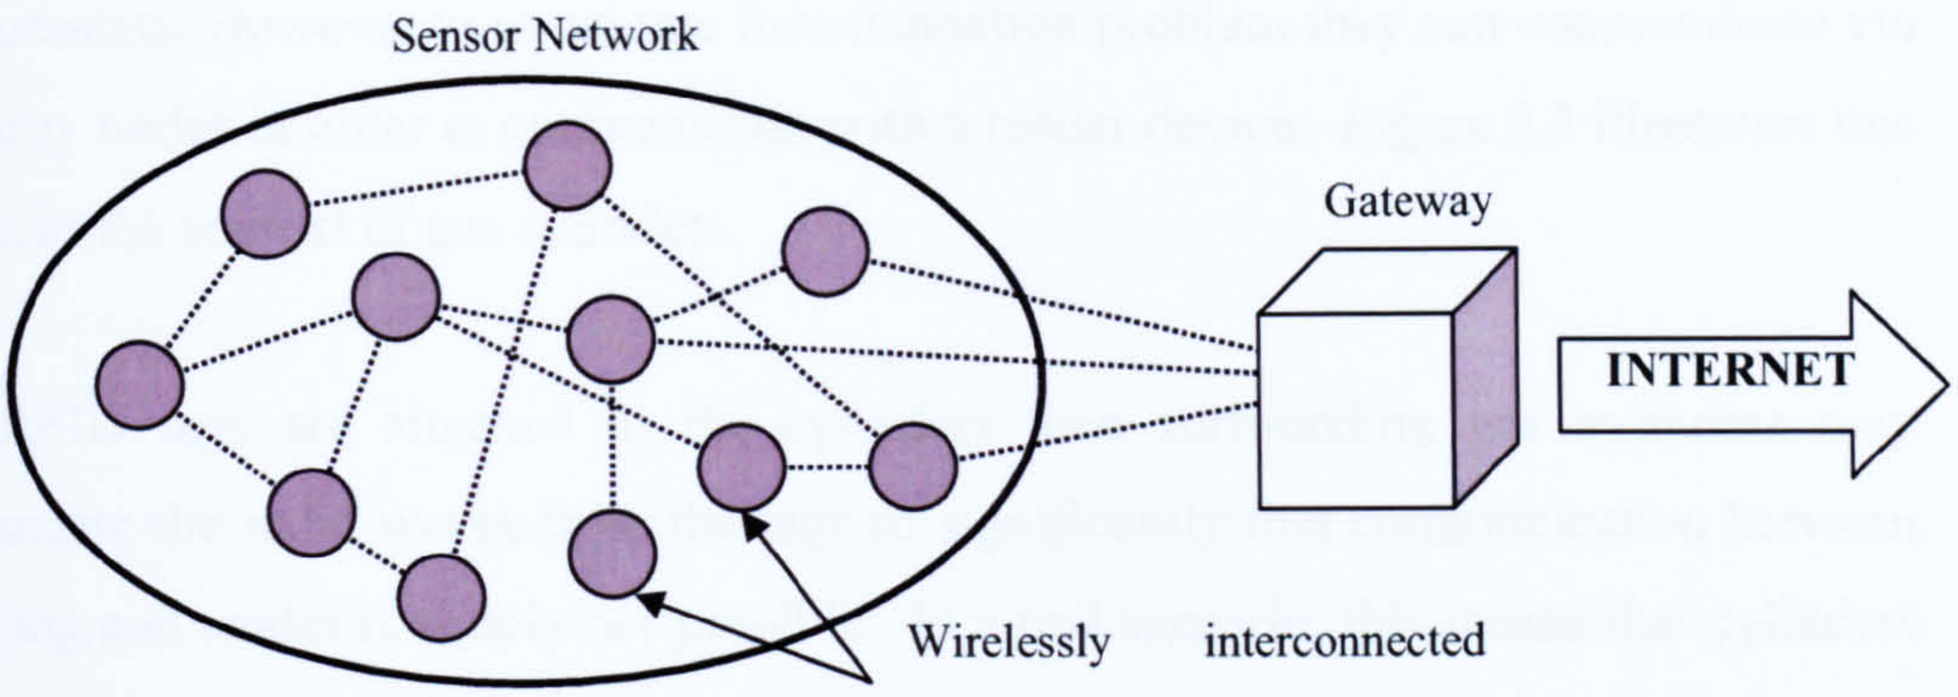

Figure 3.2 – Wireless sensor network mesh topology

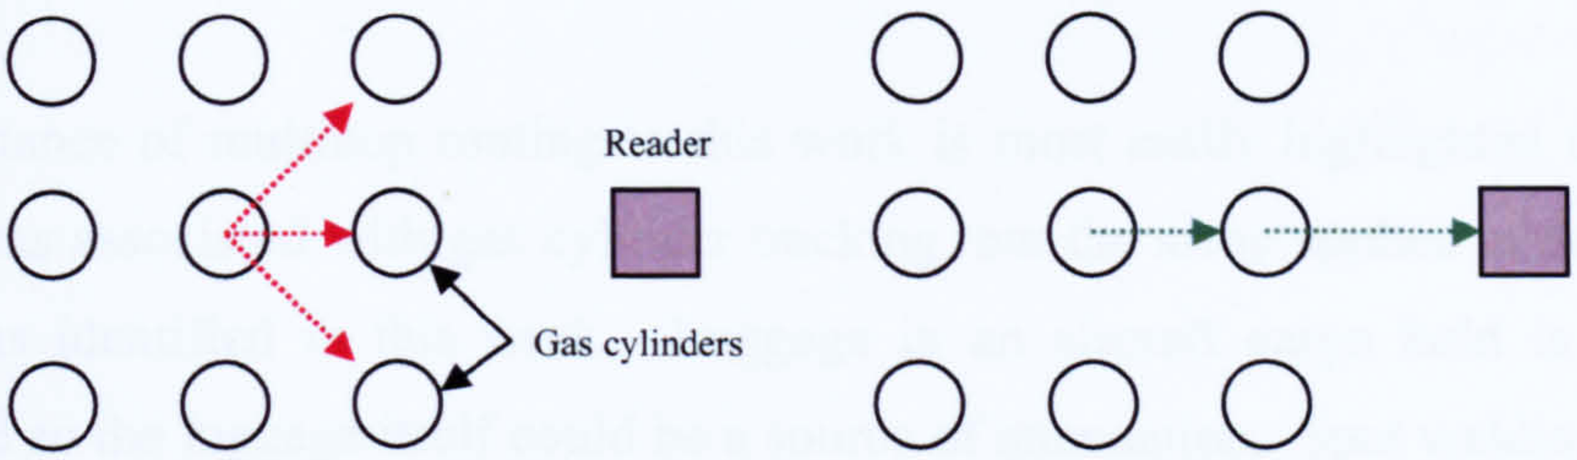

Figure 3.3 – Gas cylinders with RFID tags (left) and sensor nodes (right) attached

Connections within a mesh network may be configured and maintained throughout the life of the network, or there may be no connections at all – that is – data is sent in the hope that it will arrive at the target destination. These are described as connection orientated and connectionless networks [73] respectively. Importantly, the connections in mesh networks can self heal in the event of a failure by simply routing data around the point of failure. This increases the reliability of the network in the event of hardware failure or movement, since connectivity is preserved.

It is the multihop routing capability of WSN's which is particular important in the context of this work. In terms of asset tracking it is believed that it could be used to overcome the limitations of RFID technologies, in particular when they are placed in metallic environments. Recalling Table 2.2, the maximum range of an RFID tag in LOS conditions was approximately 3m. However, in reality, there is likely to be objects between the tag and reader meaning that this range would not be achievable. Sensor nodes use radio waves to communicate, and so are bound by similar

constraints. However to overcome the attenuation problem they can communicate via nearby nodes in order to communicate with a reader device. Figure 3.3 illustrates this idea in the context of gas cylinders.

If RFID tags are attached to the cylinders then surrounding gas cylinders may attenuate the radio waves from the tags so significantly that communication between the tag and reader is simply not possible. In a real scenario, this means that cylinders would not be identified and would lead to large discrepancies in any inventory management system. However, using WSN nodes, direct tag-to-reader communication is not necessary. Instead the nodes can utilise multihop routing and therefore negate the problem of attenuation.

The importance of multihop routing to this work is most easily highlighted through the problems associated with gas cylinder tracking, but the same applies to the other applications identified in this work. Luggage in an aircraft cargo hold is tightly packed and so the luggage itself could be a source of attenuation. Spot welding relies on large amounts of metallic machinery. In all cases, utilising short range (in the order of tens of metres) multihop communications is intended to overcome the problem of attenuation suffered in currently available tracking technologies, none of which utilise mesh networks to their advantage. The fact that data can be routed from node to node means that the effective communication range of a single node can be greatly increased [74], provided that there is sufficient node density. In asset tracking, node density is naturally assured, since assets are tightly packed to ensure the most effective use of storage space.

3.2.2. Expandability

The sensor nodes in WSN's can be expanded in one of two ways; via their software, or via additional hardware. Dealing with the software first, it was previously mentioned that a sensor node may typically include a microcontroller. These devices can be programmed, which implies that the user can define how a device will respond in a given situation to some data, be it sensor stimuli or network data.

In addition to programmability, sensor nodes can usually be expanded via hardware in many different ways. Such hardware may include sensors, such as those for detecting temperature and pressure, or user interface devices such as liquid crystal displays and keypads.

Taking the packaged gas industry as an example, WSN's may provide the basis of an effective tracking system initially, but in the long term the system may be expanded to incorporate sensor data; this is actually demonstrated in Chapter 7. Such sensor data may report the condition of the gas cylinders or the environment which they currently inhabit – this could be used to detect gas leaks, for example. A similar situation applies in the automotive industry, since sensors are being applied for the monitoring of spot welding [6,7]. Hence, a system which accommodates future expansion for such sensors is advantageous.

3.2.3. Scalability

Scalability in this context refers to the ability of some application or technology to meet the growth requirements of an organisation. In terms of WSN, this refers to the ability of the network to sustain a large, and possibly increasing, number of sensor nodes. There are a number of features of WSN's which afford them the ability to cope with larges volumes of network traffic and maintain reliable operation. Since WSN's take on a mesh topology, new devices can easily connect to the network and both organise and configure themselves appropriately [74]. In addition, the programmable nature of the sensor nodes enables protocols and networking algorithms to be written which can cope with addressing and routing data between a large numbers of devices.

With asset tracking, there may be hundreds or even thousands of items to track at any one time – as would be the case particularly for the airline and packaged gas industries.

3.2.4. Low Maintenance Cost

WSN are often designed for semi-permanent or even permanent deployment. As a result, the sensor nodes are designed to be almost completely autonomous for their entire lifespan, which can often be determined by the capacity of the nodes power supply. New nodes being introduced to the network need no human configuration, since they can configure themselves automatically - this means that there is no specialist training required in node deployment. In addition, if a system requires updating it is possible to implement in-network reprogramming, so that the devices can be updated remotely rather than having to be physically recalled.

These features all point toward an industrial critical feature – low running costs. This is important in terms of industrial applications because lower running costs may make the technology more attractive for implementation. With asset tracking, the goal really is to be able to tag an item and then just assume that it will work; in the automotive industry low maintenance is vital since the addition of sensing systems seeks to reduce the reliance on manual inspection of machinery and parts for faults.

3.2.5. Intelligence

It has been mentioned earlier in this chapter that the sensor nodes which make up WSN's can be perceived to possess some level of intelligence because they can respond to events in a user defined, and therefore intelligent, way. This is useful in a sensing scenario as it allows the nodes to respond to emergency situations themselves, such as:

- In asset tracking, a node attached to an asset might decide the asset is lost or misplaced because it has not received any network traffic for a specified period of time; therefore it may take action to attempt to rectify this situation.
- In spot weld monitoring, a poor weld being detected could trigger an alert for action to be taken; such action could be an automatic adjustment of welding parameters, or human intervention.

Sensor nodes can also operate in conjunction with one another to perform tasks (i.e. – collaboration). This can be useful in some instances in order to decide if sensor data is accurate – multiple nodes can report their data to one another and only act if they all agree that the situation requires some action.

3.3. Constraints of Wireless Sensor Networks

3.3.1. Electromagnetic Phenomena

Since WSN's make use of radio waves, a part of the electromagnetic (EM) spectrum, they are subject to the same principles. EM waves carry energy from one wireless device to another through air or any other intermediate objects. Therefore they are subject to attenuation due to their propagation medium, as well as other effects such as interference, reflection, refraction, diffraction and multi-path effects [75]. All of these principles are discussed in detail in Chapter 5, however what one must understand here is that they serve as a source of unpredictability in wireless communication; sometimes they can have a positive effect on transmission, whilst at other times they can be so detrimental that communication is simply not possible.

These effects are not limited to WSN's; they affect all types of wireless communication, including WLANs and traditional RFID, as do many of the constraints to be discussed in this Section. In WSN's the effects can possibly be exaggerated when compared to WLAN since the transceivers used in low power sensor devices have much tighter power constraints and therefore cannot use as much energy to transmit as, for example, a transceiver in a laptop. As a comparison, a typical sensor node transceiver outputs around 0dBm [76-77], whilst a typical transceiver in a WLAN can output a maximum of 20dBm, according to UK radio regulations [78]. Note that dBm refers to milliwatts (mW) being referenced to on a decibel scale, where 0dBm is equal to 1mW. Table 3.1 shows how the decibel scale impacts upon these figures; simply put, a laptop can transmit up to 100 times more power than a sensor node, giving it a larger operating range.

dBm	mW
0	1
3	2
6	4
10	10
20	100

Table 3.1 – Decibel (dB) values

3.3.2. Bandwidth

There are a number constraints placed upon wireless networks in general with regards to bandwidth. First and foremost, the bandwidth available to wireless devices is shared between all of the devices within radio range of each other because they have to share the same transmission medium. In terms of this project this is an important point, since the nature of asset tracking means that the sensing devices are likely to be densely packed into storage areas meaning that each device can actually only utilise a fraction of the available bandwidth. Further to this, most sensor nodes only have a relatively small amount of bandwidth to begin with since, as we have already seen, they are low power. This means that their bandwidth is usually in the order of kbps, as opposed to Mbps which can be utilised in typical WLANs.

An important consideration is the method of accessing the transmission medium in order to make the most effective use of the available bandwidth. The most efficient means of transmitting on a network, wireless or wired, is by allowing a node to transmit when it needs to. However, since the medium is shared, this is not possible as it would cause data collisions – therefore a Medium Access Control (MAC) protocol [79] of some design is required in order to prevent nodes from talking over one another and therefore interfering.

3.3.3. Hardware Constraints

Since a key feature of sensor nodes is autonomy, their typically finite energy resources often define their effective lifetime before they require maintenance. Therefore, an important feature in their design is reduced power consumption, a

feature already mentioned previously in this section with the sensor radio transceivers. Low power microcontrollers are used rather than full sized processors that one would find in desktop computers, which tend to include the bare minimum of everything in terms of programmable and random access memory (RAM) as well as having RISC[9] architectures. Additionally, they tend to run at relatively low clock speeds.

Power saving is two fold, and not simply related to low power hardware, since a programmer must realise how to effectively utilise the hardware through application software. With this in mind, many applications are designed to use the sensor nodes for very short periods of time, with large intervals between these bursts of processing. In addition, since memory is limited, programs must be lightweight and often sacrifice features such as complex networking protocols that maybe found as part of larger computers.

It is important to note that not all sensor nodes are limited in such ways because, as discussed later in this chapter, a number are based around typical computer components. They are however, not suited to the applications of asset tracking or spot weld monitoring due to their size, and so are not used in this work.

3.3.4. Weak Security

In any network, security is a concern – possibly more so in wireless networks since the transmission medium, air, is more accessible to third parties. Again, it is not so much an issue limited to WSN, however it is exacerbated by their low bandwidth and processing abilities. This means that the more complex a security protocol, the longer a sensor node has to process messages and the greater the overheads[10] in transmitting a data packet. Respectively, this means that the microcontroller and transceiver have to be switched on longer, therefore consuming more energy. As a result, one often has to consider a compromise between security and energy consumption. Data upon

[9] Acronym meaning Reduced Instruction Set Computer.
[10] Overheads are those data bits which encapsulate pertinent data within a packet (i.e. – its payload).

the network could be sensitive to an organisation, and therefore might need to be encrypted to make it unreadable by a third party.

3.4. Applications of Wireless Sensor Networks

As previously mentioned, there is great potential for the application of WSN in a number of areas, be it for research or industrial purposes. Much of the current work related to WSN is focused toward research applications, although more recently there have been a growing number of industrially related projects. Considering the different applications of WSN serves to emphasise and open ones mind to their potential; it is also incredibly interesting to discover the wide variation of applications, a small number of which are presented here.

3.4.1. Habitat Monitoring

On a small island off the coast of Maine in the United States a team of computer engineers from the University of California, Berkeley, conducted an experiment which applies WSN to real-world habitat monitoring [80]. Working in conjunction with biologists at the College of the Atlantic, the engineers installed 190 wireless sensors for the purpose of monitoring the habitat of the nesting petrels[11] on the island.

In the past, biologists studying the nesting behaviours of these birds had to travel to the island periodically to gather observation data. To check on the petrels they often had to disturb their burrows which often caused the birds to abandon their homes completely. As a result of the wireless sensors being installed on the island, these biologists could check on the birds remotely from thousands of miles away if they wished. Matchbox sized wireless sensors, capable of measuring variations in temperature, were left in burrows – from this data the biologists could determine the bird occupancy levels of the island.

[11] A species of bird which lives in burrows – the population is declining, hence the interest in studying them.

To facilitate the reception of live data from the sensor network, a system was developed to maintain connectivity and allow remote viewing of data (see Figure 3.4). At the lowest level of this architecture, sensor nodes act autonomously to provide the required sensor functionality. Typically the motes formed a mesh network, which means that they can utilise multihop routing in order to extend their communications range. Each sensor could then communicate with a base station, or gateway, which stored the data until it could be sent via satellite to Washington DC where it was forwarded on to servers in Berkeley. Up until recently this meant that all of the data could be viewed on the Internet via a webpage; this functionality has since been removed.

Whilst at first glance this system may appear to be a trivial matter of convenience for those who require the sensor data, the most significant benefit of this application of WSN technology was the minimisation of disturbance to the very habitat that the scientists were trying to preserve.

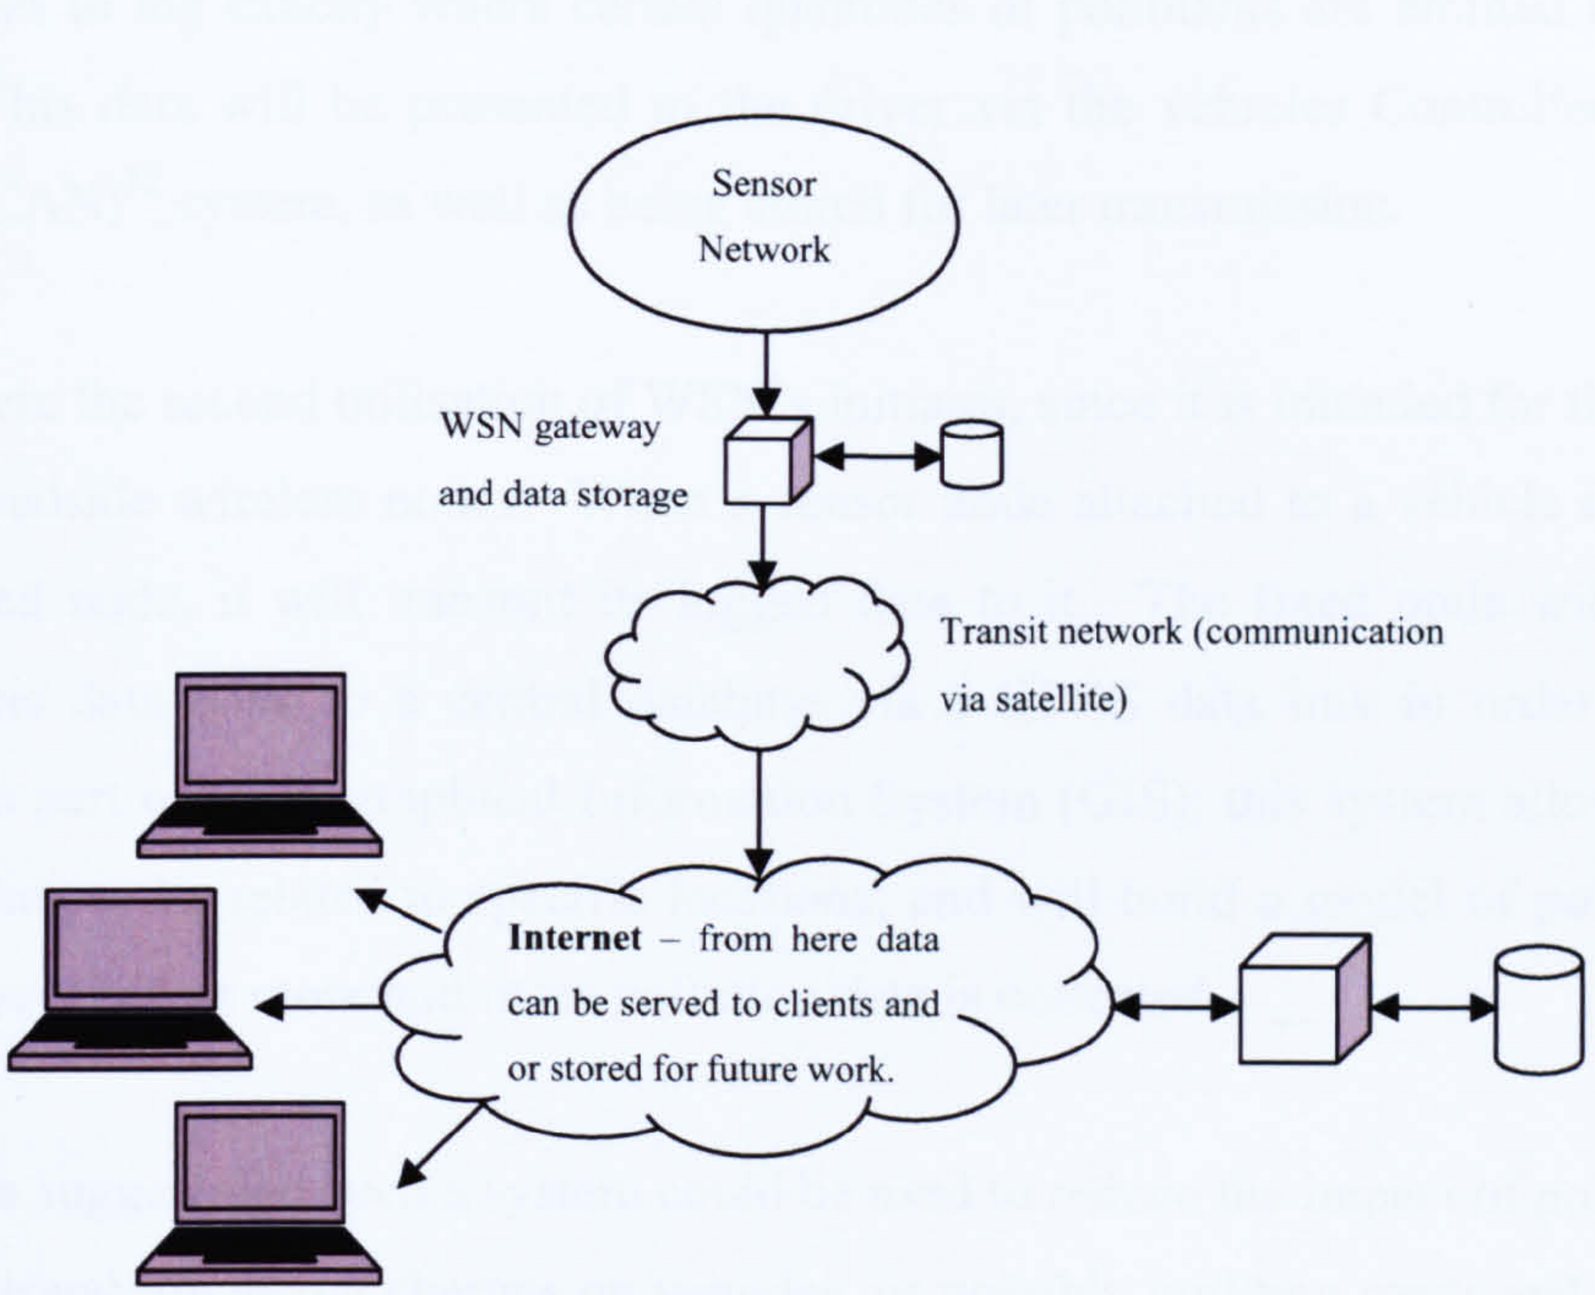

Figure 3.4 - System architecture for habitat monitoring

3.4.2. Pollution Monitoring

Currently, pollution control and monitoring is of concern due to worldwide concerns over global warming. A significant amount of the pollution which enters the Earth's atmosphere is generated from road transport; a problem which is compounded year upon year as the number vehicles on the roads increase.

Research being conducting at Liverpool John Moores University [81] has been considering ways in which the pollution from road transport could be monitored in order to identify pollution hotspots. This utilises WSN in two ways. The first way is to interface with pollutant sensors on the exhaust of a vehicle using wireless sensor nodes. The intention is that such as system should be suitable for both manufacturer fitting and also retrofitting to older vehicles. This makes wireless sensing devices ideal, since they will be small and require minimal amounts of invasive installation.

In addition to pollutant sensors a Global Positioning System (GPS) module allows the sensor nodes to log exactly where certain quantities of pollutants are emitted from a vehicle. This data will be presented to the driver via the vehicles Controller Area Network (CAN)[12] system, as well as being stored for later transmission.

This is where the second utilisation of WSN's initiates, since it is intended for there to be fixed roadside wireless nodes. When a sensor node attached to a vehicle detects such a fixed node, it will transmit its logged data to it. The fixed node will then transmit this data back to a central database via a GPRS data link in order to be analysed as part of a Geographical Information System (GIS); this system allows the pollution data to be related to specific locations, and will build a model of pollution hotspots over time as more and more pollution data is collected.

The authors suggest that such a system could be used to reduce the impact of pollution hotspots through imposing charges on vehicles, or possibly building roads and traffic

[12] CAN is a protocol and bus system which allows different systems within an automobile to communicate with one another, and is common in modern day vehicles.

controls which allow the bypassing of such locations. Measures such as these could help to reduce pollution, or could produce revenue in order to assist combating it.

3.4.3. Artificial Intelligence

At Utah State University, in the United States, research is being conducted which involves combining WSN technology with robotics. The result is given the name MAS-Net [82]. Many mobile robot platforms – MASmotes (see Figure 3.5) – which act as sensing or actuation nodes make up the mobile actuator sensor networks, or MAS-Net. Each MASmote is built using off the shelf components, which means that they are reasonably inexpensive to produce. At the heart of each MASmote is a wireless sensor node, which handles the bulk of the processing requirements as well as the communication between the robots.

The aim of creating MAS-Net was to give the robots a common task and allow them to work together in achieving the intended goal. The self-organising nature of WSN lends itself well to this aim – in addition the sensor nodes provide ample processing power whilst keeping battery consumption to a minimum.

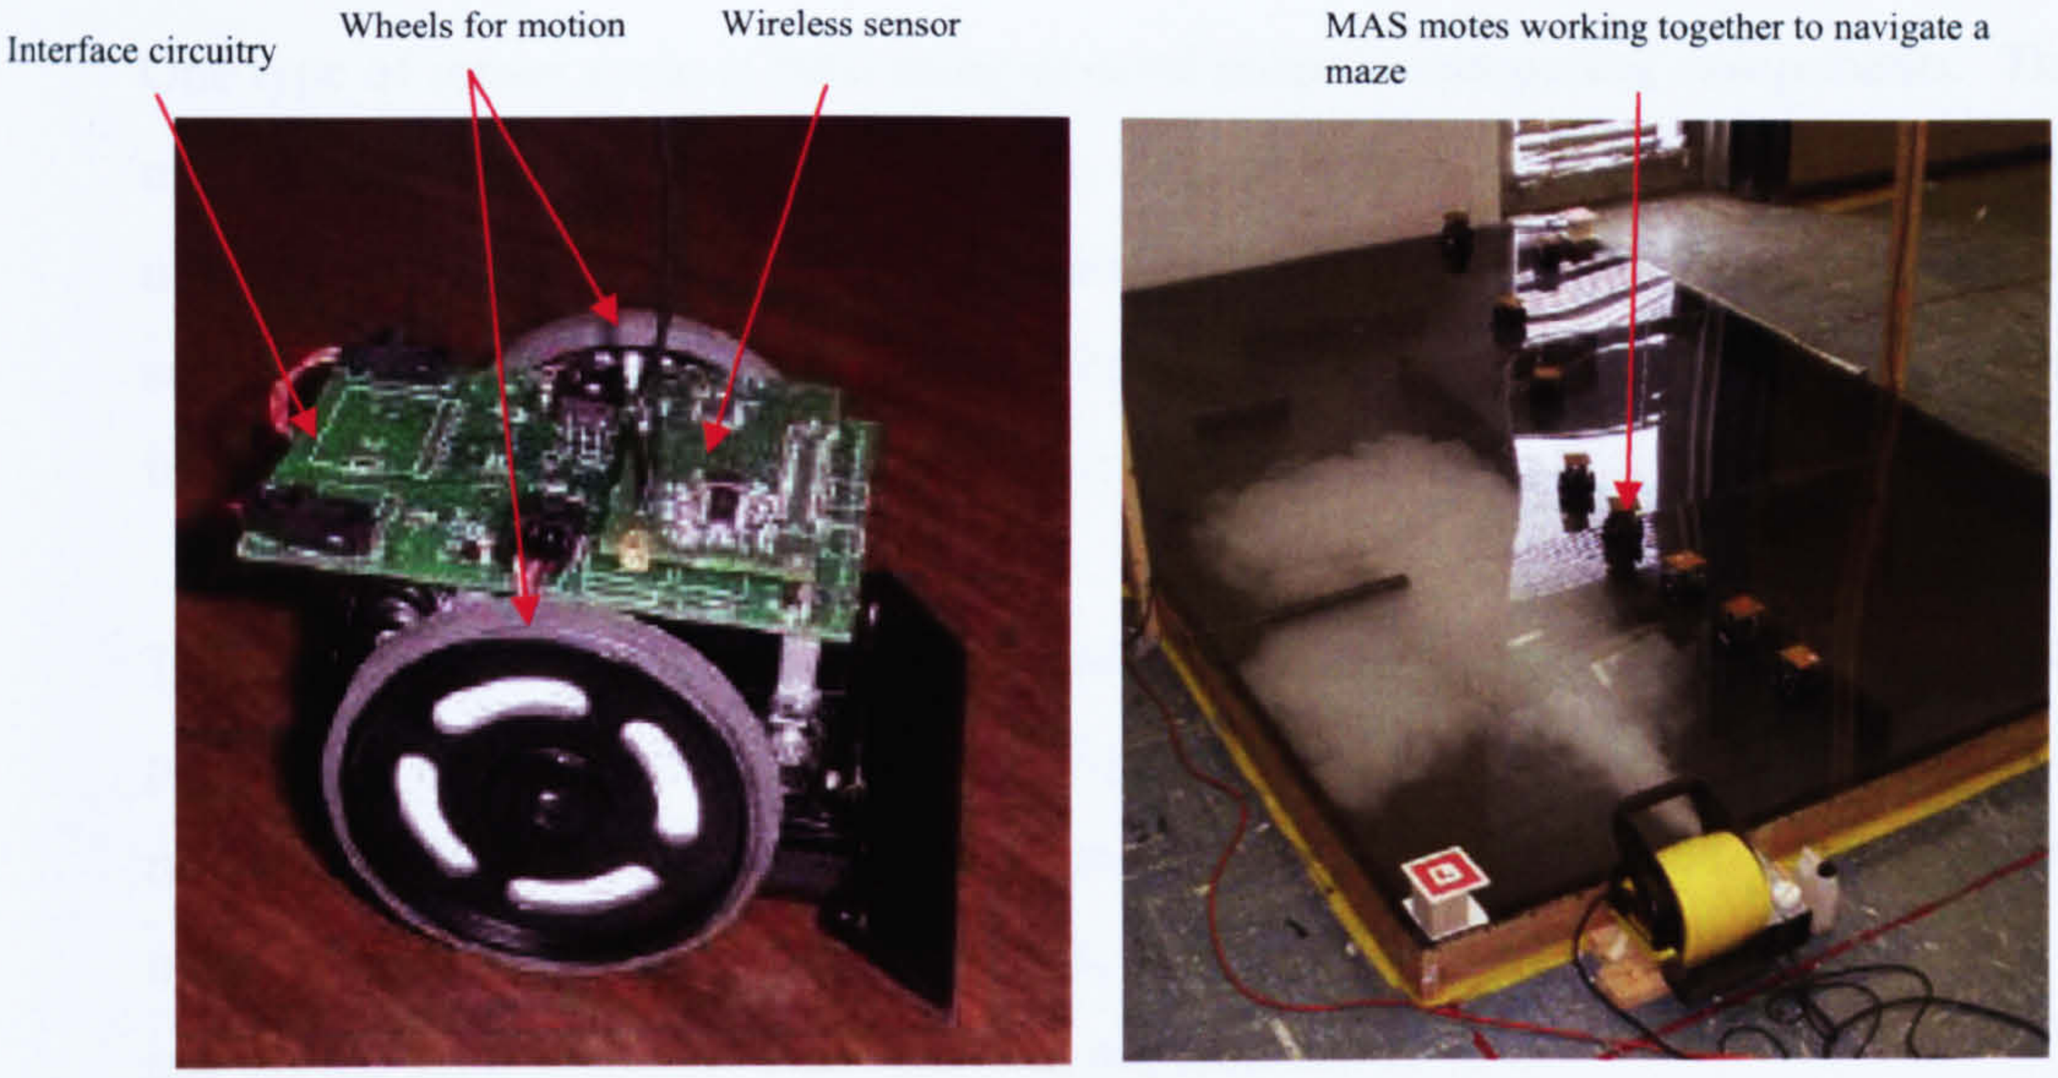

Figure 3.5 - A MASmote (left) and a group of MASmotes collaborating (right)

Currently the robots are able to move and perform tasks in formation as well as help each other navigate difficult environments, such as those which contain obstacles as well as natural phenomena such as fog. Such environments often pose a significant problem for sensors, and so the functionality of the MAS-Net as a whole is an impressive achievement.

No real-world applications have been attempted to date with MAS-Net, however the belief is that the technology could be deployed in places not suitable for humans. Their low cost makes them suitable for large scale production, as well as increasing their feasibility as an expendible device in dangerous situations for humans.

3.5. A Brief Comparison of Wireless Sensor Devices

The past decade there has been an explosion of activity aimed at developing sensing nodes suitable for a wide variety of applications. So far the focus of this document has been a general overview of WSN's, giving an idea of the advantages, drawbacks and applications of the concept itself. Now the focus must shift a little toward the actual devices themselves, firstly by looking at the different approaches that have been taken in producing them.

One type of sensor node is built using general purpose computing components. These may be custom built using, for example, ITX [83] format components, or they may utilise personal digital assistants (otherwise known as PDAs) which incorporate sensors such as those for positioning. There are also purpose built nodes such as that from Sensoria – the Sensoria WINS NG 2.0 [84] node is displayed in Figure 3.6.

These devices can typically accommodate a wide variety of sensors and other peripherals since they are built around large and relatively powerful processors. Their reliance on common computing components means that they are compatible with many well defined networking standards, such as those used in WLAN, and they can use off the shelf software, thus reducing implementation time.

Figure 3.6 – Sensor WINS NG 2.0 sensor node [84]

Offering such compatibility with common computing standards and components comes at a cost as these general purpose nodes tend to be rather power hungry; therefore they are mostly only suitable for use when power is not an issue.

Moving away from such large devices, there are the concepts of embedded and system on chip (SoC) devices, which aim to be both small and low power; fitting more accurately the description of a wireless sensor node given at the start of this chapter. These devices are often referred to as motes.

Dealing with SoC motes first, these devices are not built from off the shelf components at all, since the designers attempt to integrate every component of the mote into a single integrated circuit. An example is the Intel iMote [85], which integrates an ARM7 microcontroller, Bluetooth transceiver, RAM, flash memory and I/O ports into a single chip. The advantage of this is that devices can be much smaller in size, and potentially use less power whilst being cheaper than embedded systems if created in large volumes [86].

However, for prototyping, embedded sensor node platforms or COTS motes – so called because they use Cheap Off The Shelf components – are reasonably cheap to make in smaller quantities. This makes them ideal for research projects such as this, where the aim is to produce prototype systems for demonstration purposes which will display similar characteristics to potentially smaller future devices. Such examples are μAmp [87], BT node [88], Firefly [89] and Berkeley's range of Mica [90] motes.

Each type of mote has its own unique features and aims; however this work focuses particularly on the Mica series of motes from Berkeley since they are readily available as a commercial product, are of a reasonably small size for demonstration to industry, and have a large support community.

Berkeley motes have their roots in work that was funded by the Defence Advanced Research Projects Agency (DARPA)[13], setting the goal of "a complete sensor/communication system which can be integrated into a cubic millimetre package" [91] – this aim gave rise to a concept known as smart-dust; a fictitious mote measuring $1mm^3$ and costing just $1. Significant progress has been made in attempting to achieve this goal, although the largest challenges for most devices remain to be the size of both the battery and antenna. Possibly the most significant progress can be seen with a mote named SPEC [92], pictured in Figure 3.7; this integrates a microcontroller as well as a radio transceiver. This is a significant step from previous devices of this size which have been using optical communication, although little has been done with regarding to reducing the size of vital components such as the antenna and battery power supply.

Figure 3.7 – SPEC compared to a ballpoint pen tip

[13] DARPA is an agency of the United States Department of Defence, and is responsible for the development of new technology for military use.

3.6. The Micax Berkeley Motes

3.6.1. The Hardware

The motes used in this research work are the Berkeley Mica2, MicaZ and Mica2Dot motes [93-94] – they are referred to collectively throughout this document as the Micax motes for simplicity. The Micax motes, when considered by modern computing standards, are quite limited in terms of their processing and memory capabilities. This limitation, as mentioned previously in Section 3.3.3, is simply a result of their small size whilst utilising cheap components. Therefore one must be aware of the hardware available on the motes in order to overcome or negate these limitations in practical applications.

Figure 3.8 shows an overview of the Micax motes in terms of their major components. At the heart of a mote lies the Atmel AVR[14] ATmega128L [95] programmable microcontroller, which affords part of the expandability referred to in Section 2.2.2. It has a number of features such as eight 10-bit analogue to digital channels (ADC), 4KB RAM and 128KB of program memory. This microcontroller is aimed squarely at low power computing markets, and in particular the 128L can be run at clock speeds of up to 8MHz, allowing it to operate at minimum voltage levels of 2.7V, although lower voltages can be used if the clock speed is reduced.

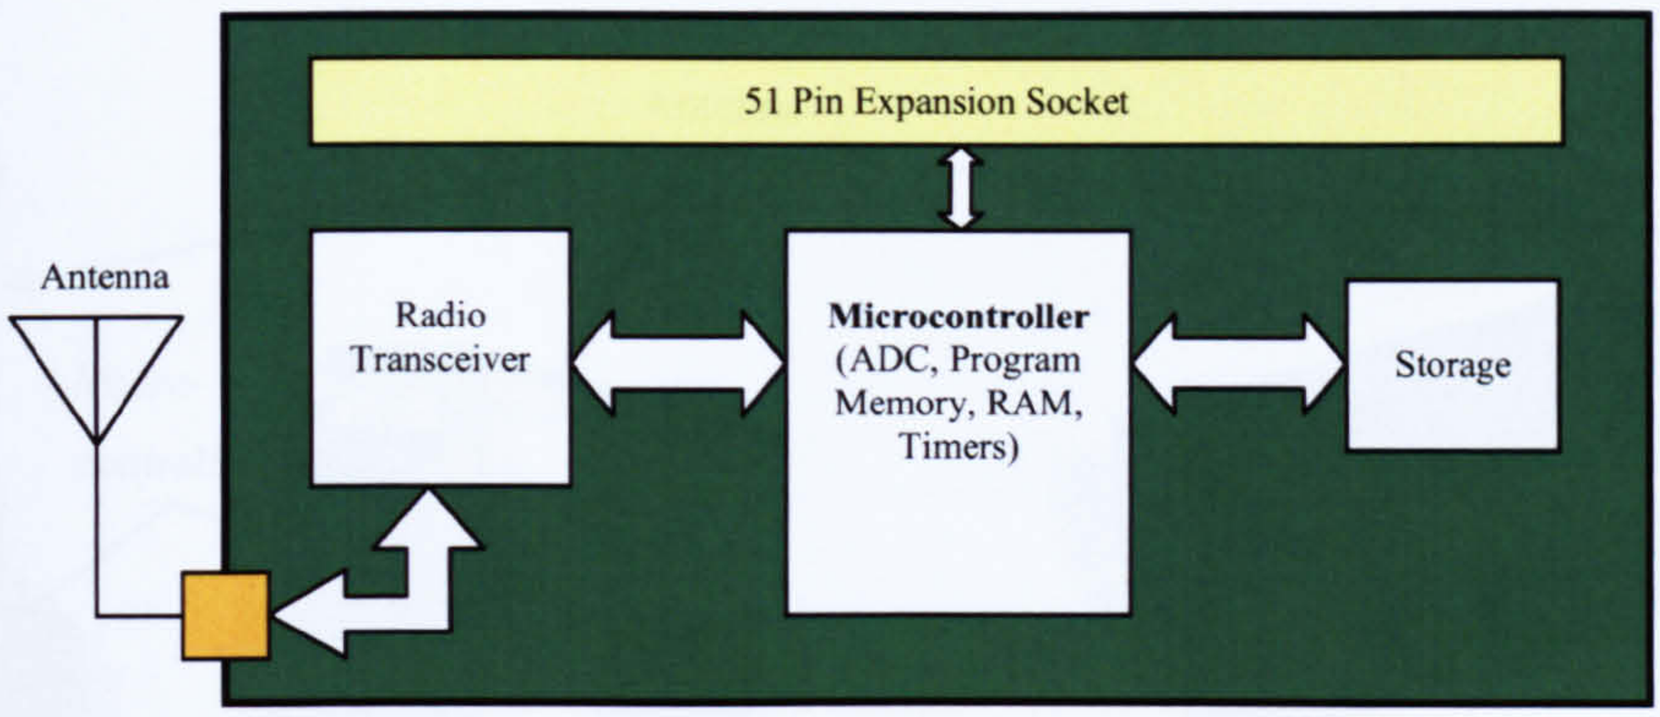

Figure 3.8 – Micax Hardware Overview

[14] Note that AVR is not an acronym, and doesn't represent anything in particular.

In contrast the ATmega128 can run at up to 16MHz, but requires a minimum of 4.5V to do this. Among its special features the ATmega128L includes a number of low power modes which range from idling (i.e. – shutting down the processor unit, but still allowing interrupts and the ADC to operate when needed), to standby mode which disables almost all the functions on the microcontroller so that it can only be awoken by an asynchronous event such as an interrupt.

The radio transceivers of the Mica2 (including the Mica2Dot) and MicaZ are the Texas Instruments CC1000 [76] and CC2420 [77] respectively. Dealing with the CC1000 first, it has a frequency range of 300-1000MHz; the motes used in this work are actually imported directly from the United States and use 915MHz since it is legal to do so there. Motes more suitable for UK use have also been produced and they typically operate at 433MHz. The CC2420 transceiver is built to comply with the Zigbee [96], hence the reason for the Mica**Z** mote, standard for 2.4GHz transceivers and therefore supports transmission speeds of up to 250kbps. Both transceivers have selectable frequency channels which range from 903-927MHz (1MHz intervals, 28 channels) for the Mica2 mote and 2405-2480MHz (5MHz intervals, 16 channels) for the MicaZ. This feature is included to avoid intereference from any existing wireless infrastructure which may operate at similar frequencies. This is particularly important in terms of the MicaZ motes which operate close to WLAN frequencies; there are at least four channels available which do not overlap however.

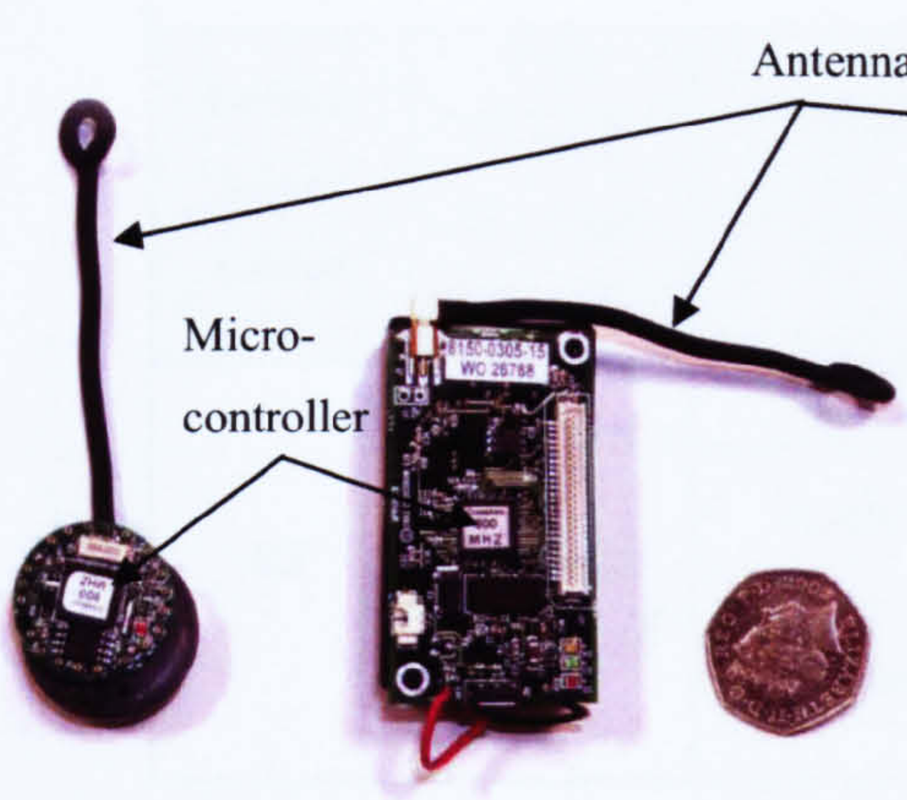

Figure 3.9 – Mica2Dot (left) and Mica2 (right)

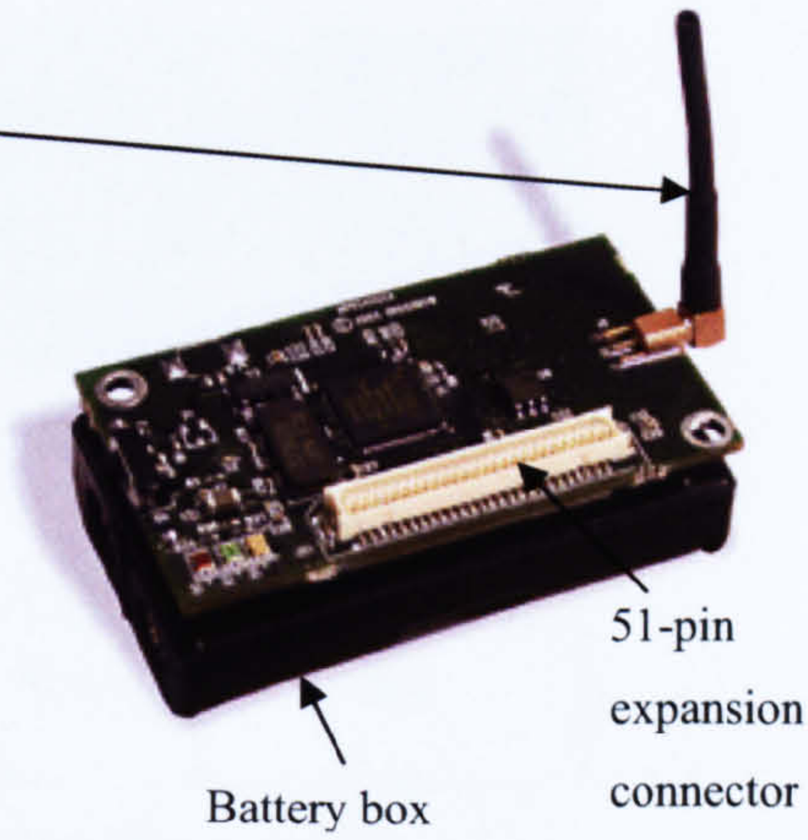

Figure 3.10 – MicaZ mote

In order to allow for data logging, 512KB of non-volatile storage is included on the Micax motes in the form of flash memory (Atmel AT45DB [97]). Other features include LED indicators, a battery compartment, MMCX (micro-miniature coaxial) antenna socket in addition to a surface mounted 51-pin socket [98] which enables the connection of a variety of peripheral devices, enhancing the motes expandability.

The Micax motes are illustrated in Figures 3.9 and 3.10. The Mica2Dot mote is almost identical to the Mica2 in terms of hardware, yet it is much smaller as can be seen from Figure 3.9. Throughout this work the larger Mica2 and MicaZ motes are used since their size makes them more convenient to work with on a daily basis, as well as them accepting much higher capacity (and cheaper) batteries. Table 3.2 serves to give an overview and comparison of these different motes in terms of their specification.

Mote Type		**MicaZ**	**Mica2**	**Mica2Dot**
MCU	Chip	ATmega128L		
	Type	8MHz		4MHz
	Program memory (KB)	128		
	RAM (KB)	4		
Non-volatile storage	Chip	AT45DB014B		
	Size (KB)	512		
Default power source	Type	2xAA		Coin Cell
	Typical capacity (mAh)	2850		1000
RF	Chip	CC2420	CC1000	
	Radio frequency	~2.45GHz	~915 or 433MHz	
	Raw speed (kbps)	250	78.4	38.4

Table 3.2 - Hardware specifications of Micax motes

3.6.2. Commercially Available Sensors

There are a number of ready made sensor boards [99] for the Mica2 and MicaZ motes. Figure 3.11 shows the MTS310 sensor board which includes a piezoelectric speaker, a thermistor for measuring temperature, a light sensor, a microphone, as well as an accelerometer for measuring orientation and a magnetometer for measuring magnetic fields. Figure 3.12 shows a further sensor board, the MTS420, which has a combined humidity and temperature sensor, a barometric pressure sensor, and a GPS module for positioning in outdoor locations. In addition, it also includes an accelerometer and light sensor.

Additional prototyping boards are available which allow the attachment of custom sensors and devices to the motes; these are shown in Figures 3.13 (MDA100) and 3.14 (MDA300).

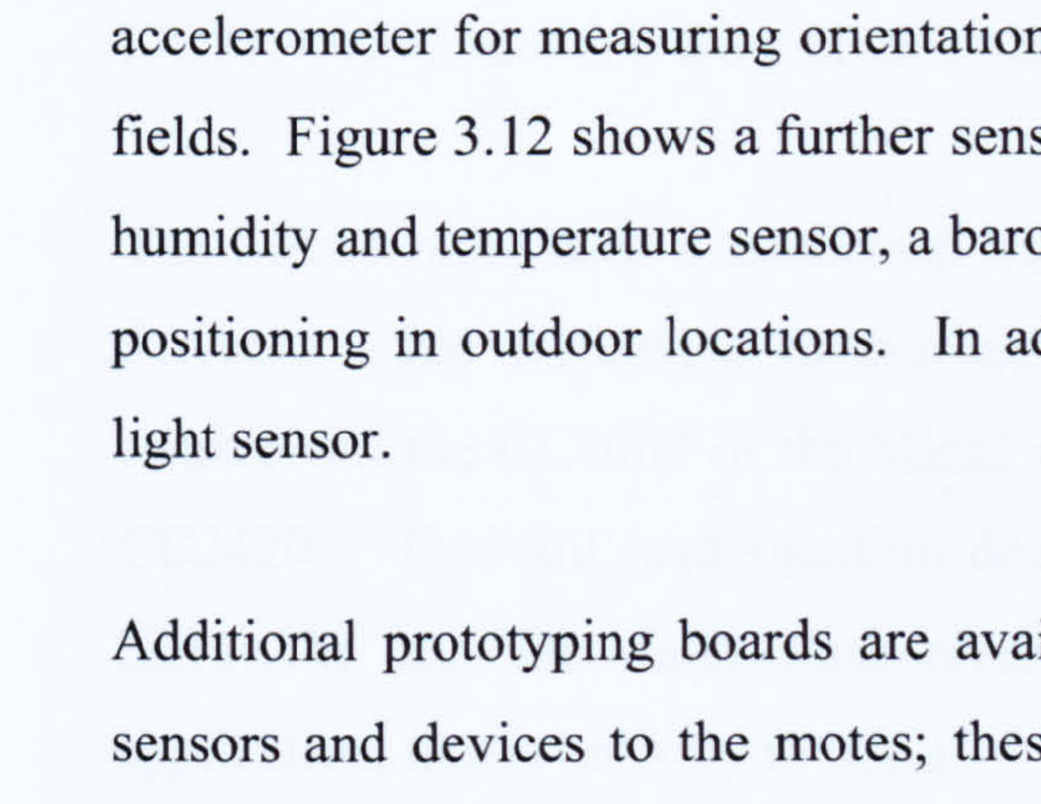

Figure 3.11 – MTS310

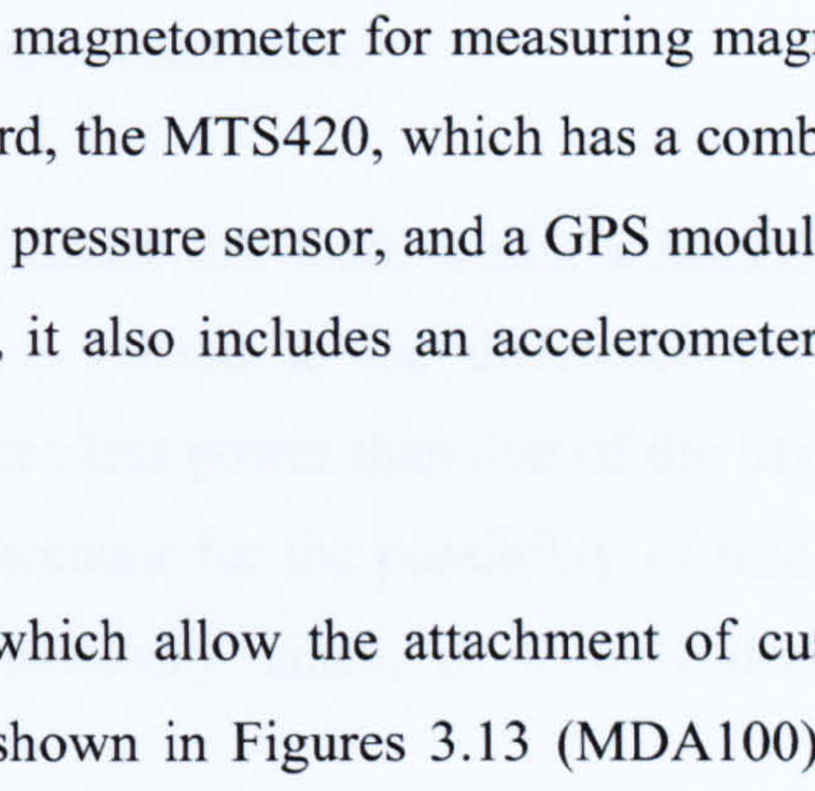

Figure 3.12 – MTS420

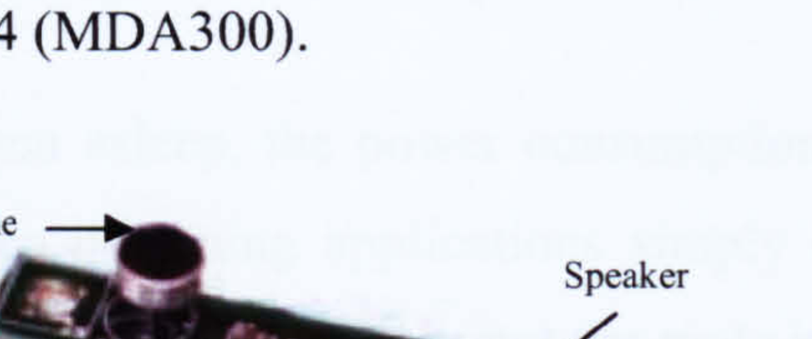

Figure 3.13 – MDA100

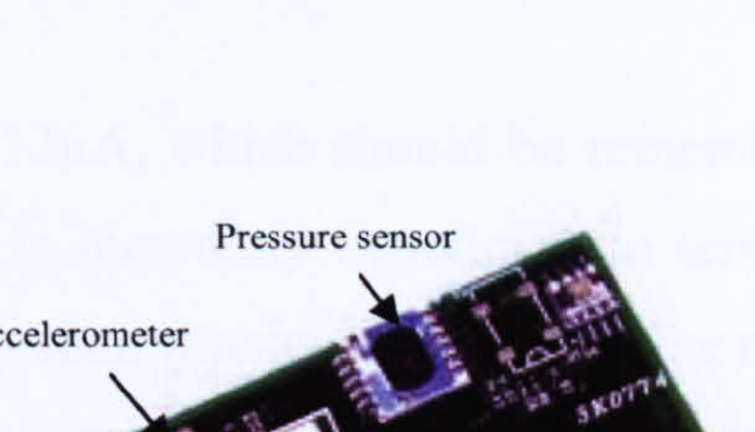

Figure 3.14 – MDA300

3.6.3. Power Consumption

The Micax motes are impressively economical in terms of their power usage, and are estimated by their manufacturer to have a battery life of 1.5 years at a 1% duty cycle. Table 3.4 gives some typical figures, taken from relevant data sheets and manuals [76,77,93,95], showing the current consumption of each mote.

The Mica2 mote uses a maximum of 22mA under full load, whilst a MicaZ uses 31.7mA. The difference in current usage is related to the difference in radio transceiver; the CC1000 of the Mica2 consumes less power than that of the MicaZ's CC2420. This full load situation does not account for the possibility of using the EEPROM and the radio transceiver simultaneously since, in a well designed application, this is unlikely to occur.

When asleep, the power consumption drops to 13μA, which should be remembered when designing applications simply because it is incredibly beneficial, in terms of energy conservation, to put the mote into its sleep state whenever possible. The motes can then wake up again at a set time interval or maybe upon detection of wireless activity.

Device or Condition	**Operating Current (mA)**	
	Mica2	**MicaZ**
CPU at full speed	12	
CPU in sleep mode	0.010	
Radio, receive	7	19.7
Radio, transmit (0 dBm power)	10	17
Radio, sleep	0.001	0.001
Serial flash memory, write	15	
Serial flash memory, read	4	
Serial flash memory, sleep	0.002	
Maximum power usage (full load)	22	31.7
Minimum power usage (sleep)	0.013	

Table 3.3 - Typical power consumption of Micax motes

Transmission Power (dBm)	Power Consumption (mA)	
	Mica2	**MicaZ**
0	16.8	17.4
-1	15.8	16.5
-3	14.5	15.2
-5	13.8	13.9
-7	10.8	12.5
-10	10.1	11.2
-15	9.3	9.9

Table 3.4 – Power consumption at various transmission power levels

It is worthy of mention here also that the motes can transmit at less than 0dBm output power, as well as transmitting at higher power too. The implication of this is that one can save energy by using as little power as possible to transmit. Alteration of the transmission power is possible by setting a register on each transceiver; a list of some of the possible values along with their respective energy consumption is shown in Table 3.4.

3.6.4. Hardware Interface

The hardware interface provided for programming and interfacing the Micax motes with a desktop computer comes in the form of various programming boards. Throughout this project the MIB510 and MIB520 [93] are used, with the major difference being that the MIB520 uses a Universal Serial Bus (USB) interface, whilst the MIB510 simply uses RS232. Obviously for more modern computers USB is essential since RS232 communications ports are not all that common nowadays. Both boards are pictured in Figure 3.15.

Connectivity to the motes is provided via their 51-pin socket, with the boards providing power to the motes whilst they are connected. The typical use for these boards, apart from programming, is collecting data from motes which have formed a WSN, as well as transmitting data out onto the network for the motes to act upon (i.e. – they can act as a gateway device).

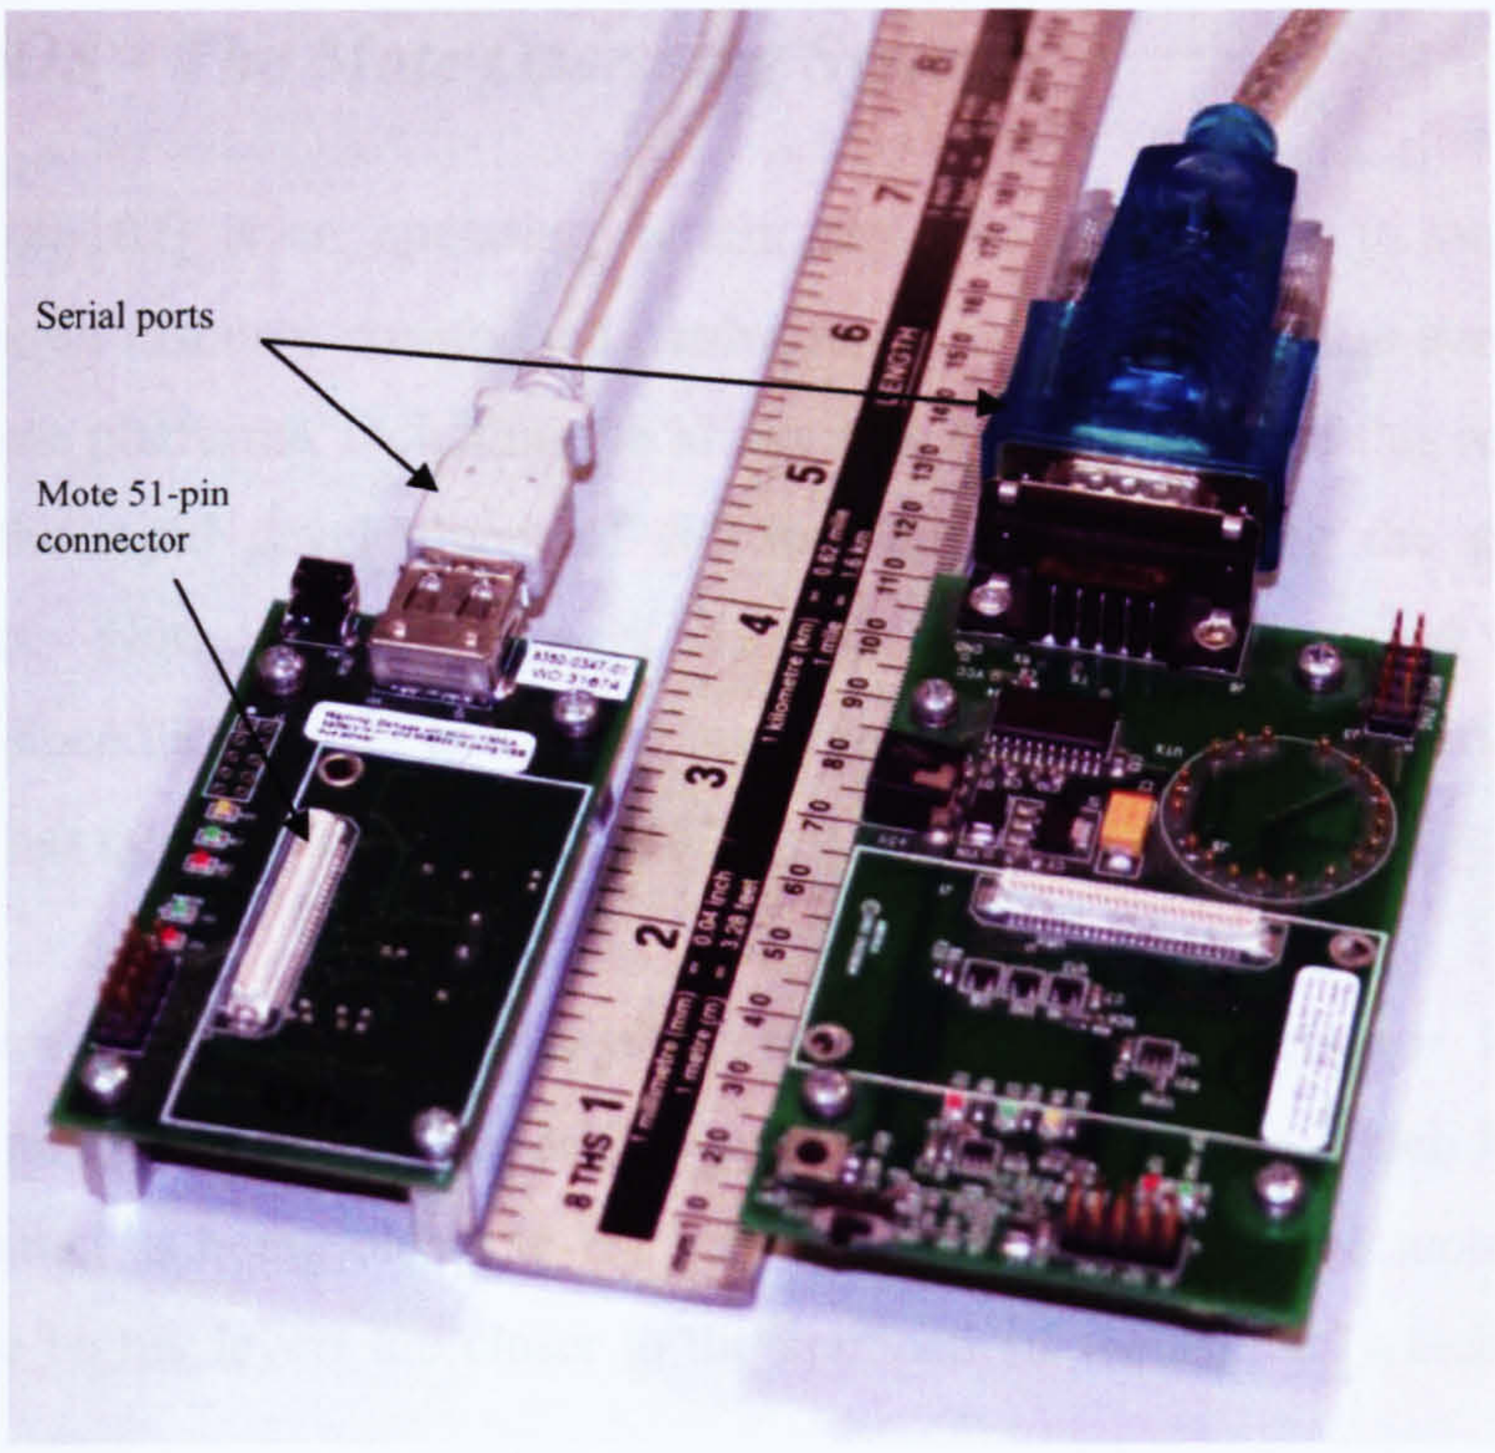

Figure 3.15 – MIB520 (left) and MIB510 (right) programming boards

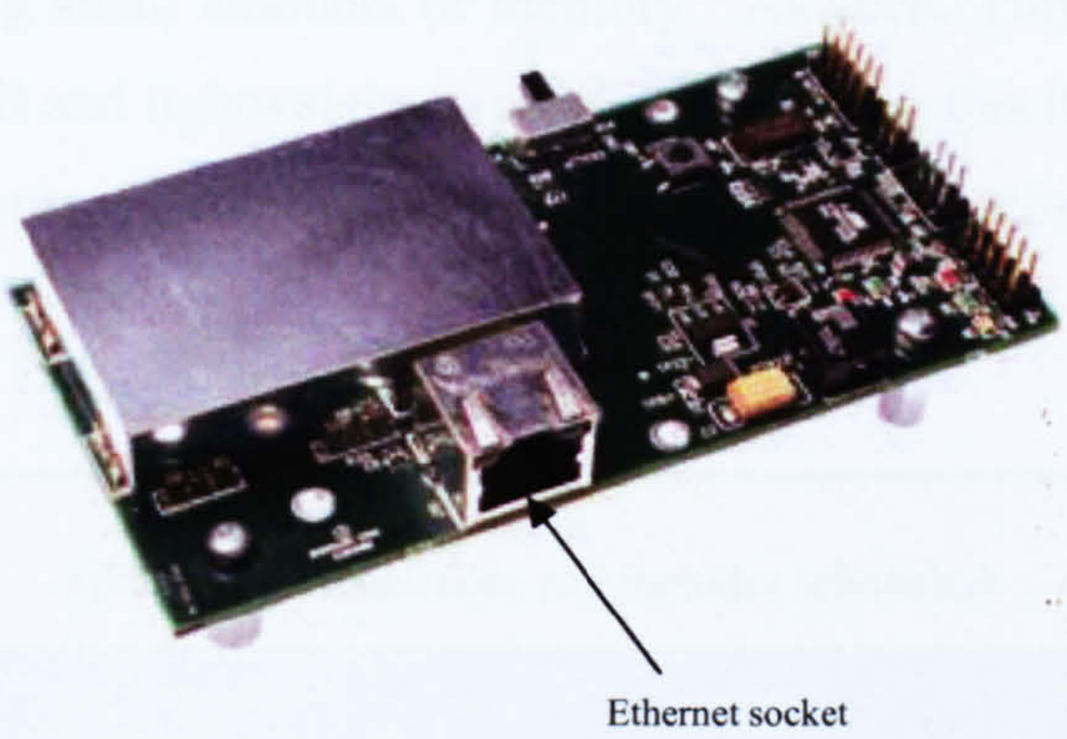

Figure 3.16 – MIB600 Ethernet interface board

In addition to the MIB510/520, there is also an MIB600 programming board, pictured in Figure 3.16, which provides a true Ethernet interface. This means that the board can be plugged directly into an Ethernet network socket and act directly as a gateway to a network or the Internet.

3.7. TinyOS – The Mote Operating System

TinyOS [100-102] is an operating system designed specifically to support WSN applications on resource constrained hardware. It is suitable for use across a broad range of mote platforms, including the Micax motes used as part of this research, and assists in the rapid development of software aimed at realising the goals of the programmer. Note that in this work, TinyOS version 1.0 is referred to – this is significant since there has been a recent major release of the operating system (version 2.0) which has numerous differences [103].

Many operating systems organise their components into layers (or levels), and TinyOS is no exception. Figure 3.17 shows its layer structure in which lower layers are represented as being closer to the physical hardware (i.e. – the mote platform), whereas the higher layers are closer to the applications themselves which run on the mote.

As discussed previously in this chapter, the motes have a number of hardware constraints, including small amounts of memory. As such, TinyOS implements the aim of being as small and lightweight as possible. In doing this it has had to sacrifice many features that one might find in a standard computer operating system.

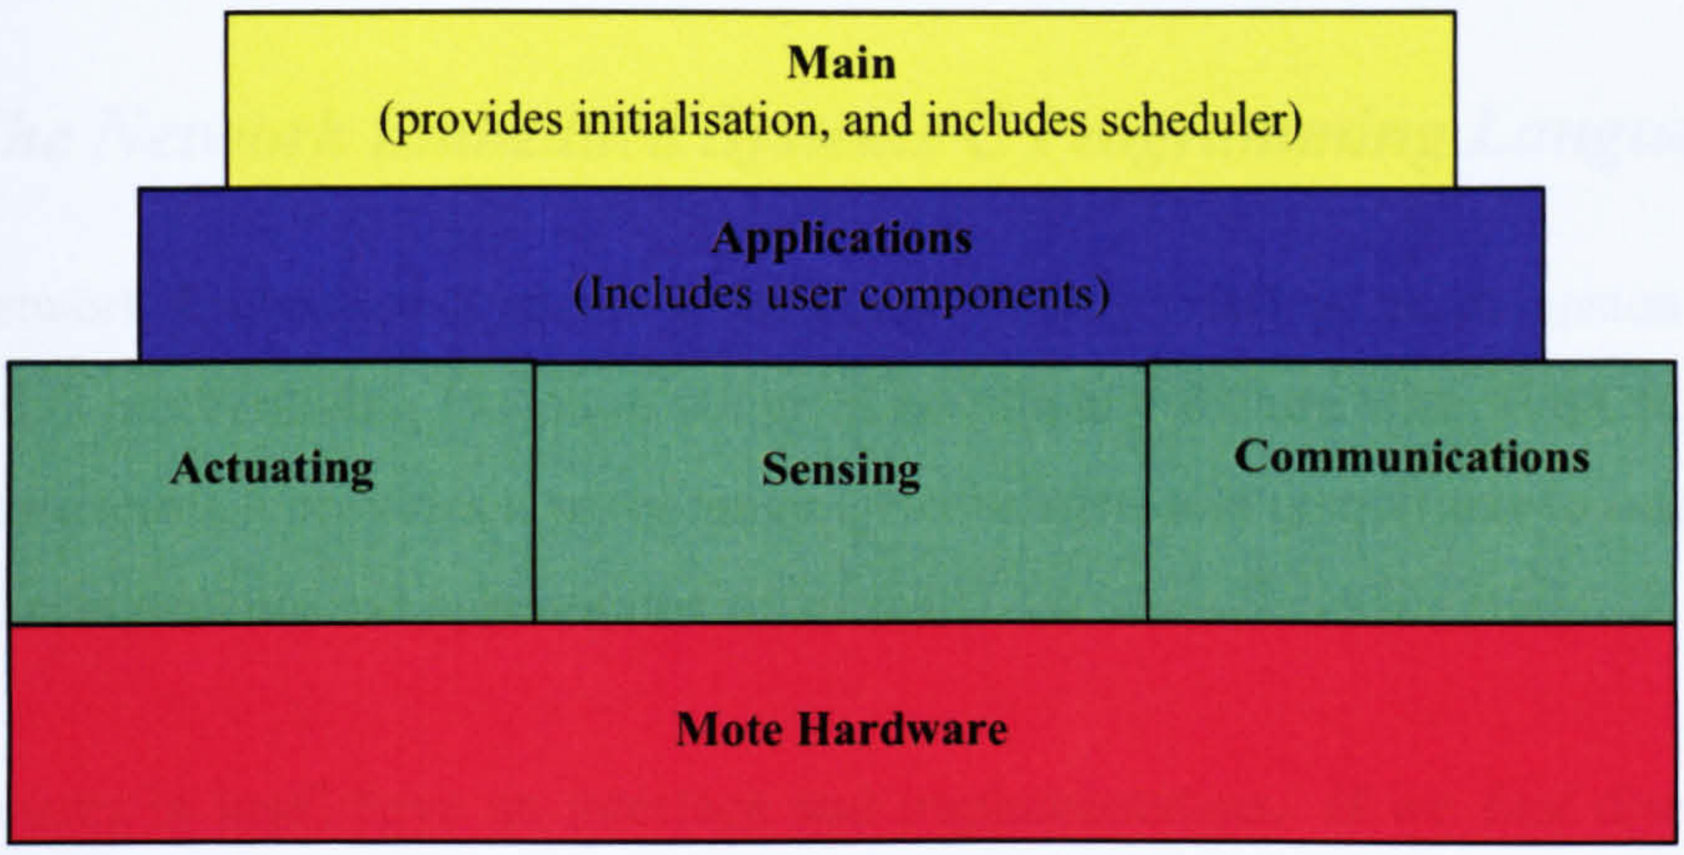

Figure 3.17 - TinyOS layer structure

Such sacrifices include there being no file system or even dynamic memory allocation which would be beneficial to manage memory more effectively. However the implementation of these features would be costly in terms of both complexity and size which in turn could imply increased processing time and memory requirements, negating any potentially positive effect in most cases.

A simple task model is implemented in TinyOS; however it is still able to meet the demands of mote applications in terms of concurrency. Concurrency refers to the ability of motes to be able to handle multiple flows of information simultaneously, something which motes may have to do when they are sensing and communicating. The task scheduler is able to handle two sources of concurrency; from tasks and events. Tasks are sections of code which can be posted to the scheduler, and always run to completion. They can, however, be pre-empted by events (which also run to completion) such as hardware interrupts which might be triggered by incoming communications or sensors. When there are no tasks to execute, the task scheduler puts the microcontroller to sleep in order to conserve energy.

There is also minimal device and networking abstraction, as well as a component based architecture being used. The later entails only the required parts of the operating system being compiled into each application. This assists in making resulting applications as compact as possible.

3.8. The Network Embedded Systems C Programming Language

The Network Embedded Systems C language (nesC) [104] is an extension of the standard C programming language designed specifically for use with TinyOS. In the most basic terms it provides a set of language constructs and restrictions to implement TinyOS components and applications.

Components in nesC have an interface and implementation. If we first discuss the idea of interfaces, they are classified in terms of `provides` or `uses`, which reflects the TinyOS layer model.

`Provides` interfaces deal with calls which can be made by applications at higher layers of this model, whilst a `uses` interface refers to calling components in a lower layer. Further to this, nesC distinguishes the directionality of an interface call between layers using the `event` and `command` keywords. An `event` call is a method call from a low layer component to a higher one; `command` achieves the opposite. There are many standard interfaces included with TinyOS, which promotes reusability and aids in keeping memory usage to a minimum.

Moving finally to the implementation of components, there are two types: modules and configurations. Modules are implemented by application code written with C-like syntax. Configurations on the other hand are implemented by connecting interfaces of different components. There may be several instances of a single component interface used through the declaration of abstract components. This allows, for example, multiple timers to be used in the same application despite there only being one actual hardware clock. It is important to remember that TinyOS does not support dynamic memory allocation so all components are statically constructed at compile time.

3.9. Wireless Sensor Networks vs. Existing Technologies

In order to show the capabilities of WSN in comparison to those technologies discussed in Chapter 2, Table 3.5 compares them in different ways. This table shows that, at the moment, WSN's are mostly on par with technologies based around WLAN and UWB. One should keep in mind that the statistics given for WSN's are based upon the Micax motes being used. The aims of the technology are given as a reminder that the Micax represents a convenient research platform, but in reality the technology to support WSN's is still in its infancy. The future for the technology holds much more, a glimpse of which has been shown through the SPEC mote. The future aim of $1 per node would actually put WSN quite close in terms of cost to passive RFID technology whilst allowing it to offer significantly more in the way of features. A cost such as this would be attractive to those interested in asset tracking as well as those in the automotive industry. In addition, the aim of a size of $1mm^3$ would make the technology the smallest of all available which could be particularly beneficial for asset tracking.

	Cost	Range	LOS?	Size (mm)	Battery Life	Reusable?
Bar codes	Much less than £0.01	Up to 3m	Yes	35×25	n/a	No
RFID (Passive)	£0.03+	Up to 3m	No	Varies greatly	n/a	No
WLAN	£50	Over 100m	No	62×40×17	1-3 years	Yes
UWB	£100	Approx 40m	No	58×92×10	1 year	Yes
RuBee	£5	5-10m	No	25×25×4	Up to 2 years	Yes
WSN	£55; aim of £0.50 / $1	20-30m; unlimited	No	57×32×15; aim of 1×1×1	1.5 years with typical AA	Yes

Table 3.5 – Technology comparison

Table 3.5 also mentions reusability, a subject which has not been touched on before in this work. Bar codes and RFID tags are designed to be placed on an asset, and disposed of when the asset is finished with. As a result they are usually embedded in or on the asset, or attached via some means which makes reusing them impractical. The added cost of WSN hardware may encourage industry to find ways of reusing the technology far beyond the life of an asset. In luggage handling, a mote could be placed on a bag at check-in and removed when the luggage is claimed at the arrival airport by the passenger. With gas cylinders, a mote could be placed on a cylinder as soon as it is manufactured and removed when it is found to be no longer serviceable. The same also applies to the automotive industry, and it means that the actual cost of the technology may not be as much as it first seems since the organisations involved may not have to keep constantly replacing a disposable technology.

Battery life is important, and it is clear that all the battery operated technologies here have a comparable life. It is worth noting that the numbers quoted are taken from manufacturer specifications so it is difficult to know the extensiveness of testing which allows these predictions, or the duty cycles that they can be associated with. Chapter 7 looks at the battery life of the Micax motes in order to show their feasibility for industrial use.

The final, and possibly most important point, is that WSN have a potentially unlimited range through the use of multihop communication. This sets WSN and motes apart from other technologies which effectively have a limited range; from tag to reader.

4. Inventory Management System

4.1. Introduction

Whilst having an idea of how to improve asset tracking through the use of an alternative concept (i.e. – WSN's) is good first step, it is important to carry this idea on into practise. This chapter looks at the design and implementation of a prototype inventory management system, utilising the Micax motes, which is demonstrable to industry.

Asset management in the airline and packaged gas industries has requirements common to both. For example they both have:

- Moving assets, whether it be luggage moving through airports or gas cylinders being moved from one place to another.
- Situations where assets do not move at all – for luggage this would occur in an aircraft cargo hold and for gas cylinders when they are stored at a storage centre.

As a result, it was decided that a generic prototype system would be advantageous. A generic system can be shown to any industry regardless of their area of interest; it can give them an idea of the capability of the technology, and demonstrate it working. Not only was this thought to be useful in terms of saving time within the project itself, it meant that future demonstrations could be possible to industries which have not yet been identified as having an interest in this inventory management method. Of course, as will be seen in Chapter 7, application specific features are important and so the prototype was designed with extendibility in mind so that features specifically applicable to the airline and packaged gas industries could be demonstrated.

A key restriction of this prototype is the scale of its demonstration. For the purposes of this project there are five Mica2 and three MicaZ motes available for use; the Mica2 and MicaZ motes are not compatible for use together since they operate at

different frequencies. As a result, the concept of the system is proven in practise with just five Mica2 motes. Despite this limitation, the design of the prototype takes into account many features which would be vital to a large scale deployment.

4.2. System Overview

The prototype system is made up of a number of different key components; ID tags and relay nodes, a Graphical User Interface (GUI) with an associated base station, and a Hand Held Unit (HHU). An overview of each of these components is given in this section, and Figure 4.1 gives an idea of how they are linked together to form this prototype inventory management system.

Each asset requires identification and so should be individually tagged. The tags in this prototype system are Mica2 motes, which are programmed with a TinyOS application allowing the unique identification of any asset to which they are attached. Each tag can communicate both directly and indirectly (i.e. – via multihop communication) with its neighbouring tags, the GUI and the HHU. Relay nodes are also included with the purpose of providing additional network coverage between the GUI and the tags. In practise the relay nodes are identical to the tags in this system.

A HHU has been created to complete the system. This portable device allows a user to interact with the tags remotely – it has controls, a Liquid Crystal Display (LCD) and its own power supply.

Finally, a GUI is included to give the user a clear representation of the assets; it allows the tracking of assets as they move, as well as keeping a record of the current inventory at each appropriate location. The GUI instigates interrogation of the tags, as well as allowing communication with the HHU. The GUI interfaces with the tags and the HHU via a base station. The two are inseperable, as the base station serves as the interface (or gateway) between the GUI and the other components. As a result, the remainder of this document refers to data as being sent or received by the GUI for simplicity.

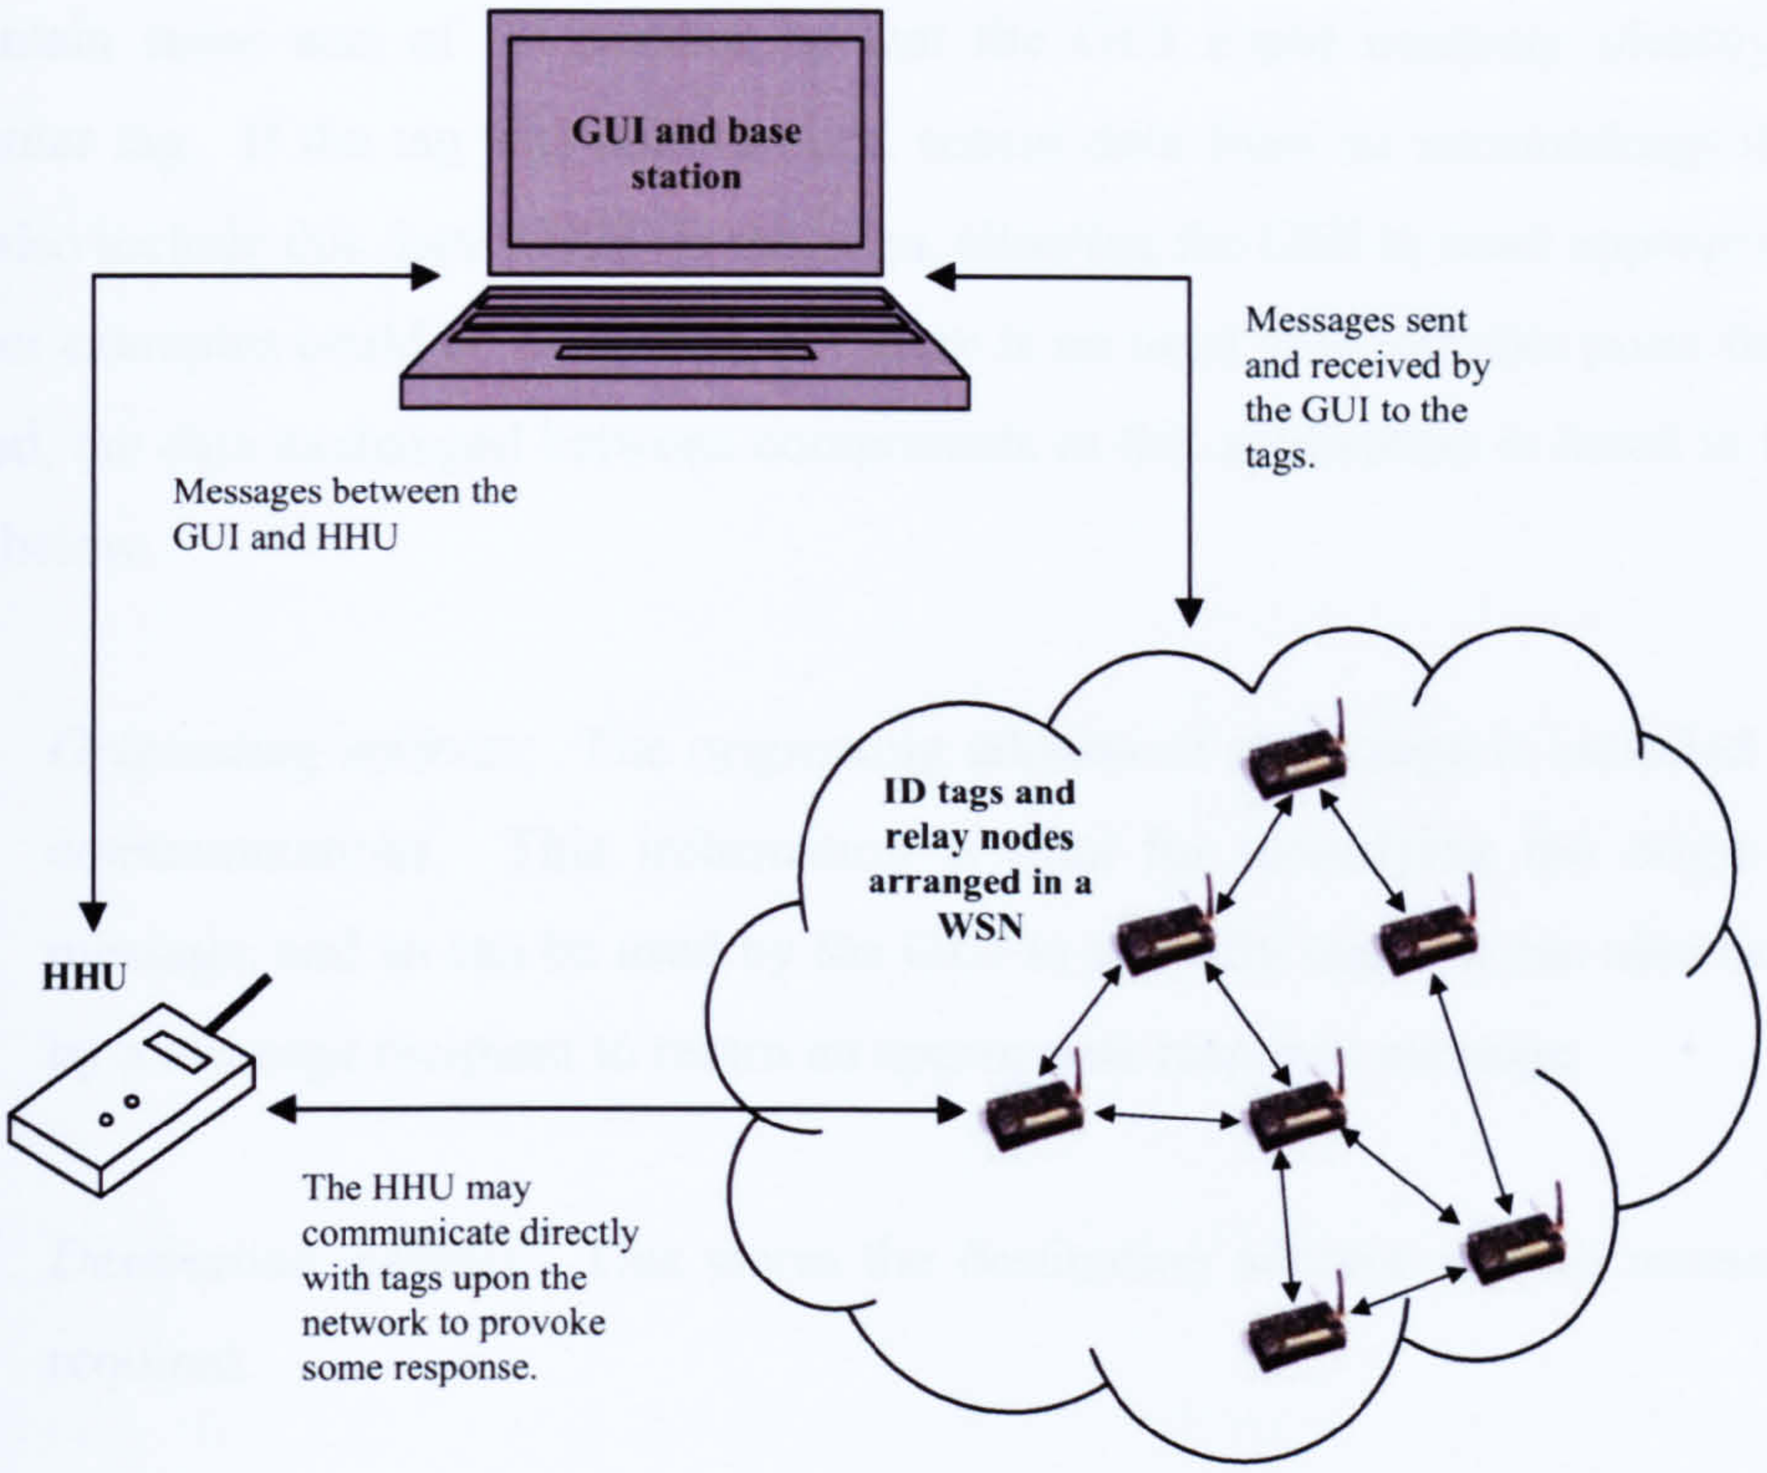

Figure 4.1 – Prototype inventory management system overview

Each component has its own unique role in this prototype system. It is the intention of this chapter to look at how each is implemented. In addition to detailing the methods of implementation, justification of the design decisions are discussed wherever appropriate. Before this, however, the communication protocols and data structures are described. The reason for doing this is to highlight the actual data communicated between each component, as well as to prevent duplication of this information throughout the chapter.

4.3. Inter-Device Communications

4.3.1. Message Data

Communication between the different system components is facilitated through the exchange of messages, which are formatted into data structures known as packets. These messages contain important data which allows the components to perform certain operations. For example, a message from a tag to the GUI would be required

to contain some sort of ID number so that the GUI could uniquely identify that particular tag. If the tag was able to elicit sensor data from its surroundings then it may also include this data within its message, allowing the GUI to react appropriately. Further examples could be explained, but there is no need to labour this point further. Instead, the data exchanged between components in this application is listed in bullet form below.

- *Originating address*: The originating address of a message is included in all communications. This information is vital for identifying the origin of a message, and so can be used by the GUI to identifty tags – it can also be used by a message recipient to return an appropriate response message.

- *Destination address*: This stores the destination address of each message, if required.

- *Intermediate node address*: Since messages must sometimes propagate via intermediate nodes, intermediate destination and origin addresses are required; these are stored in this field whenever appropriate.

- *Message type*: This stores the actual message type, and is discussed further in Section 4.3.2. This information is used to define what a component should do when it receives a message.

- *Battery voltage*: This information is specifically present for the tags so that they can report the status of their battery power supply. This is important, as tags with failing battery supplies will require maintenance to ensure that they do not fail.

- *Other data*: A section of each message is reserved for any additional data which may be exchanged. With the prototype system, it is used for sensor data.

4.3.2. Message Types

As mentioned in the previous section, a message contains some information regarding its type. This is important to the successful operation of the prototype system because each message type causes a component to perform a different operation. Table 4.1 defines these different message types. There is only one message type which deals with inter-tag communication, and that is `TAG_WAKEUP`. This message type is received by a sleeping tag to allow it to awake; it then retransmits this message to wake up its neighbouring tags. The reason for having this message type is discussed later in Section 4.4.3.

Messages which deal with communication between the GUI and tags are prefixed with `GUI_*`. A `GUI_REQUEST` message type is sent by the GUI to all tags in order to provoke them to send a `GUI_RESPONSE`, which allows the GUI to collect information regarding the tags which are presently available. In addition to this, a `GUI_PING` type message is included, which allows the GUI to send a message destined for a particular tag. All of these message types are allowed to use multi-hop communication methods to propagate.

`type` field value	Meaning
`TAG_WAKEUP`	Message sent by a tag which will wake up all surrounding tags.
`GUI_REQUEST`	Request from GUI for tags to respond so those present can be identified.
`GUI_RESPONSE`	Response from tags, triggered by `GUI_RESPONSE`.
`GUI_PING`	GUI sends this to a specific node to force it to perform some action.
`HHU_REQUEST`	HHU sends this as a request to a specific tag to respond.
`HHU_RESPONSE`	Tag response to `HHU_REQUEST`.
`HHU_UPDATE`	Message from the GUI to the HHU to update some data.
`HHU_UPDATE_ACK`	HHU acknowledgement of `HHU_UPDATE` to the GUI.

Table 4.1 – `type` field values

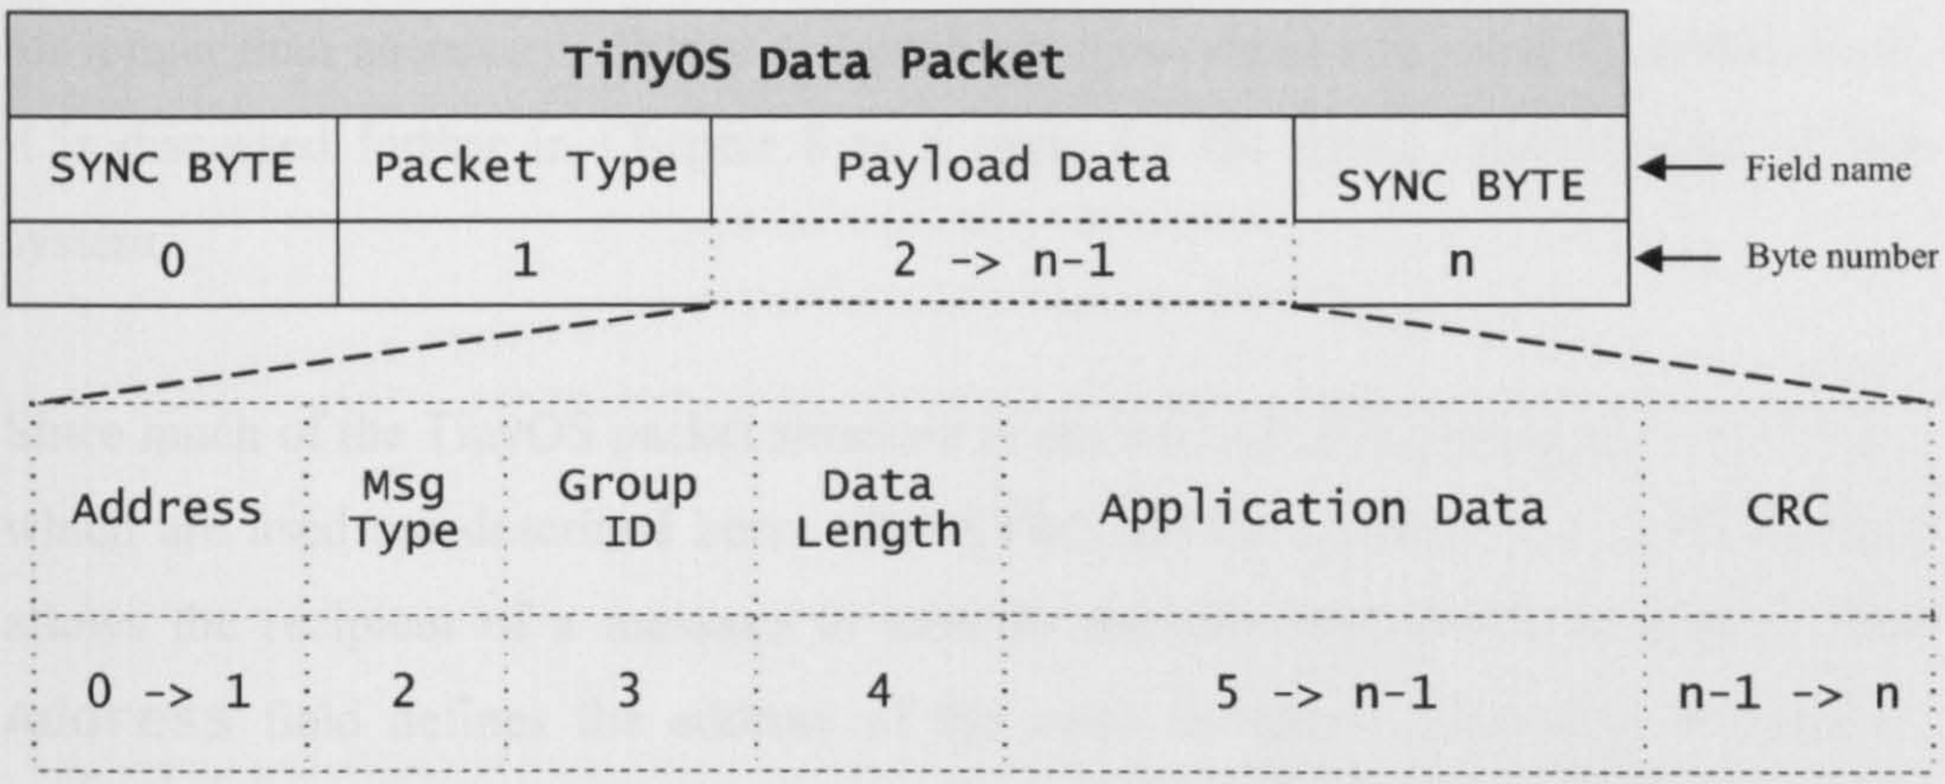

Figure 4.2 – TinyOS packet structure

The other message types listed in Table 4.1 are specific to handling messages relating in some way to the HHU, and are given the prefix HHU_*. The HHU_REQUEST message is sent from the HHU to a particular tag (not to all, like the GUI) to provoke it to send back a HHU_RESPONSE. The HHU_UPDATE and HHU_UPDATE_ACK type messages allow data to be transferred between the GUI and HHU; HHU_UPDATE messages are directed from the GUI and contain some tag data for storage by the HHU. The HHU_UPDATE_ACK messages propagate in the opposite direction and serve to allow the HHU to acknowledge the receipt of HHU_UPDATE message types. None of these messages are allowed to propagate by multihop methods since they are intended to be sent directly between the HHU and a tag or the GUI.

4.3.3. Packet Structure

Since the content of each message has been discussed, as well as the different types of message, it now remains to present the packet structure which is used. The TinyOS operating system has its own generic packet structure [105] which is outlined in Figure 4.2. The TinyOS packet structure allows for great flexibility when prototyping as different applications have different data requirements which means that data encapsulated within the TinyOS packet may have a different structure and size. This flexibility does come at a cost however, as the TinyOS structure is 10 bytes in size without further application data being added. These additional bytes represent wasted energy and bandwidth since they force the mote transceiver to be actively transmitting

for longer than necessary. Whilst this problem is not dealt with directly in this work, it is discussed further in Chapter 8 as a topic for the future improvement of this system.

Since much of the TinyOS packet structure is not utilised in this prototype, only those which are used are described here. The `SYNC BYTE` is always set to `0x7E`, and allows the recipient of a message to identify the start and end of a packet. The `Address` field defines the address of the mote to which a message is targeted. Whilst this is not used directly by this application to address data to specific nodes, it is used to send messages either via the transceiver or the serial port. This involves setting the field to either `0xFFFF` or `0x007E` respectively. The later is used by the base station to forward messages to the GUI, whilst the former allows the broadcasting of messages wirelessly. The final data from the TinyOS structure which is used is the `CRC` (Cyclic Redundancy Check) field. This allows TinyOS to check the validity of incoming data by applying an algorithm [105] which generates a 16-bit number. This number is placed in the CRC field by the sender; the receiver then generates its own number from the received packet and compares this with the one in the packet. If they match, the packet is assumed to be valid; if they do not match then the packet is discarded.

The message data described in Section 4.3.1 is placed in the `Application Data` field of the TinyOS packet; its format is shown in Figure 4.3, where:

- `origin` = Originating address
- `dest` = Final destination address
- `inter` = Intermediate node address
- `type` = Message type
- `battery` = Battery voltage
- `data1` = Other data

Application Data					
origin	dest	inter	type	battery	data1
0 -> 7	8 -> 15	16 -> 23	24	25 -> 26	27 -> 28

Figure 4.3 – Application data structure

Messages incorporating this packet structure are typically 39 bytes in length; 29 bytes of this represent the application data. The key word typically is used here since a TinyOS packet can include escape characters for reserved bytes. The SYNC BYTE is an example of such a reserved byte because it is required for packet synchronisation by the receiver. As a result, if the payload data contains this byte it is altered by XOR'ing the payload byte (0x7E) with 0x20 and is then appended to 0x7D, to give 0x7D5E. This allows it to be identified by the receiver. It also means that one should avoid the use of reserved TinyOS characters if possible, in order to reduce the likelihood of packets becoming larger than necessary.

4.3.4. Message Propagation

All GUI_* and TAG_* message types are allowed to propagate via a multihop method. HHU_* messages use direct communication, which means that the HHU and the tag must be in direct radio communication range. Hence, if a tag, for example, receives a HHU message that is not addressed specifically to it then the message will simply be discarded and no further action will be taken.

The main aim of this section is to look specifically at the multihop communication technique used. Many network systems make use of routing algorithms in order to facilitate the propagation of data within a network. The responsibility of the routing algorithm is to select intermediate nodes which can act to relay messages between the source and the destination nodes. The problem is that many routing algorithms insist upon prior knowledge about the network in order to select routes. Proactive methods, such as Destination-Sequenced Distance Vector (DSDV) routing [106], seek to ensure that every node knows a route to every other node on a network. This is useful for low latency networks because if a node wishes to communicate with another then it

already knows the route along which a message must propagate. However, the memory requirements for storing all of these routes can become excessive, as too can the overheads required to maintain routes since nodes must exchange messages simply to provide up-to-date route information. Therefore, proactive methods are unsuitable for low memory devices, such as the Micax motes, and would waste energy through the need to provide route updates. In addition, there is no guarantee that route information would be up-to-date if a tag was to be moved which is to be expected when it is attached to an asset.

An alternative is reactive, or on-demand, routing. An example of this is Ad-hoc On-Demand Distance Vector (AODV) routing [107]. This takes an alternative approach to routing because it does not require nodes to store routing information. Instead, when a path is required it is discovered. It can then be discarded once used. This reduces the need for nodes to have large amounts of memory and also removes the need for nodes to be frequently exchanging messages containing updated route information. This concept was thought to be ideal for this work, since it means that if no data is required then the tags can lay dormant and save their finite energy supply.

Since this work borrows a little from the AODV routing concept, now is an appropriate time to be a little more specific about how AODV implements route discovery. A route is the path through a network from a source node to a destination node, most likely via some intermediate nodes. To initiate route discovery the source node broadcasts a message onto the network. This broadcasted message is destined for a particular node (i.e. – the destination node) but, since it is a broadcast, it may be received by all nodes within communication range. Each node will decide whether it was the intended recipient of the broadcasted message and record the address of the node from which the message was received. If the node is not the intended recipient it will rebroadcast the message; if it is the intended recipient it will send a response message – not a broadcast – back to the node from which the broadcast was received. That node will then record the address of the sending node and, provided it is not the source, will send the response to the node from which it received the original broadcast message; this process continues until the response message reaches the source node.

During this process, two routes are being set up. The first is a reverse route, since the nodes which receive the original broadcast message store the address of the node from which it was received – that is – the node one hop[15] closer to the source than they are. In doing this, a route from the destination to the source is established. When the response message from the destination traverses this reverse route, the forward route is set up (from source to destination); nodes are now recording the addresses of nodes closer to the destination. Hence, when the response message is received by the source both forward and reverse routes exist between the source and destination which allows them to begin communicating.

Despite AODV routing having the potential to operate on nodes with small memory capacities and finite energy supplies due to its on-demand nature, it is still inefficient in the context of this work because there is no need for a forward route. This can be highlighted by an example scenario whereby the GUI broadcasts a `GUI_REQUEST` message. As this message propagates through the network reaching every available tag, either directly or via intermediate tags, a reverse route can be constructed as is the case in AODV routing. The `GUI_RESPONSE` message from each tag can then propagate along this reverse route and back to the GUI where the inventory can be analysed and stored. There is no immediate need for any further communication – therefore no forward route is required. The reverse routing described in this section is illustrated in Figure 4.4 and pseudo code is provided in Figure 4.5

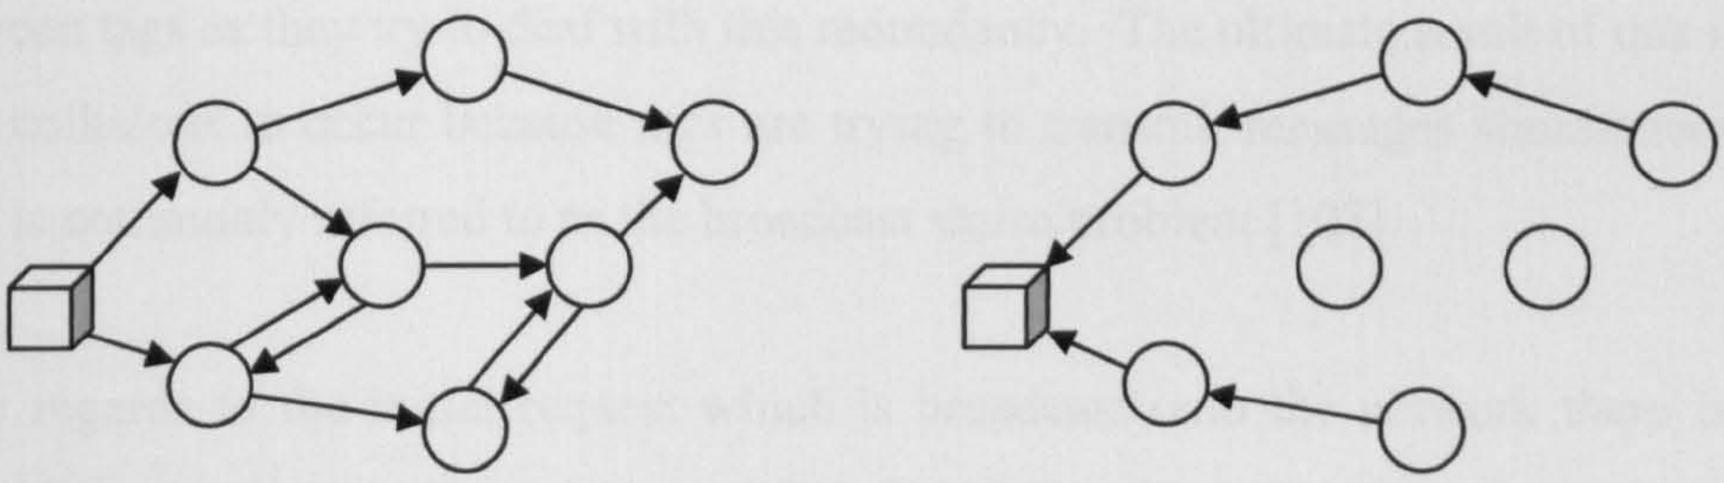

Request broadcast out to tags from GUI

Response propagation using reverse route

Figure 4.4 – Request and response message propagation

[15] A single hop in a network route is the step from one node to the next.

```
if received(msg) = FALSE then     // loop prevention
begin
  received(msg) := TRUE;          // loop prevention
  if msg->type = GUI_REQUEST then
  begin
    reverse_route := msg->inter;
    send_msg(GUI_REQUEST, broadcast);
    send_msg(GUI_RESPONSE, reverse_route);
  end
  else if msg->type = GUI_RESPONSE then
  begin
    if msg->final_dest <> local_id then
      send_msg(GUI_RESPONSE, reverse_route);
    else
      process(msg);
  end;
end;
```

Figure 4.5 – Reverse multihop routing pseudo code

One can see from Figure 4.4 that broadcasting causes a significant number of messages to be transmitted. Figure 4.6 serves to further highlight this issue by simply looking at the case were a single message is sent out on a network were all the nodes are exactly one hop from the source node, and can all communicate with each other, error free. With a network consisting of 10 nodes, broadcasting generates almost fives times more messages than the reverse routing technique used in this system; once 50 nodes is reached this increases to nearly 25 times more sent messages! This indicates that broadcasting is acceptable for transmitting a single on-demand request to many tags, but allowing the tags to broadcast their response causes a significant number of messages to be generated for propagation through the network. This leads to a high level of redundant packets on the network which would increase contention between tags as they try to deal with this redundancy. The ultimate result of this is for data collisions to occur because tags are trying to transmit messages simultaneously. This is commonly referred to as the broadcast storm problem [108].

With regards to the initial request which is broadcast onto the network there is the possibility that it can be retransmitted indefinitely by the tags unless they take some preventative measures. As a result, this system allows each tag to a keep a record of its most recently received messages so that it can disregard those which have already been rebroadcast. This is demonstrated in Figure 4.5 – see the loop prevention comments.

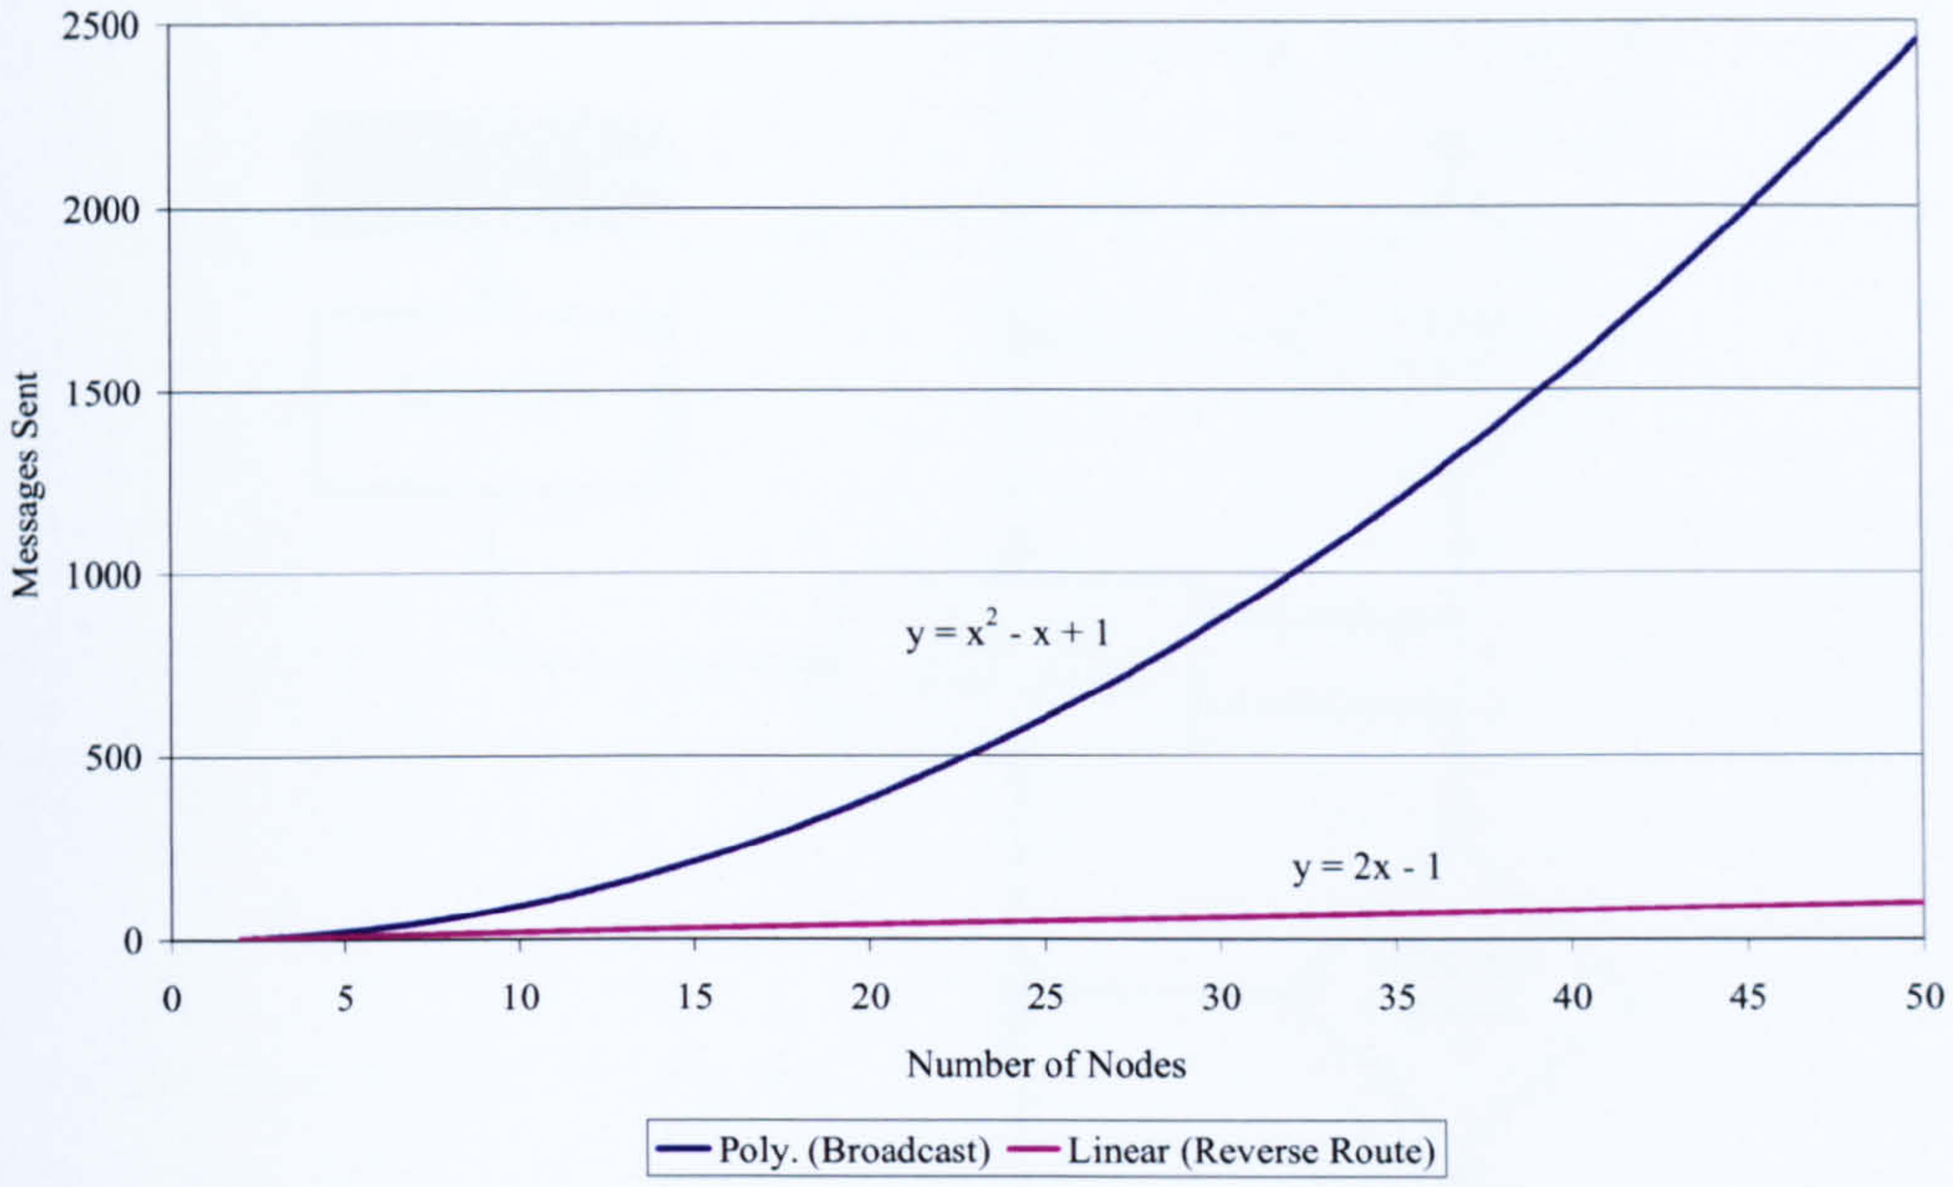

Figure 4.6 – Broadcasting vs reverse routing in terms of messages sent

The use of reverse routing in this system attempts to minimise the number of messages generated through nodes issuing responses, and therefore reduces the energy consumption as well as attempting to overcome the broadcast storm problem. Further issues discussed in Section 3.3, such as limited bandwidth, are also dealt with since fewer messages mean reduced bandwidth consumption. In addition, tags placed on mobile assets should be catered for adequately since the time between a reverse route being created and a response being issued should be minimal.

4.4. ID Tags

4.4.1. Overview

Now that both the data and the mechanism for messaging between each component of the prototype system has been defined, their implementation can be described, starting with the ID tags. The hardware for the ID tags is the Mica2 mote; this section looks specifically at the implementation of a TinyOS application for these devices. An overview of the TinyOS application is given in Figure 4.7 in the form of a brief flow diagram.

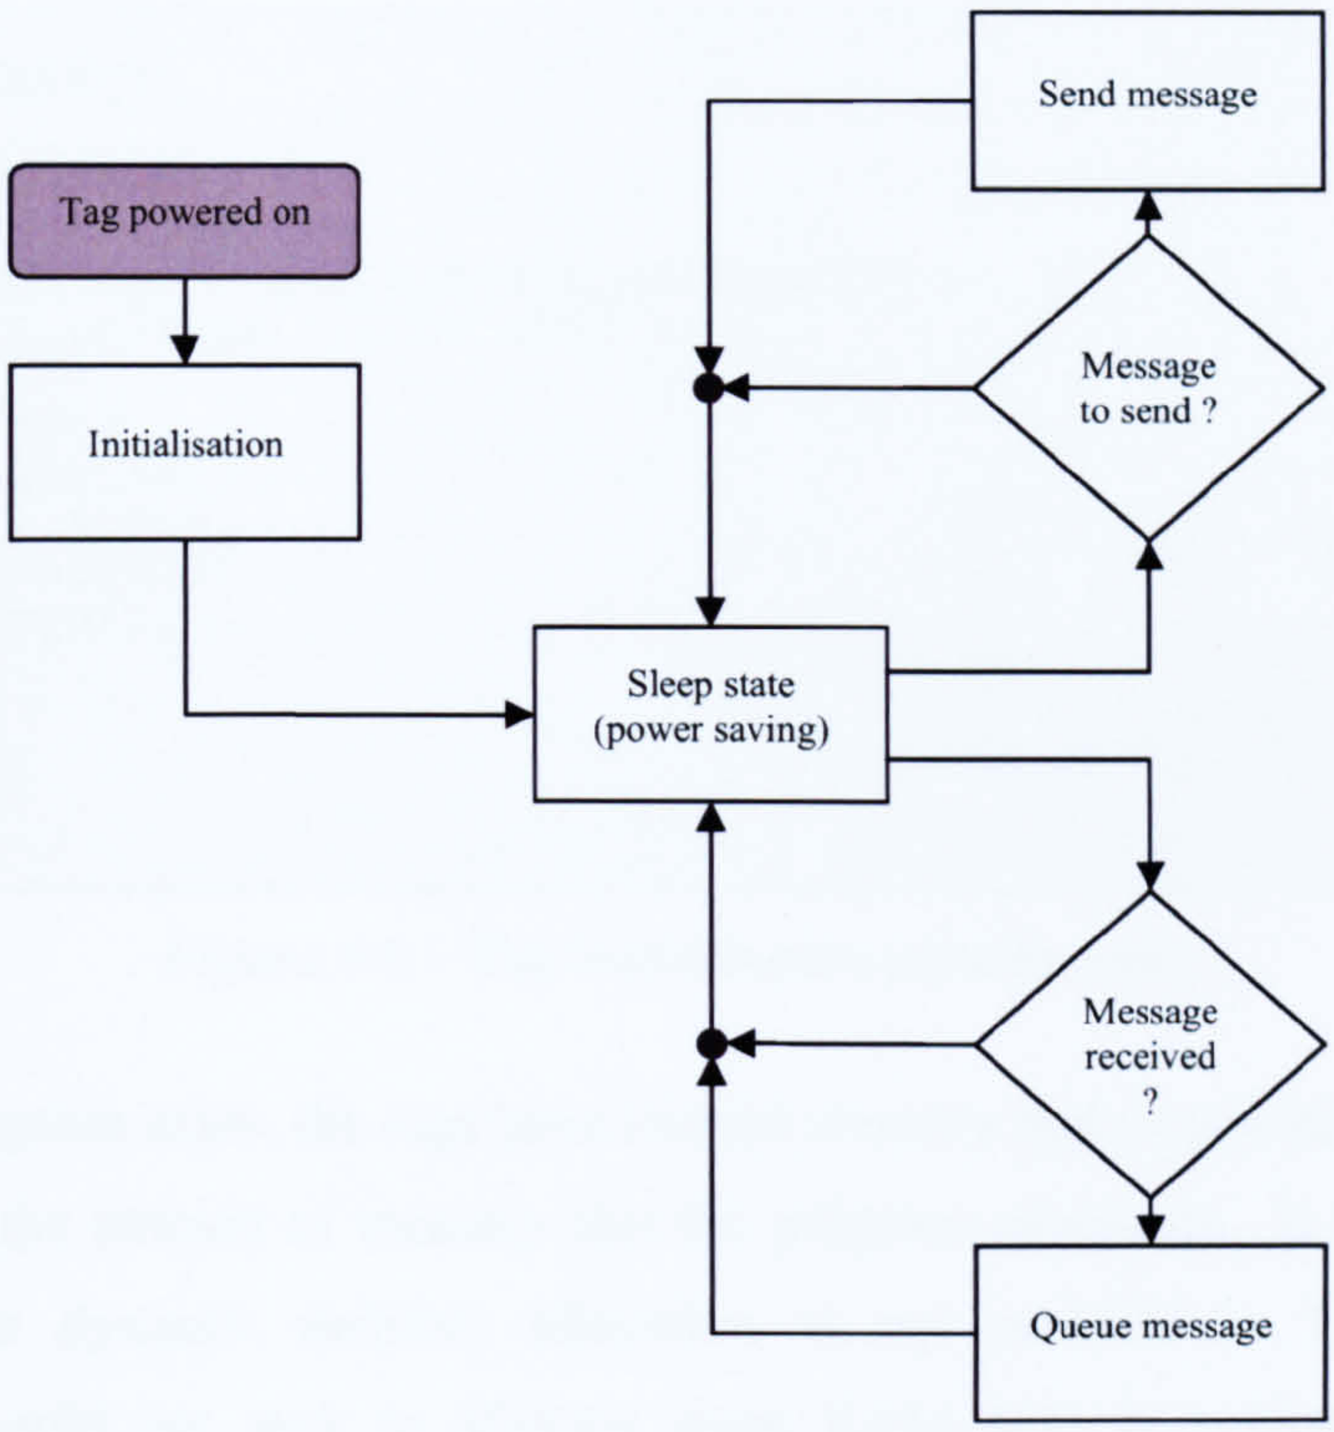

Figure 4.7 – TinyOS application flow diagram

It was mentioned in the overview of this chapter that there are relay nodes included in the system. These are intended to extend the range of the network, whilst operating in an almost identical way to the tags; therefore one should assume that the design of the tag software also applies to the relay nodes. It is largely assumed that the relay nodes will be mains powered, however since they need to support the power saving features of the tags, they also include their own power saving features.

A number of key concerns are considered with the creation of the tag software. One of the foremost concerns is that the tag lifespan should be as long as possible which means that power consumption must be minimised. It is assumed (from information obtained via the Micax manual [93]) that a tag will consume most power when it is transmitting, and to a lesser degree when it is actively processing data. As a result, an aim of this TinyOS application is to minimise the time that tags spend actively processing and transmitting data. The former concern is discussed shortly, whilst the later concern is largely accommodated through the use of the reverse routing mechanism, as well as the rather small message size; both of which have been previously discussed.

```
procedure init();
begin
  init_variables();
  init_components();
  battery_voltage = get_battery_voltage();
  if battery_voltage > 2.5V then
  begin
    light_led();
    get_unique_id();
    state := SLEEP;
    enable_radio();
    power_down();
  end
  else
    error();
end;
```

Figure 4.8 – Tag initialisation pseudo code

Power consumption aside, the tags have limited memory and so a watchful eye needs to be kept on the amount of memory that the program consumes. It is important to remember that dynamic memory allocation is not possible in TinyOS so the application should not seek to allocate more RAM than is available (i.e. 4kB). Furthermore, the application needs to fit into the 128kB of ROM available to the tag. In many ways this points toward trying to keep the application as simple as possible so as to keep its size to a minimum; doing so assists the notion of power saving since the simpler and smaller an application, the smaller the amount time spent processing complex tasks.

Obviously, whilst trying to produce a lightweight application, one must keep in mind the task at hand. Ultimately the application must be able to perform the task of enabling a mote to act like a tag and to perform all of the functions that are required of it in a reliable fashion.

4.4.2. Initialisation

When a tag is powered on, its first operation should be to initialise. This process involves the initialisation of local variables, as well as any TinyOS components which are used within the application. In Section 3.8 it was briefly mentioned that TinyOS features a component based architecture. This means that many of the common

hardware functions that one would want to utilise on a mote can be handled through calls to existing components.

Once variables and components are initialised, further operations are undertaken to make the tag ready to operate – pseudo code which summarises these operations is shown in Figure 4.8. The first operation is to check whether the battery voltage of a tag is sufficiently high – that is – above 2.5V. Note that it is probable that a Mica2 mote will operate at as little as 2.1V, as this is the minimum voltage required by the CC1000 [76] transceiver. However, it is unlikely that it will provide much in the way of an operational lifetime if the voltage is already below 2.5V when first powered on.

Obtaining the battery voltage requires access to the mote's ADC channels. The TinyOS component which handles ADC control is called `ADCC`. Like many operations in TinyOS, retrieving data from the ADC is a split phase operation. This means that one makes a call to get the ADC data, and then some time later an event occurs which allows the capture of this data. The reason for doing this is that capturing data from an ADC channel can be lengthy process, and so waiting for the ADC circuitry to become ready would cause significant code blocking. This would cause the mote to waste energy as it would be powered up whilst waiting. Therefore, split phase operations are advantageous because they allow the mote to continue with other tasks while the hardware is making itself ready to invoke the original call to use it; quite a number of TinyOS components operate in this manner.

Each call to an ADC channel returns a 16-bit integer. Only the least significant 10-bits are utilised, since this is the limit of the microcontrollers ADC precision. Computing the battery voltage (V_{batt}) from this value is accomplished using Equation 4-1, where V_{ref} is equal to 1.223V, *ADC_LEVELS* represents the number of levels of precision available (i.e. 2^{10} or 1024) and the *ADC_VALUE* is the value returned by a call to read an ADC channel. If the battery voltage is above 2.5V then the initialisation process can proceed; this application specifically lights a LED for a short period of time as a visual indicator.

DS2401 Memory Map		
CRC	Serial Number	Family Code (0x01)
0	1-6	7

Figure 4.9 – 64-bit silicon serial number memory map

$$V_{batt} = V_{ref} \times \frac{ADC\_LEVELS}{ADC\_VALUE} \qquad \text{(Equation 4-1)}$$

The next task is to obtain the unique ID for the tag. A key feature of any inventory management method is its ability to uniquely identify the asset it is associated with. Bar codes and RFID tags do this through globally unique serial numbers, the assignment of which is largely overseen by GS1. The Micax motes include a silicon serial number IC (the Dallas Semiconductor DS2401 [109]), which stores, in ROM, a 64-bit number. The format of this number is illustrated in Figure 4.9.

According to the manufacturer datasheet it is guaranteed that no two parts will have the same serial number, which is likely since there are at least 2^{48} unique ID combinations available. Therefore, this IC is quite useful in terms of asset identification. Furthermore, the manufacturer allows organisations to choose their own 12-bit identifier which is placed in the most significant 12-bits of the 48-bit serial number. In a similar way to bar code numbers this could allow both the manufacturer and the product to be identified – this could be useful for assets which do not necessarily stay with their parent organisation throughout their movement in a supply chain. As an example, this could be useful in the situation of lost luggage as it could mean that luggage is expedited to the originating airline (which could have its own 12-bit identifier) who can then deal with returning the luggage to their customer.

TinyOS includes a component, called `SerialId`, which can retrieve the serial number from the silicon serial number IC. Rather than retrieving the serial number each time it is required, it is more effective to retrieve it when the tag boots and then store it in memory. This saves processing time in the long term, as well as saving energy since the serial IC only uses power when it is accessed. In addition, the developer of the `SerialId` component notes that, although the DS2401 is connected

directly to the Micax microcontroller, accessing it at the same time as other IC's, such as the radio transceiver, can cause the operation to fail.

Moving on from this, TinyOS has a component named `GenericComm` which allows messaging via both the radio transceiver and the microcontroller UART (Universal Asynchronous Receiver and Transmitter) communication device. The later is useful for serial communications between a mote and a base station, such as the MIB510 programming board used in this work. The `GenericComm` component provides interfaces which allow for both sending and receiving messages. Enabling this component is the penultimate operation of initialisation; before this point the tag should not respond to messages because it is simply not ready to. The final operation is to put the tag into its sleeping state which is detailed further in the next section.

4.4.3. Sleep State

The reason for the sleep state not being discussed in the previous section is that it is central to the operation of a tag. If one looks back at Figure 4.7, every process leads back to the sleep state. This is because, for the majority of its time, the tag should be consuming minimal amounts of power in order to extend its effective lifespan. TinyOS facilitates this in two ways: powering down the microcontroller, and also altering the power consumed when listening for incoming messages.

TinyOS provides the component `HPLPowerManagementM` which allows the TinyOS task scheduler to put the microcontroller into a sleep state. The Atmel 128L supports six different sleep states, four of which `HPLPowerManagementM` can choose to implement; these are listed in Table 4.2. The chosen sleep state depends upon a number of factors such as whether the task queue is empty, and whether the radio transceiver and ADC are enabled. Obviously, if the radio transceiver has to be disabled for power management to operate then it is of no use, since a tag must always be able to respond to messages that it receives – if the transceiver is turned off then it will not receive any!

Sleep state	Description
Idle	This mode stops the CPU of the microcontroller, but allows other components to carry on operating, including all interrupts, SPI (Serial Peripheral Interface) and Timers. TinyOS will go into this state if any external timers are running, or if the UART component is enabled.
ADC Noise Reduction	If external timers and the UART are disabled, then `HPLPowerManagementM` checks to see if the ADC is enabled. If it is then this state is selected. It is similar to the idle state, but stops two of the three timers included in the microcontroller package.
Power Save	If this mode is entered, the ADC is disabled and all other conditions are met. It allows a timer to run, which means that the microcontroller can be woken after a set period. In addition, external interrupts are still enabled, although they need to be kept at the interrupt level for a while so that they are reliably recognised.
Power Down	This is the same as the power save state, except no timer is left running. The microcontroller can only be woken via an external interrupt and represents the lowest power mode supported under current versions of `HPLPowerManagementM`.

Table 4.2 – HPLPowerManagementM sleep states

To deal with this issue, a component responsible for controlling the CC1000 transceiver can be utilised. It is called `CC1000RadioIntM`. This component allows the listening mode of the transceiver to be set such that it powers on only for extremely brief and infrequent intervals to check if there is any wireless activity. The minimal power consumption setting turns on the transceiver once every 100ms. If it detects activity then it can trigger the microcontroller to wake and process the incoming message; otherwise the transceiver can go back to sleep.

The `TAG_WAKEUP` message type is included in order to accommodate low power listening. When a tag sends a message whilst it is fully powered on, it will only be transmitting for a short period of time; in the order of tens of milliseconds. As a result, there is a high probability that the receipt of a message would be missed if the transceiver is only on briefly every 100ms. To deal with this issue,

`CC1000RadioIntM` includes commands for setting both the listening and transmitting mode. In the low power transmission mode the tag sends a long preamble, in the order of hundreds of milliseconds long, which allows a receiving tag to wake up reliably. Hence, when a tag wakes up it will change its listening mode to full power and then broadcast a `TAG_WAKEUP` message using the low power transmitting mode. It can then set its own transmission mode to high power, and wait a little while for its neighbours to wake before sending further messages. The tag will not resume using low power listening until an ample time has passed with no transceiver activity.

A short experiment was carried out in order to highlight the benefit of enabling power saving and to ensure that it was functioning correctly. Using the well known relationship between voltage (V), resistance (R), and current (I) shown in Equation 4-2, placing a 1Ω resistor in series with the positive mote power terminal and its battery allows the determination of current consumption through monitoring the voltage across the resistor. An Instek GDS-810S [110] digital oscilloscope was used to perform this task, whilst a mote was set to transmit a message periodically – with and without power saving enabled. A picture of this setup is shown in Figure 4.10, and the results of the test are shown in Figure 4.11.

$$V = I \times R \qquad \text{(Equation 4-2)}$$

Despite the GDS-810S oscilloscope not having ample resolution to show the actual current used by the mote using power saving (minimum detectable voltage is 1mV), Figure 4.11 does show the quite remarkable benefits of using power saving. This chart shows that a mote without power saving enabled will consume at least 14mA – it is presumed at this stage that a mote with power saving enabled consumes tens of μA per hour based upon data shown in Table 3.3 in Chapter 3. The pseudo code for enabling and disabling power saving is detailed in Figure 4.12. Note that the transmission power (see Table 3.4 in Chapter 3) is also set to its lowest possible value when the mote is saving power, since the CC1000 data sheet [76] recommends this to avoid current leakage.

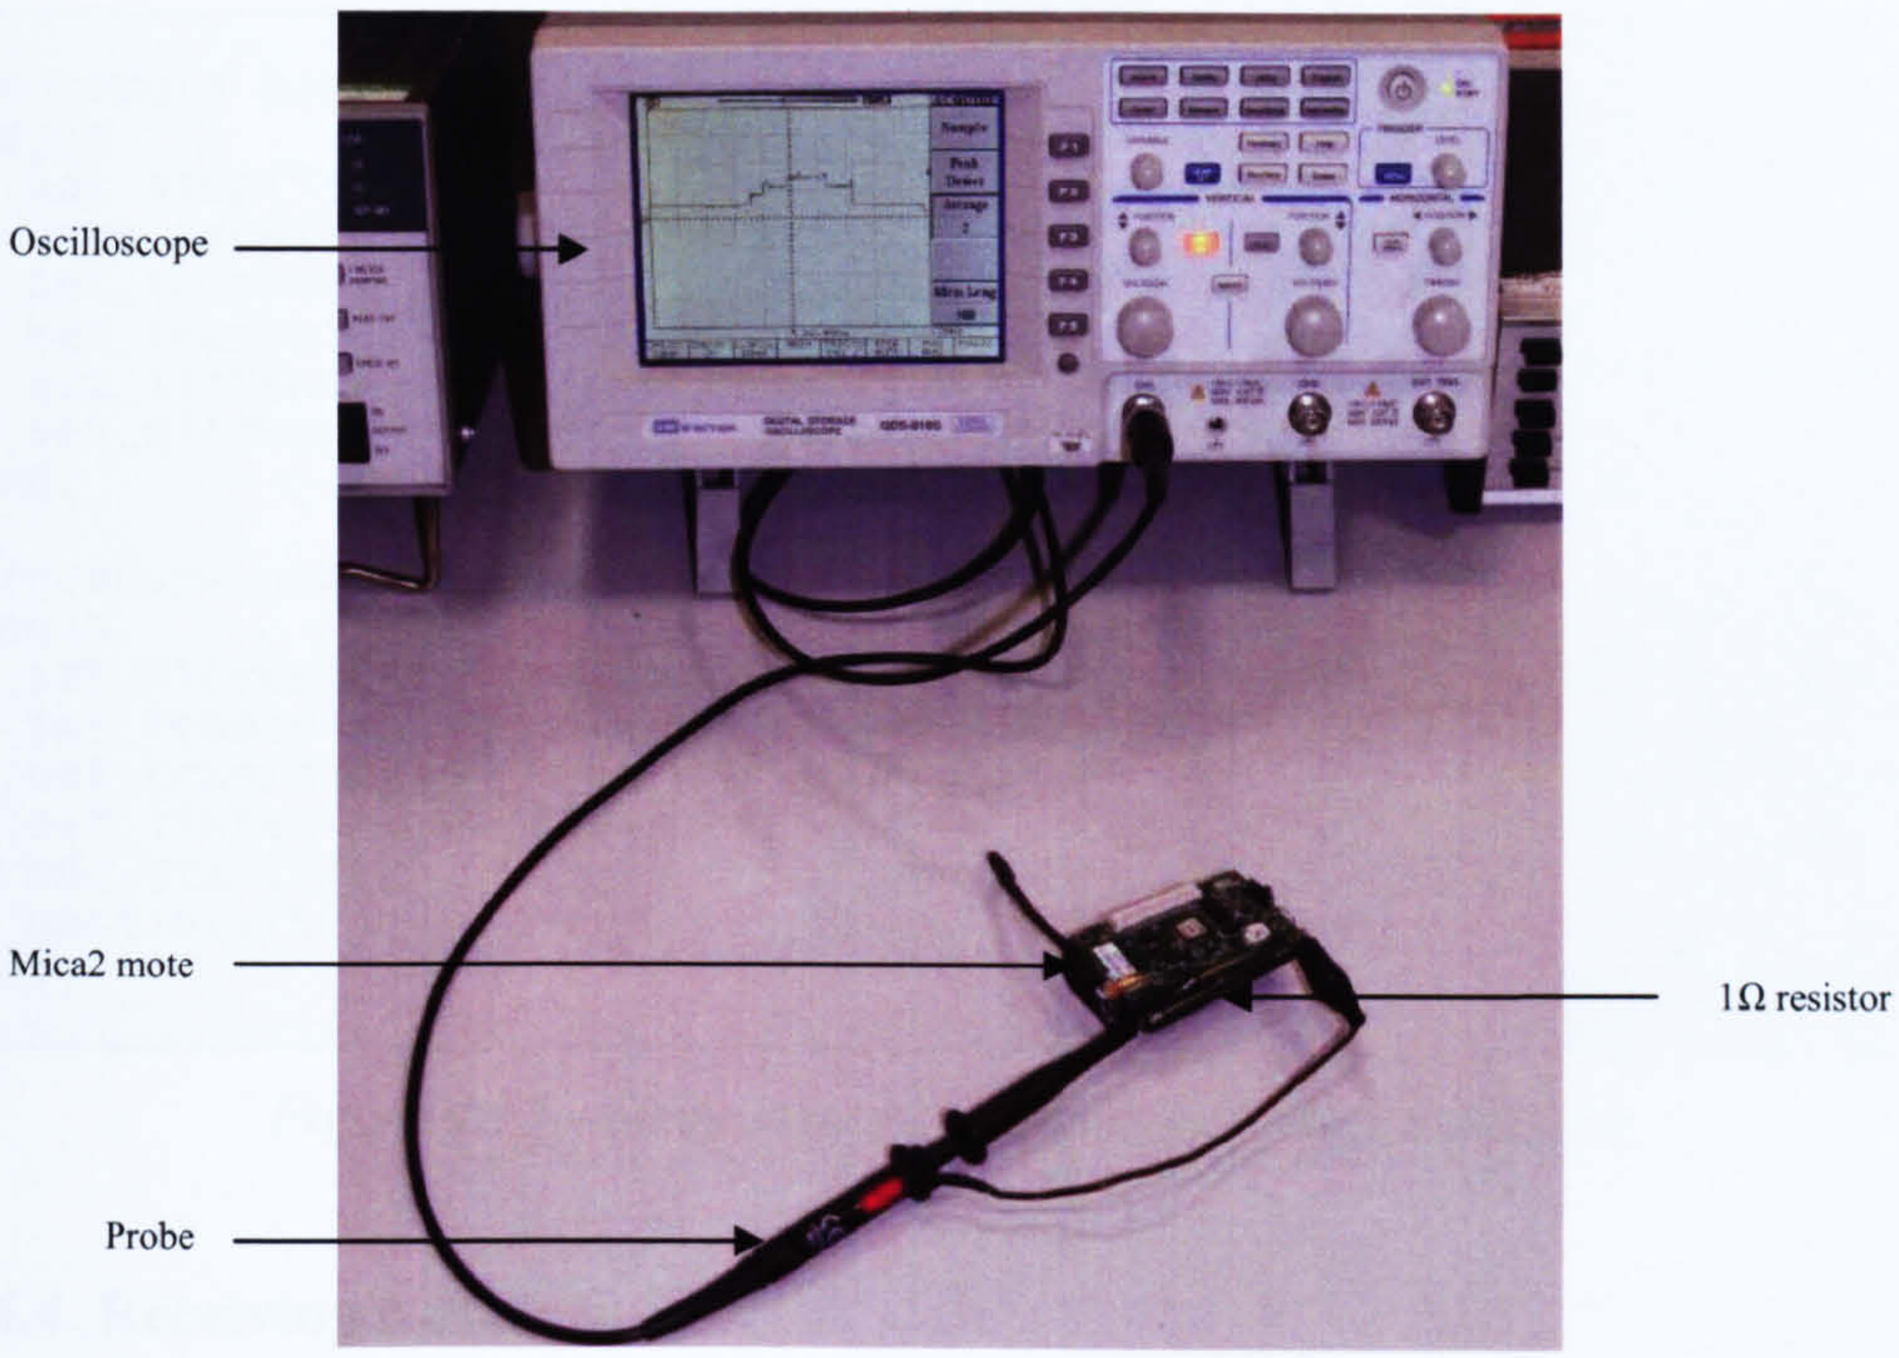

Figure 4.10 – Powering saving experimental setup

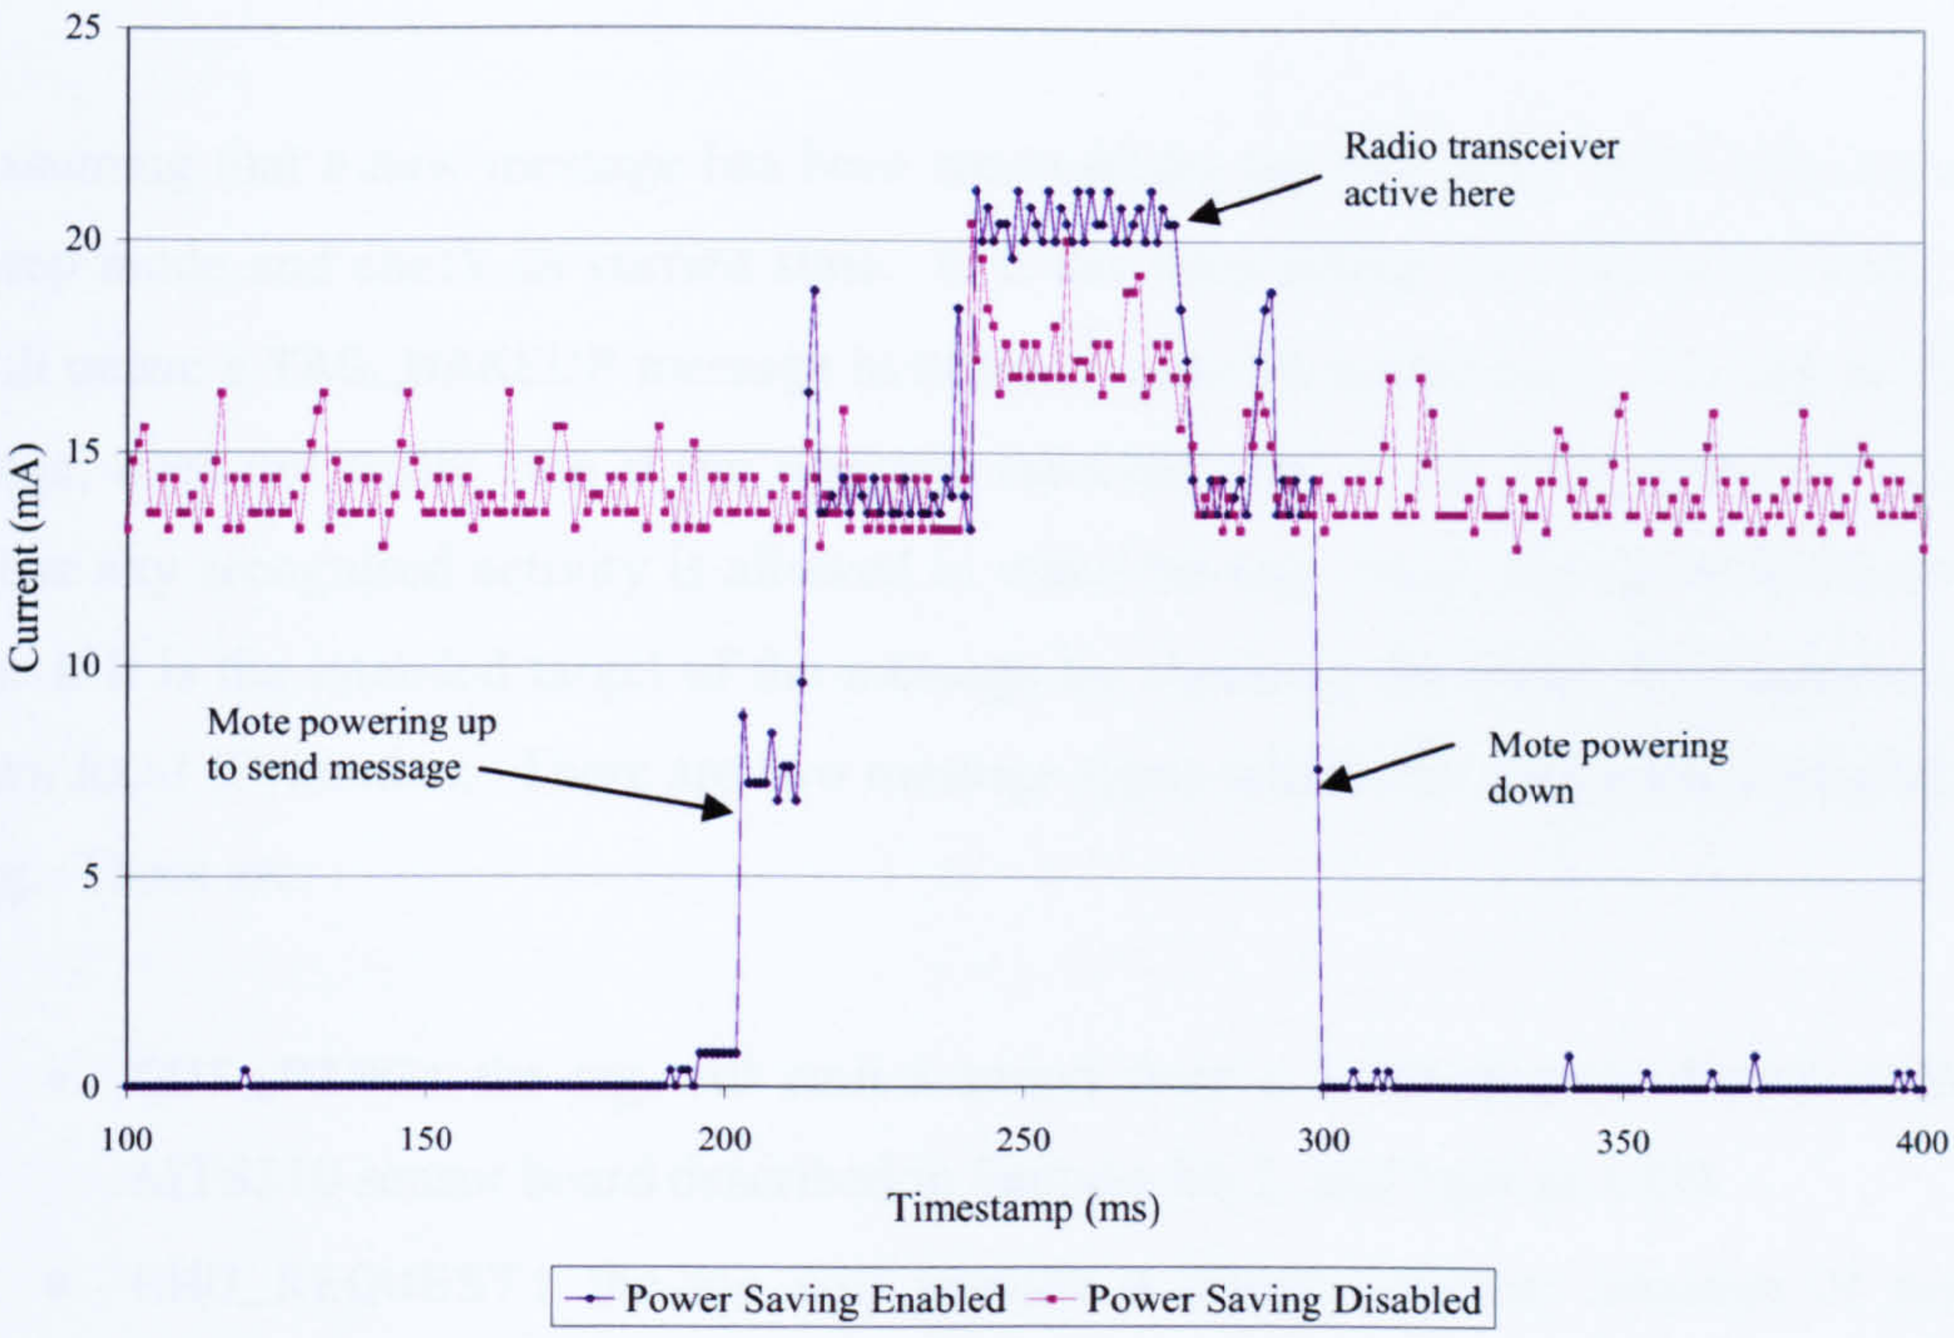

Figure 4.11 – The effect of enabling TinyOS power management

```
procedure power_down();
begin
  adc.stop();
  uart.stop();
  set_transmission_power(min_power);
  set_transmission_mode(min_power);
  set_listening_mode(min_power);
  set_microcontroller(sleep_state);
end;

procedure power_up();
begin
  set_microcontroller(awake_state);
  set_transmission_power(max_power);
  set_transmission_mode(max_power);
  set_listening_mode(max_power);
  adc.start();
  uart.start();
end;
```

Figure 4.12 – Sleep state power management pseudo code

4.4.4. Receiving a Message

When a tag receives a message, a `receive` event is triggered. The tag will check to see if it has received this message previously; if it has then it can simply go back to sleep and ignore the message.

Assuming that a new message has been received the tag will bring itself fully out of sleep mode and check its current state. If it has been asleep for a period of time, it will queue a TAG_WAKEUP message in order to wake its neighbours. The tag, at this stage, does not really care if the message received was of the TAG_WAKEUP type, since any recognised activity is allowed to wake the tag. Next, the tag will check to see if it is the intended target of the message by checking the `dest` field against its own local ID number. There are two message types which can define the destination tag. These are:

- GUI_PING: the tag will emit a sound from a small speaker, found on the MTS310 sensor board described in Section 3.6.2, and light an LED.
- HHU_REQUEST: the tag will provide a direct response message of type HHU_RESPONSE.

If the message type GUI_PING is received by a tag to which it is not addressed then the tag will rebroadcast this message exactly once. GUI_REQUEST messages are treated similarly. The reverse routing mechanism discussed in Section 4.3.4 is implemented here for the GUI_RESPONSE message type, which is the only message type requiring it.

At this point it is worth noting that a tag can receive messages from multiple sources almost simultaneously, so it may not be able to service one message before the next arrives. As a result, a message queuing system is implemented to temporarily store messages so that they are not simply lost. This queuing system is separated into two queues; one for general purpose messages (general_queue) and another for messages which must be given priority (priority_queue). Dealing with the general queue first, this is intended to store messages of the types GUI_REQUEST and GUI_RESPONSE. These packets, whilst important, are not time critical – all that matters is that they are received by the tags (GUI_REQUEST) and GUI (GUI_RESPONSE) in order to allow a tag to be identified.

TAG_WAKEUP messages are placed in the priority queue because it is important that these messages are serviced before any others to ensure that neighbouring tags are awake to receive subsequent messages. HHU_* message types (see Table 4.1) are also placed here – by giving priority to its messages the hope is to reduce the time spent waiting for the HHU features to operate.

After a message is appropriately queued, two timers are started. The first timer is for sending a message at some random time interval after message receipt. The second timer, the timeout timer, is started (or restarted if already running) to ensure that a tag conserves as much power as possible. This timer will fire if no message activity occurs for a period of time. It will allow the tag to go from its idle mode to the power down sleep mode (see Table 4.2), therefore saving more power. In addition, it removes any reverse routes which have been stored, since this information cannot be guaranteed to be accurate the next time a tag wakes up. The pseudo code for all of these operations is shown in Figure 4.13.

```
event receive(msg);
begin
  if received(msg) = FALSE then
  begin
    received(msg) := TRUE;
    power_up();
    if state = SLEEP then priority_queue(TAG_WAKEUP);
    state := AWAKE;
    if msg->dest = local_id then
      if msg->type = GUI_PING then
        enable_action()
      else if msg->type = HHU_REQUEST then
        priority_queue(HHU_RESPONSE, msg->origin)
    else if (msg->inter = local_id)
                                and (msg->type = GUI_RESPONSE) then
      general_queue(msg->type, reverse_route)
    else if msg->type = GUI_REQUEST then
    begin
      reverse_route := msg->inter;
      general_queue(msg->type, broadcast);
    end
    else if msg->type = GUI_PING then
       general_queue(msg->type, broadcast);
    timers.start();
  end;
end;
```

Figure 4.13 – Receive event pseudo code

```
event send_timer.fired();
begin
  power_up();
  if priority_queue <> empty then
  begin
    if priority_queue_head->type = TAG_WAKEUP then
      set_transmission_mode(min_power);
    send_msg(priority_queue_head)
  end;
  else if general_queue <> empty then
    send_msg(general_queue_head);
  if (priority_queue and general_queue) <> empty then
    timers.start();
end;
```

Figure 4.14 – Send timer pseudo code

4.4.5. Sending a Message

As mentioned in the previous section, messages are sent on a timer interval. This timer is given a random offset to prevent all nodes simultaneously responding to a message. MAC is provided by TinyOS through channel assessment which checks that no other nodes are transmitting before sending a message. Once the tag is powered up, the priority queue is serviced first; if this is empty then the general queue is serviced. Should both queues be empty then the timeout timer will not be restarted.

If no further activity occurs then the tag will go to sleep. Otherwise the send timer will be restarted, as will the timeout timer.

It is noteworthy that the TinyOS code for the tags, included in Appendix A.1, extends the send timer interval if the message which has just been sent is of the type `TAG_WAKEUP`. The reason for doing this is that there is the need to allow time for neighbouring nodes to wake up before transmitting further data to them. The queuing system allows nodes to receive traffic whilst this waiting period elapses. In addition, `TAG_WAKEUP` messages are sent using the lowest power transmission mode available in `CC1000RadioIntM` so that the message preamble is sufficiently long. The pseudo code for the send timer event is included in Figure 4.14.

4.4.6. Memory Consumption

The memory consumed by the application is shown in Table 4.3, and is largely dependant on the size of the queues implemented as message buffers. The maximum combined queue size (i.e. the length of the general and priority queues combined) is 96 messages, as this consumes 3.88kB of the available 4kB RAM; using more causes a mote to cease operating. In terms of ROM, the application uses approximately 25% of the available capacity. This information was obtained from the nesC compiler at compile time, an example of which is included in the Appendix A.1.4.

Queue size (combined)	RAM (kB)	ROM (kB)
0	1.16	31.82
20	1.72	33.04
40	2.29	33.04
60	2.86	33.04
80	3.42	33.04
96 (Max)	3.88	33.04

Table 4.3 – TinyOS application memory consumption

4.5. Handheld Unit

4.5.1. Overview

A handheld unit (HHU) was created as part of this work to help support the inventory management system, and to give an example of how the hardware supporting WSN is flexible enough to allow expansion. It was thought that by demonstrating such capabilities the technology may become a more attractive investment prospect for industry.

The HHU is designed in a modular way, mainly because some parts of it are designed to be reusable. In particular the mote interface board displays this characteristic as it is used extensively in this work, as well as in other research projects [81,111-112]. Figure 4.15 shows a block diagram of the HHU.

In this particular implementation, the HHU is being used as a positioning device – that is – its purpose is to attempt to determine the position of tags relative to a human operator. There are other possible uses of this device, which are discussed in Chapter 8, but this was deemed appropriate in the context of this project.

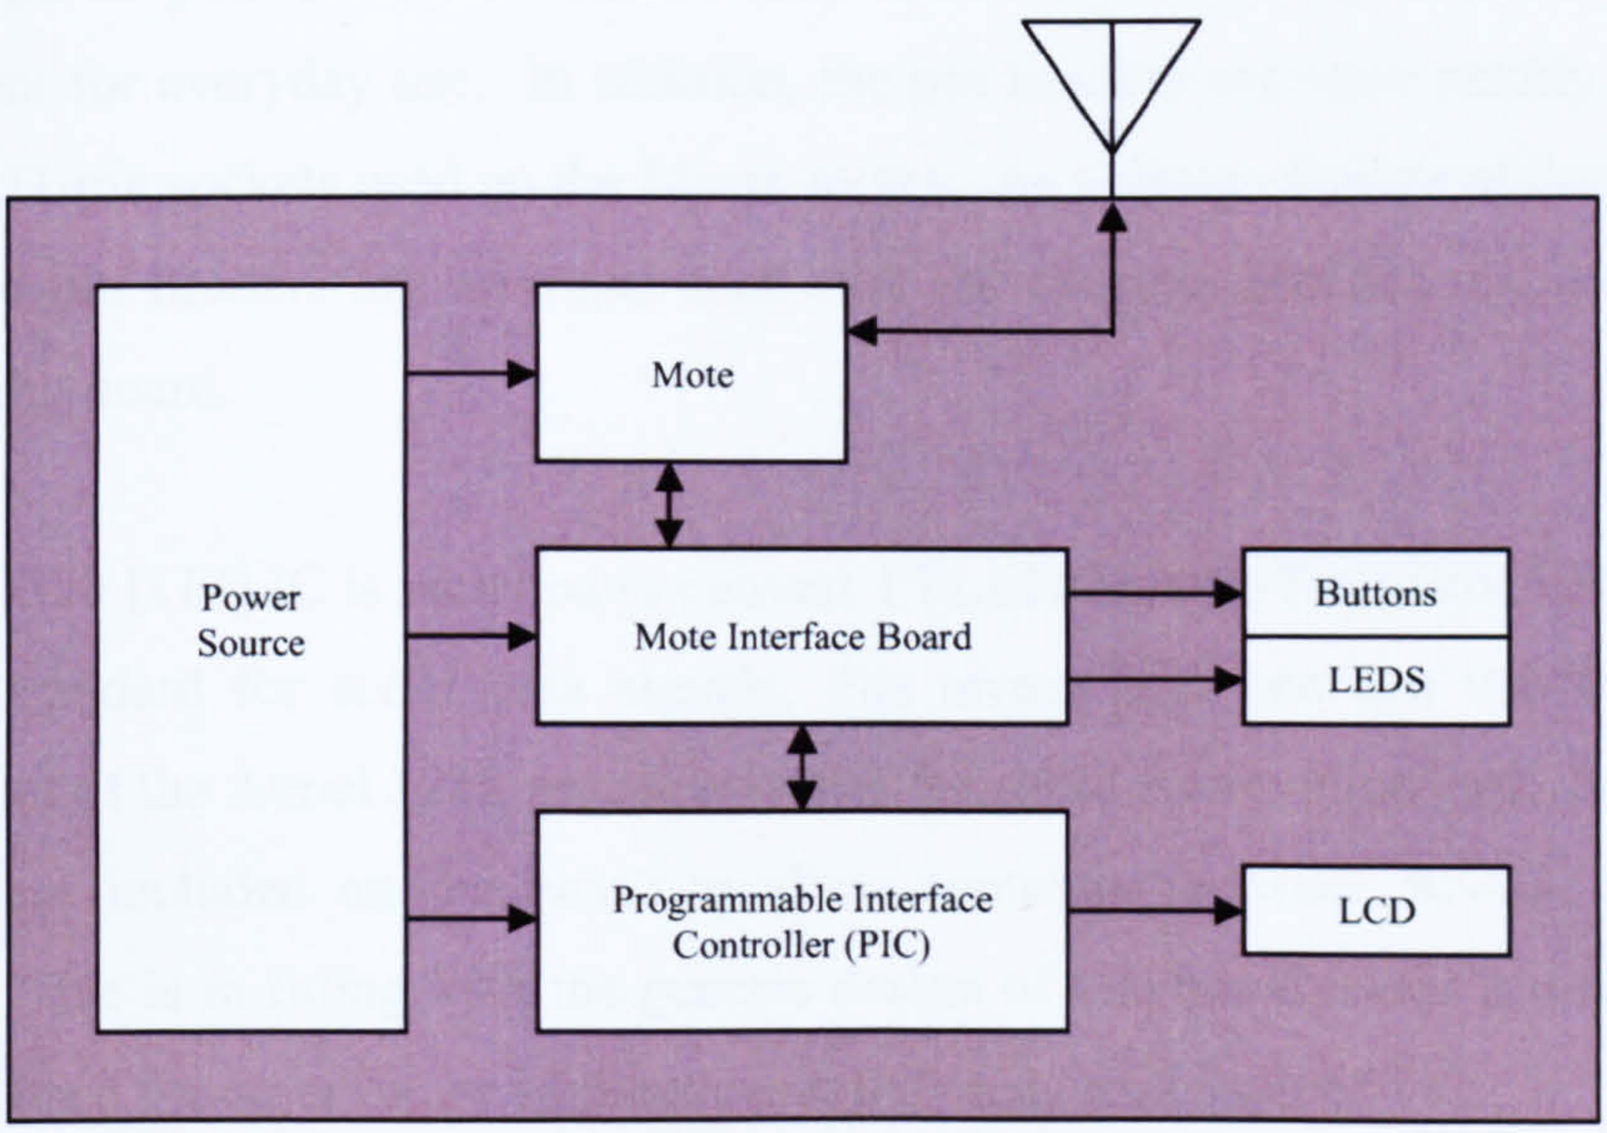

Figure 4.15 – Hand held unit hardware block diagram

The Mica2 motes have a Received Signal Strength Indicator (RSSI), which, in theory, should allow the computation of the range between a sending and receiving device. This section describes the hardware and software implementations used in attempting to achieve this goal.

4.5.2. Hardware Implementation

The HHU is encased in a 191×110×61 mm ABS box which provides adequate protection for the circuitry inside, whilst being a reasonable size for a hand held device. Power is provided to the box from a 9V PP3 type battery. Figures 4.16 and 4.17 show the outside and inside of the unit, respectively.

At the heart of the HHU is a Mica2 mote which handles data processing and communication with both the tags and GUI. The mote requires connection to the buttons, the LCD (via the LCD interface board) as well as the indicator LEDs and power source. To achieve this, a general purpose interface board (see Figure 4.18) was created which provides break out connections to all 51 pins of the surface mount mote interface socket [97] via two rows of dual in-line pin headers arbitrarily named SKT1 and SKT2. Mappings of these connections are included in Appendix B.1. These type of pins are used because they are easily connected to, making them convenient for everyday use. In addition, the pin headers are more readily available than the 51-pin sockets used on the Micax motes. As a design feature of this interface board the pin headers are arranged such that the spacing matches that of 2.54mm prototyping board.

A MAX3223 [113] IC is included to convert TTL (Transistor-Transistor Logic) to the RS-232 standard for serial data signals; this means that one can use the UART component of the Atmel 128L microcontroller for serial communications. Jumper pin headers are included on the board to allow switching between RS-232 and TTL outputs. This is in fitting with the generic design of this board, since it means that it could be used for other future applications which may make use of TTL.

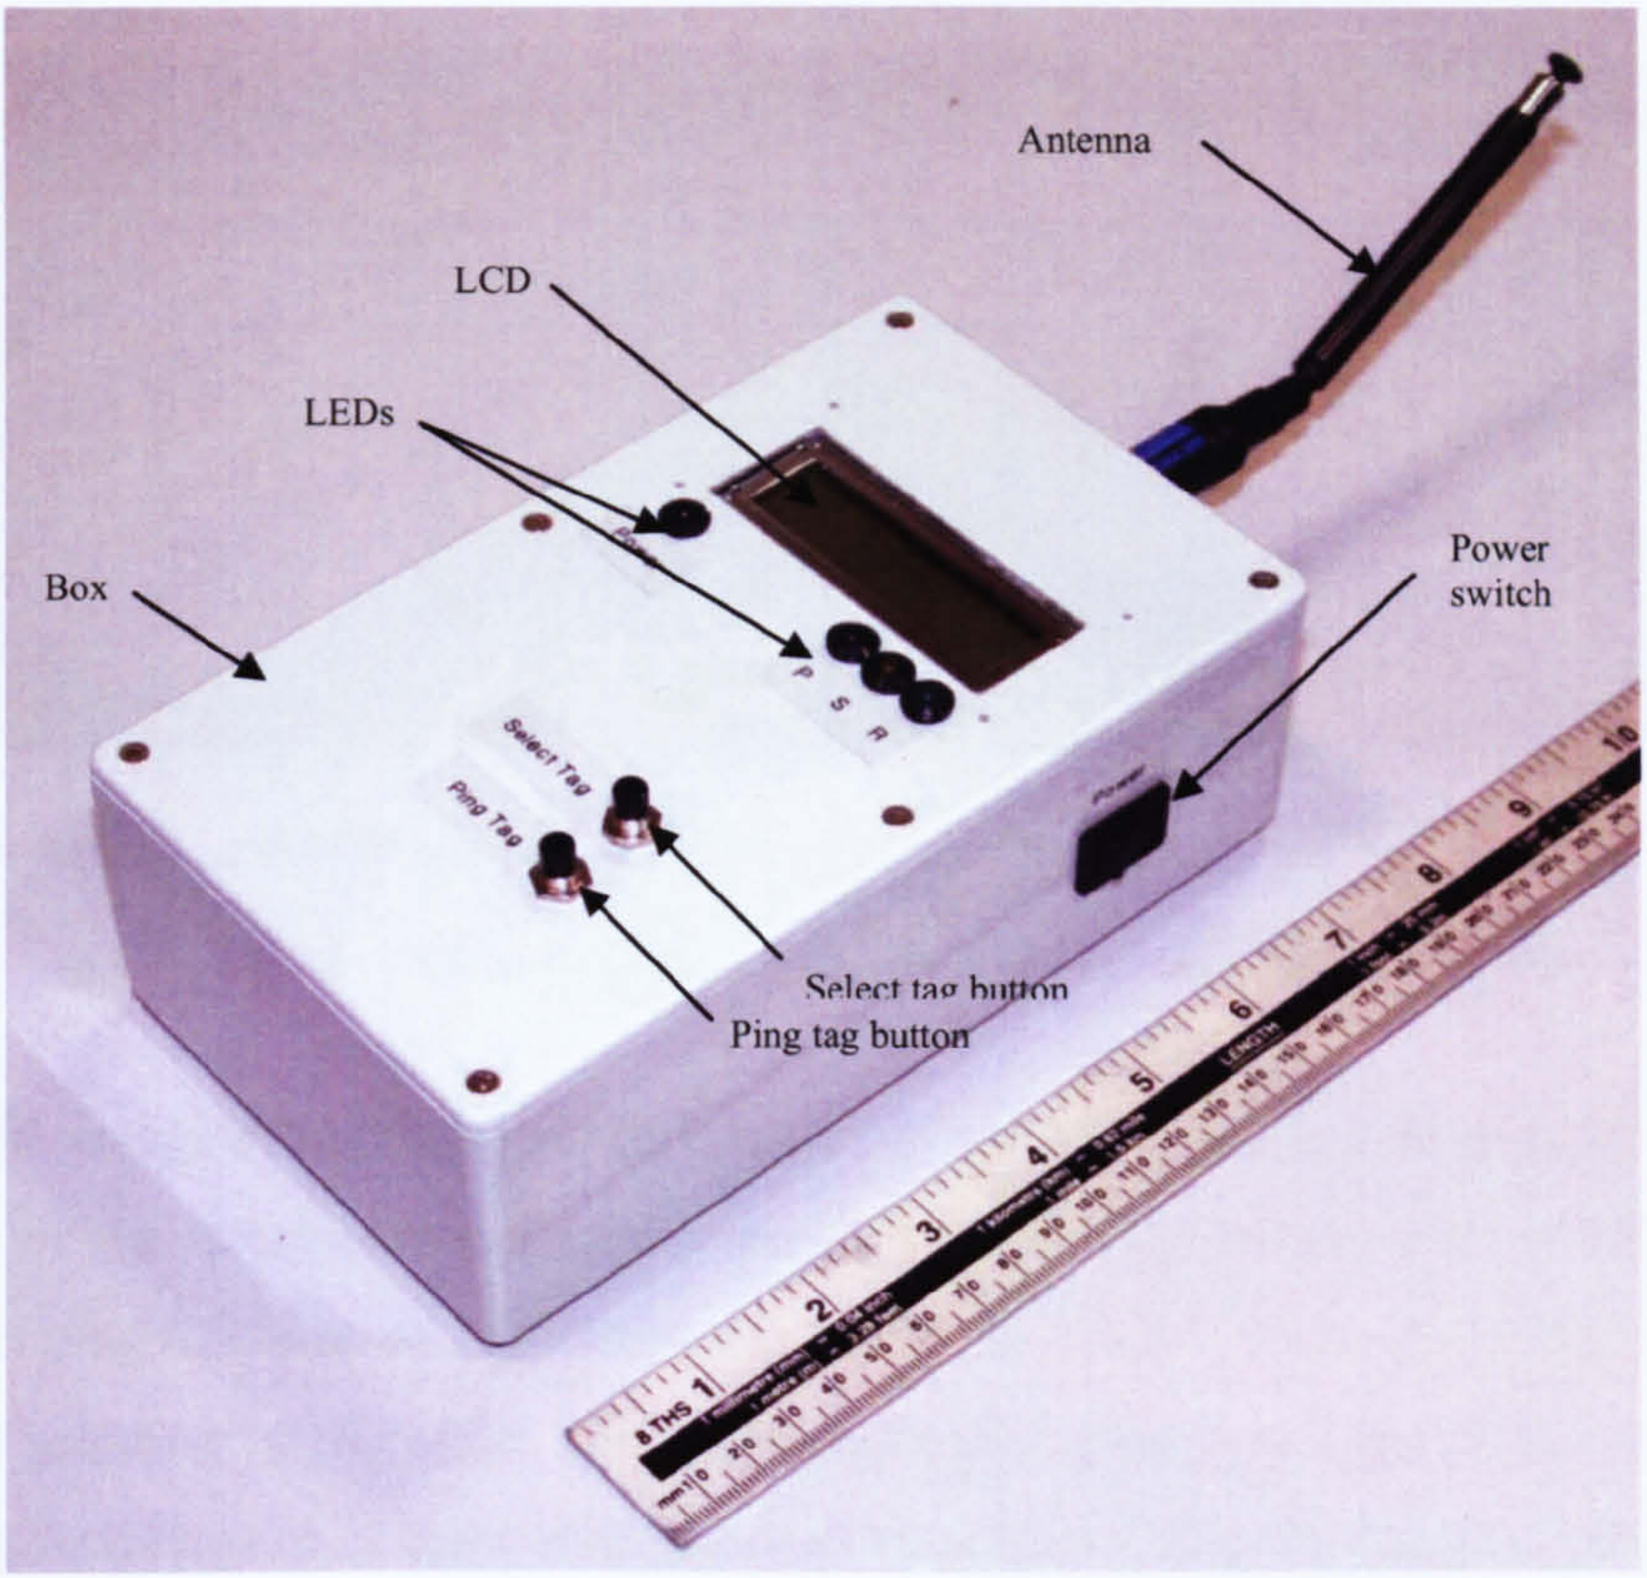

Figure 4.16 – Outside the handheld unit

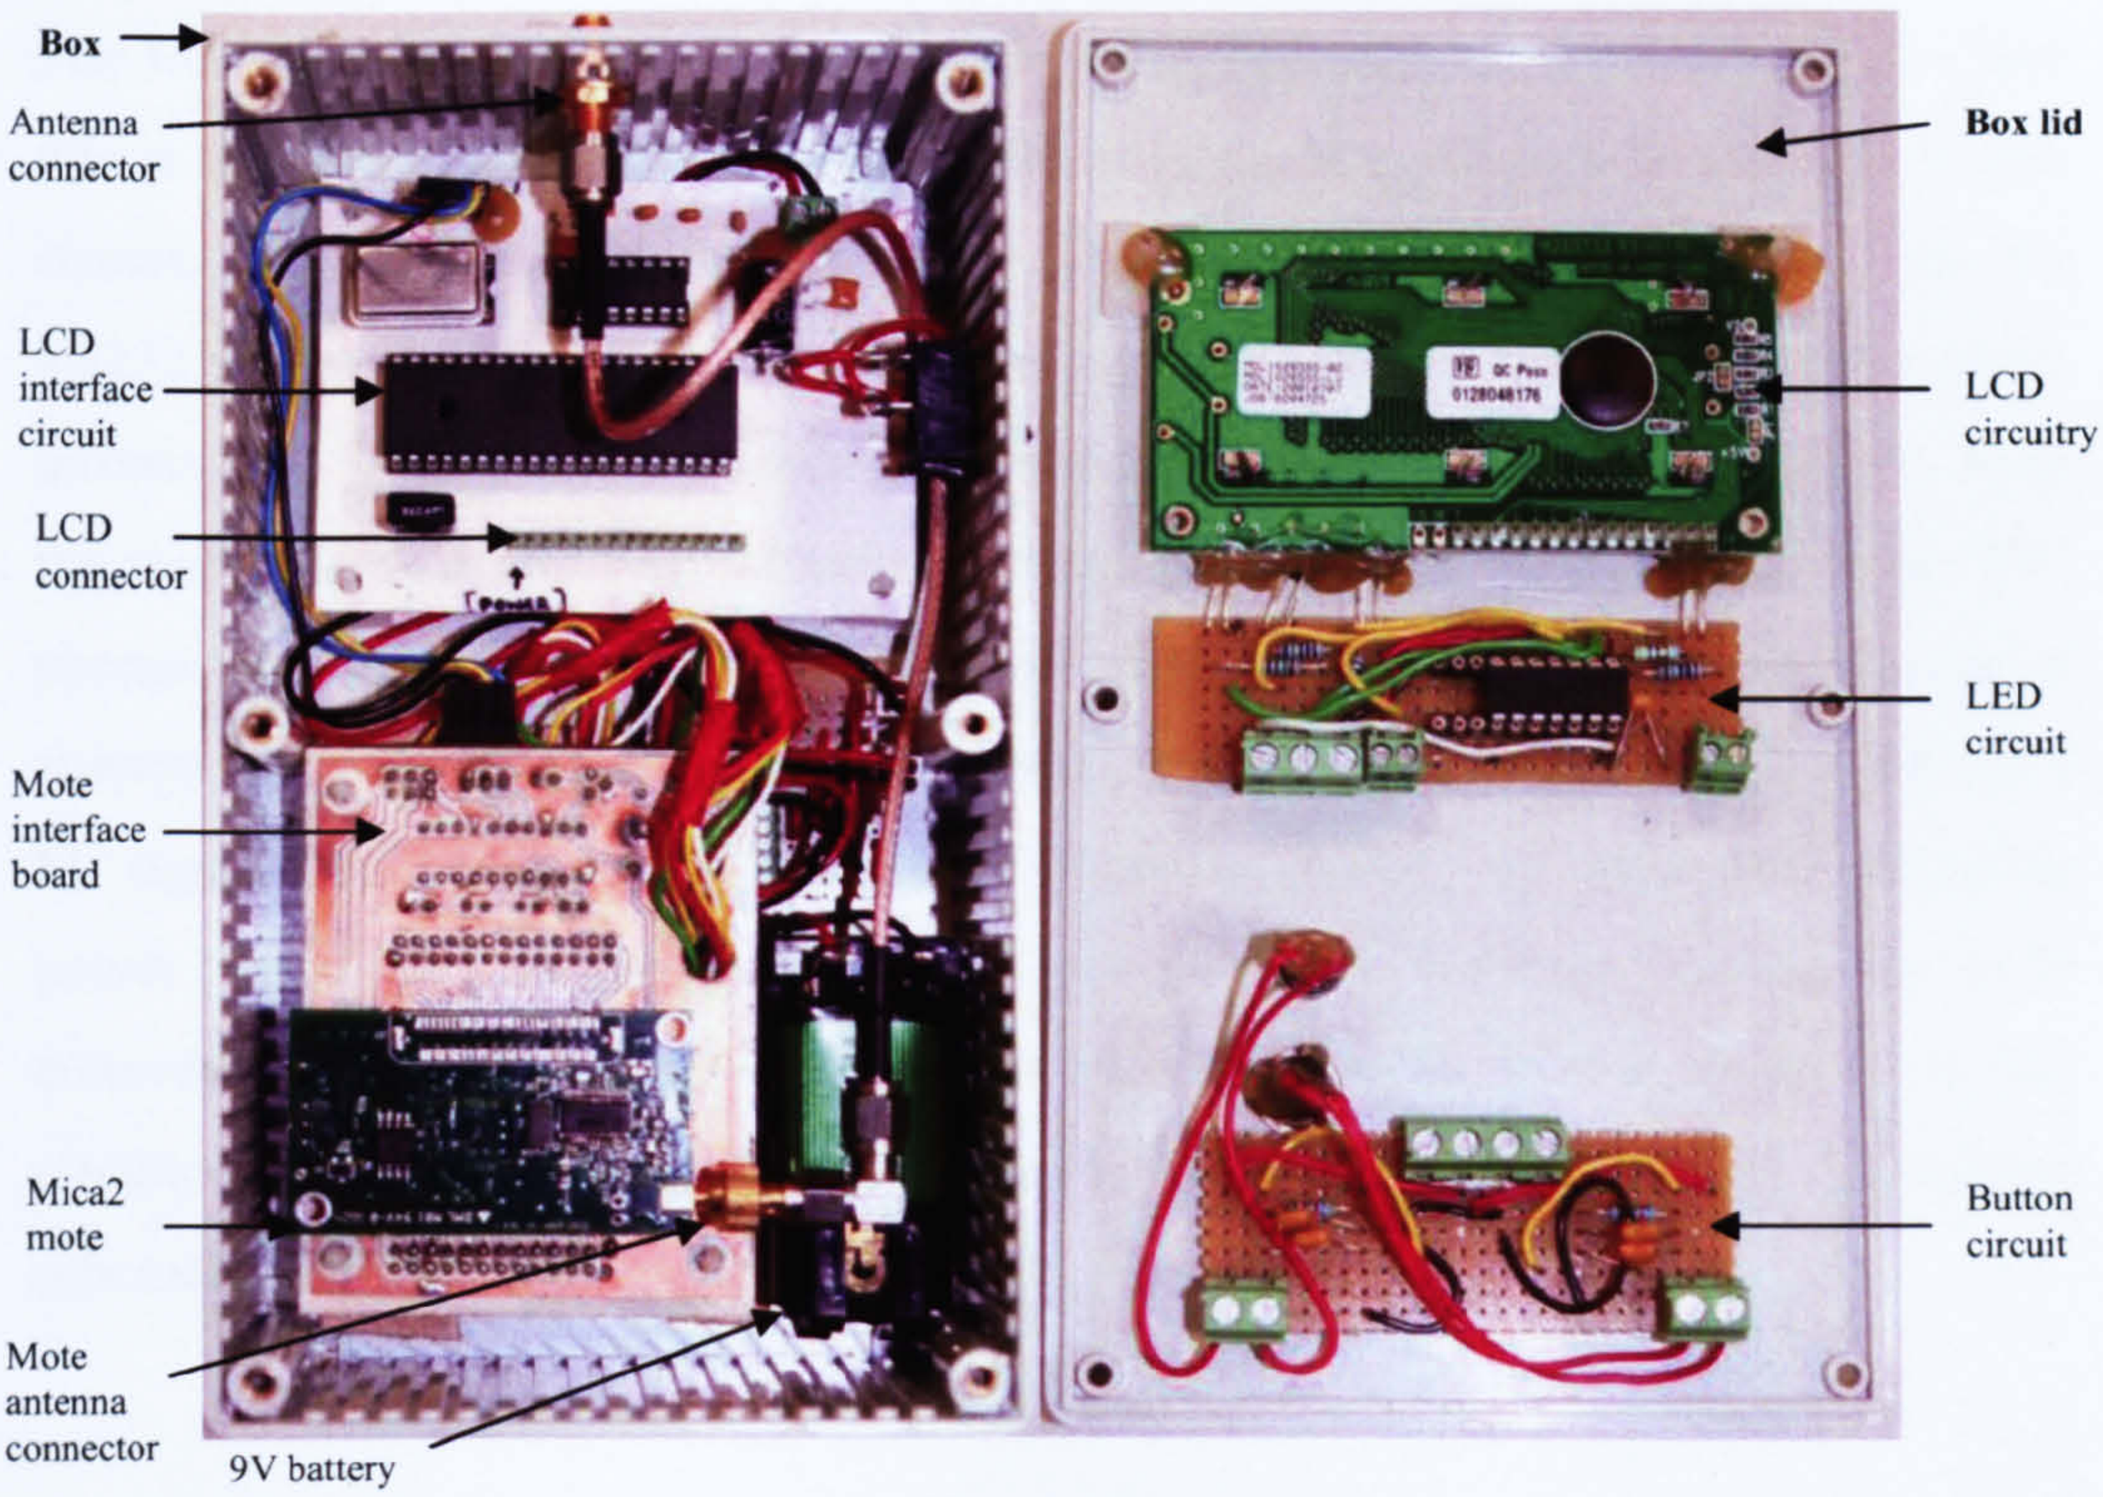

Figure 4.17 – Inside the handheld unit

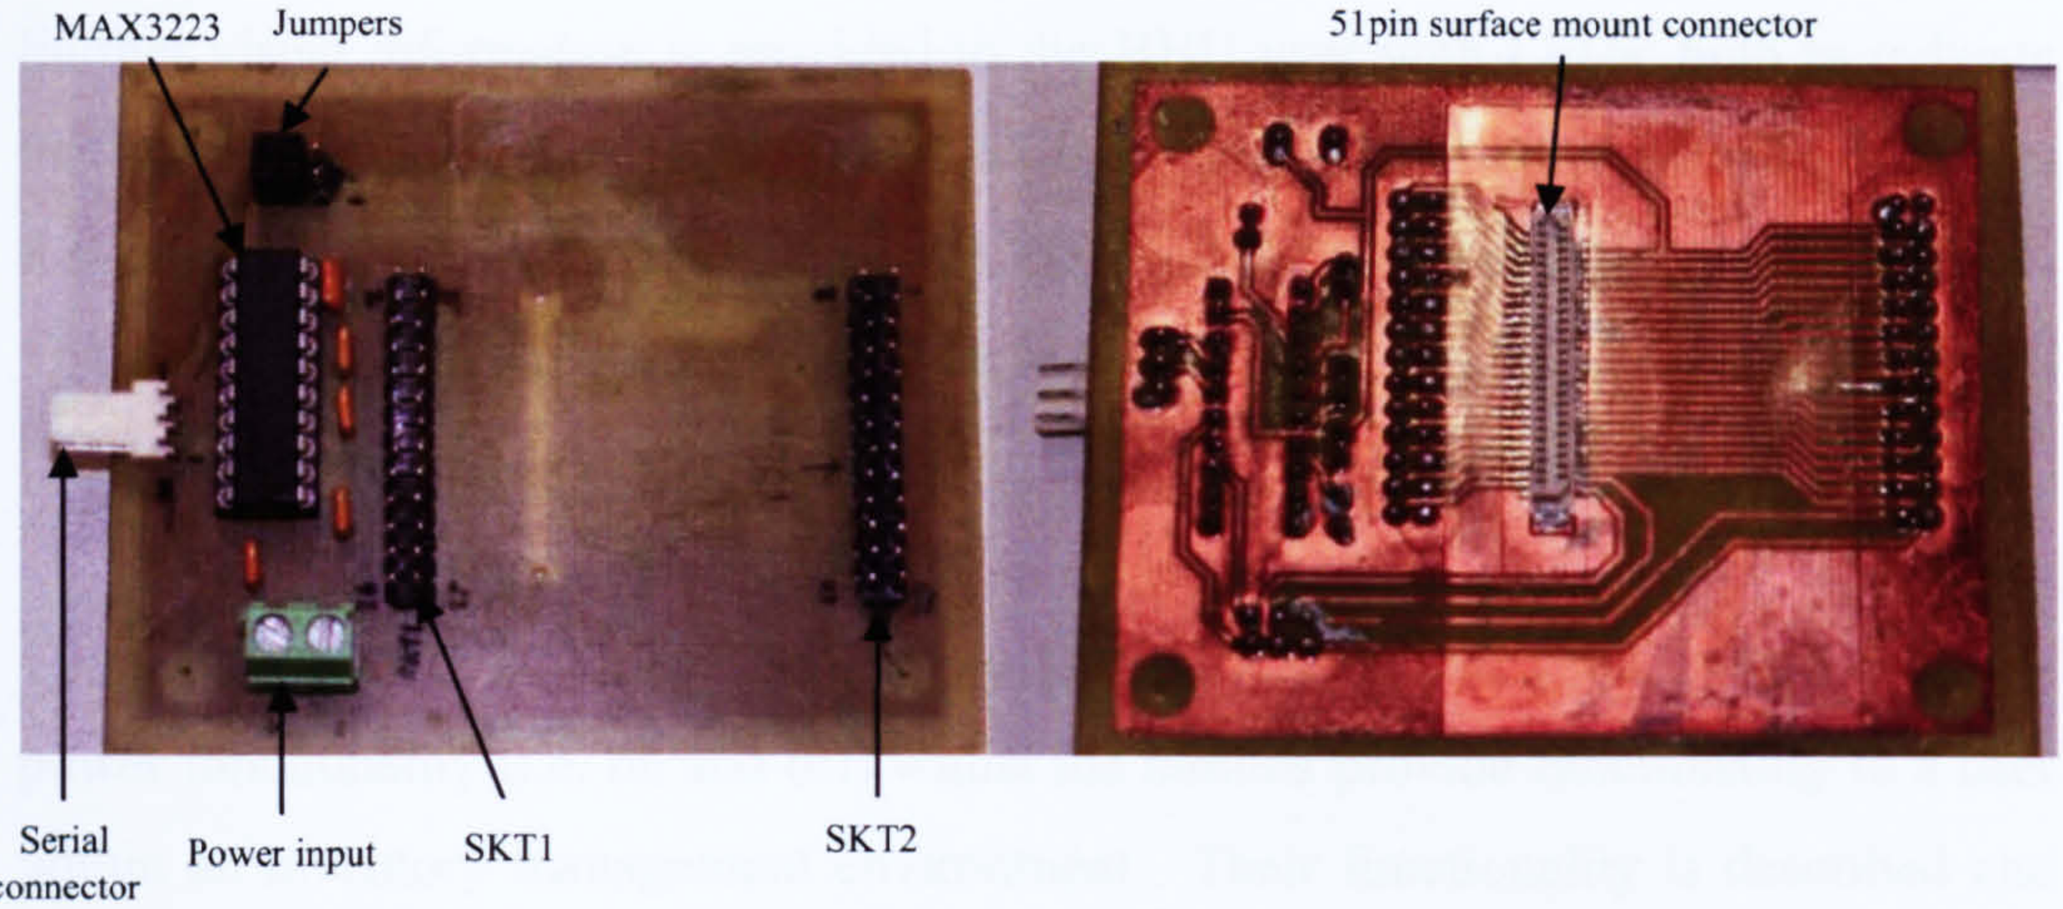

Figure 4.18 – Mote interface board top (left) and bottom (right)

The mote is powered via a power connector on the mote interface board. Later versions of the interface board (not pictured here) include an LM1117 [114] voltage regulator, which is used to reduce the voltage of the HHU 9V battery power source to the 3.3V required by the mote. Full schematics of the latest revision of this board are shown in Appendix B.2; these schematics were created using the Eagle CAD software [115], designed specifically for PCB creation.

The LCD interface circuit allows the mote to display data on the LCD [116] – how this is achieved is described briefly in the next section. The mote connects to this circuit, which features a Programmable Interface Controller (PIC) [117], via the RS232 connection provided on the mote interface board. The PIC, like the mote, utilises TTL for serial communication, and so a further MAX3223 IC is used. Both the mote and PIC are TTL capable but the mote operates at 3.3V, whilst the PIC operates at 5V; this means that a high output from the mote may actually be too low in voltage for the PIC to utilise. The PIC requires a 5V power supply which is fed via a 5V regulator (LM7805 [118]) from the 9V power source. In addition to providing power to the PIC, the 5V regulator also powers the LCD screen through the LCD connector shown in Figure 4.16. A variable resistor (4.7kΩ) is included in the circuitry so that the contrast of the LCD can be altered for better visibility. The circuit schematics for the LCD interface circuit are shown in Appendix B.3.

Further visual information is provided to the HHU user with LEDs; both to indicate the power status of the HHU as well as the status of the mote contained within. Referring to Figure 4.15, the left most LED signifies that that HHU is powered on. The rightmost LEDs indicate that the mote is receiving, sending and/or powered (in that order from right to left).

Controlling the HHU is possible via two buttons and a switch; the switch provides power functionality (i.e. on and off) whilst the buttons provide functionality to a user within an inventory management environment. Their functionality is described and illustrated later in this chapter.

Rather than use the standard Mica2 antenna, which would most likely be restricted inside the HHU casing, an external solution was sought in the form of a Watson WSMA-889 [119]. The antenna is extendable and operable for radio frequencies of 25-1900MHz. This makes it suitable for the Mica2, but not the MicaZ. A female SMA connector is provided on the exterior of the HHU to attach the antenna; this could allow the MicaZ to be used if the antenna was swapped.

4.5.3. PIC Software

In order to operate effectively as an interface between the mote and LCD, the PIC requires programming. The MPLAB compiler software developed by Microchip [120] allows the programming of PICs using assembly language. Whilst it is possible to obtain compilers which allow the use of high level languages, such as C, MPLAB is free to use. The down side is that development can take significantly longer using assembly code because it is less intuitive.

The LCD used in this work is based on the Hitachi LM016L [121], which is very common, easy to obtain from many sources, and inexpensive. It is capable of displaying two lines of characters, with 16 characters per line. There are 14 pins in total used to provide power, control and to write data to the screen – the later of these are connected to the PIC as shown in the schematic in Appendix B.3.

LCD Packet				
START (0x40)	LCD LINE 1	MIDDLE (0x2C)	LCD LINE 2	END (0x26)
0	1-16	17	18-33	34

Figure 4.19 – LCD packet structure

Before the LCD can be used for displaying information it must be initialised (see Appendix A.2 for full assembly code). Prior to this, a long delay is required before writing any instructions to the LCD so that one can guarantee that all components in the HHU have powered up correctly. In addition, each instruction sent to the LCD requires a short delay in order to allow the LCD to respond appropriately.

Once the screen is initialised, the PIC may enter an idle state where it simply waits for an interrupt to occur – in this case, the interrupt will be the mote transmitting some data serially to it. Upon detecting data, the PIC will jump to its interrupt service routine which checks the data for reserved characters. These reserved characters are `0x40`, `0x26` and `0x2C` (or '@', '&' and ',' in ASCII[16]) which respectively mark the start, end and middle of a data packet. Each packet, which is sent in ASCII format, is 35 bytes in length. This is depicted in Figure 4.19. It is not encapsulated in the TinyOS packet structure, unlike other message types previously mentioned.

The PIC will first buffer a full packet, using the start and end characters as a guide, and then write the appropriate information to the screen. When in the write phase, the PIC will use the `0x2C` character to initiate an instruction to move the LCD cursor to the second line. After a full message has been written to the screen, the cursor moves back to the home position – that is – the first character on the first line. At this point the LCD is ready to be written to again in order to display further data. The LCD can be written to at a rate of at least 20Hz; it is probable that one could increase this but there has been no need to in this application.

[16] American Standard Code for Information Interchange character encoding scheme.

4.5.4. Mote Software

In order to support the functionality of the HHU, the implementation of TinyOS application code is required. Since the network logic used for the tags is the same as that used for the HHU, the code referred to in Section 4.4 has been updated rather than writing a completely separate application. This code is shown in Appendix A.1.

Beginning first with the LCD screen, the mote needs to send data to the LCD interface circuit, as described in the previous section. The data in this case is not required to be encapsulated in the TinyOS packet structure and so direct access to the microcontroller UART component is required; TinyOS provides this via the HPLUARTC component. This allows one to write data to the UART one byte at a time. Updates to the LCD screen occur when the HHU is:

- First powered on (see Figure 4.20) to tell the user that the HHU requires an update, which is possible from the GUI.
- Updated (see Figure 4.21), via the HHU_UPDATE message type or the select tag button is pressed.
- Receiving a HHU_RESPONSE message (see Figure 4.22), whereby the RSSI is calculated and displayed.

The buttons (see Figure 4.16) for selecting and pinging tags are wired to the mote interface board and act as interrupts to the Atmel 128L. The select tag button allows the user to select the tag they wish to locate from the HHU memory; these tags are stored in memory as a result of the HHU receiving HHU_UPDATE messages from the GUI. Subsequently pressing the ping tag button will initiate the sending of HHU_REQUEST messages to the selected tag. When the tag responds, the mote inside the HHU can sample the strength of the received signal using an ADC channel, meaning that the RSSI can be calculated using Equation 4-3 [93]. It is then displayed on the screen as shown in Figure 4.22. Chapter 7 looks further at the results of actually utilising this hand held unit as a positioning tool.

$$RSSI(dBm) = -50 \times \left(V_{batt} \times \frac{ADC\_VALUE}{ADC\_LEVELS} \right) - 45.5 \qquad \text{(Equation 4-3)}$$

Figure 4.20 – HHU display when first powered on

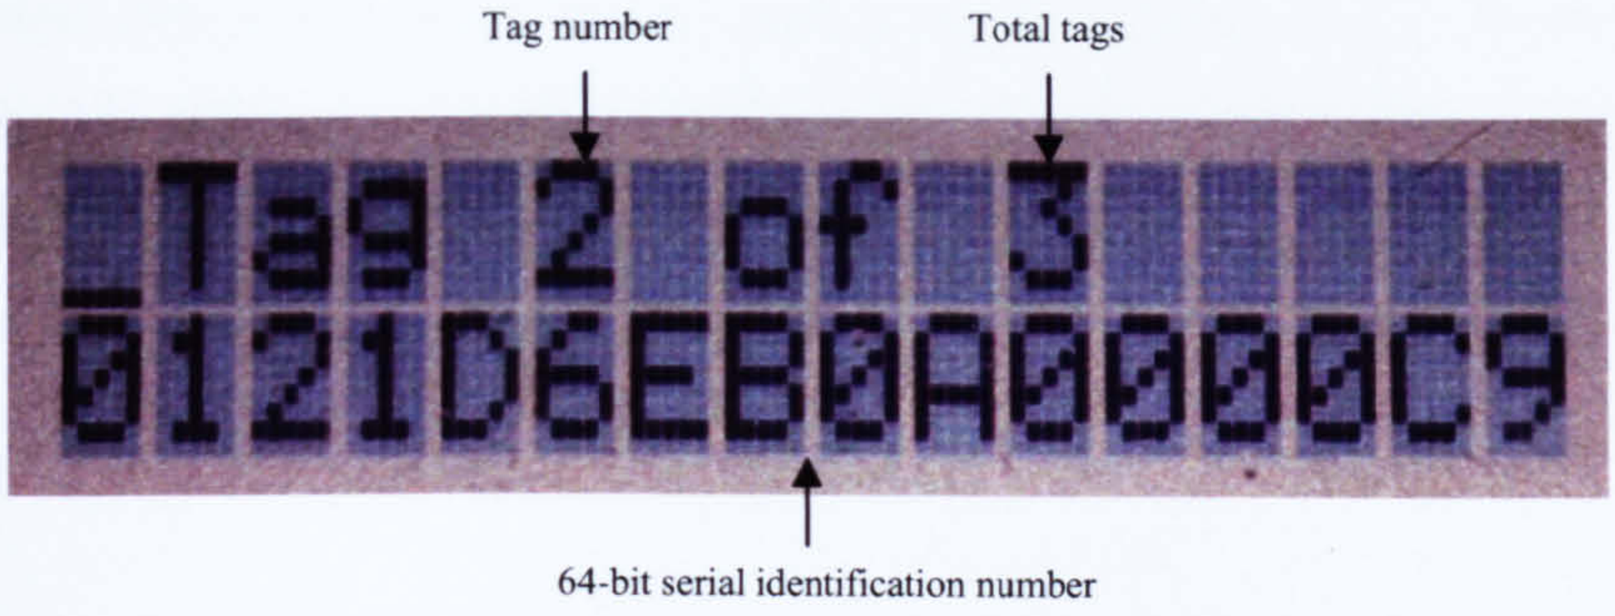

Figure 4.21 – HHU display after receiving a HHU_UPDATE message

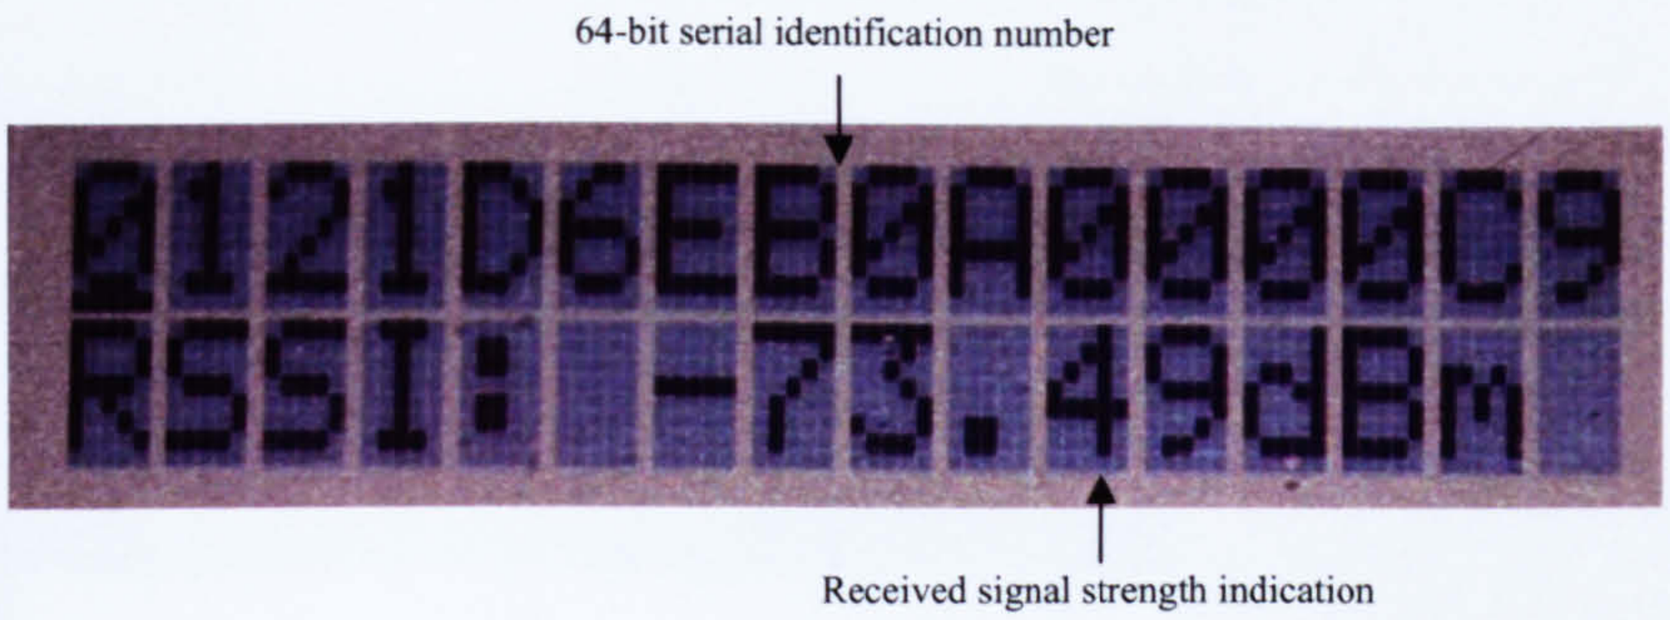

Figure 4.22 – HHU display after receiving a HHU_RESPONSE message

4.6. Graphical User Interface

The GUI ties together the tag and HHU components of the system – that is – one can use it discover tags, as well as communicate with them and the HHU. As with the other components, the GUI has been kept generic so that it can be shown to representatives from any industry and demonstrated as part of this inventory management system. The software is written using the Delphi [122] programming language which is an object orientated programming language for rapid development of applications under the Microsoft Windows operating system. This language was chosen primarily because of the portability that it affords, since, unlike other languages, it compiles everything required into the executable program. This means that no extra files are required for the application to operate on any Windows based computer. In addition, Delphi includes components which allow access to the Windows API (Application Programming Interface), providing functions such as serial port access. Such features make it appropriate for an application which requires rapid prototype development in addition to portability for demonstration purposes.

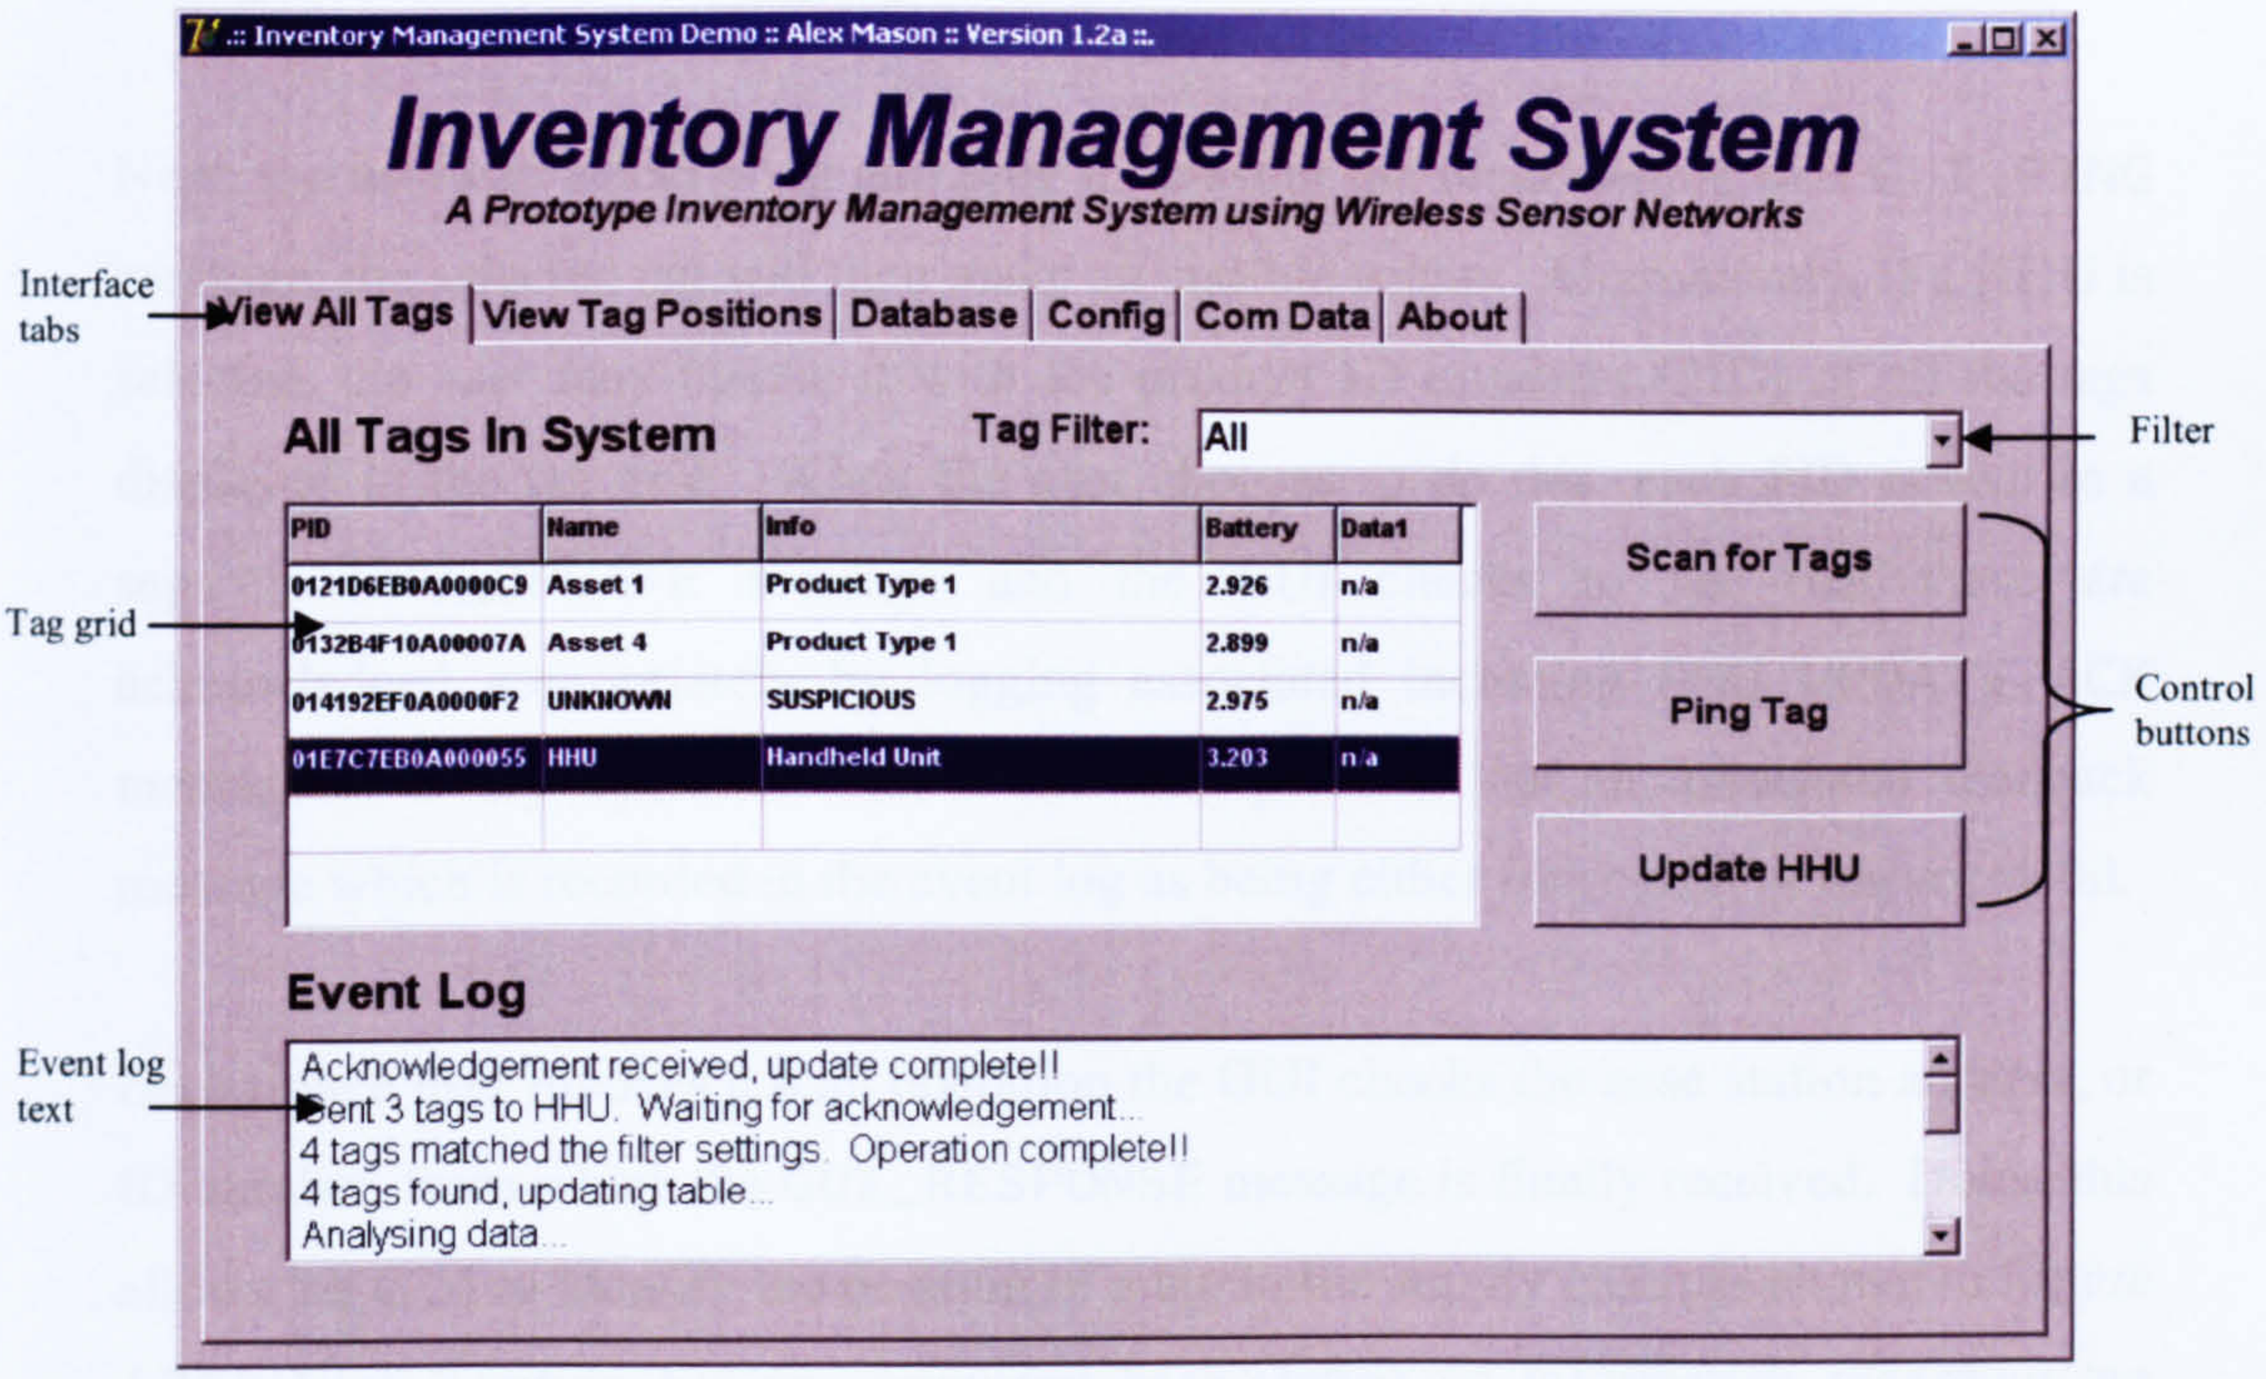

Figure 4.23 – Inventory interface

A number of interfaces make up the GUI allowing a user to perform a variety of operations pertinent to the concept of inventory management, as well as defining program parameters and viewing incoming and outgoing messages. Each interface can be viewed via a separate tab, as shown in Figure 4.23. This figure also shows the first interface that a user views when executing the application, allowing a user to:

- Scan for tags which are present in the inventory management system
- Ping discovered tags.
- Update the HHU if it is present.

Each of the above functions is invoked by the appropriately labelled button on the GUI itself. A user will typically scan for tags initially; doing so triggers the GUI to broadcast a `GUI_REQUEST` message type, and then it listens for `GUI_RESPONSE` messages. After a period of time it will then update the tag grid with any devices which have been discovered; if a filter is selected then only those tags with `info` fields matching the filter will be displayed. The `info` field for each tag is stored in a database (discussed shortly) – tags which are not present in the database are marked as suspicious.

Next, the user can select a tag and ping it, causing the broadcasting of a `GUI_PING` message; the selected tag will then make an audible sound. Alternatively, if a HHU is selected, the user may update it with the product ID numbers (PID) of all the tags displayed in the tag grid. When the user chooses to do this, each PID is sent in a separate `HHU_UPDATE` message and the GUI checks to see that these are acknowledged appropriately by logging associated incoming `HHU_UPDATE_ACK` messages. Every operation that a user can perform has an associated feedback message which is recorded in the event log as being either successful or unsuccessful.

Each time a user invokes a scan operation the GUI checks the base station address, or ID number, from which the `GUI_RESPONSE` message is finally received. Doing this allows the GUI to identify the position of a tag in the supply chain as shown in Figure 4.24. Each position has an associated base station so information regarding tag positions can be listed appropriately.

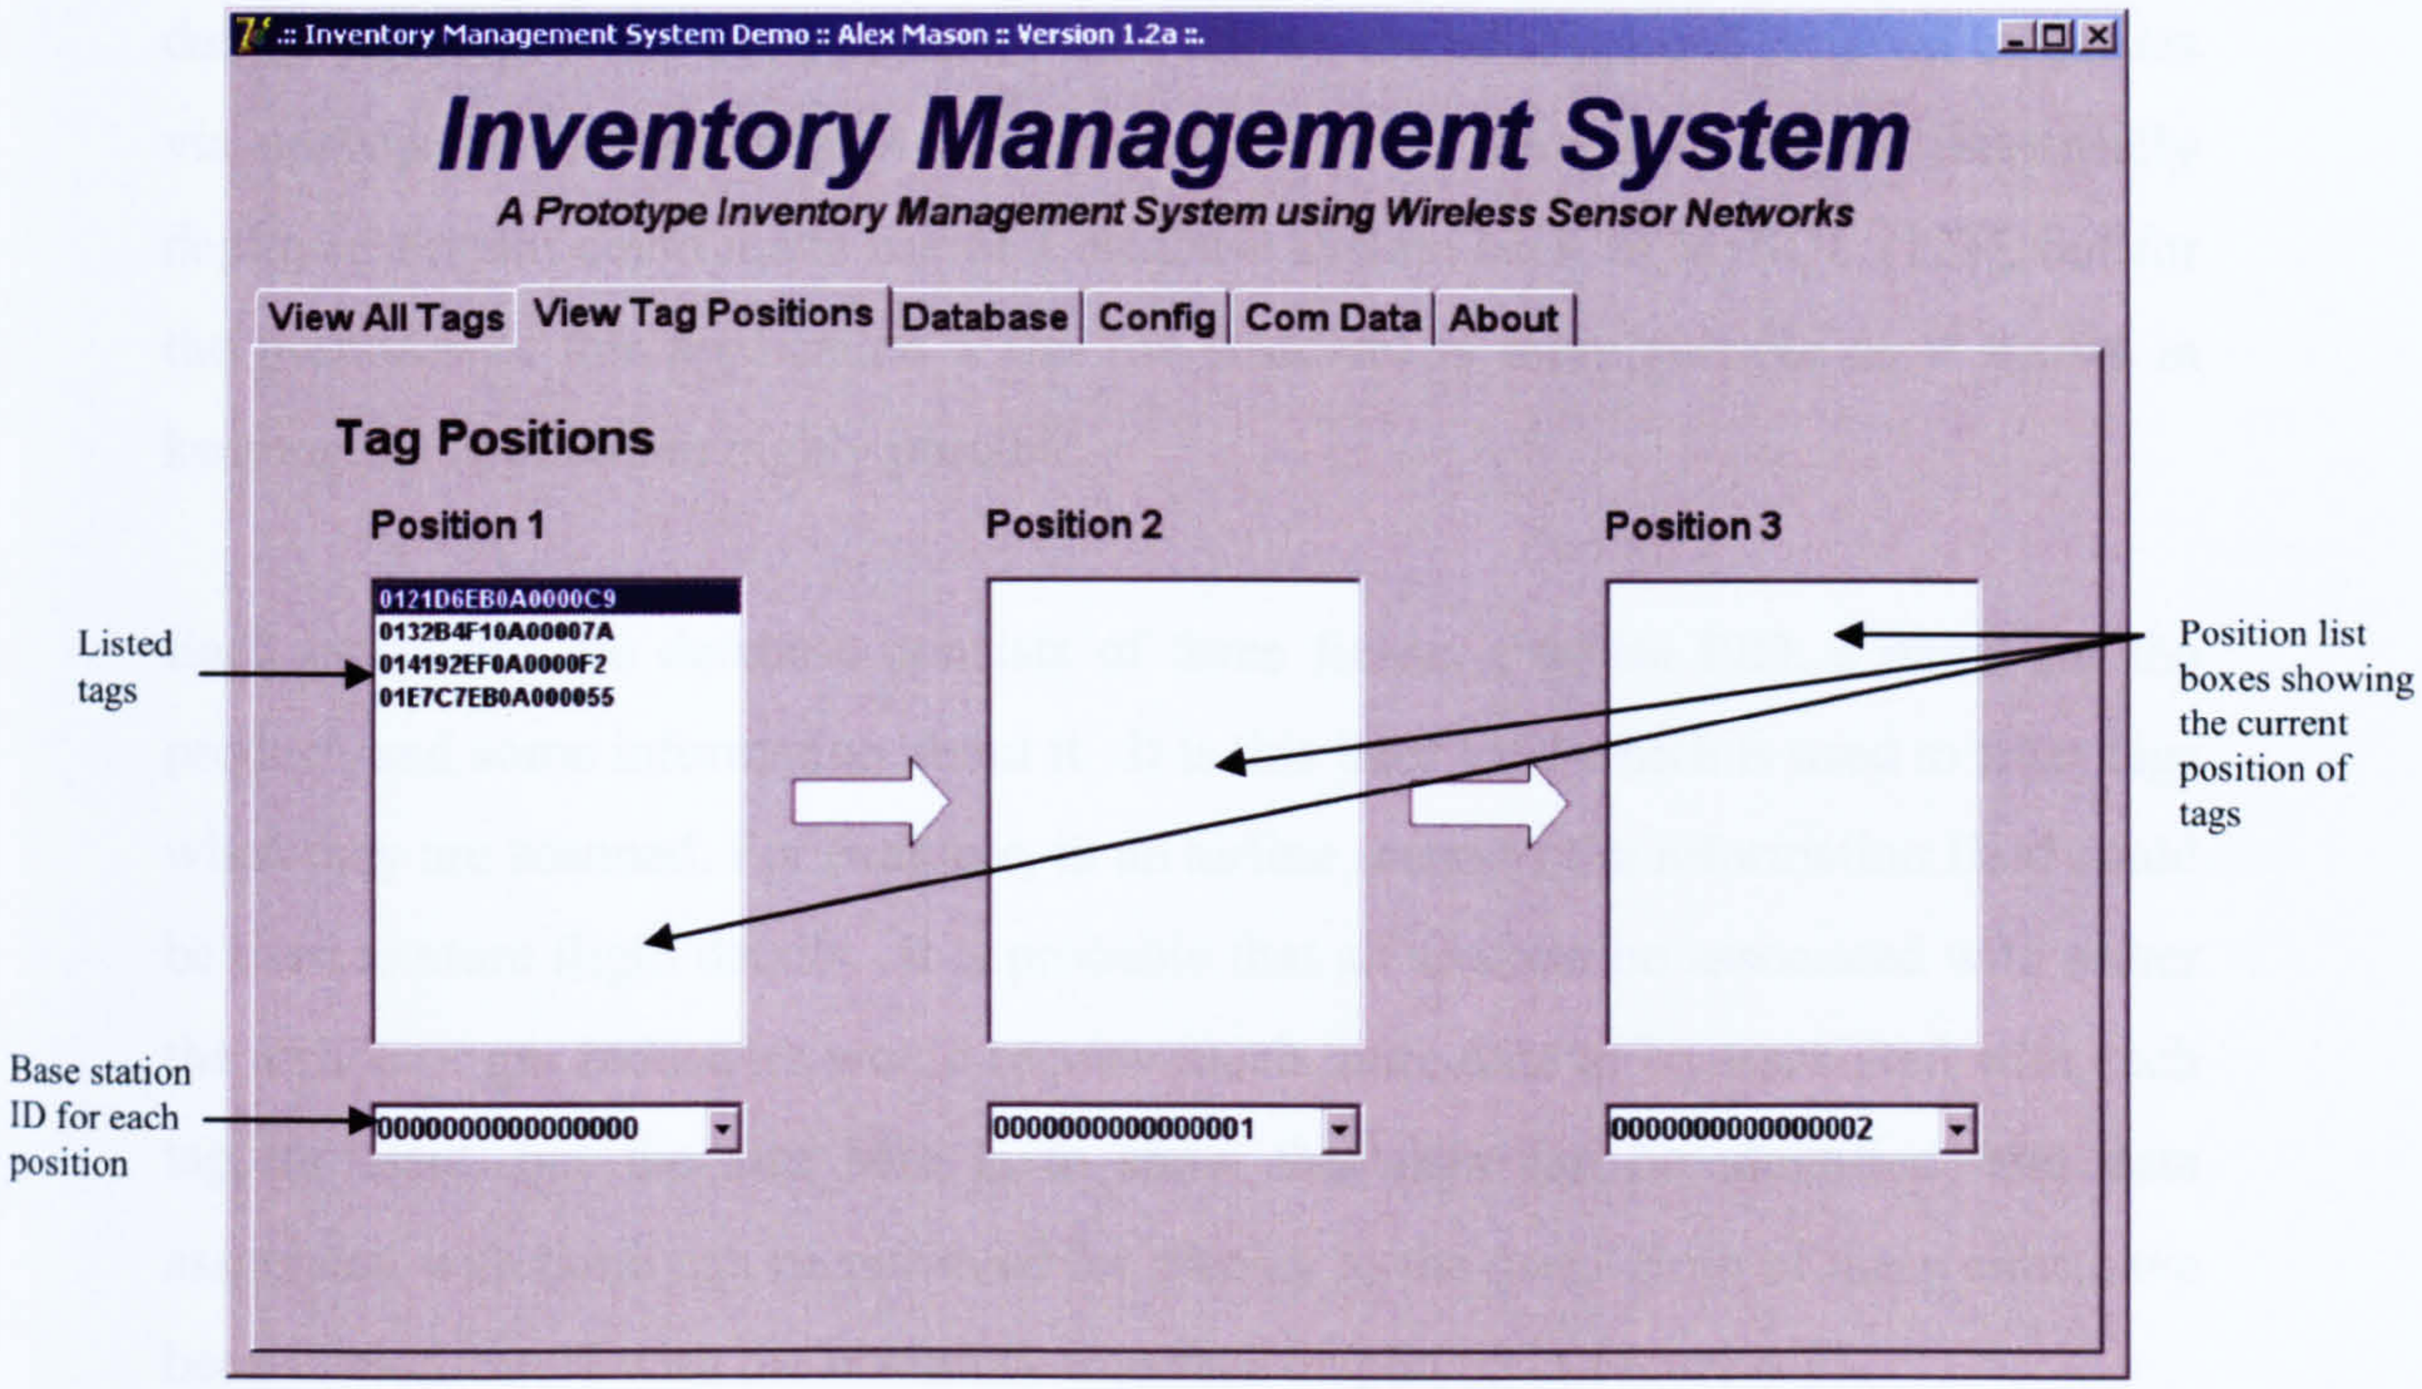

Figure 4.24 – Tracking interface

This tracking method is difficult to demonstrate with just 5 Mica2 motes but the idea is that at least one base station could be located at each point of interest in a supply chain. Taking an airport as an example a base station could be located at:

- Check in, where the passenger hands over their luggage
- The departure gate to check that luggage has traversed the luggage handling system successfully
- The aircraft cargo hold to ensure has luggage safely arrived at the aircraft

Taking these examples, they could be simply labelled position 1, 2 and 3 respectively; this idea is shown in the Figure 4.24 to preserve the generic nature of the application. One could increase the resolution of the system by simply increasing the number of points of interest and thereby adding more base stations.

Moving on, this application is supported by a database which was previously alluded to. This is important because it allows information to be associated with tags detected by the GUI, and means that the tags themselves do not have to store this data. The

data is stored in a flat file structure which allows the addition and removal of entries via appropriate buttons shown in Figure 4.25. It is imagined that an industrially deployed system could make use of a database system such as MySQL [123], but for the purposes of this application a flat file structure is more suitable as it assists in keeping the application highly portable.

Each entry into the database consists of three fields; a 64-bit PID, a name for the product, and some information about it. It is this later field which is used to filter tags when they are scanned. For example, in an airline scenario the information field could be used to store flight details. It is probable that an application associated with either the airline or gas industries would require much more data to be associated with each tag, or asset, but the aim here is to show that tags can be identified, and data associated with them can be retrieved for display to the user. Both of these aims have been demonstrated with the inventory interface displayed in Figure 4.23.

The application implements the TinyOS message structure [105] including CRC checking, which was previously described in Section 4.33 – the code for this implementation can be viewed in the Appendix, A.3.2. Figure 4.26 shows incoming and outgoing messages whilst the GUI is under typical operating conditions. In order to facilitate communications, a few configuration options are also included (see Figure 4.27). The configuration of the communications port is rather self explanatory; however the network configuration options are not. It is possible to set the number of iterations of each message sent as this can sometimes improve the reliability of the prototype through the provision of data redundancy. To assist this, a time out is used between message iterations to allow a tag to re-enter its sleep state, and therefore clear its memory of received messages. This allows it to then receive duplicate messages. The number of iterations should be kept to a minimum however, as whilst it may improve system reliability, it also increases the energy consumption of deployed tags. This is discussed further in Chapter 7, where the reasoning behind adding such redundancy is considered further.

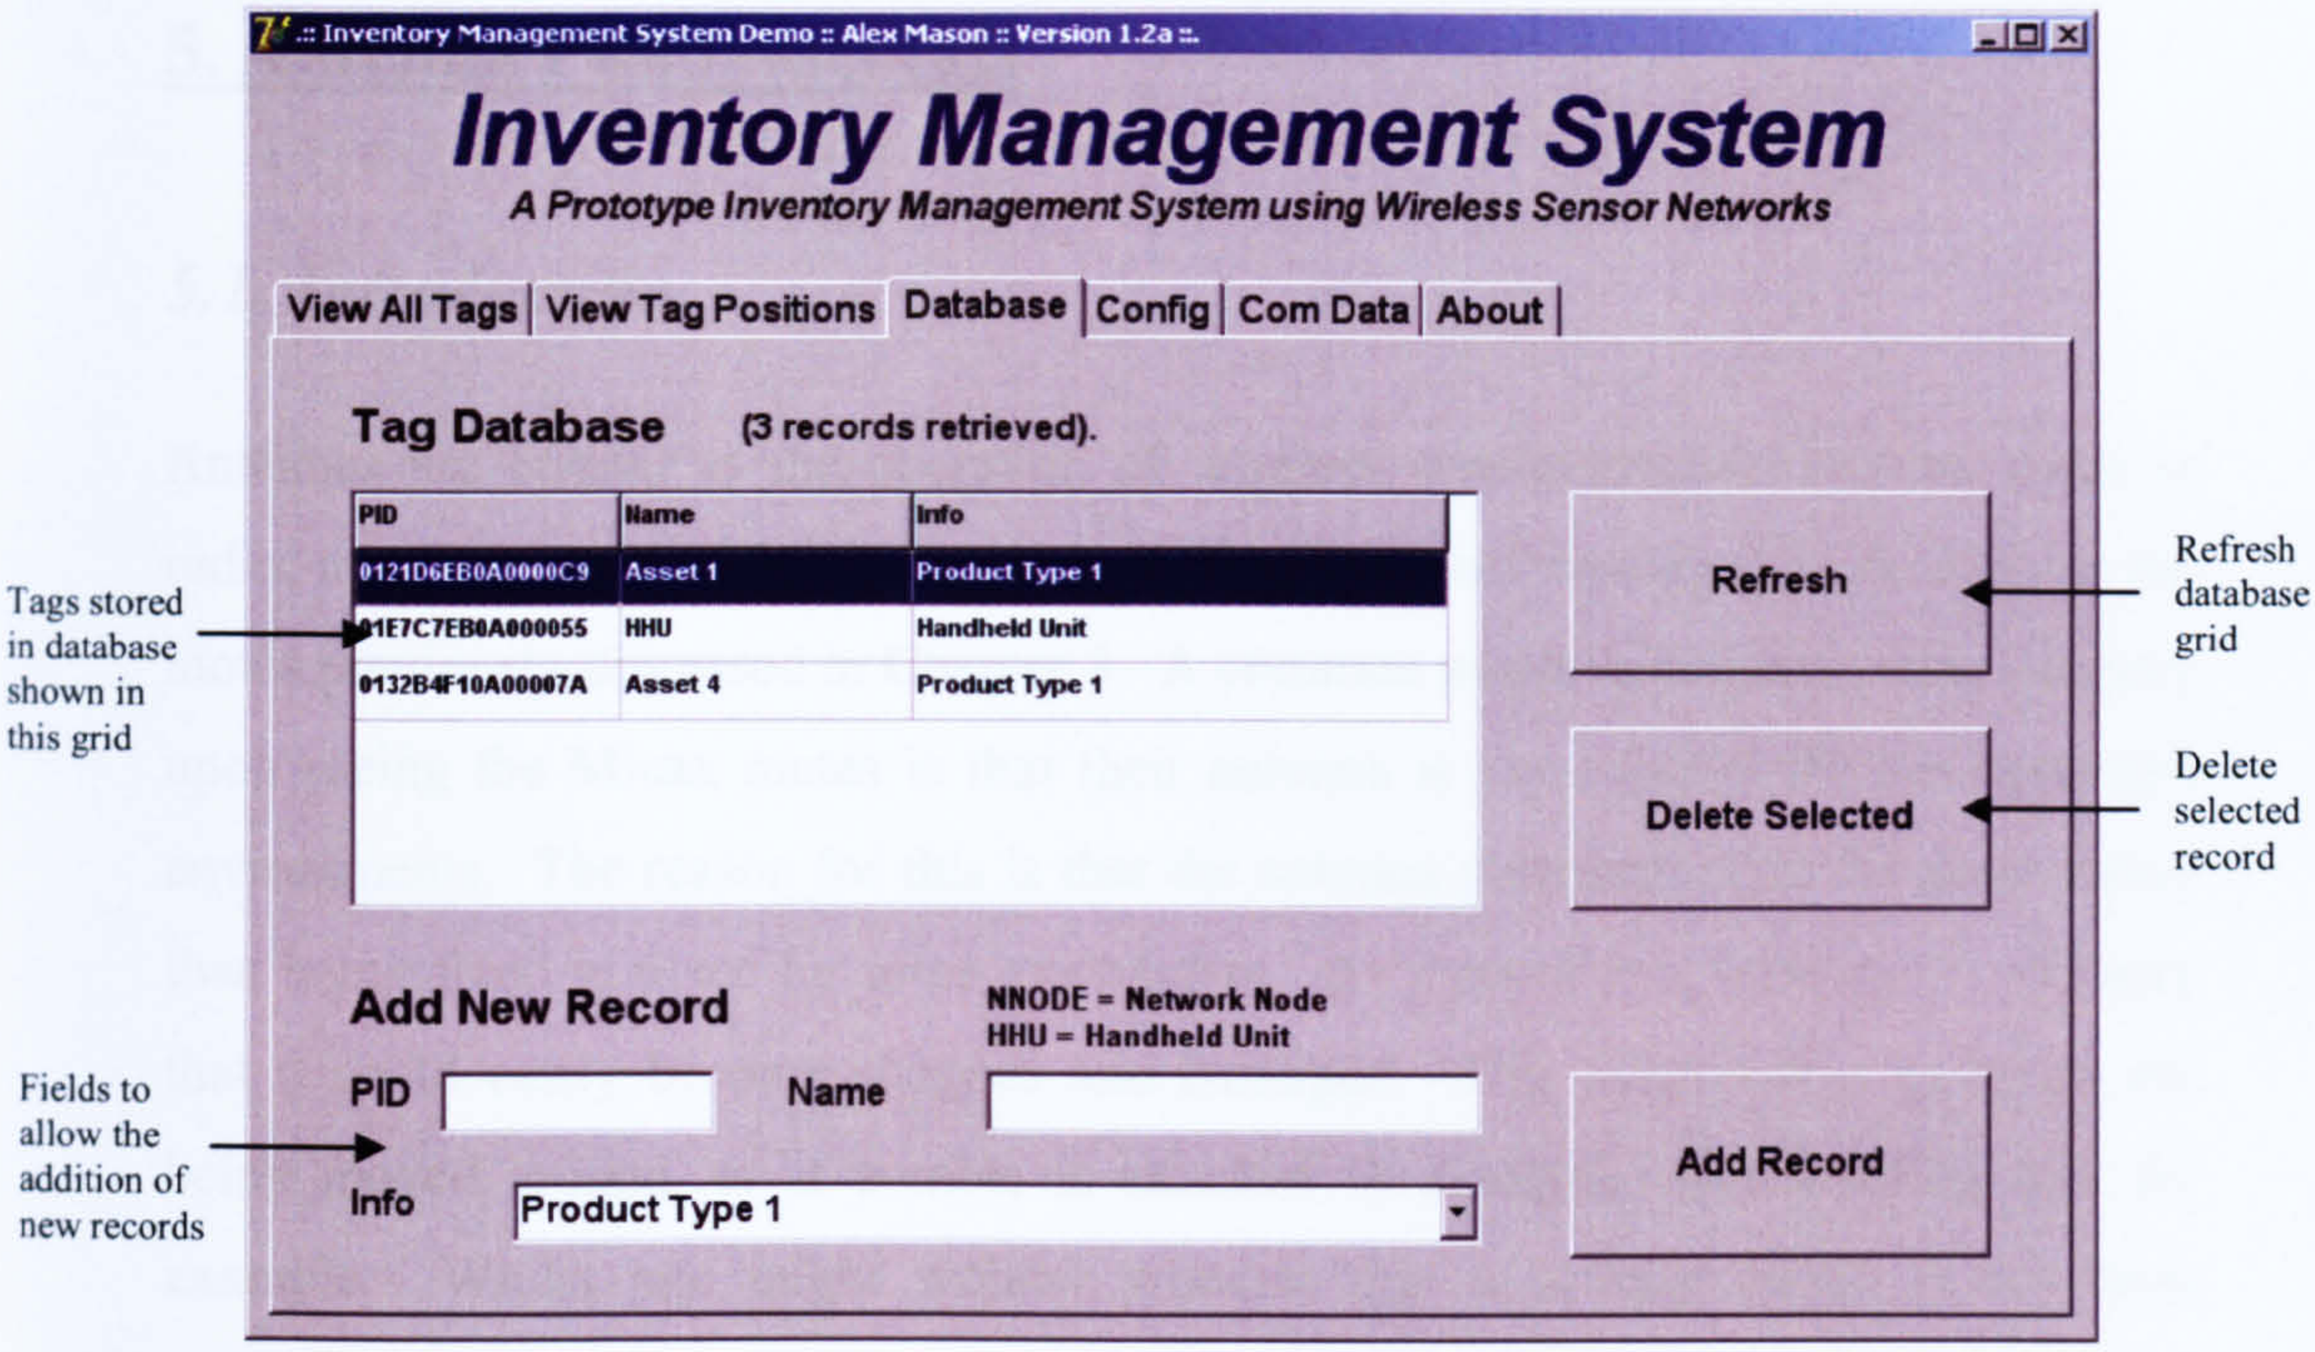

Figure 4.25 – Database interface

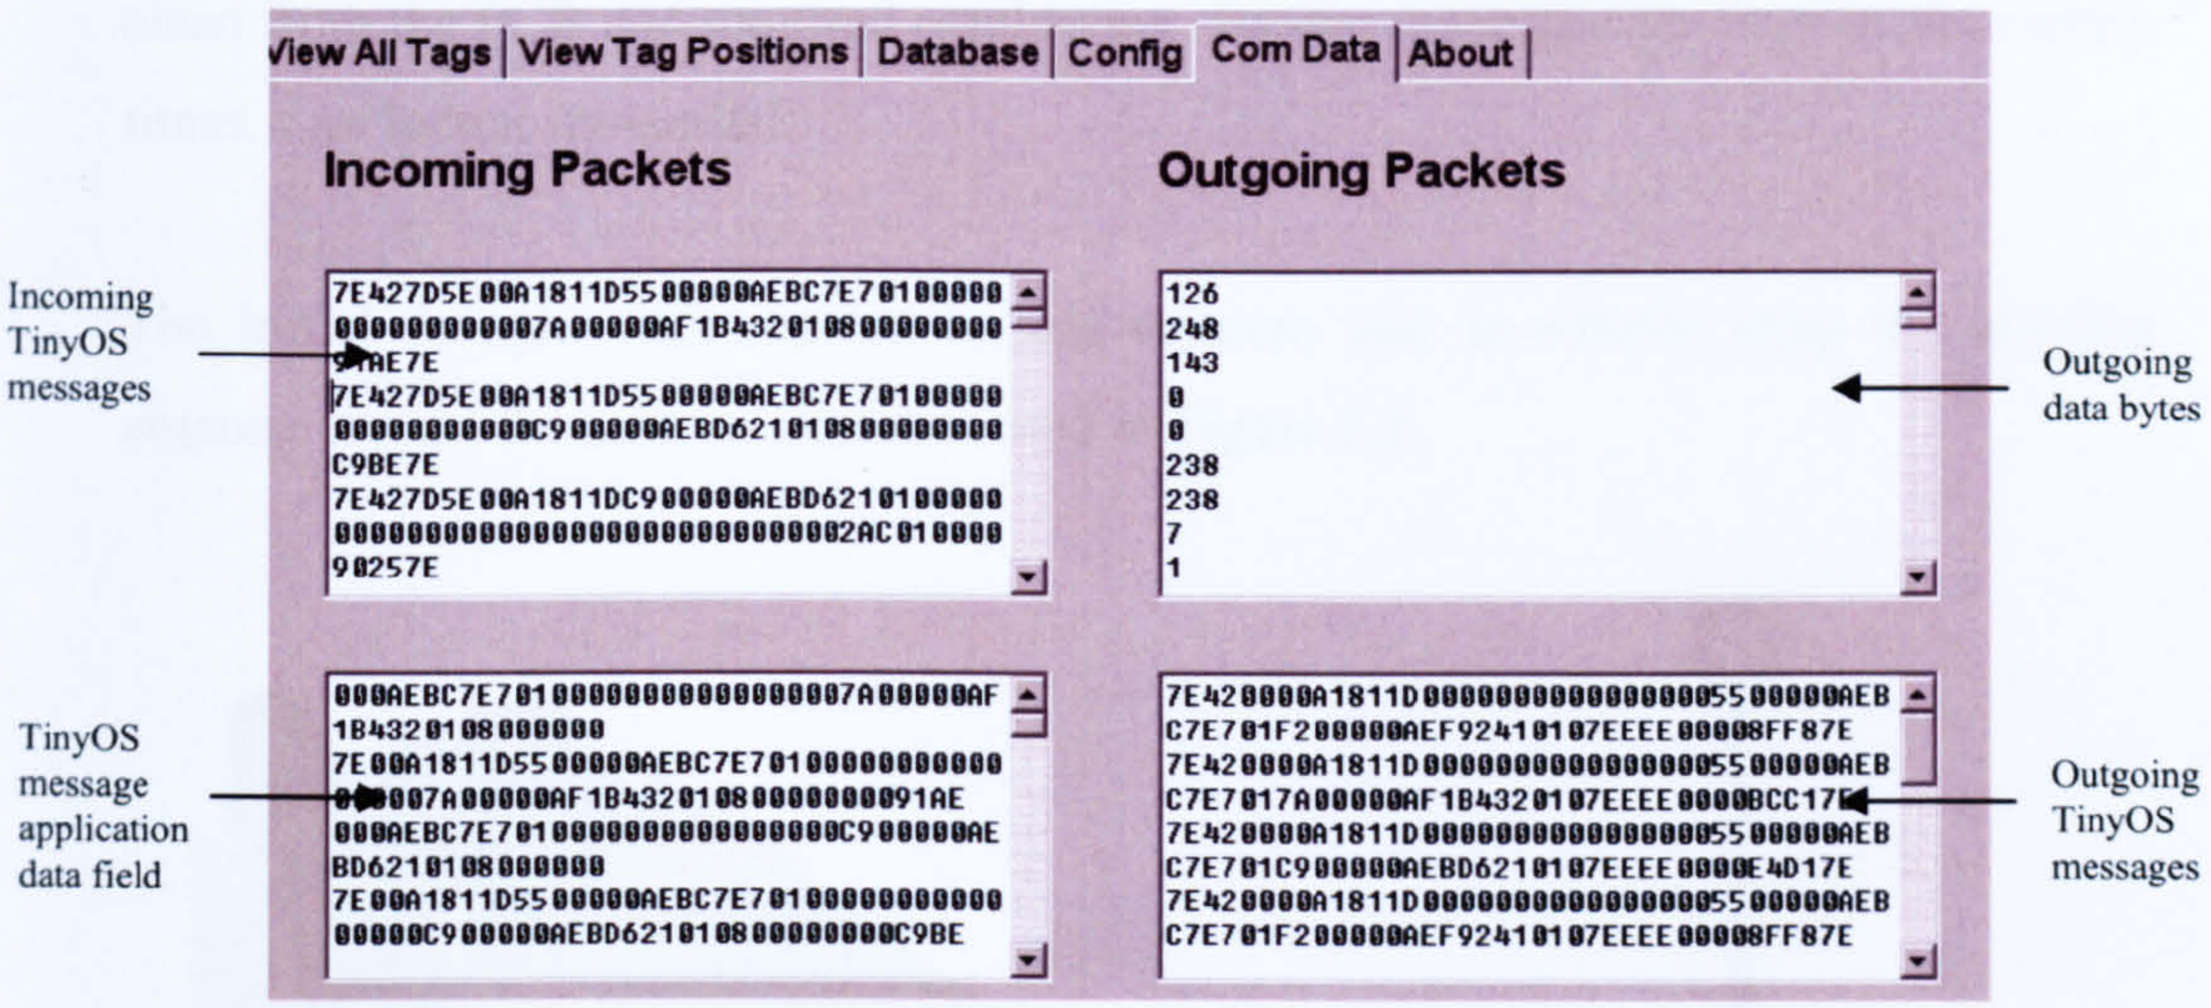

Figure 4.26 – Message interface

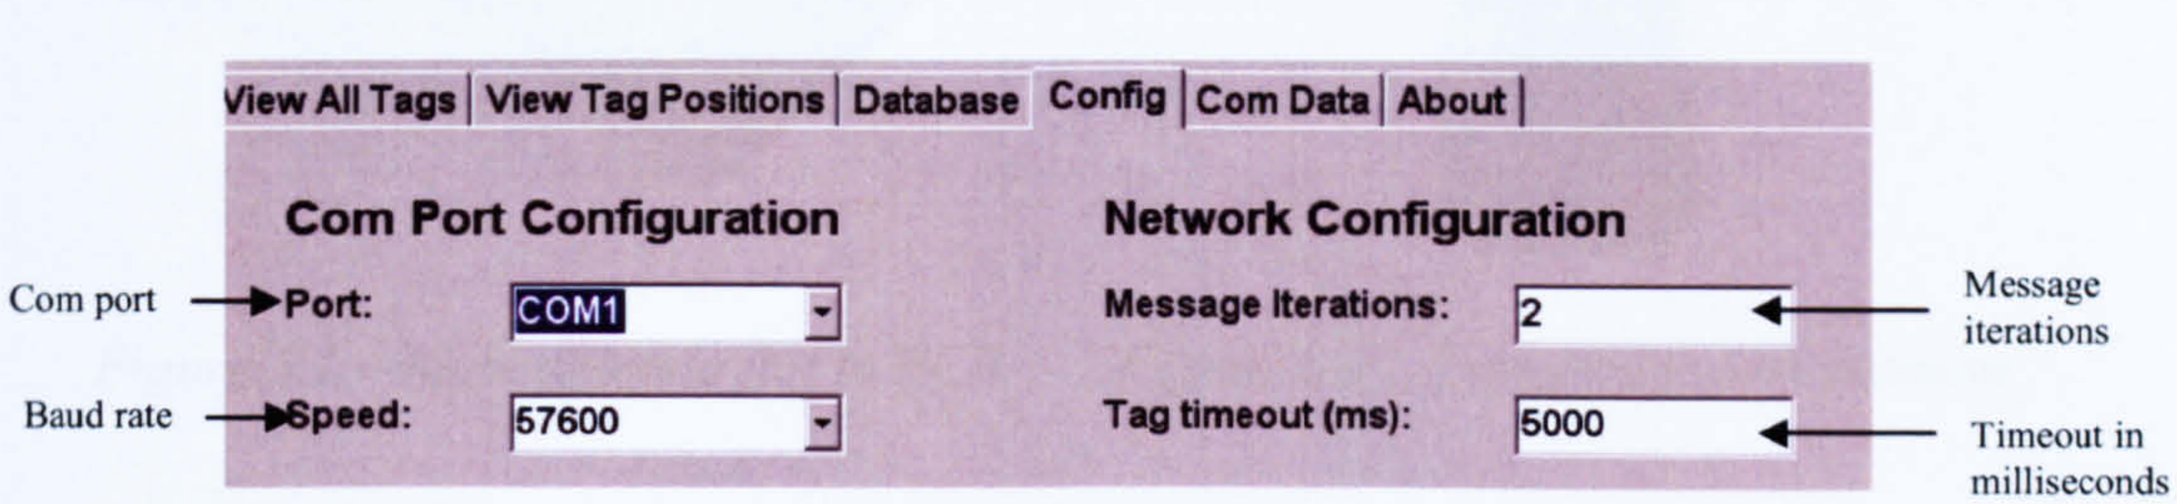

Figure 4.27 – Configuration interface

5. Antenna Fundamentals

5.1. Introduction

Antennas are critical to the operation of wireless communication systems such as radio, television, and mobile phones, as well as low cost sensing devices such as the motes previously discussed in Chapter 3. A common negative response from industry upon seeing the Micax motes is that their antenna is not suitable for the suggested environments. The reason for this is that the antenna protrudes from the mote rather than being fixed in place by some mechanism. As a result it is believed by industry that it could easily become snagged and damaged. This could occur as assets are being moved around, or if a mote is attached to a robotic spot welding arm for example. Whilst one might wonder whether this is a valid point, it has been experienced first hand during this project that the mote antenna and associated MMCX socket is quite fragile. On numerous occasions the socket has been ripped clean from the PCB and required resoldering; for one mote this has happened so many times it no longer transmits!

The initial thought with regards to this concern was to simply wrap the existing antenna around the mote as demonstrated in Figure 5.1.

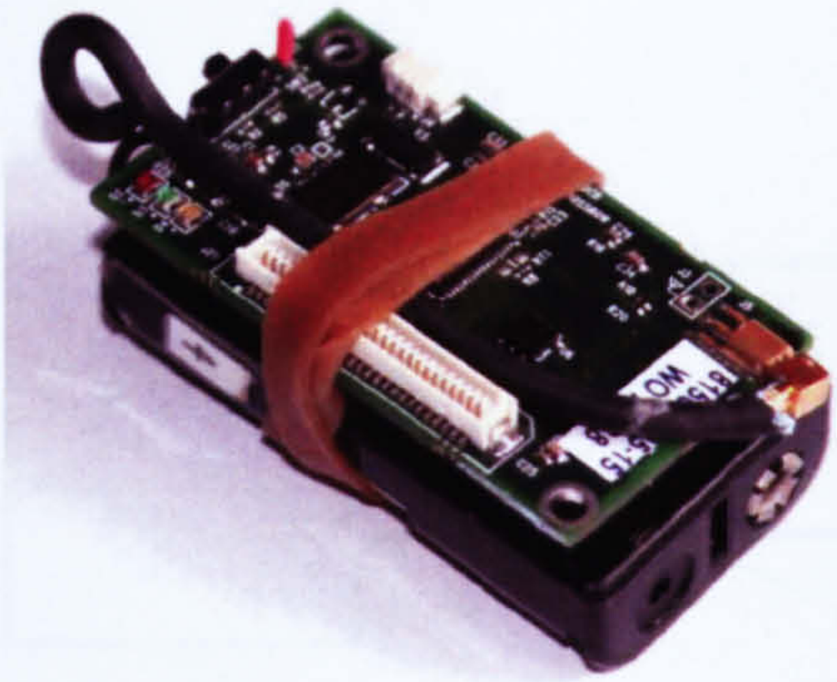

Figure 5.1 – Mote antenna flat to PCB

Figure 5.2 – Typical antenna position

A simple experiment, using one mote as a transmitter and another as a receiver, demonstrated that wrapping the antenna caused significant loss in signal strength. At a distance of 1m between the transmitter and receiver, the signal strength dropped by 10dBm when the antenna was wrapped around the mote compared to when it was normally orientated, as in Figure 5.2. As a result it was thought that designing a new antenna for the motes would be a more effective solution.

5.2. The Antenna System

Rather than just thinking of an antenna as a standalone component it can be useful to consider it as part of a system of components. Such a system would be made of three components, as listed below.

- *Source:* This is responsible for generating the signal which will be radiated out into space. In the case of the Mica2 motes, the CC1000 [76] transceiver performs this task.
- *Transmission line:* This carries the generated signal between the source and antenna.
- *Antenna:* This is responsible for ensuring that signals are effectively transmitted and received.

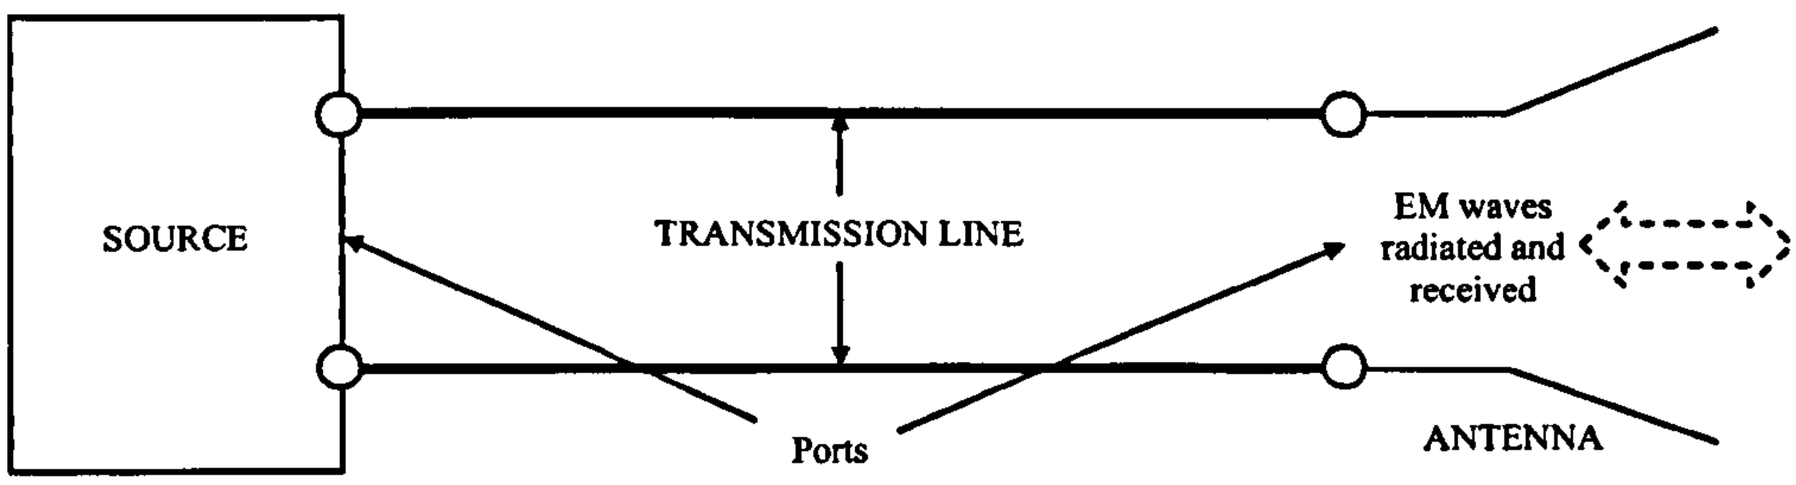

Figure 5.3 – A basic antenna system overview representation

An overview of a typical antenna system is shown in Figure 5.3; where appropriate, the components are shown to be connected by ports. Before delving into the proprosed solution to this antenna problem, it is important to give a brief discussion of how energy is conveyed between wireless devices.

5.3. Electromagnetic Waves

Electromagnetic (EM) waves transfer energy between components, as well as between two or more antennas which, in turn, allows wireless communication. EM waves are made up of electric (**E**) and magnetic (**H**) fields, both of which are vector quantities. If neither field varies then they can be considered as static and will not produce EM radiation. However, when the fields vary dynamically they become coupled such that a time varying electric field creates a time varying magnetic field and vice versa. It is this coupling which results in EM radiation. The **E** and **H** fields are taken to be perpendicular to one another and are always in phase, as shown in Figure 5.4.

Whilst Figure 5.4 represents the two component fields of an EM plane wave, when considering their propagation in space it can sometimes be simpler to represent EM waves in terms of rays and wavefronts. This is particularly useful when considering the behavioural characteristics of EM waves, a topic which will be discussed a little later in this chapter. A ray is a line which is drawn along the direction of propagation of a wave, such as the lines drawn in Figure 5.5 labelled R_A, R_B and R_C. Taking a look back at Figure 5.4, the *z* axis may represent a ray, as would any line which was drawn parallel to it. A wave front is a surface perpendicular to the direction of wave propagation. If EM waves are emitted from a source which radiates equally in all directions (i.e. – is isotropic) then the wavefronts will appear to be spherical, as shown in Figure 5.5 by the wavefronts A, B and C. Whilst the isotropic source is useful in theory, it is worth noting that it cannot be created in practise.

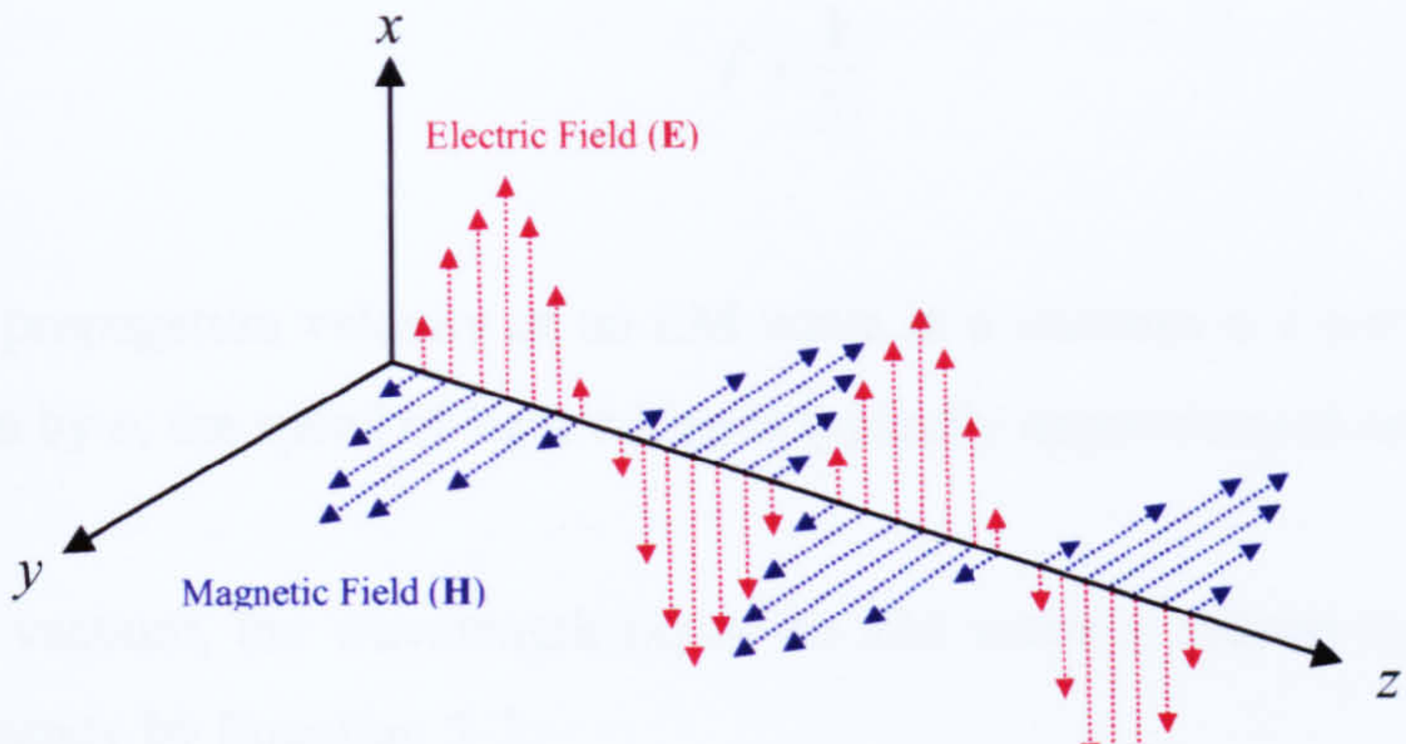

Figure 5.4 – Electromagnetic travelling wave visual representation

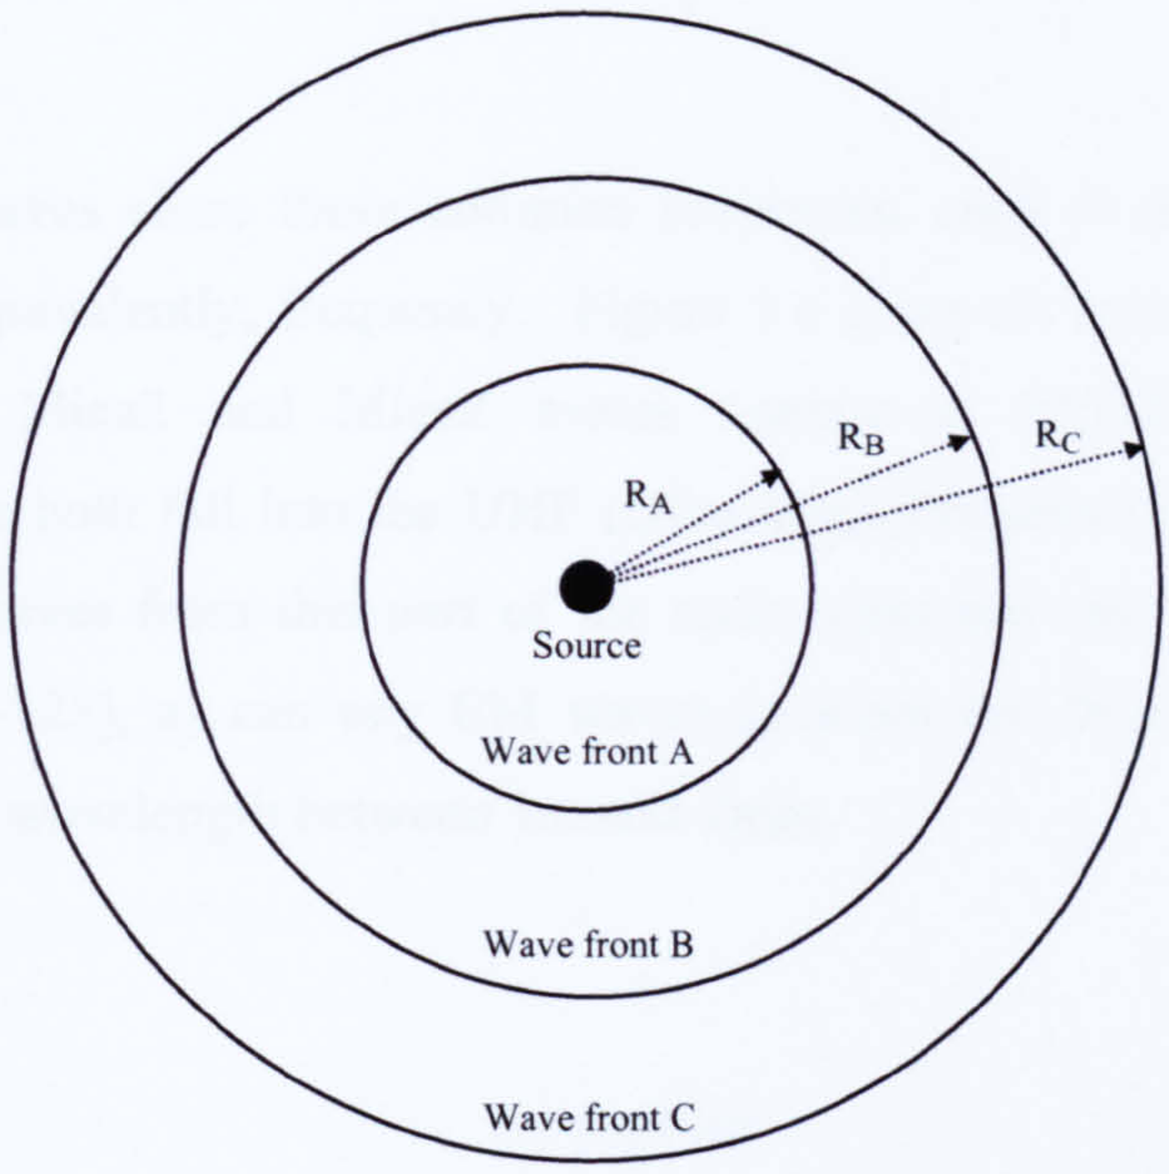

Figure 5.5 –Spherical wave in free space

There are many different types of EM wave, all of which belong to a family of waves called the EM spectrum. Members of this family include visible light, x-rays, gamma rays, as well as radio waves. All EM waves share a set of basic properties:

- An EM wave consists of electric and magnetic field intensities, as previously described, which oscillate at the same frequency (f). Frequency is inversely proportional to the period (T) as shown in equation 5-1.

$$f = \frac{1}{T} \qquad \text{(Equation 5-1)}$$

- The propagation velocity of an EM wave in a vacuum is a universal constant given by c, the speed of light which is typically approximated to $3 \times 10^8\ ms^{-1}$.

- In a vacuum, the wavelength (λ) of an EM wave is related to its oscillation frequency by Equation 5-2.

$$\lambda = \frac{c}{f} \qquad \text{(Equation 5-2)}$$

Whilst all EM waves share these common properties, each is defined by its own wavelength or, equivalently, frequency. Figure 5.6 gives an overview of the entire spectrum. The Mica2 and MicaZ motes operate at 915MHz and 2.45GHz respectively; these both fall into the UHF (Ultra High Frequency) band of the radio spectrum. EM waves from this part of the radio spectrum can also be classed as microwaves [124-125], as can any EM waves between the frequencies 0.3GHz to 300GHz or with a wavelength between 1m and 1mm.

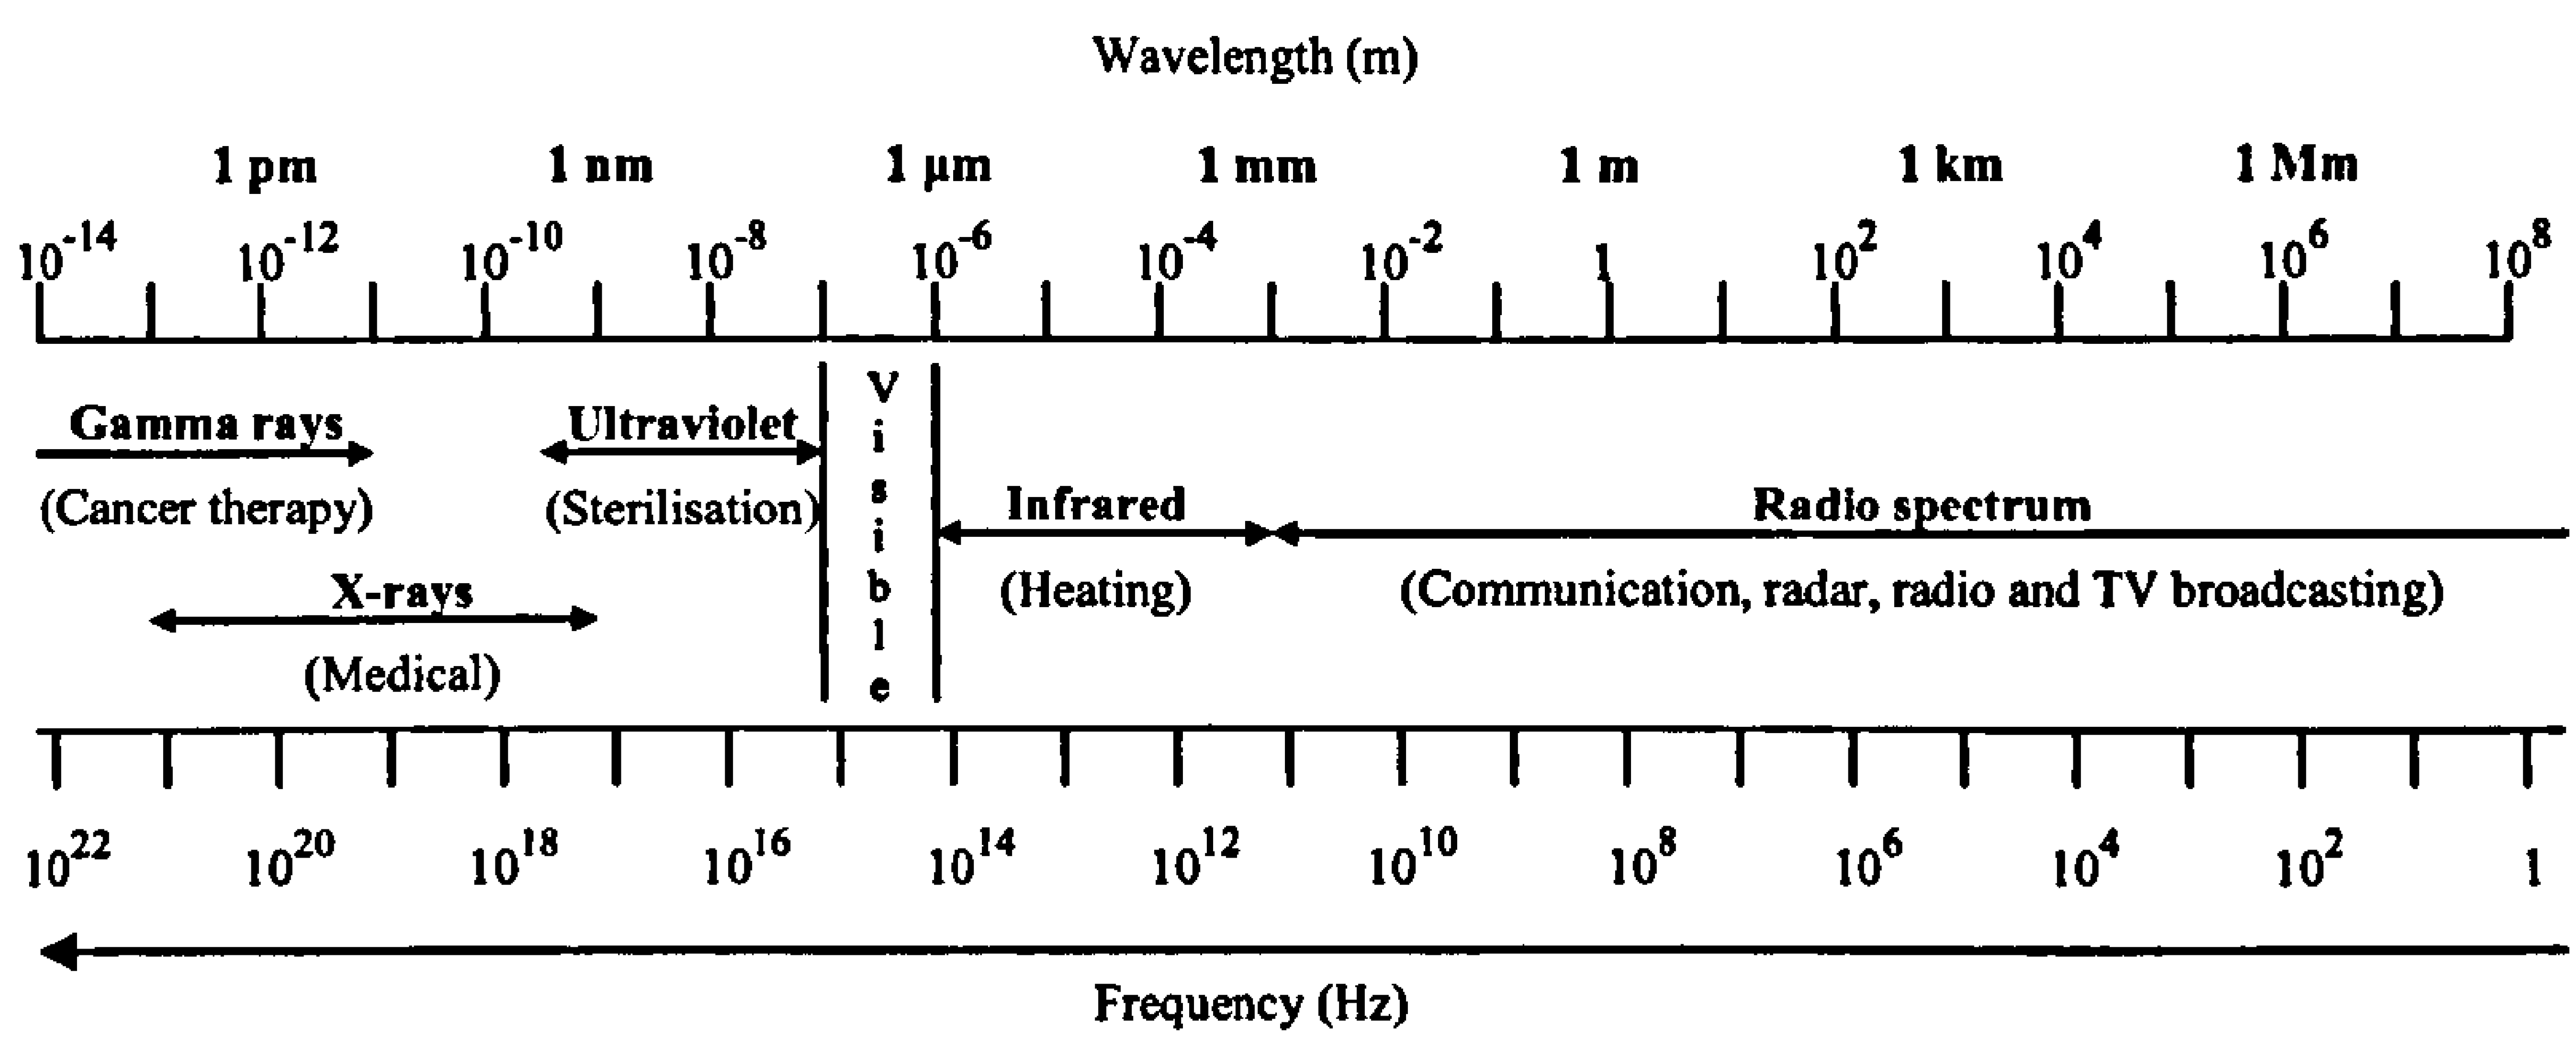

Figure 5.6 – The electromagnetic spectrum

5.4. Factors Effecting Electromagnetic Wave Propagation

5.4.1. Polarisation

An EM wave propagating in space has an associated polarisation which is taken to be in the direction of the **E** field. As an example, Figure 5.4 can be said to be linearly polarised since the E field has a particular direction in space for all values of z – that is, it is always directed in the x direction. If x is vertical in respect to the Earth it can also be said to be vertically polarised, whilst if x were horizontal with respect to the Earth then it would be horizontally polarised.

EM waves are not generally linearly polarised, they can, for example, be circularly polarised too. In this case the **E** field rotates around the z-axis such that the wave advances through space in a screw-like motion making a single rotation for every wavelength it advances. Circularly polarised waves are created when two linearly polarised waves are combined, usually because they have been launched from the same source antenna and in the same direction. For circular polarisation to be achieved one wave has to be polarised at 90° to the other and also be 90° out of phase. For true circular polarisation the amplitude of each linearly polarised component must be equal, or else elliptical polarisation results.

5.4.2. Attenuation

The attenuation of an EM wave refers to a reduction in its magnitude as it moves through a medium (e.g. – air, wood, metal). Attenuation may be due to the medium itself absorbing the energy of the EM wave as well as due to the spreading of energy as it propragates away from its source.

When an EM wave is radiated into space from a source, it will spread to some degree depending how the source focuses this energy; this topic is dealt with later in relation to antennas. For the moment though it is easier to consider the case of an isotropic source which produces perfectly spherical wavefronts, such as that shown previously in Figure 5.5. Considering these wavefronts, if no energy is lost as a result of

propagation in the medium then each wavefront will represent an equal amount of energy. However, as the distance (R) from the source increases the energy per unit area, or power density (P), will decrease. This relationship is known as the inverse square law of radiation, and can be expressed as in Equation 5-3 [126].

$$P = \frac{1}{R^2} \qquad \text{(Equation 5-3)}$$

In some literature this relationship is referred to as free space loss, although this is a little misleading if one is familiar with the law of conservation of energy [75]. There is no energy conversion however, just a decrease in its density. It is this feature of EM wave propagation that the HHU, discussed in Chapter 4, seeks to utilise for positioning via the mote RSSI feature. However, as will be discussed in this and other chapters, the relationship between distance and signal strength is not as simple as that defined by the inverse-square law.

The second cause of EM wave attenuation may occur due to the absorption of energy by the propagation medium. A quantity known as the skin depth (δ) serves to characterise how well an EM wave can penetrate a medium. It defines the penetration distance of an EM wave before it decays to e^{-1} (36.8%) of its original magnitude [127] and can be expressed generally for any propagation medium by Equation 5-4, where α is the attenuation constant of a medium. For a conductor, α can be defined by Equation 5-5 where f is the propagating frequency, μ is the magnetic permeability of the medium and σ is its conductivity. The concept of skin depth arises from related phenomena known as the skin effect, which describes the tendency of an alternating current (AC) to distribute itself most densely at the surface of a conductor rather than at its core.

$$\delta = \frac{1}{\alpha} \qquad \text{(Equation 5-4)}$$

$$\alpha = \sqrt{\pi f \mu \sigma} \qquad \text{(Equation 5-5)}$$

Material	Skin Depth (m)		
	131kHz	**915MHz**	**2.45GHz**
Copper	2.05×10^{-7}	2.45×10^{-9}	1.50×10^{-9}
Iron	6.97×10^{-9}	8.34×10^{-11}	5.10×10^{-11}
Tin	5.20×10^{-7}	6.22×10^{-9}	3.80×10^{-9}
Aluminium	2.63×10^{-7}	3.15×10^{-9}	1.93×10^{-9}
Steel	1.49×10^{-6}	1.78×10^{-8}	1.09×10^{-8}

Table 5.1 – Skin depth for various conducting metals at different frequencies

Whilst α can be defined for many materials, conducting materials (i.e. – metals) are of particular relevance to this work since they are of concern when considering using radio waves in metallic environments. A perfect conductor has an infinite conductivity ($\sigma = \infty$, hence $\alpha = \infty$) and therefore zero skin depth ($\delta = 0$). For a perfect dielectric on the other hand, $\sigma = 0$, $\alpha = 0$ and hence $\delta = \infty$; therefore there is no absorption. Table 5.1 shows the skin depth for various conducting metals which has been calculated using a combination of Equations 5-4 and 5-5.

Taking the packaged gas industry as an example, the UK Health and Safety Executive recommends a specification for steel gas cylinders. This specification defines the minimum cylinder wall thickness [128] (t) of a cylinder, as shown in Equation 5-6, where P_1 is the is the test pressure (measured in bar), f_e is the maximum equivalent stress (measured in Nmm^{-2}) and D_0 is the height of the cylinder.

$$t = \frac{0.3 P_1 D_0}{7 f_e - 0.4 P_1} \qquad \text{(Equation 5-6)}$$

Taking a typical cylinder pressure of 300bar, P_1 is assumed to be 1.5 times this value. f_e is given as $702Nmm^{-2}$. Assuming D_0 is 500mm, this would indicate that t would have to be at least 14.3mm. It is possible to see from Table 5.1 that the skin depth for steel at all of the frequencies listed is in the order of μm to tens of nm. Therefore, one wall of a 500mm tall gas cylinder would be in the order of almost ten thousand skin depths thick. To pass completely through such a gas cylinder, an EM wave would actually have to pass through two such walls, and the gas contained within!

Passive RFID systems which utilise backscatter coupling may only transmit in the order of μW of power [129] and so only small amounts of metallic material are required in order to attenuate the signal to a point where it is indistinguishable from background noise present in an environment. The Micax motes can transmit at power levels in order of mW, a factor of a thousand greater than passive RFID, due to their power supply and therefore have a significant advantage in terms of absorbtion. By coupling this with the multihop routing approach described in Chapter 4, it is hoped that active transmission over short distances will overcome the issues experienced by passive RFID systems.

The 131kHz frequency is also included in the Table 5.1 in order to consider the claims made by the developers of the RuBee system [55] mentioned in Chapter 2. Specifically, they say that the system suffers little problem with metallic surroundings. However the theory, despite agreeing that lower frequencies do fair better in metallic environments, suggests that the skin depth at 131kHz is still only approximately 1.5μm for metals such as steel. As a result, one has to wonder if RuBee holds much practical advantage in the real world solely based on its operating frequency, particularly since they present little evidence to support their claims.

5.4.3. Optical Priniciples

In Section 5.2 it was mentioned that EM waves from different parts of the EM spectrum share a number of common properties. As a result it is common to find radio wave phenomena explained in a way similar to that of the behaviour of light in terms of refraction, reflection, interference and diffraction. Knowledge of these phenomena is important because it offers some insight into how radio waves propagate and interact with objects in the real world, despite not being able to physically see them.

Beginning with refraction, this refers to when an EM wave comes across a change in medium density and therefore changes direction due to a change in speed – the amount of refraction which takes place is related to the ratio of the two medium densities. This can be observed in practice by passing light at an incident angle of less

than 90° into a glass block. Since the glass block is denser than the air surrounding it refraction, or bending, of light occurs.

When a wave strikes a denser medium it may not be totally refracted. Instead some of the power may be reflected. Reflection refers to the creation of a new wave that travels away from the point at which there is a change in medium density; therefore the incident wave can be observed to split into two waves – one which is refracted and travels through a medium boundary and also one which is reflected away from that boundary; this is illustrated in Figure 5.7. The ratio of reflected power to incident power is referred to as the reflection coefficient (Γ), and is expressed mathematically in Equation 5-7.

$$\Gamma = \frac{reflected}{incident} \qquad \text{(Equation 5-7)}$$

Γ is a complex number which takes into account the magnitude and phase of the reflected and incident signals. Neglecting phase (i.e. – the imaginary part of Γ), a reflection coefficient of 1 would represent a perfect open circuit whilst -1 would represent a perfect short circuit. A value of 0 represents a perfect match, were no reflection occurs; this is important in terms of transmission lines and antennas, two topics which will be discussed later in this chapter.

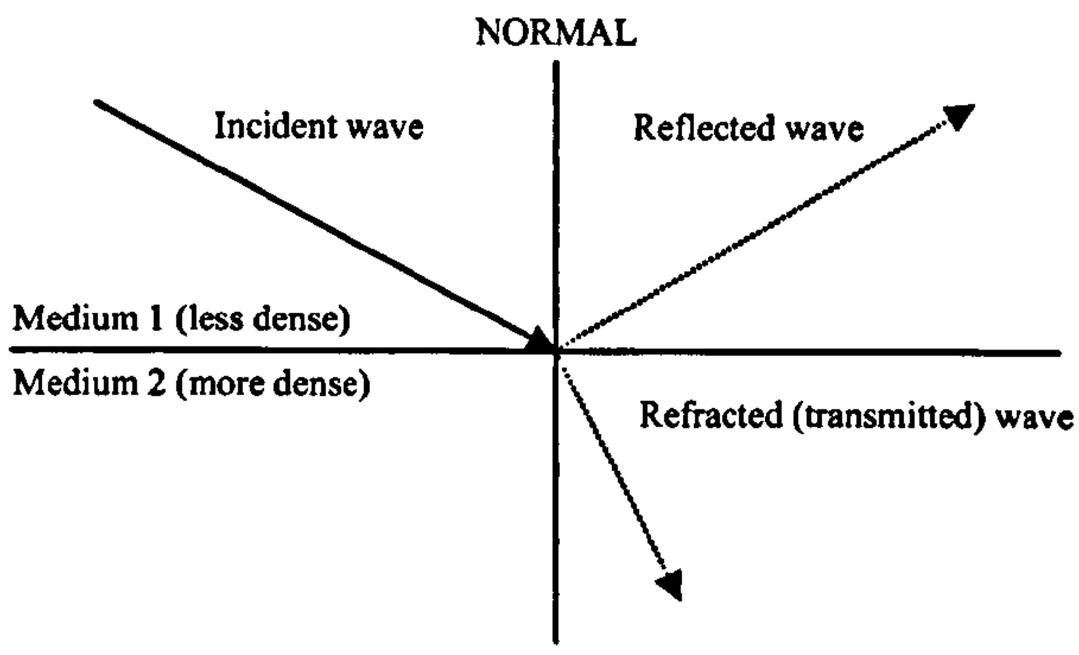

Figure 5.7 – Waves being reflected and refracted

If one assumes that no energy is lost from the incident wave, it follows that the sum of the reflected and refracted power is equal to the incident power. The fraction of the total incident power refracted, or transmitted, through the media boundary is given by the power transmission coefficient *(T)*, which is calculated as in Equation 5-8.

$$T = 1 - |\Gamma|^2 \qquad \text{(Equation 5-8)}$$

Moving on from reflection and refraction, diffraction is another important behavioural property of EM waves, which refers to the deviation of wave propagation from its usual straight-line path. Diffraction accounts for the apparent bending of waves around the edges of objects. Figure 5.8 shows a plane wave incident to a solid barrier which has an opening. In part (a), the size of the gap is very large in comparison to the wavelength, and therefore, apart from some small side effects, the wave emerges and continues to propagate in a straight line. In parts (b) and (c), the gap size is equal to and less than the wavelength, respectively, and so significant diffraction occurs; the smaller the gap the greater the effect of diffraction.

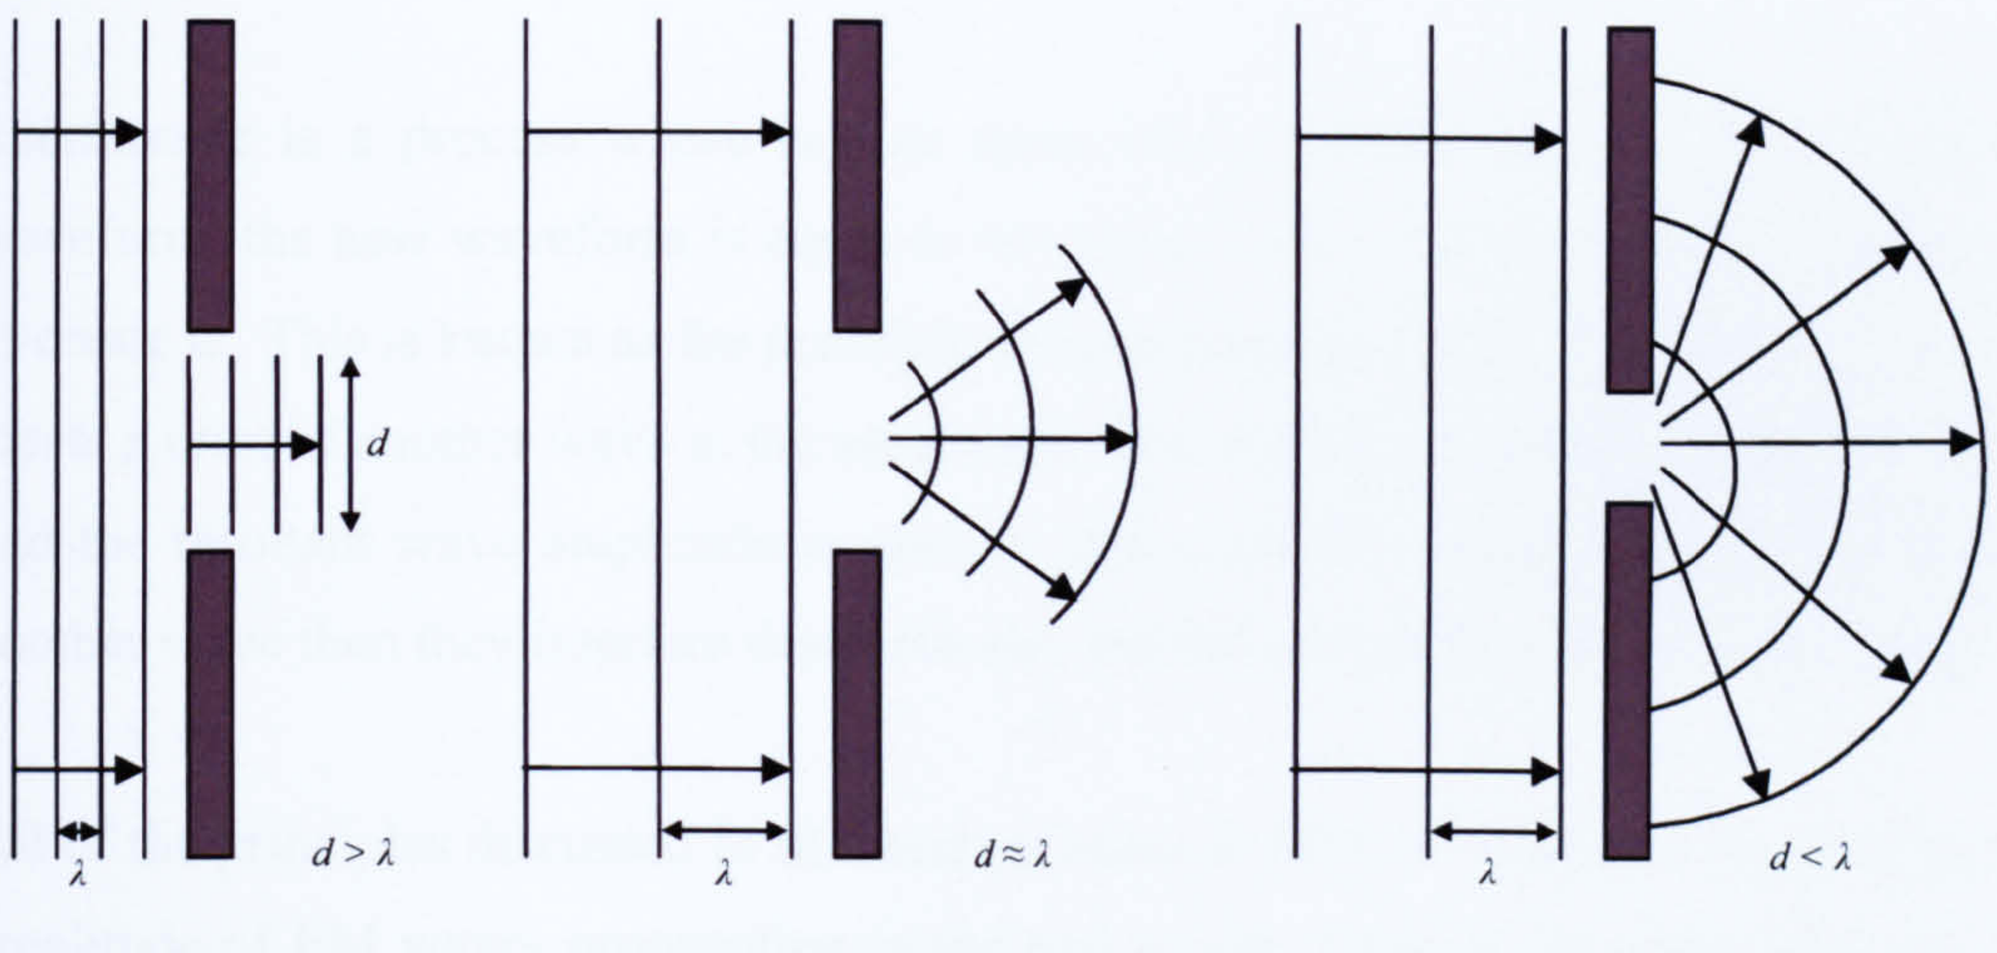

Figure 5.8 – Diffraction of EM waves

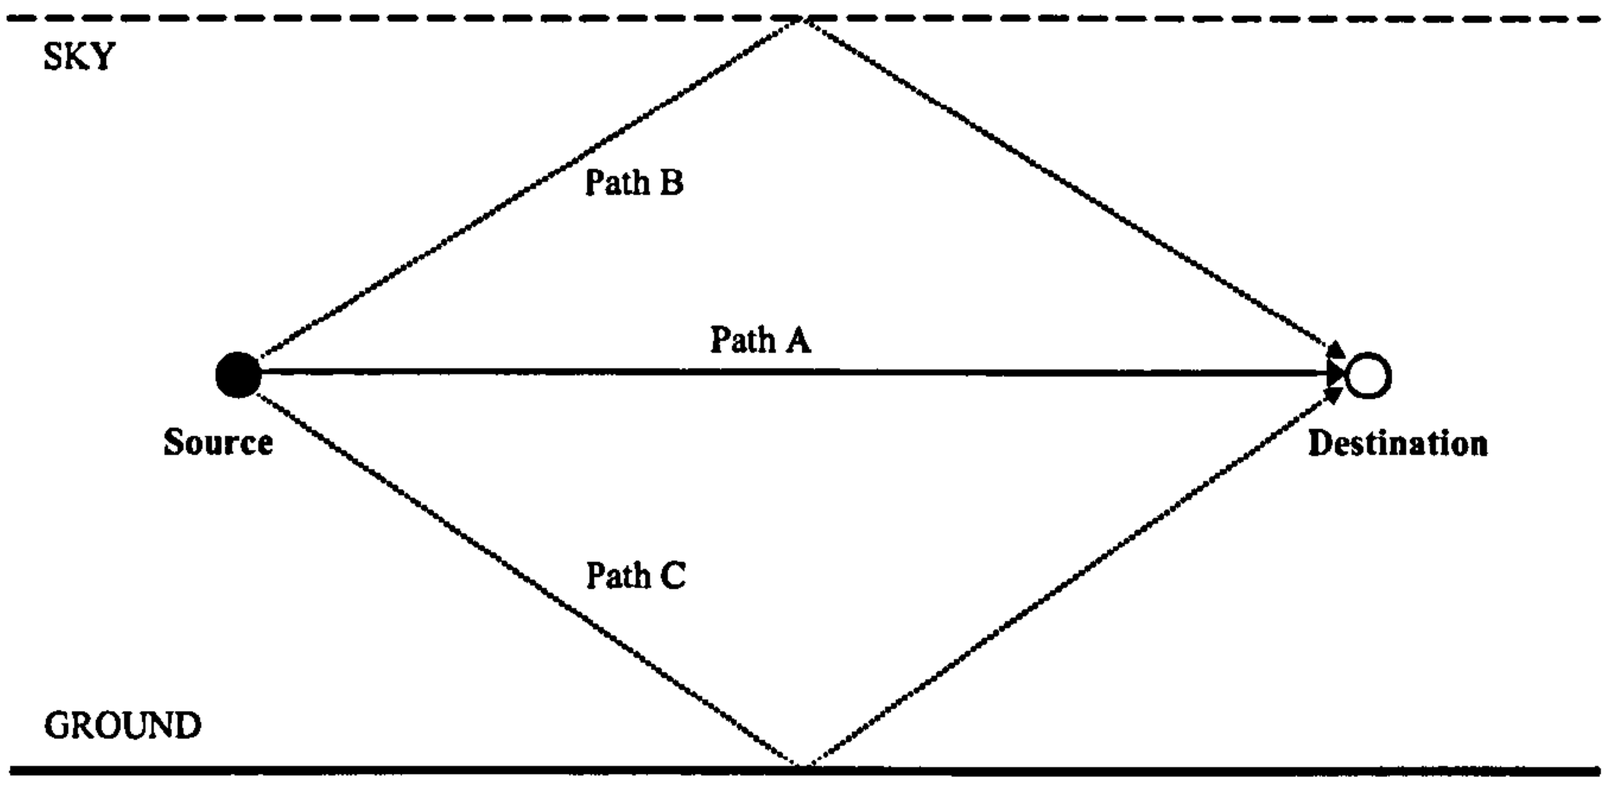

Figure 5.9 – Multipath fading

Reflection, refraction and diffraction are particularly important as they are major contributors to a phenomenon known as multipath fading [130-131]. Fading is a reduction in the amplitude of a wave, which results in received power loss. Unlike attenuation due to the spreading or absorption of EM energy, as discussed in Section 5.3.2, this occurs due to waves from a source taking many different paths to reach a destination. Since there is no guarantee that each path will be of the same length, waves may reach their intended destination out of phase and therefore interfere with one another. A simple case of what may happen if only reflection occurs is shown in Figure 5.9.

Interference is a process where two or more waves superimpose to create a new waveform; the new waveform is equal to the vector sum of the waves which interact to create it. This is known as the principle of superposition [132]. If a crest of a wave meets a crest of another wave at the same point then the crests interfere constructively and the resultant wave amplitude is greater. If a crest of a wave meets a trough of another wave then they interfere destructively, and the overall amplitude is decreased.

All of the principles discussed in this section make it difficult to predict the path and amplitude of EM waves propagating in the real world. Objects encountered along a propagation path can be made up of many different materials, each with its own effect on signal attenuation, refraction and reflection. In addition, openings and gaps (such as windows and doors) can cause diffraction. This unpredictability is compounded further as waves can travel along different paths and superimpose to cause

interference. In some instances where it is constructive, this might assist transmission, but similarly it may be destructive and restrict transmission. Therefore, an awareness of these factors is essential in the application of antennas as it helps to build an understanding of energy transfer from one to another.

5.5. Transmission Lines

5.5.1. Purpose and Construction

The general purpose of a transmission line in the context of this work is to convey EM energy from the source to the antenna or from the antenna to the receiver. The theory of how transmission lines operate is particularly useful because it shares much in common with the operation of antennas themselves.

There are a number of different types of transmission line, but generically they consist of a pair of conductors separated by a dielectric medium. The EM energy travels in the space between the conductors rather than being radiated into space. Two common types of transmission line are shown in Figure 5.10; the simplest is made up of two wires separated by an air dielectric. This type of line can suffer from electrical noise since it is not shielded, hence the reason for noise sensitive systems making use of coaxial type lines which do include shielding.

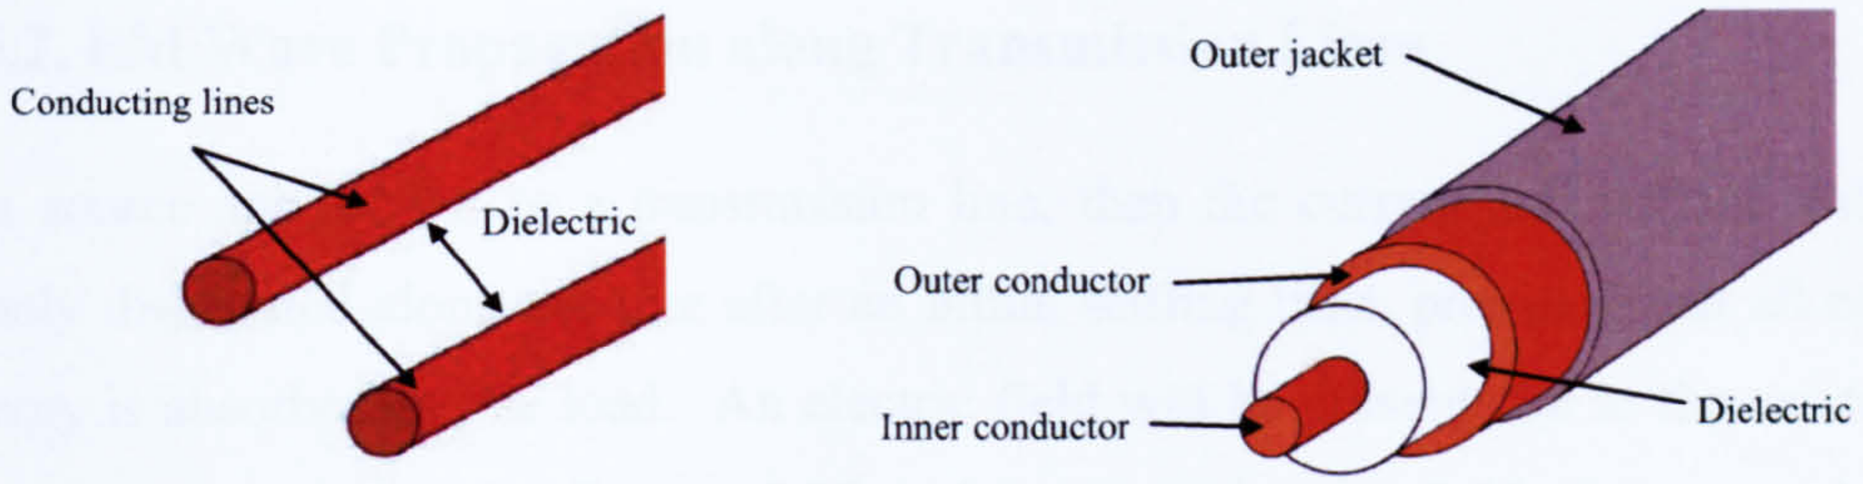

Figure 5.10 – Two pair open (left) and coaxial (right) transmission lines

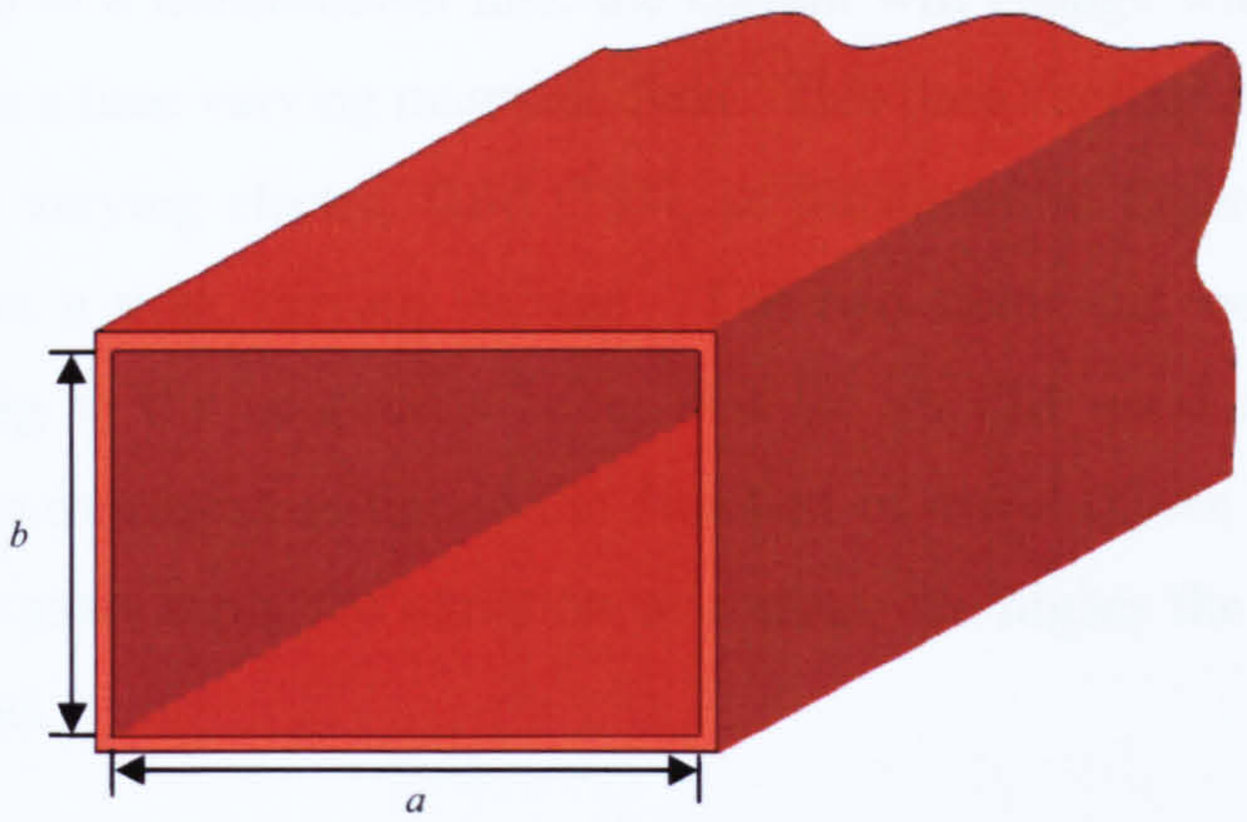

Figure 5.11 – Waveguide structure

At high frequencies even coaxial lines can be impractical due to large losses and so hollow metal structures known as waveguides are employed. Waveguides are typically rectangular in their cross section (although not necessarily so); an example is shown in Figure 5.11. As opposed to wire transmission lines which can be defined by their maximum practical operating frequency, wave guides are defined by their minimum operating frequency, or cutoff frequency (f_c). This is determined by the dimension of the waveguide using the expression given in Equation 5-9.

$$f_c = \frac{c}{2a} \qquad \text{(Equation 5-9)}$$

5.5.2. EM Wave Propagation along Transmission Lines

If a source applies DC to a transmission line, then the current and voltage will be evenly distributed along the line after an initial settling time, provided that all of the energy is absorbed by the load. An electric field will be present due to the existence of charge as shown in Figure 5.12, where the electric field lines are drawn perpendicular to the direction of travel. A magnetic field will also exist since charge is flowing, however the charge distribution will not change. This means that the electric field will not change direction at any point, nor will the magnetic field; both will remain static. Therefore, as described in Section 5.1, the conditions for EM wave propagation are not satisfied since there are no time varying fields.

If AC is applied to a transmission line, the current will change with respect to time and will result in a time varying magnetic field. This time varying magnetic field will result in a time varying electric field [133], as illustrated in Figure 5.13 where the waveform shows a time varying voltage. The two fields can now be said to be coupled resulting in the conditions being met for an EM wave to be created for propagation on a transmission line in the direction of travel shown in Figure 5.13. It follows that the more rapid the variation of current, the higher the frequency of the EM wave created.

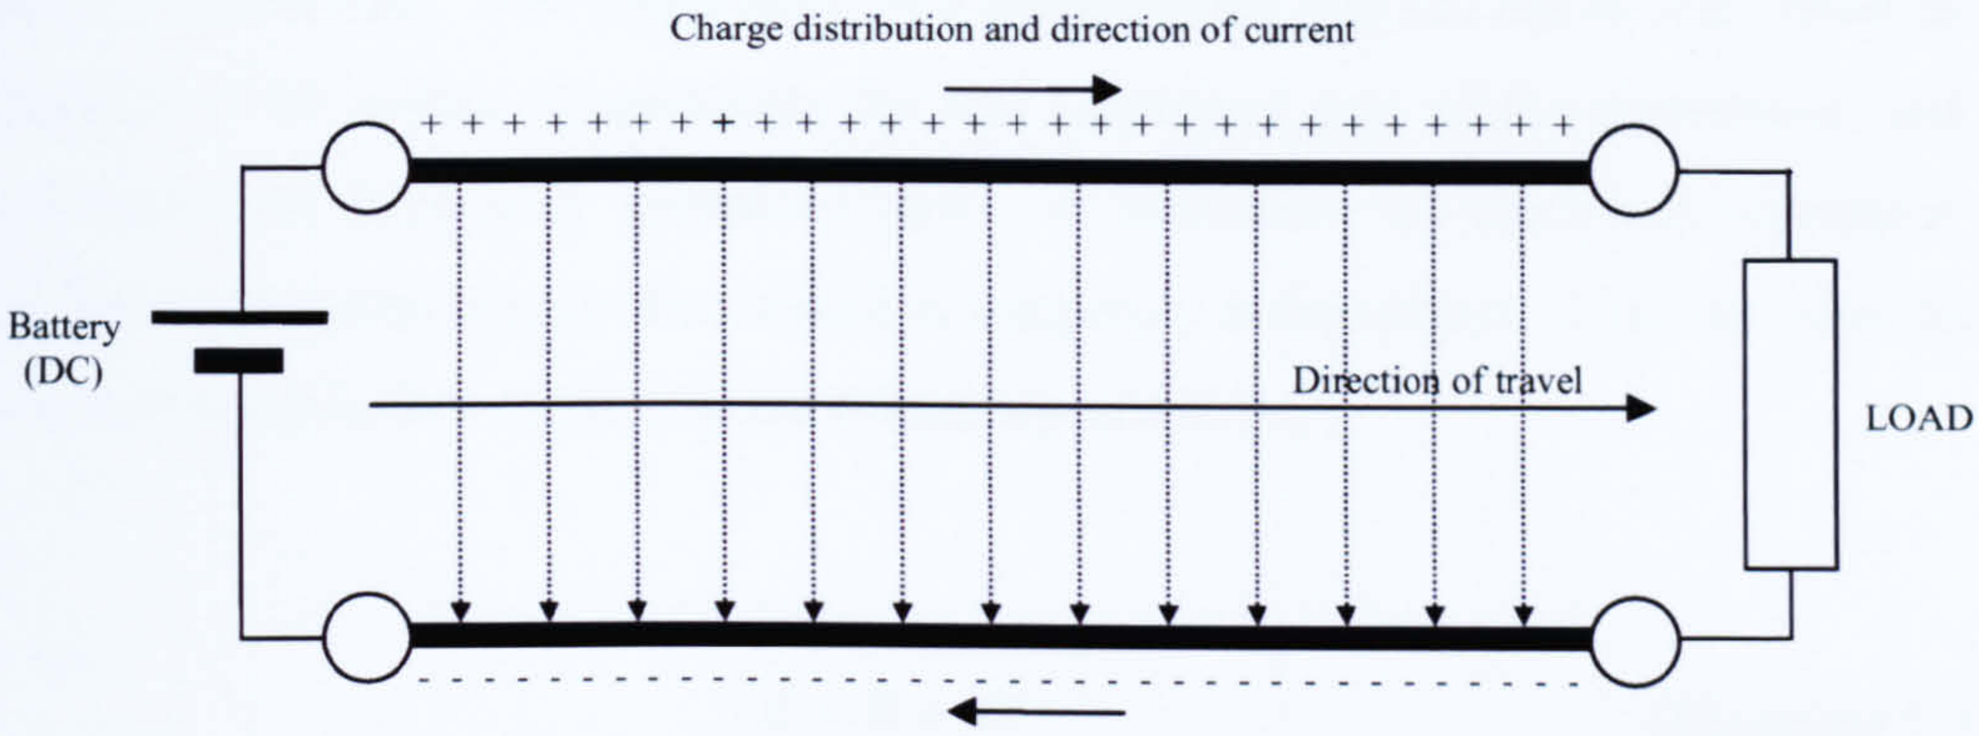

Figure 5.12 – DC on a transmission line

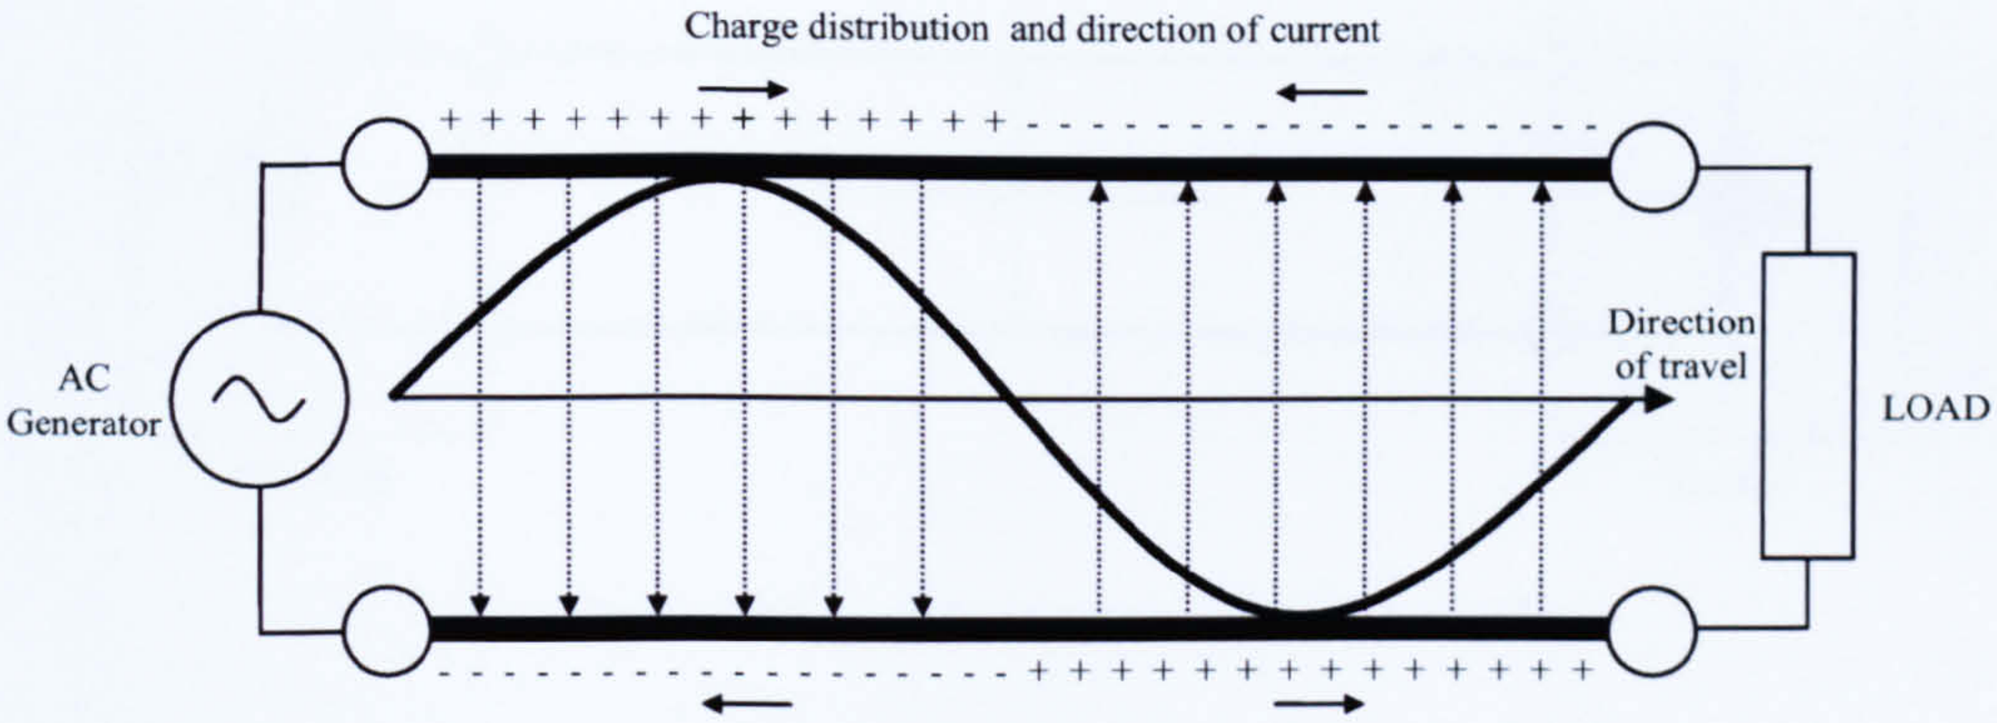

Figure 5.13 – AC on a transmission line

5.5.3. Losses

In order to demonstrate the losses present on a transmission line it is possible to represent the antenna system as a Thevenin equivalent circuit consisting of an input voltage source (V_S) and a resistance (R_S). The circuit connected to the second port represents the load, or antenna, which is said to have a resistance (R_L). This setup is represented in Figure 5.14.

Since EM waves have amplitude and phase it is more appropriate to talk of resistance as complex impedance. Doing this means that R_S and R_L are better represented as Z_S and Z_L respectively. The impedance of a transmission line (Z) can be expressed as in Equation 5-10, where R represents the real (resistive) part of the impedance, and X represents the imaginary (reactive) part. R represents an electrical resistance to current on the transmission line which is frequency independent. X on the other hand is reactive and is determined by the frequency of the AC.

$$Z = R + jX \qquad \text{(Equation 5-10)}$$

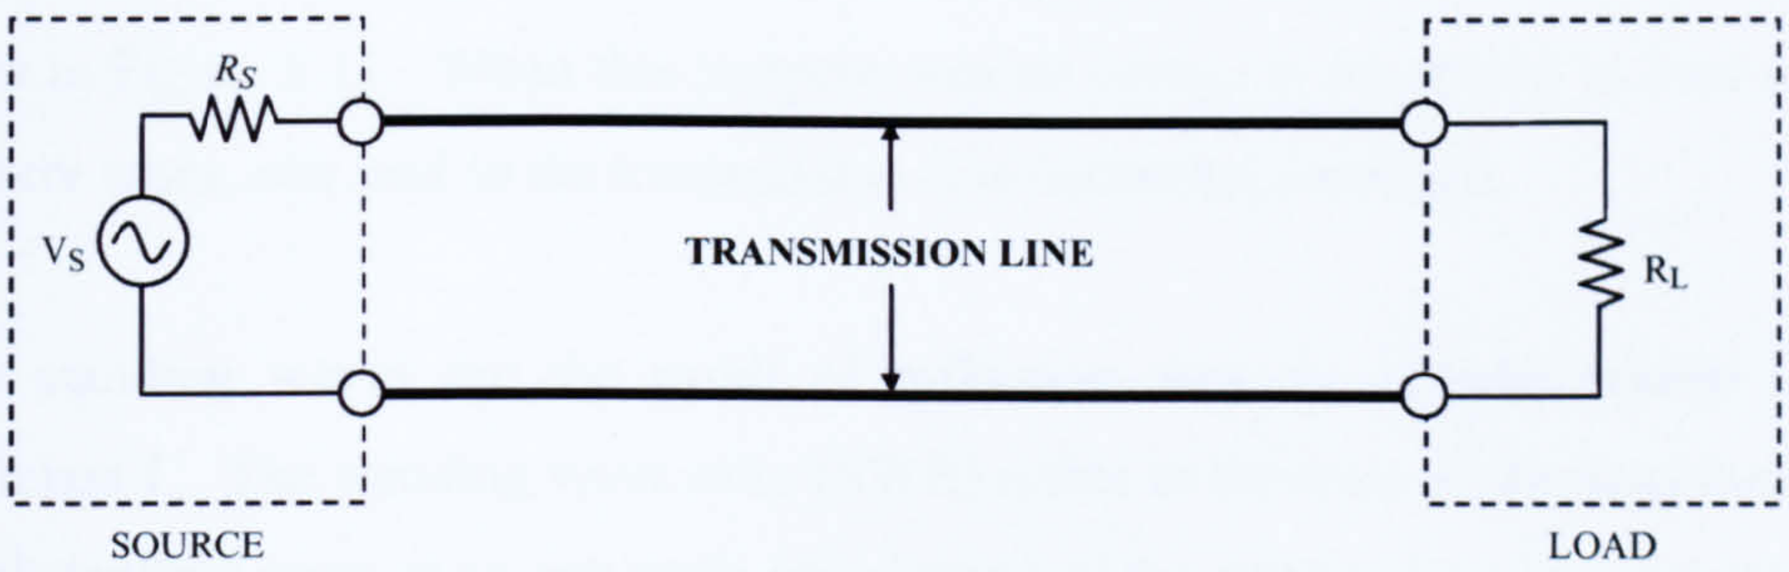

Figure 5.14– A basic transmission line

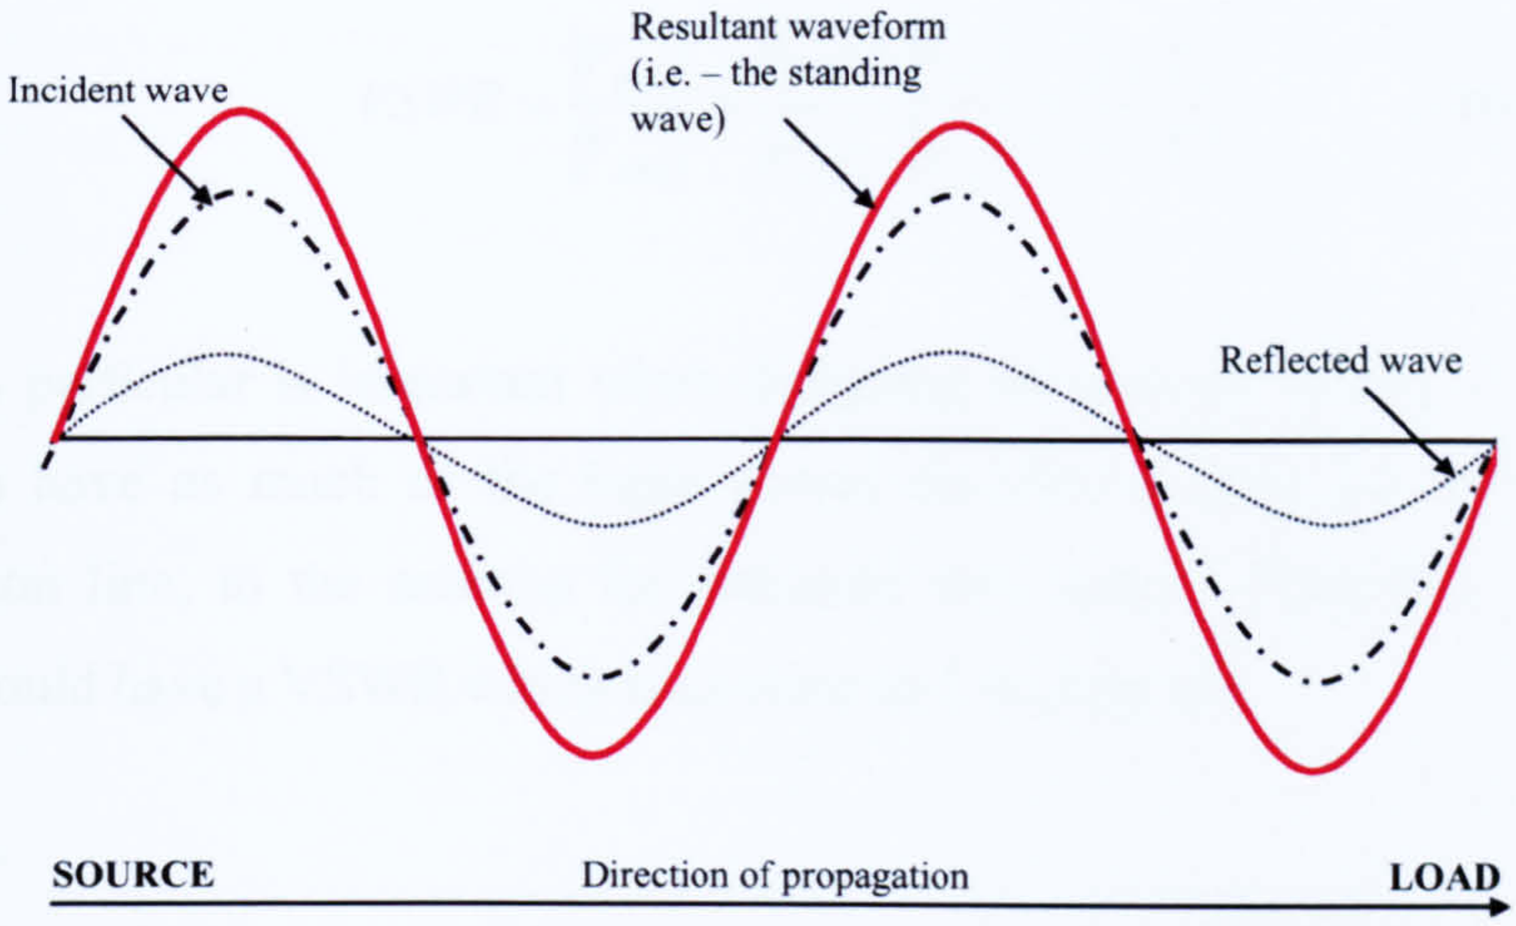

Figure 5.15 – The creation of a standing wave

For the most efficient transfer of energy, the impedance of a source, transmission line and load must match. A transmission line is said to have certain characteristic impedance (Z_0). If a finite transmission line is termintated by $R_L = Z_0$, then to the source the line appears to be infinitely long. Typical values for Z_0 in radio systems are 50Ω and 75Ω [134].

If the impedance of a load does not match that of the transmission line then incident waves will be reflected. Reflected waves have an important impact on transmission line efficiency as they can result in standing waves, whereby the incident and reflected waves superimpose and therefore interfere as discussed in Section 5.3.3 and shown in Figure 5.15. When this happens wasted energy is dissipated as heat which, in severe cases, can lead to the transmission line becoming damaged.

Since standing waves are the result of reflection they are closely related to the coefficient Γ. The standing wave ratio (SWR) refers to the ratio of the amplitude of a partial standing wave at an anti-node (maximum) to the amplitude at an adjacent node (minimum). This can be utilised to determine how well matched a transmission line is, and therefore how efficient the energy transfer is between a transmission line and its load. A SWR of 1 represents a perfect match and as the mismatch increases, so does this number. SWR is typically given in terms of voltage, and is known as the voltage standing wave ratio (VSWR) which is expressed mathematically in Equation 5-11.

$$VSWR = \frac{|V_{\max}|}{|V_{\min}|} = \frac{|V_{(i)}| + |V_{(r)}|}{|V_{(i)}| - |V_{(r)}|} = \left|\frac{1+\Gamma}{1-\Gamma}\right| \qquad \text{(Equation 5-11)}$$

VSWR in particular is important when designing an antenna sytem, as the general goal is to have as much of the input power transferred from the source, via the transmission line, to the antenna for radiation into space. Therefore, the antenna system should have a VSWR which is as close to 1 as possible.

5.6. Antennas

5.6.1. Types of Antenna

There are a large amount of different antennas available, ranging from simple straight wires to complex arrays; from the very small to the very large. Figure 5.16 shows this contrast, picturing a radio telescope measuring 25m in diameter and over 40m tall alongside a small RFID tag with integrated antenna which measures approximately a centimetre in length.

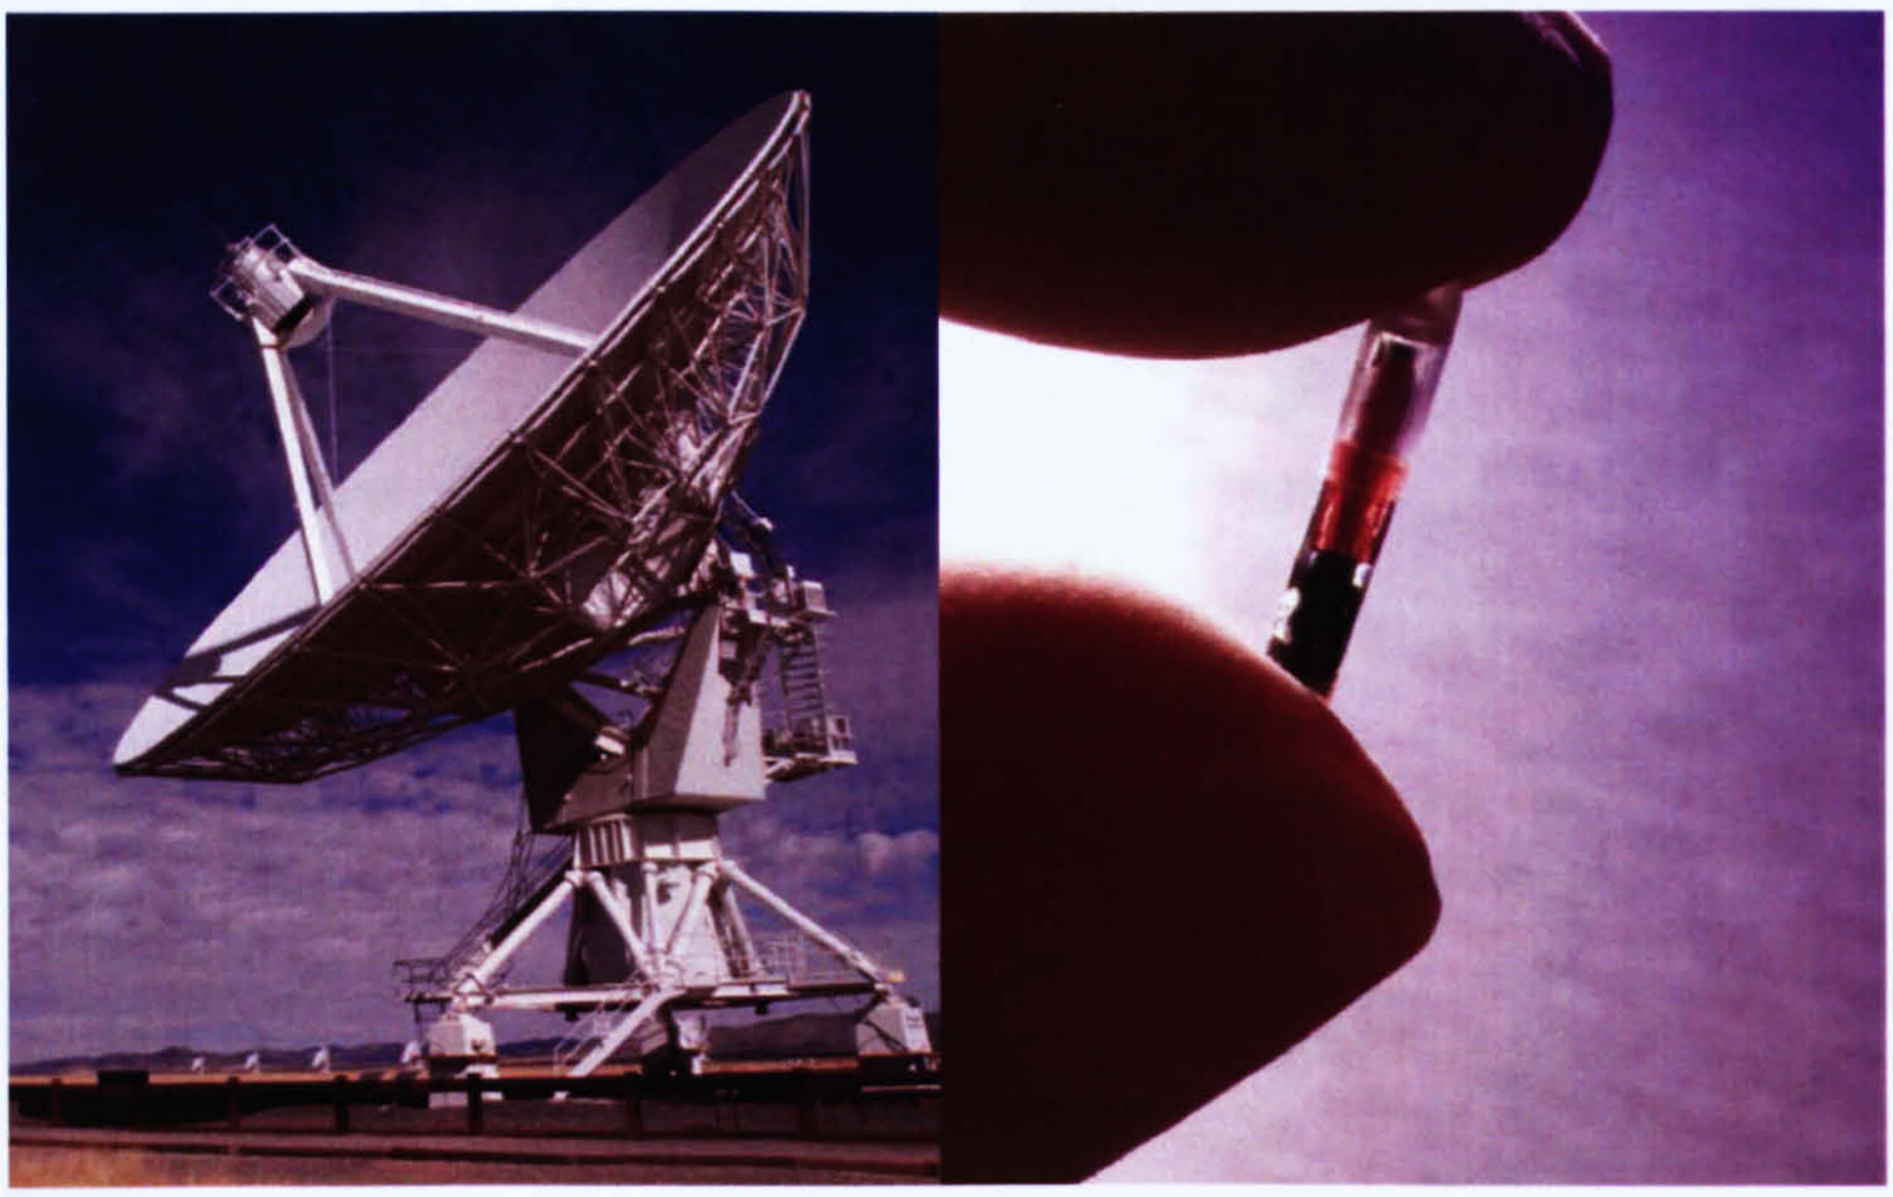

Figure 5.16 – Comparing antenna sizes

Practical antennas are typically created with a particular application or purpose in mind, and there is no one size fits all solution. Due to the vast amount of antennas available, this work initially began by considering the antennas which are currently used by the motes, or could be used in this work for the purposes of creating a new antenna. This Section deals with five such types of antenna which are illustrated in Figure 5.17; whilst these may appear quite simple, complex designs are often based upon such antennas.

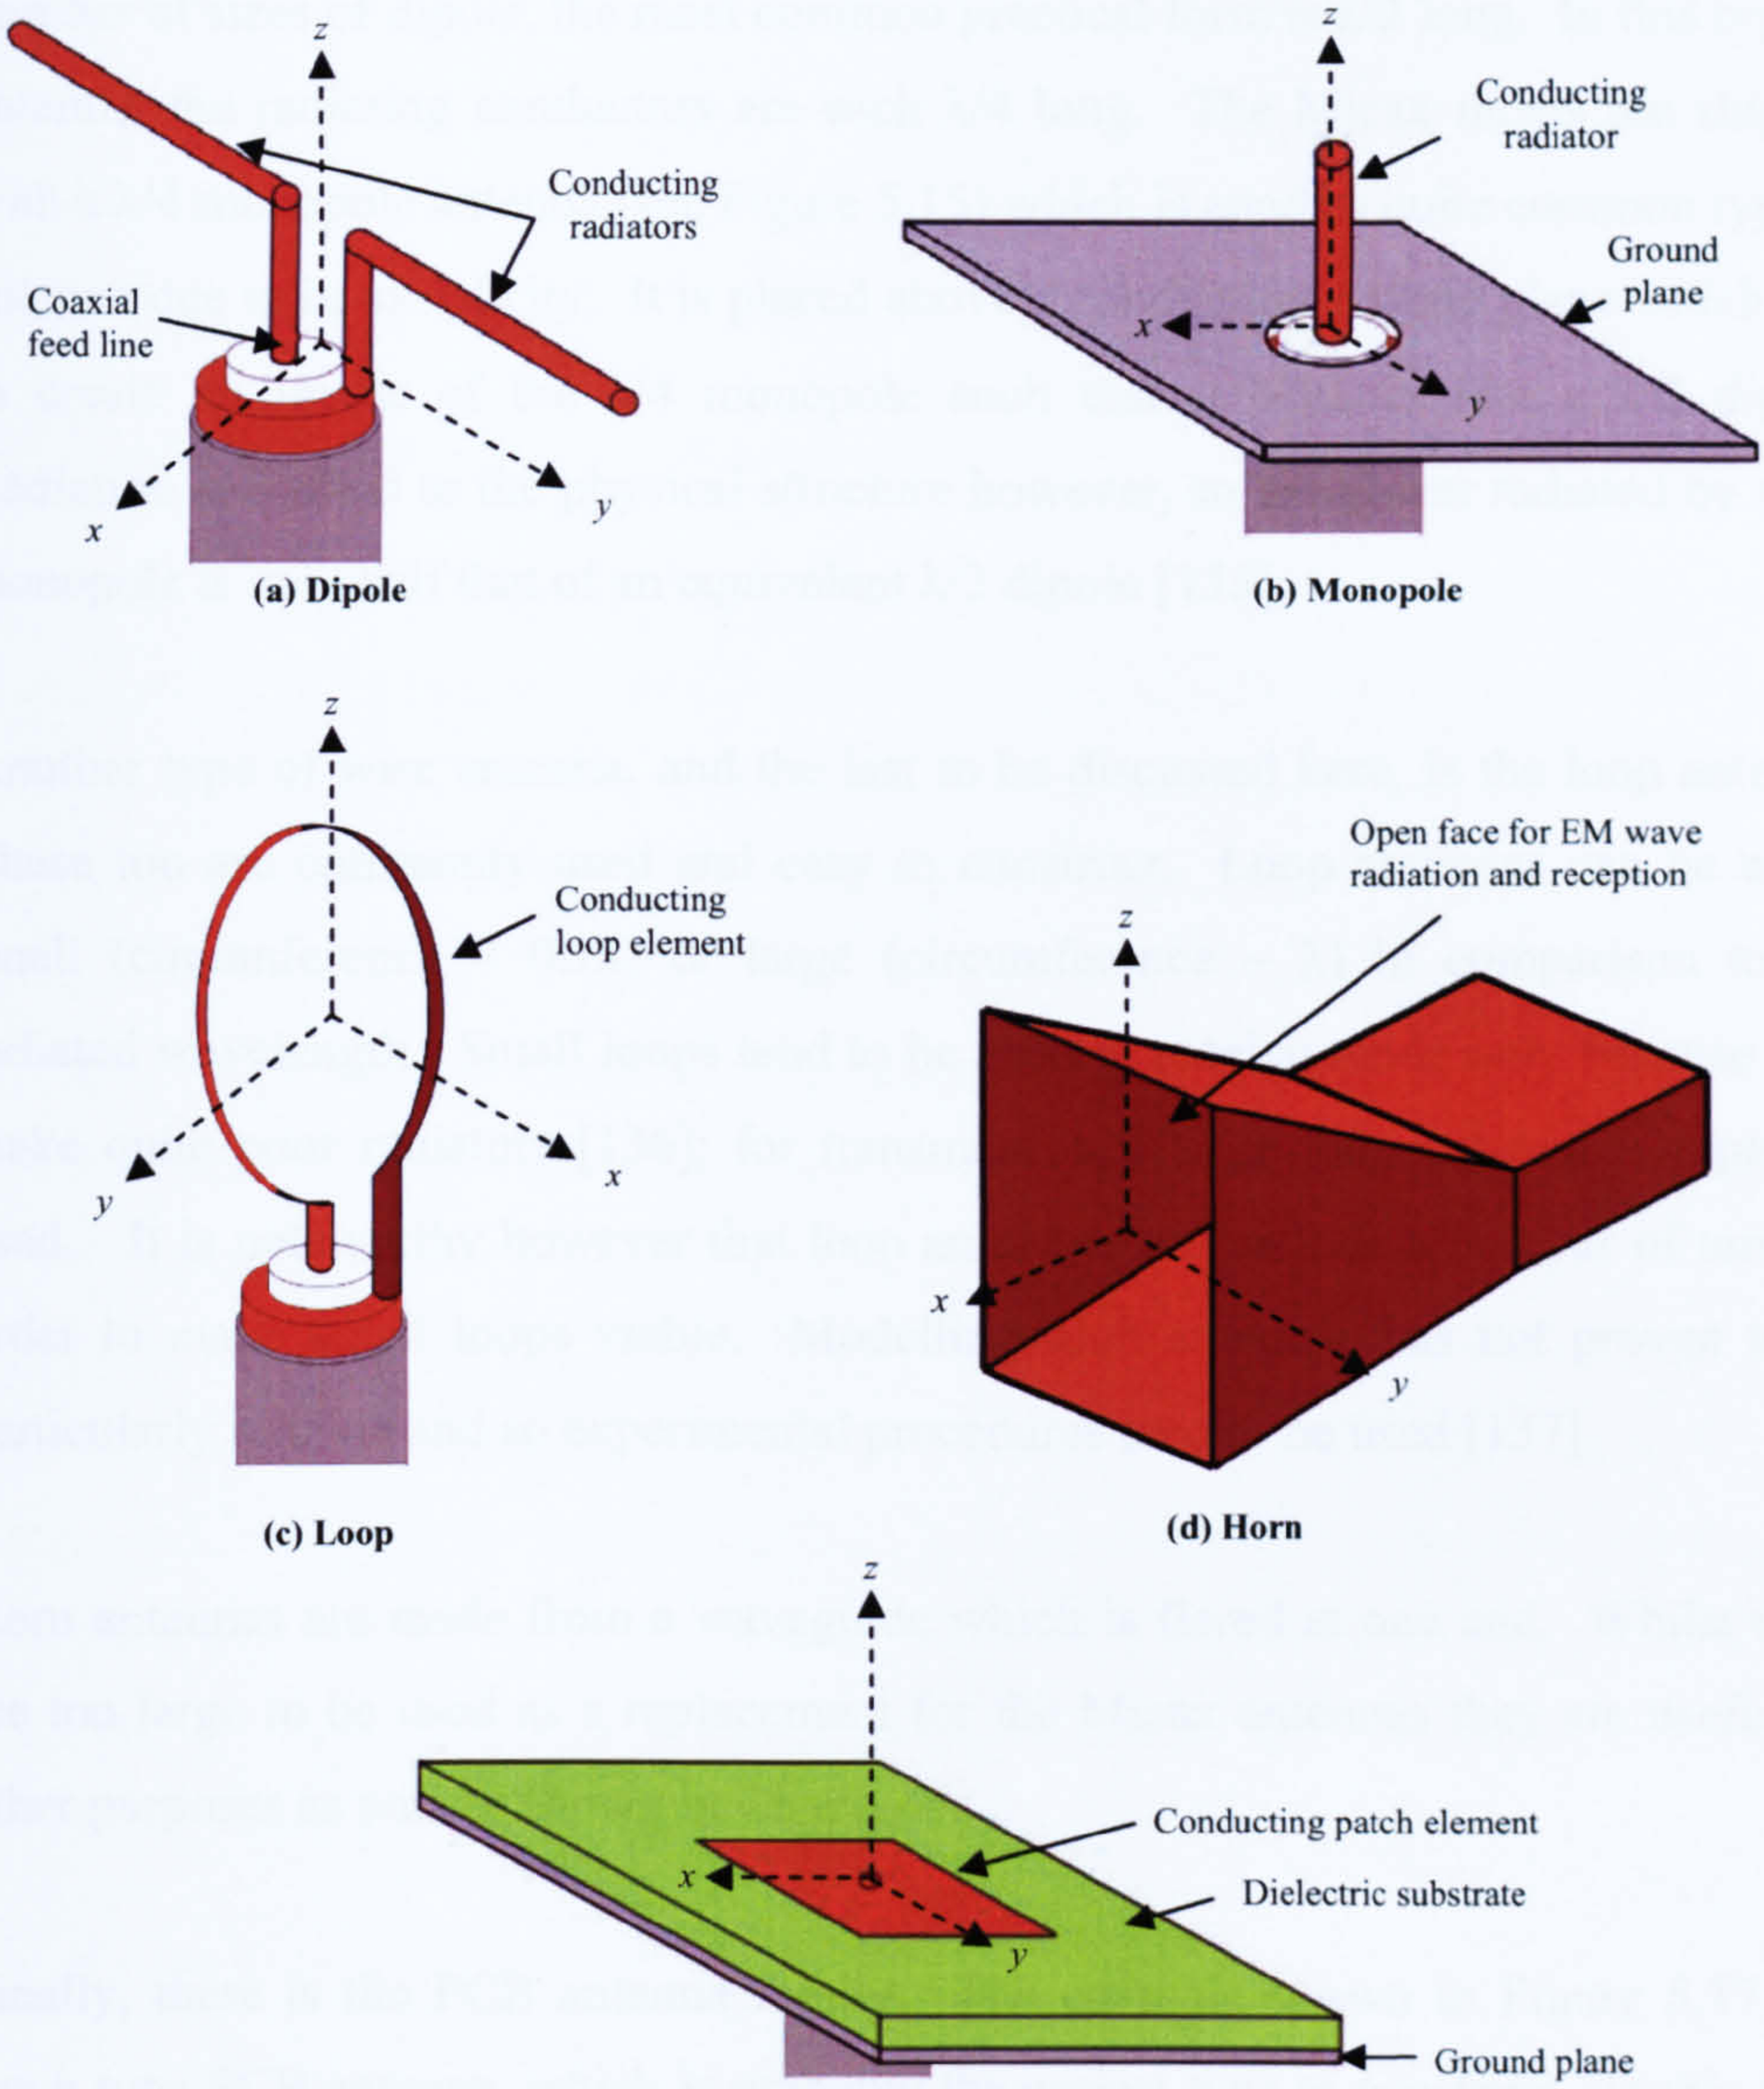

Figure 5.17 – A selection of different antennas

Firstly let us deal with a few popular types of wire antenna. Wire antennas form a large family of antennas and include those such as dipoles, monopoles (such as that used on the Micax motes) and loops. These are some of the oldest and simplest antennas available, yet it seems that for many applications they also offer quite practical solutions. Beginning with the dipole antenna, it is constructed by placing two conducting radiators back to back. In terms of a dipole being connected to a coaxial cable, such as that shown in Figure 5.10, one radiator will be connected to the inner conductor, and the other to the outer conductor. Whilst it is possible to have a number of sizes of dipole, the most common practical form is $\lambda/2$ long. In this type of antenna, the radiating conductors are each $\lambda/4$ long. The Micax motes are shipped with a $\lambda/4$ monopole antenna (see Figure 5.15) which is another quite common type of antenna due to its simplicity. It is placed above a conducting ground plane which acts to create an image of the $\lambda/4$ monopole such that it behaves like a $\lambda/2$ dipole. Radiation is limited to the physical structure however, so the power radiated by a $\lambda/4$ monopole is only half that of an equivalent $\lambda/2$ dipole [135].

Another type of wire antenna, and the last to be discussed here, is the loop antenna. These too are commonly used and easy to construct. Loop antennas can be either small (circumference $< 0.1\lambda$) or large (circumference $\sim \lambda$) in comparison to the radiated wavelength. Small loops tend to be used in receive mode only because they make quite poor radiators [136]; for transmission a large loop antenna is typically used. It is noteworthy however that loop antennas can include a number of turns in order to make small loops viable. Modelling such antennas has not proven to be particularly reliable and so experimental procedures tend to be used [137].

Horn antennas are made from a waveguide which is flared at one end. Whilst these are too large to be used as a replacement for the Micax antennas they are useful for other purposes as will be shown in Chapter 6.

Finally, there is the PCB antenna family. The example shown in Figure 5.17 is a patch type PCB antenna, which is probably the easiest type to construct. In this form they consist of a conducting patch element which is placed a short distance (less than 0.05λ) above a ground plane – the two are separated by a dielectric substrate. These

types of antenna have been popular in many different and fairly recent applications, such as in aircraft communications as well as in mobile wireless communication devices. The reason for this is that they are low profile devices which can be made inexpensively and can even be integrated into the circuitry of devices in some cases. Feeding such an antenna is possible either via a strip of copper leading to the patch, or via a probe feed; the later is shown in Figure 5.17.

The remainder of this chapter considers various properties of antennas to show how they can be characterised. Where appropriate, the antennas illustrated in Figure 5.17 will be discussed further, however the first task is to consider how antennas actually radiate EM waves.

5.6.2. Antenna Radiation Mechanism

The antenna can be thought of as an extension of a transmission line. EM waves propagate along the transmission line and continue into the antenna; the difference is that the antenna does not seek to shield the EM energy but attempts to radiate it out into space. All antennas radiate in very much the same way although some are more difficult to explain in terms of how charge is distributed across them, and so illustrating how electric field lines form around them becomes complex. Therefore a description of how radiation occurs from a simple $\lambda/2$ dipole antenna will be given, as shown in Figure 5.18.

This example assumes that it takes a time T for a full wave of length λ to form. It is assumed that initially the dipole is at the point just before it receives its first full EM wave, and therefore there is no charge upon it. At time $T/4$ the charge on the dipole will have reached a maximum and the electric fields lines will have reached a distance of $\lambda/4$ away from the dipole.

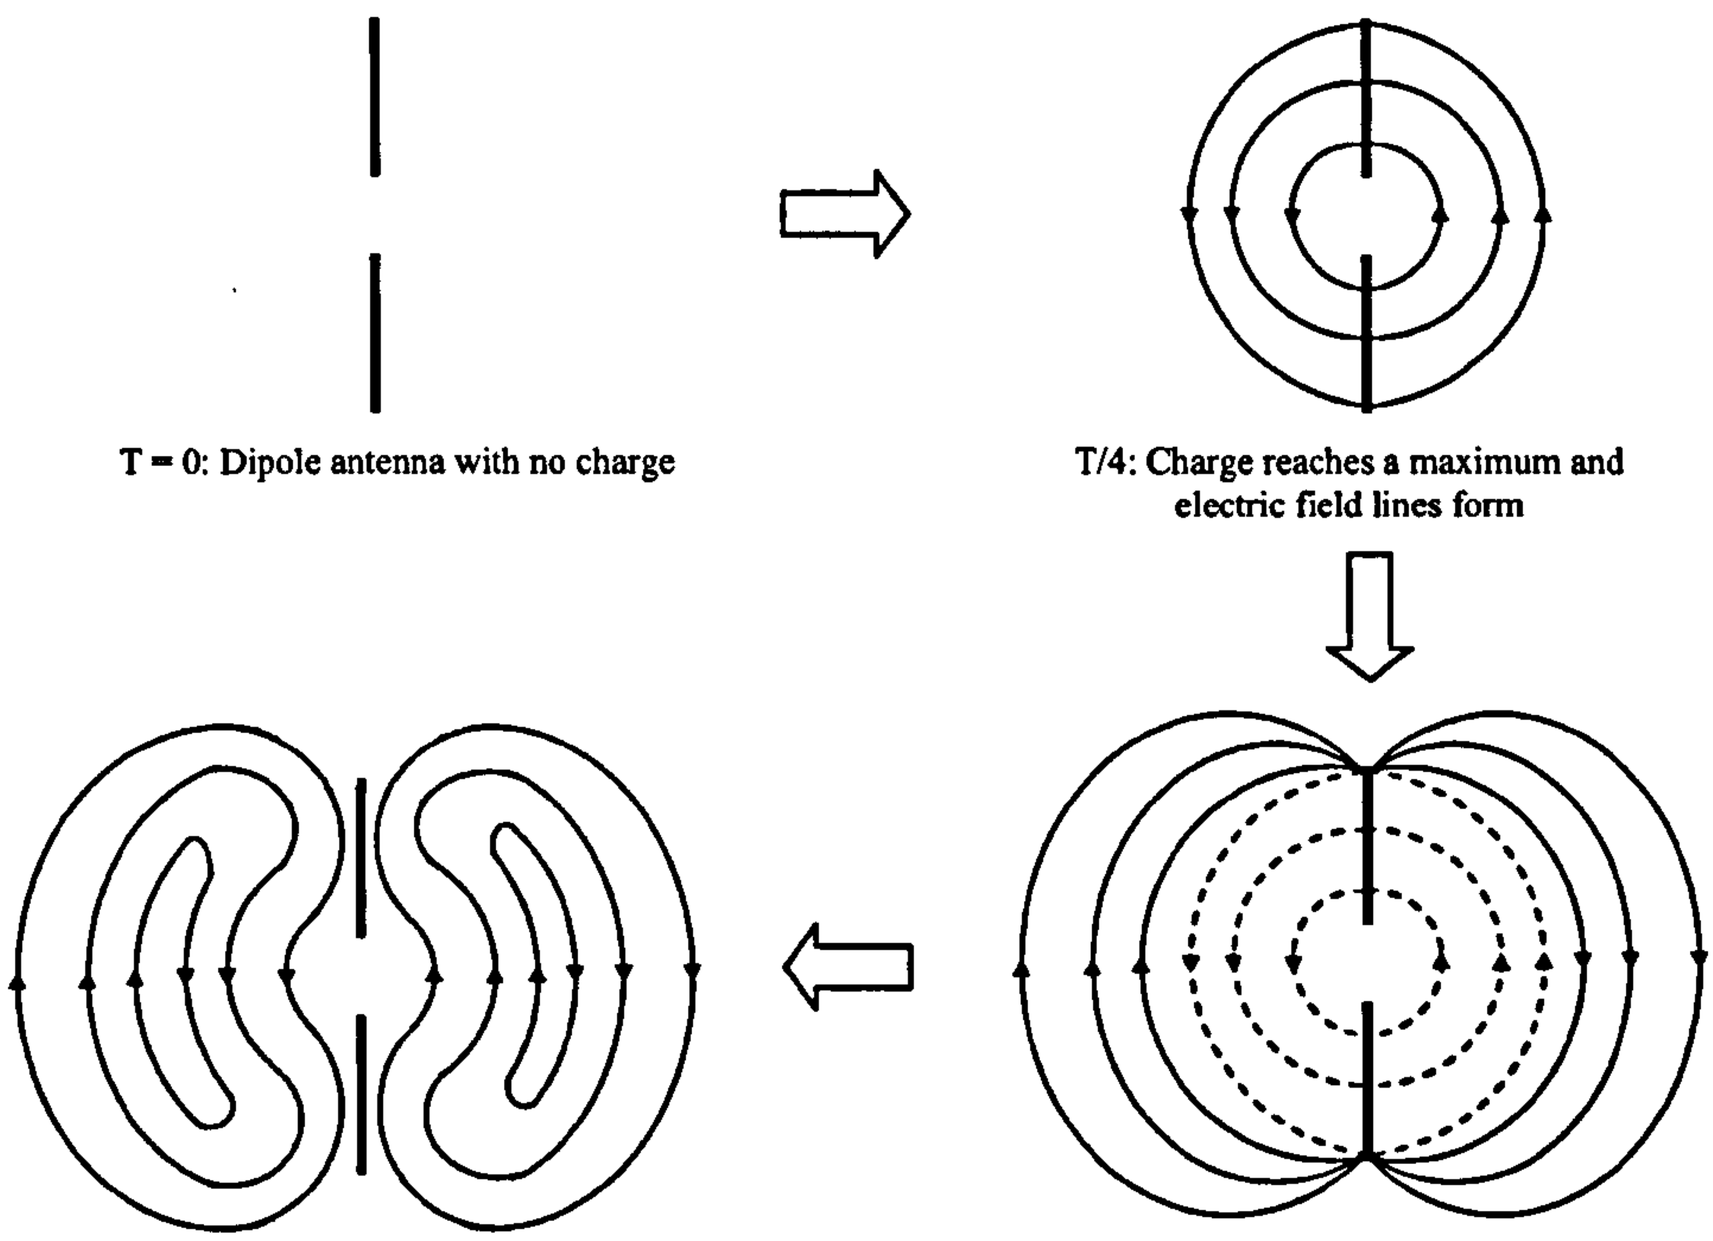

Figure 5.18 – EM radiation from a half wave dipole antenna

Between the times $T/4$ and $T/2$, the charge on the dipole will be reduced such that it is completely neutralised at time $T/2$. The neutralisation can be thought of as charges being introduced which have electric fields lines that act in the opposite direction to those created up to the time $T/4$. These field lines can be considered to be $\lambda/4$ away from the dipole, whilst the original lines are $\lambda/2$ away.

Since there is no charge on the dipole at time $T/2$, the electric field lines have no reason to stay attached to the dipole structure and can therefore form closed loops. These closed loops can then propagate into space and can be likened to ripples in water. Whilst one may create ripples in water with a stone, the action does not need repeating since the ripples will continue to spread out from the point of disturbance. Similarly, once an EM wave becomes separated from an antenna it does not require further electric charges; it is self sustaining.

5.6.3. Field Regions

The space surrounding an antenna can be split into two general regions; the near field and the far field. The near field, as the name suggests, is close to the antenna and contains more energy than the far field because of the electric and magnetic fields which are close to the antenna. The near field diminishes quickly however into the far field; this is generally agreed to start at a minimum distance R from the antenna as shown in Equation 5-12 [138-140], where D is the largest dimension of the antenna and λ is the wavelength in question. In work relating to antennas it is the far field that is of most significance since it is concerned with the radiation rather than storage of energy.

$$R = \frac{2D^2}{\lambda} \qquad \text{(Equation 5-12)}$$

5.6.4 Reciprocity

Whilst the previous section focused on how an antenna may cause the radiation of EM waves into free space – that is – transmission, this is not the only role of an antenna. Any transmitted energy will typically require reception at some point and so an antenna will need to perform this operation. Fortunately, reception is the exact opposite of transmission; this feature is known as reciprocity [141]. This means that whilst time varying current results in radiation, an EM field applied to an antenna causes current to flow in it. Reciprocity is an important feature of antennas because its means that one can typically assume that various characteristics, such as those discussed shortly, are applicable to both transmission and reception.

5.6.5 Radiation Patterns

The radiation pattern of an antenna shows how the EM energy it radiates is distributed at a fixed distance away – this distance is usually in the antennas far field region. Radiation patterns are one of the most fundamental properties of any antenna, and often allow other antenna characteristics to be calculated. Whilst most literature uses

the term radiation pattern, it is worth noting that reciprocity leads to the assumption that the reception and radiation pattern of an antenna are identical.

In order to plot radiation patterns, it is necessary to have an appropriate co-ordinate system. Since the antenna pattern is 3D, a 3D co-ordinate system is required which takes into account both cartesian (*x,y,z*) and spherical (r,θ,φ) co-ordinates; the use of these is shown in Figure 5.19. This system is useful when displaying the measured radiation pattern of an antenna through the use of 2D plots, since the co-ordinate system allows the reader to determine the radiation plane that such a plot represents.

Appendix C gives the radiation patterns one might expect from the antennas shown in Figure 5.17; these radiation patterns are relative to the cartesian axes shown in Figure 5.17. It is worth noting that the radiation patterns represent what one might theoretically expect and are therefore reasonable approximations based on literature rather than measured results.

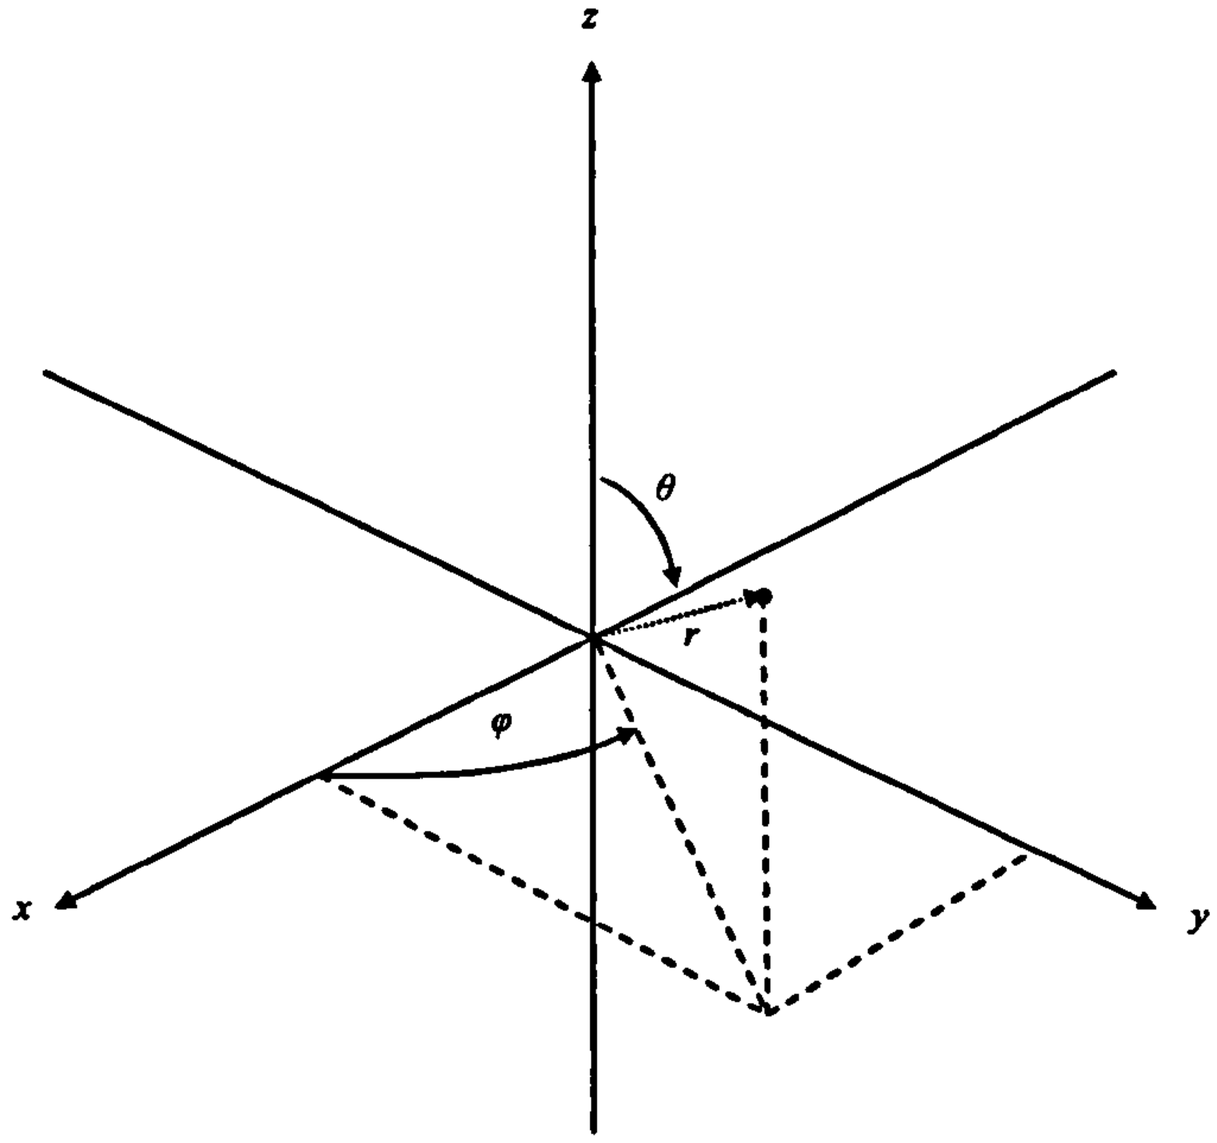

Figure 5.19 – 3D co-ordinate system

5.6.6. Antenna Polarisation

Earlier in this chapter (Section 5.4.1), the polarisation of EM waves was discussed briefly. In general, an antenna will not be orientated such that its polarisation matches that of an incoming EM wave and there will be a polarisation mismatch. Antennas are usually polarised in the direction which their elements are facing – this is easy to determine with some antennas (such as horns and monopoles), but less obvious with others (patch).

If the polarisation of an antenna is such that it is at 90° to that of an incoming linear wave then theory suggests that the antenna would not receive anything. In practise however, signals will still be received albeit severely attenuated. If an incoming wave is circularly polarised, then there should be no attenuation, regardless of an antennas orientation. One should assume, unless otherwise stated, that radiation patterns shown in this document are such that there is no polarisation mismatch.

5.6.7. Gain and Directivity

Antenna gain is a concept which relates to the direction in which radiated power is focused. An ideal isotropic source is assumed to radiate power in all directions uniformly, so at a fixed distance away from such an antenna the measured power density would be the same regardless of direction. In this case, the gain can be said to equal zero because there is no direction in which more power is radiated. As isotropic sources cannot be created physically, real antennas have a gain above zero, which is measured on the decibel scale. The unit typically used is dBi because each antenna is given a gain referenced against an isotropic antenna radiating the same power, where an isotropic antenna has a gain of 0dBi. Sometimes the figure is also expressed as a natural number, where 0dBi = 1. Literature often refers to directivity, rather than directive gain; this is defined as the maximum directive gain.

The antennas mentioned in Section 5.5.1, and illustrated in Figure 5.17, all have a directive gain associated with them which could be derived from the radiation plots which are shown in Appendix C. Whilst such a derivation will be not be performed

here, it is plain to see that the dipole and loop antennas give a similar radiation pattern. Such antennas have a low directive gain (maybe less than 2dBi) because in general much of the radiated power is evenly distributed. Next, the monopole and patch antenna give reasonably similar results in the *xz* and *yz* planes, with the ground plane of both antennas acting, in theory, to prevent radiation in the negative *z* direction. However for the patch antenna, in the *xy* plane (i.e. – parallel to the PCB substrate) there is no radiation. This means that the monopole has a lower directivity than the patch. Finally, the horn antenna directs much of its energy in the direction of the horn opening; this means that it has a high directive gain (maybe as high as 10dBi).

Moving on, there is another type of gain referred to as power gain. This is different to directive gain because instead of giving a value for gain which compares a practical antenna against an isotropic antenna which is radiating the same power, it gives a value for gain where it is the input power that is the same. This is an important distinction because there could be losses in the material of the antenna due to heating effects, for example. This could mean that not all of the input power is radiated. Directive gain (D) and power gain (G) are related by a factor called radiation efficiency (ξ), as shown in Equation 5-13 [142].

$$G = \xi D \qquad \text{(Equation 5-13)}$$

The radiation efficiency is the ratio of power radiated by an antenna to the power input to it; if there are no losses then this value will be 1 and $G = D$.

5.6.8. Bandwidth

The bandwidth of an antenna refers to the range of frequencies over which it will operate to a satisfactory level. It was mentioned in Chapter 3 that the Micax transceivers have a selectable frequency range; in addition to this they require a small range of frequencies in order to transmit data. As a result, an antenna must not just operate at one specific frequency, but rather over a range of frequencies suitable for the particular application.

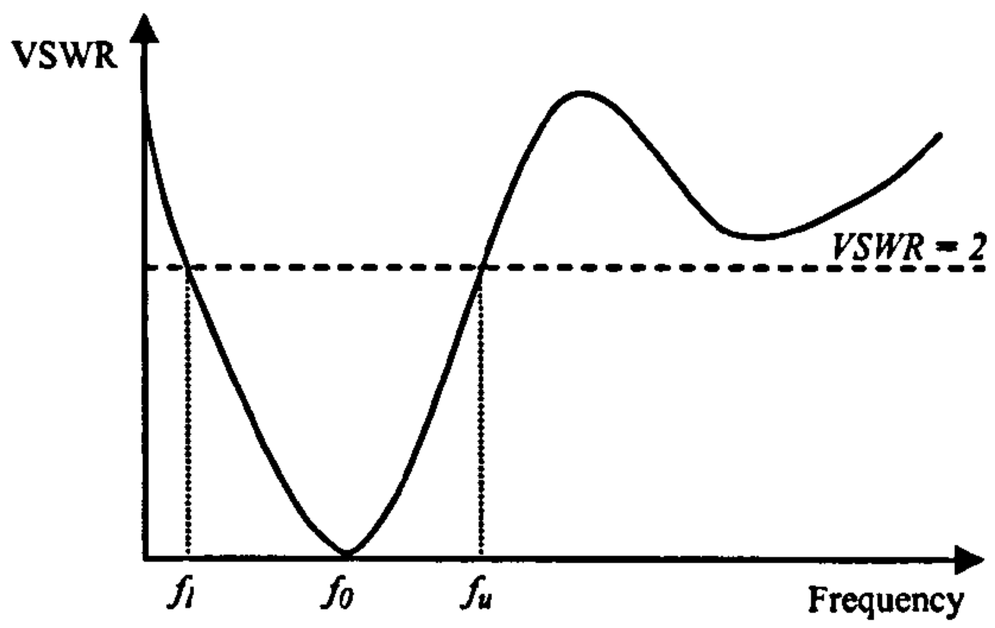

Figure 5.20 – Calculating antenna bandwidth

$$Bandwidth(\%) = \left(\frac{f_u - f_l}{f_0} \right) \times 100 \qquad \text{(Equation 5-14)}$$

Balanis [143] notes that there is no unique way to characterise the bandwidth of an antenna; many however use what is known as impedance bandwidth. This can be found using the VSWR of an antenna such that the operable frequency range is where the VSWR ≤ 2 [144] as shown in Figure 5.20. The upper and lower frequencies (f_u and f_l respectively) can be used in conjunction with the centre frequency (f_0) to express the bandwith as a percentage (see Equation 5-14).

5.7. Simulation Software

Designing antennas is a non-trivial task. Trial and error methods of design can be tedious, time consuming and waste materials. For this reason there are a number of modelling packages available. NEC (Numerical Electromagnetics Code) packages [145-146] are popular amongst amateur radio groups, but are limited generally to a text based input interface and antennas made up of many wire elements.

More advanced CAD packages are available which provide the facility to create complex structures. Such packages are far better suited to modelling antennas with complex geometries, such as those which will be presented in Chapter 6. The particular package chosen was the HFSS (High Frequency Structure Simulator) package from Ansoft [147]; its use is not strictly limited to antenna simulation but needless to say it provides many powerful features for assisting in this task. In

addition it requires no scripting input from the user, unlike packages such as Feko [148], which allows novice users the opportunity to become competent with the software without such a steep initial learning curve.

To assist in giving a brief overview of the HFSS package, a number of illustrations have been included. Figure 5.21 shows the HFSS main window. Models can be constructed here using combinations of the various geometric shapes available on the toolbars; the example shown is actually a simplified model of a $\lambda/4$ monopole antenna similar to that used on the Mica2 mote. The antenna is modelled as a thin conductor, and is placed within a cube assigned as an air boundary which essentially tells HFSS that any EM waves reaching this boundary would be radiated out into free space. It is possible to assign other types of boundaries too; in Figure 5.21 a perfect infinite ground plane is assigned to the bottom face of the air box in order to complete the monopole. In addition to vacuum and perfect electrical conductors (PEC), HFSS also provides a comprehensive material library (see Figure 5.22). Using real materials can be necessary for accuracy, but they can slow simulations significantly.

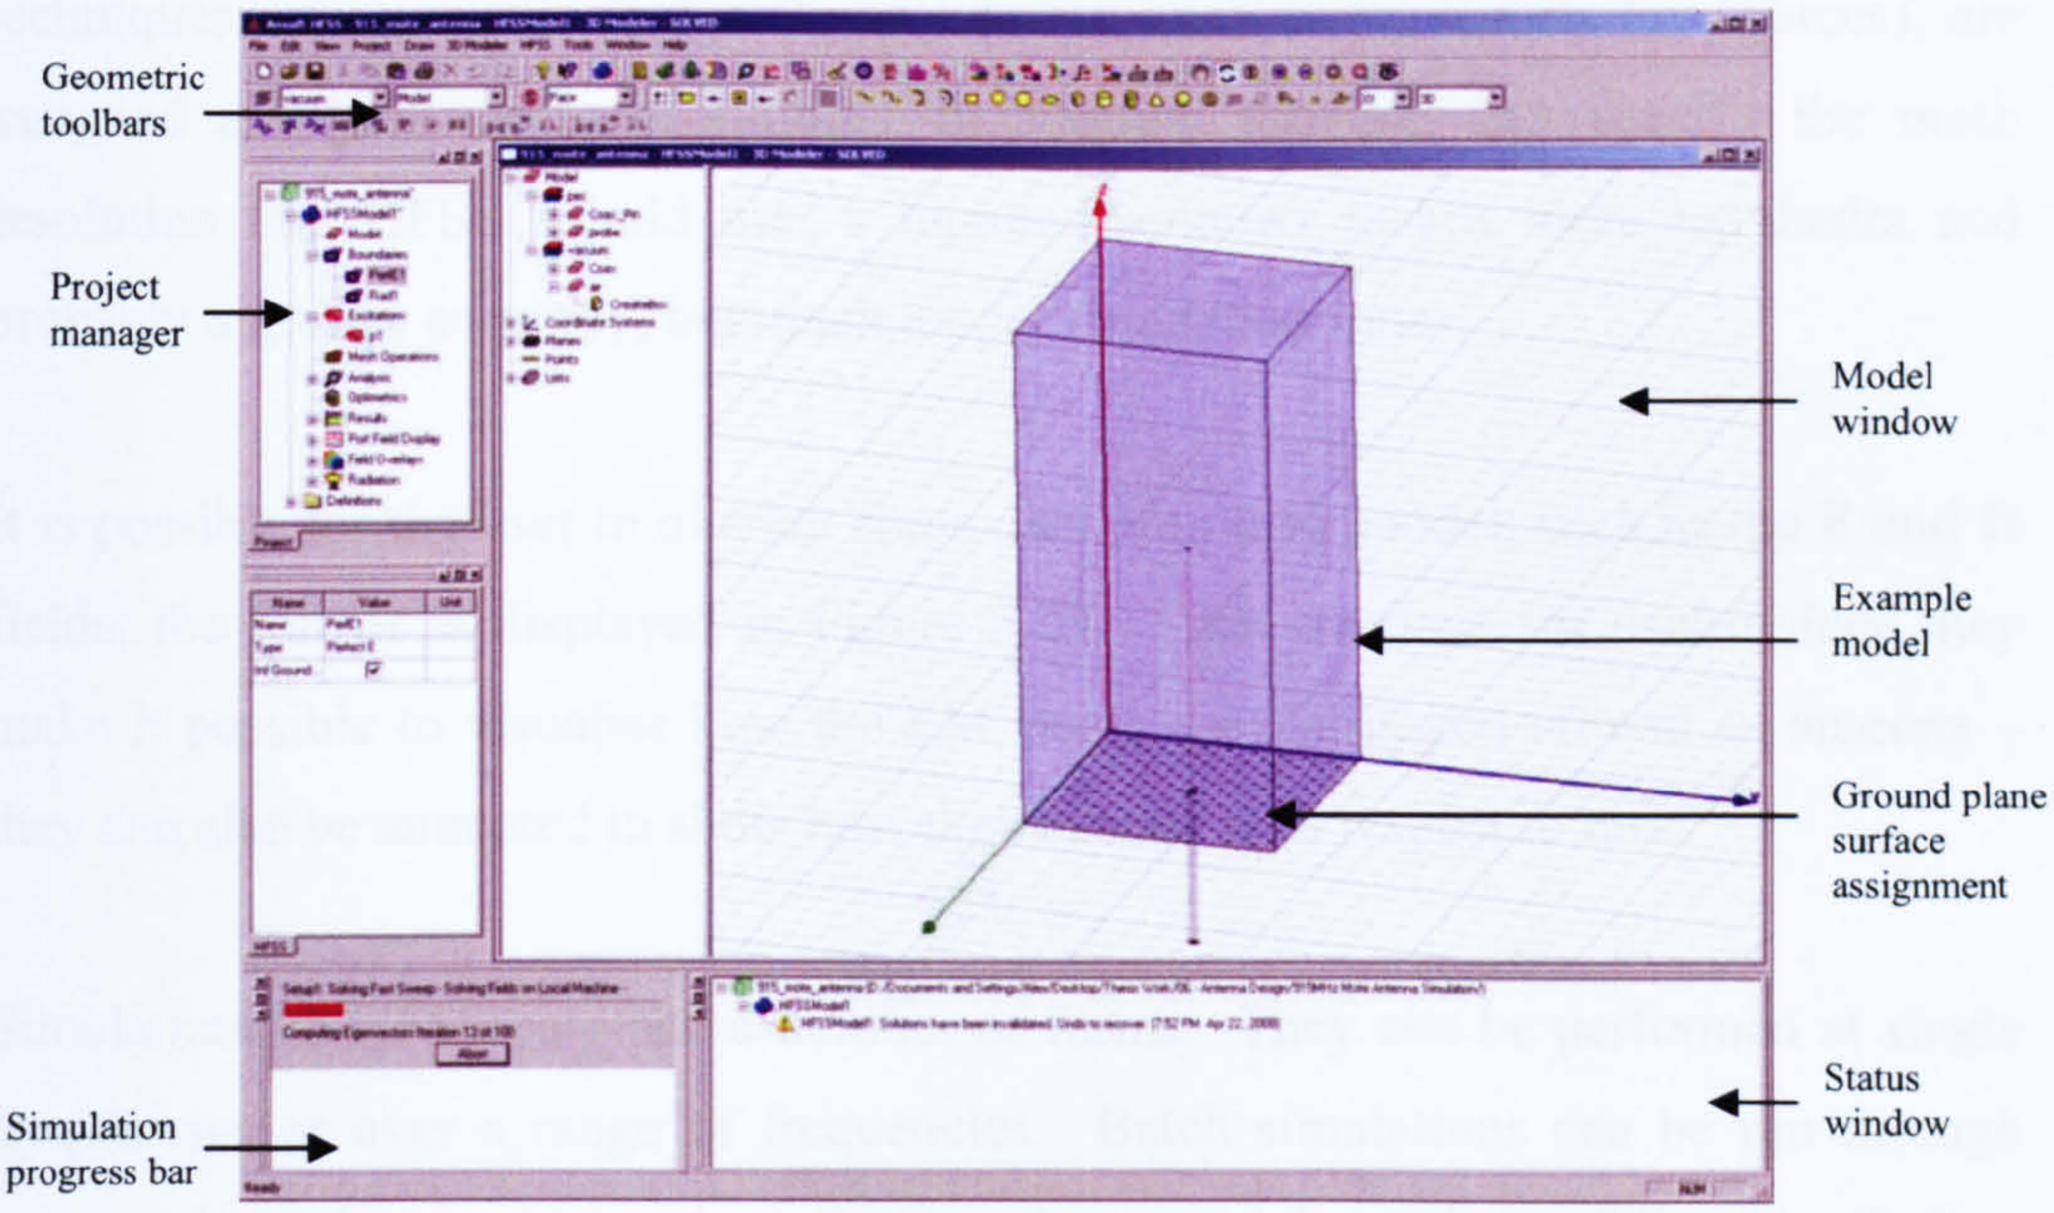

Figure 5.21 – Main HFSS window

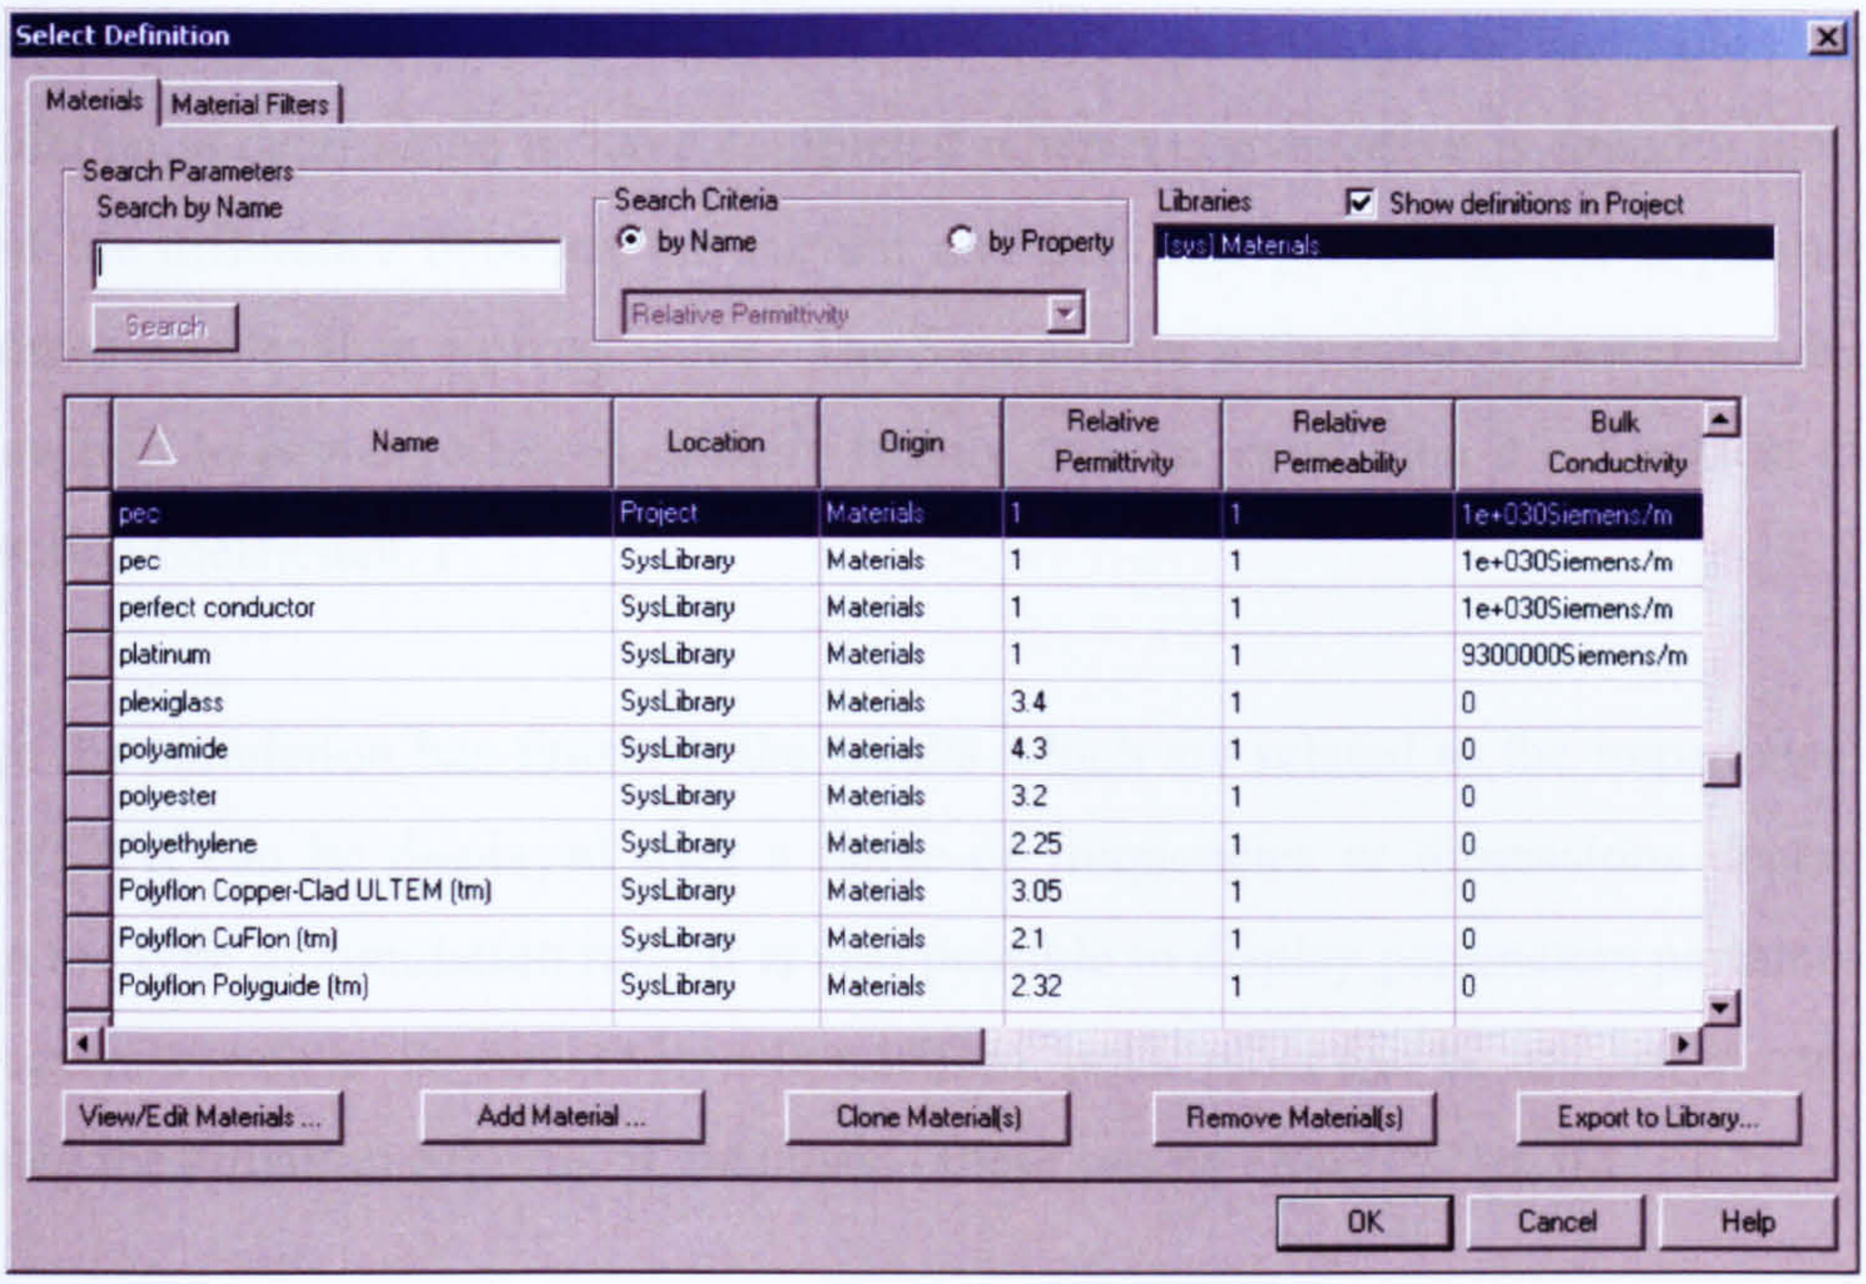

Figure 5.22 – Materials selection in HFSS

Simulations in HFSS employ a FEM (Finite Element Model) approach whereby HFSS divides its models into many different smaller elements called tetrahedra which can each be solved individually. The elements are created through adaptive meshing techniques which ensure that areas of interest, such as waveports (or sources), are assigned a higher mesh resolution. In addition the user can specify the mesh resolution that HFSS should use; a higher resolution means more tetrahedra and probably a greater accuracy, but much longer simulation times.

It is possible for the user to overlay some data onto their model, such as the **E** and **H** fields; the former is displayed in Figure 5.23. Such overlays are useful since they make it possible to visualise how the EM fields are distributed around an antenna – they can also be animated to show how they change with respect to time.

Simulations in HFSS can take a number of forms. They can be performed at single frequencies or over a range of frequencies. Batch simulations can be run through parametric analysis, which automatically adjusts model variables. When simulations are run HFSS takes a number of passes to calculate the required simulation data, starting first with a relatively small number of tetrahedral and gradually increasing

them with each pass. HFSS calculates scatter (S) parameters for each pass, and a simulation is determined to have completed when a convergence is found – that is – when the difference between the current and previous passes for the S-parameters becomes smaller than a preset value. The S parameter is the ratio of power incident to a waveport to power reflected; if there is only one waveport then it is identical to the reflection coefficient, Γ.

Once the simulation has finished, the results which are related to the impedance and the VSWR can be displayed over a range of frequencies or dimensions depending upon the type of simulation run. It is also possible to display parameters pertaining to the antenna such as its directivity. In addition, polar plots can be created in order to display the radiation patterns of antennas – these can be either 2D or 3D.

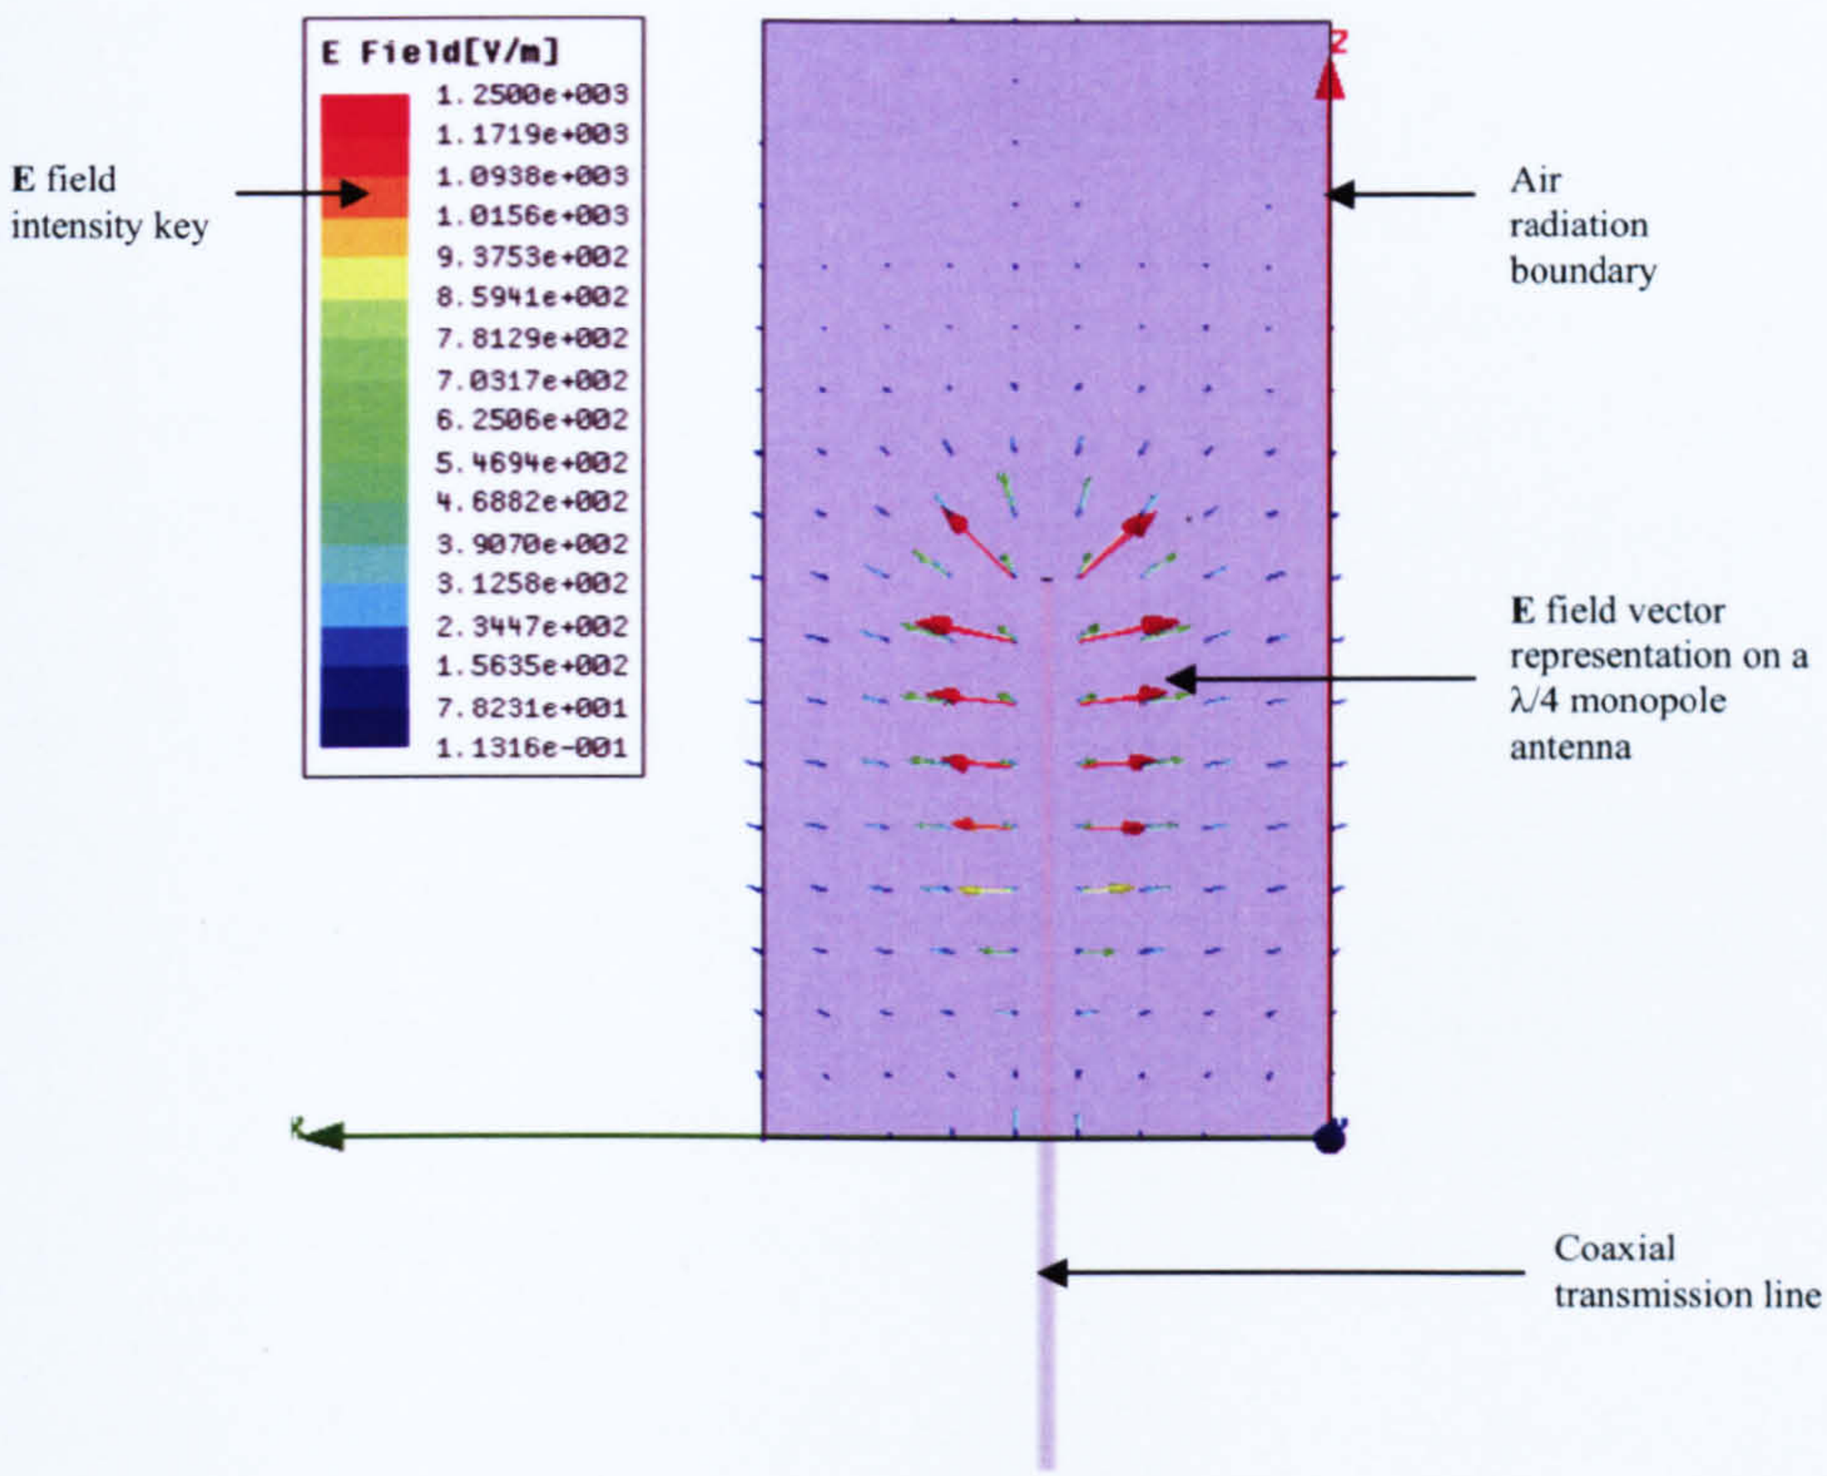

Figure 5.23 – Electric vector field overlay in HFSS

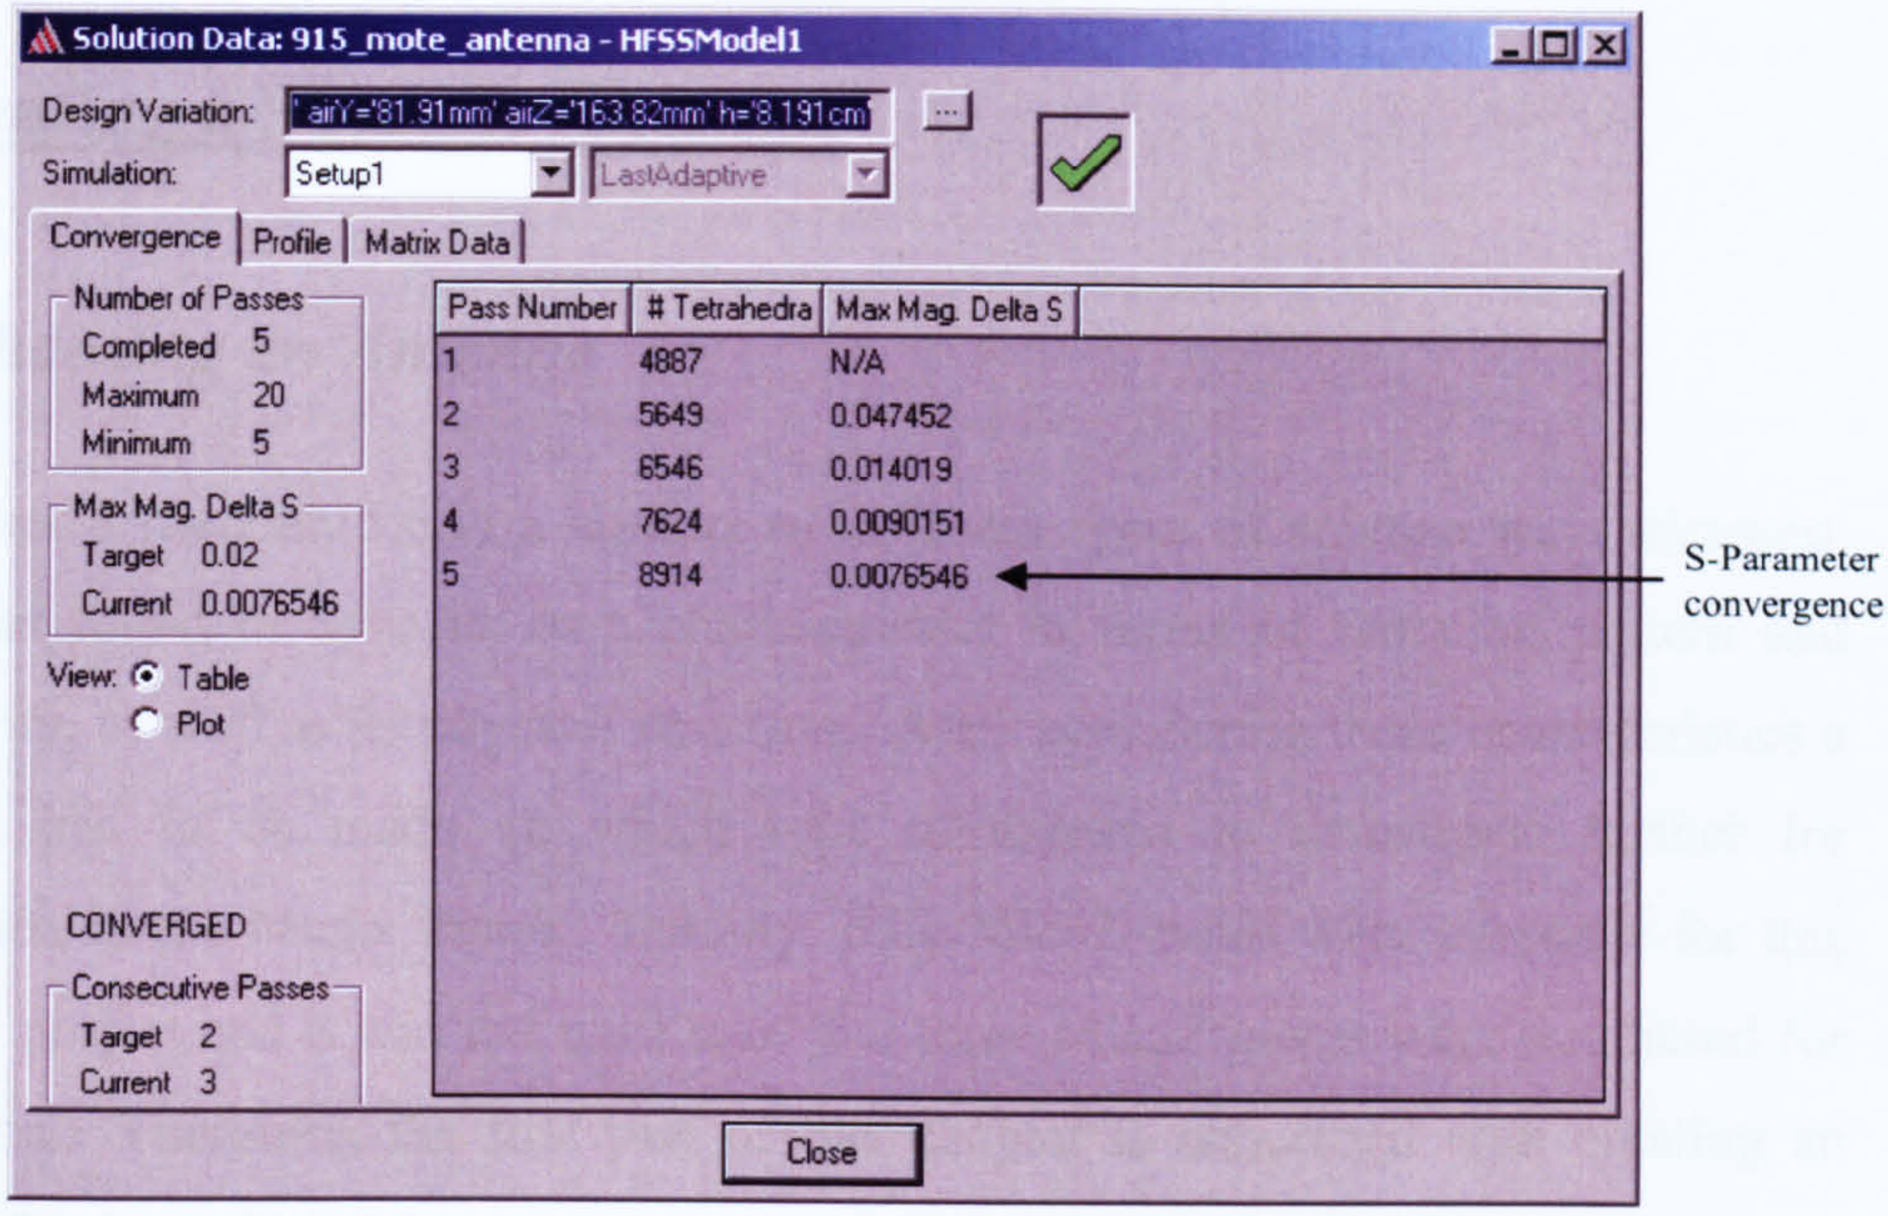

Figure 5.24 – HFSS simulation convergence

6. Creating a Suitable Antenna

6.1. Choosing an Antenna

In Chapter 5 (Section 5.6.1) a number of different types of antenna were dicussed. Each was found to have its own characteristics in terms of radiation pattern and direcitivity, as well as its physical structure. After considering these characteristics a decision had to be made on which type of antenna to investigate further for application to the Micax motes. Initially, only Mica2 motes were available for this research project and it was not until later that three MicaZ motes were purchased for evaluation. Therefore, the first part of this chapter is concerned with creating an antenna for use at 915MHz.

Industry disliked the idea of the existing Mica2 $\lambda/4$ monopole antenna because of its potential to snag. Dipole antennas are typically in the order of $\lambda/2$ in length, and loop antennas often have a circumference equal to λ. Smaller sizes are possible for both loops and dipoles but they do not make effective radiators [136,149] so for the same reason as the monopole both designs were discarded. Horn antennas were also discarded due to their size, although they have been used for evaluation purposes.

This left the option of a PCB antenna. With size and practicality being major concerns of industry it seemed that a low profile PCB antenna would be ideal. It was thought that such an antenna would be suitable for retrofitting to the current Mica2 motes, and in the future could possibly be integrated with the mote circuitry in a combined PCB design. In order to facilitate this, an aim was set of creating an antenna no greater than the size of Mica2 PCB (i.e. - 57mm × 32mm).

Although PCB antennas do have their advantages, it is noted in literature that they tend to suffer from a narrow bandwidth, quite often in the order of just a few percent [150]. It was determined that the Mica2 motes required a bandwidth such that f_0 = 915MHz, f_l = 898MHz and f_u = 933MHz. Note that the values of f_l and f_u include an extra 5MHz allowance for the FSK (Frequency Shift Keying) [151] modulation

scheme used by the Mica2 motes. Putting these figures into Equation 5-14 results in a calculated bandwidth requirement of 3.83%.

6.2. The Patch Antenna

6.2.1. Design and Simulation Setup

Chapter 5 talked about a specific type of a PCB antenna: the patch. These are considered to be the simplest of the PCB antennas because they consist of a rectangular piece of copper above a ground plane. The two are separated by a substrate. Balanis offers a partial guide to designing a probe fed patch antenna [152]. This guide is included in Appendix D along with the appropriate equations used for dimension calculations.

From the guide it was possible to calculate the dimensions of the copper patch itself when the operating frequency (915MHz), height of substrate (h) and dielectric constant (ε_r) were known quantities. Table 6.1 shows the dimensions used for this antenna, and Figure 6.1 shows an overview of the model in HFSS. FR4 epoxy board ($\varepsilon_r = 4$) was used as this is common for printed circuits.

h	*pcbX*	*pcbY*	*W*	*L*
1.6mm	100mm	155mm	99.7mm	77.9mm

Table 6.1 – Patch antenna dimensions and parameters

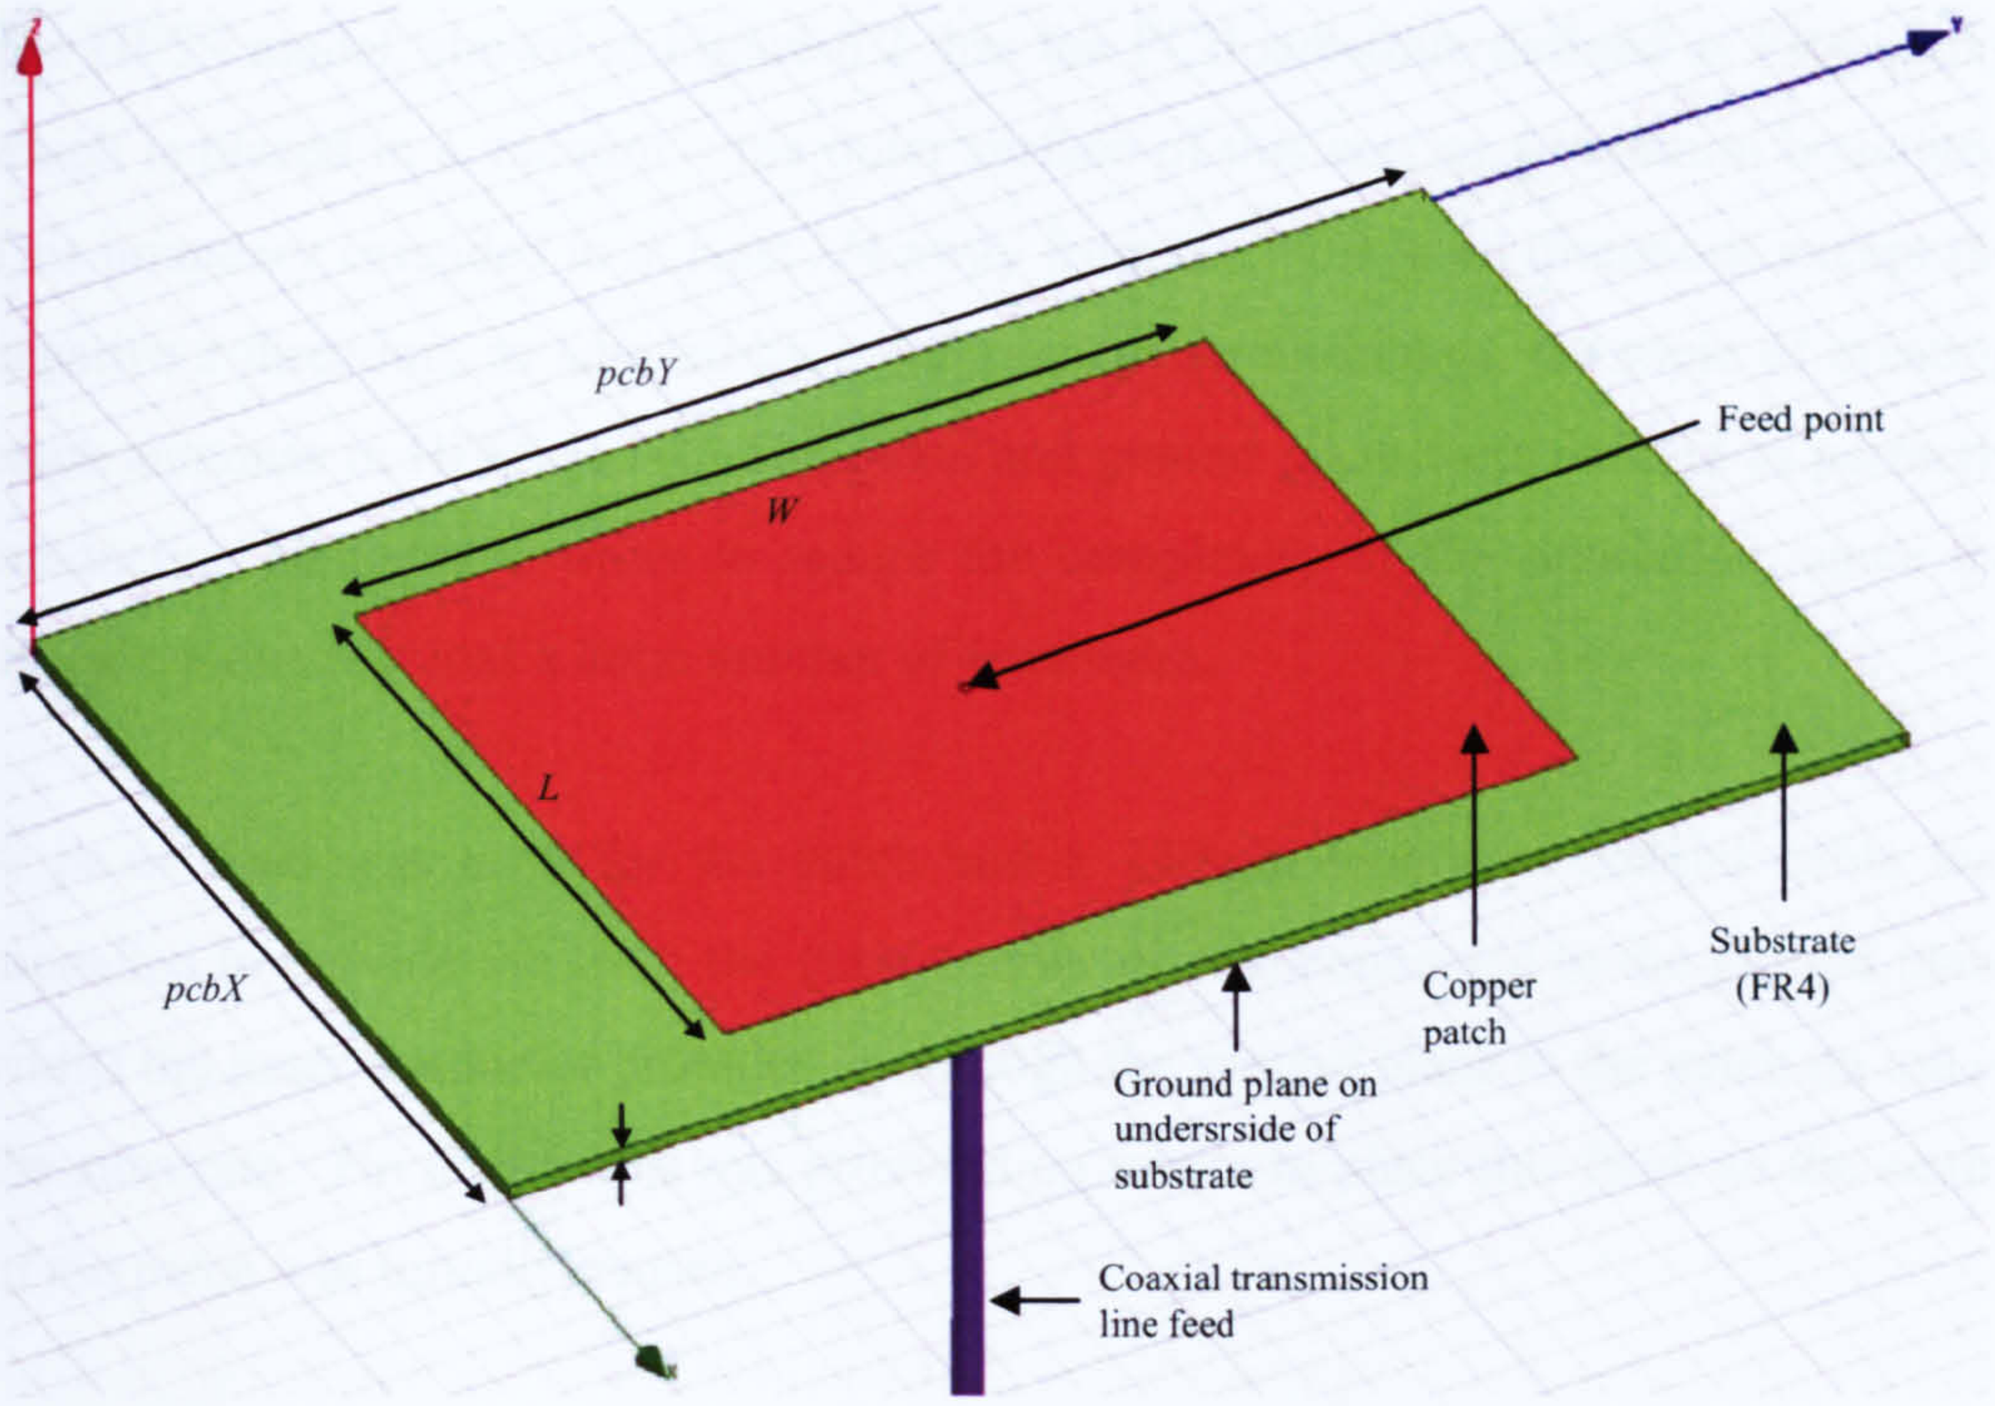

Figure 6.1 – Patch antenna simulation model

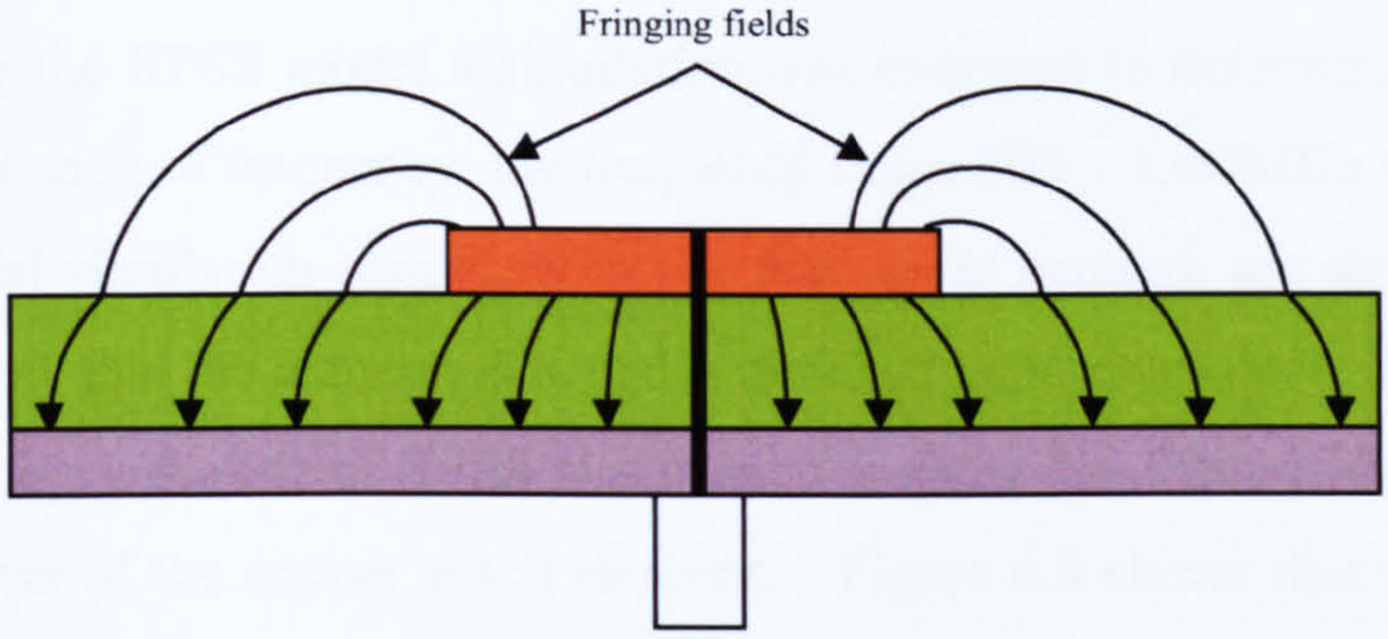

Figure 6.2 – Fringing fields on a patch antenna

It is possible to have the PCB and the ground plane equal in size to the radiating patch, but this does not allow fringing fields to form [152-153]. These are a result of the electric field lines between the top edge of the patch and the ground plane, as shown in Figure 6.2. Having these fields increases the radiated power and so they are desirable, although the additional size of the ground plane is not. There is no agreed suitable size for a ground plane, and so it was decided that it should be the same size as the PCB for simplicity.

The HFSS model, shown in Figure 6.2, has the PCB substrate defined as FR4 epoxy which is placed in a vacuum – the outer surface of this acts as a radiation boundary. This boundary extended to at least $\lambda/4$ away from the PCB in all directions except the negative z direction, as advised by a comprehensive tutorial on the topic of creating patch antennas in HFSS [154]. The patch and ground plane were defined as perfectly conducting elements in order to reduce the complexity of the simulation, since the vacuum alone required a large number of tetrahedra.

A probe feed was used for the patch and a 250mm length of coaxial cable was modelled to provide the feed; the outer conductor was connected to the ground plane whilst the inner conductor protudes up through the ground plane to the patch on top of the substrate. No instruction was provided on where to place this feed, so the centre of the patch was initially chosen.

6.2.2. Simulation Results

After creating the HFSS model a simulation was executed to determine the reflected power as a function of frequency; the frequency range 800 – 1500MHz was arbitrarily chosen. Initial simulation results, with the feed point centred, are shown in Figure 6.3. They show that the antenna is actually matched at approximately 1.42GHz. The parametric analysis feature of HFSS was used to to show the effect of moving the feed toward a corner of the copper patch element. Figure 6.3 shows that when the feed offset is set to 55mm the antenna becomes well matched to 915MHz. Figure 6.4 shows that the minimum reflected power at 915MHz occurred when the feed offset was 55mm. This approach of minimising the reflected power was adopted since it follows that very little reflected power (or VSWR close to 1) represents a good impedance match which means that one does not have to be concerned with complex impedances. This strategy was adopted throughout this work.

Figure 6.3 does not display the VSWR (as a result of it having a huge range), but an equivalent of the VSWR = 2 line exists at 11.1% reflected power; this was determined from Equations 5-11 and 5-8. Using this line in conjunction with Equation 5-14 the

predicted bandwidth of this antenna is only 2.19%, 1.64% less than that required by the Mica2 in order to utilise all of the transceiver channels.

Figure 6.5 shows the simulated radiation pattern of the patch antenna with the modified feed point. Much of the radiated power is focused perpendicular to the radiating patch element in the positive *z* direction (i.e. 0° in the *xz* and *yz* planes); here the power radiated is 9.5dB (or almost 10 times) more than any direction in the *xy* plane. This shows some agreement with the predicted radiation pattern in Appendix C.5. Since so much power is radiated in the positive *z* direction this antenna has a high directivity, which HFSS calculates to be 7.3dBi.

There are, however, some noteable differences between the simulated results and those predicted in Appendix C.5. In theoretical literature many authors consider the case of an infinite ground plane which would allow no radiation between 90° and 270° in the *xz* and *yz* planes as well as there being no radiation at all in the *xy* plane; the simulated results show some radiation does occur here however. In practise ground planes are finite and often as small as possible. This means that EM waves can diffract at the edges of the ground plane – something which is not taken into account in the infinite case [153]. Therefore, although the theory predicts that a ground plane prevents transmission in a particular direction, the reality is that it causes significant attenuation but not a complete null.

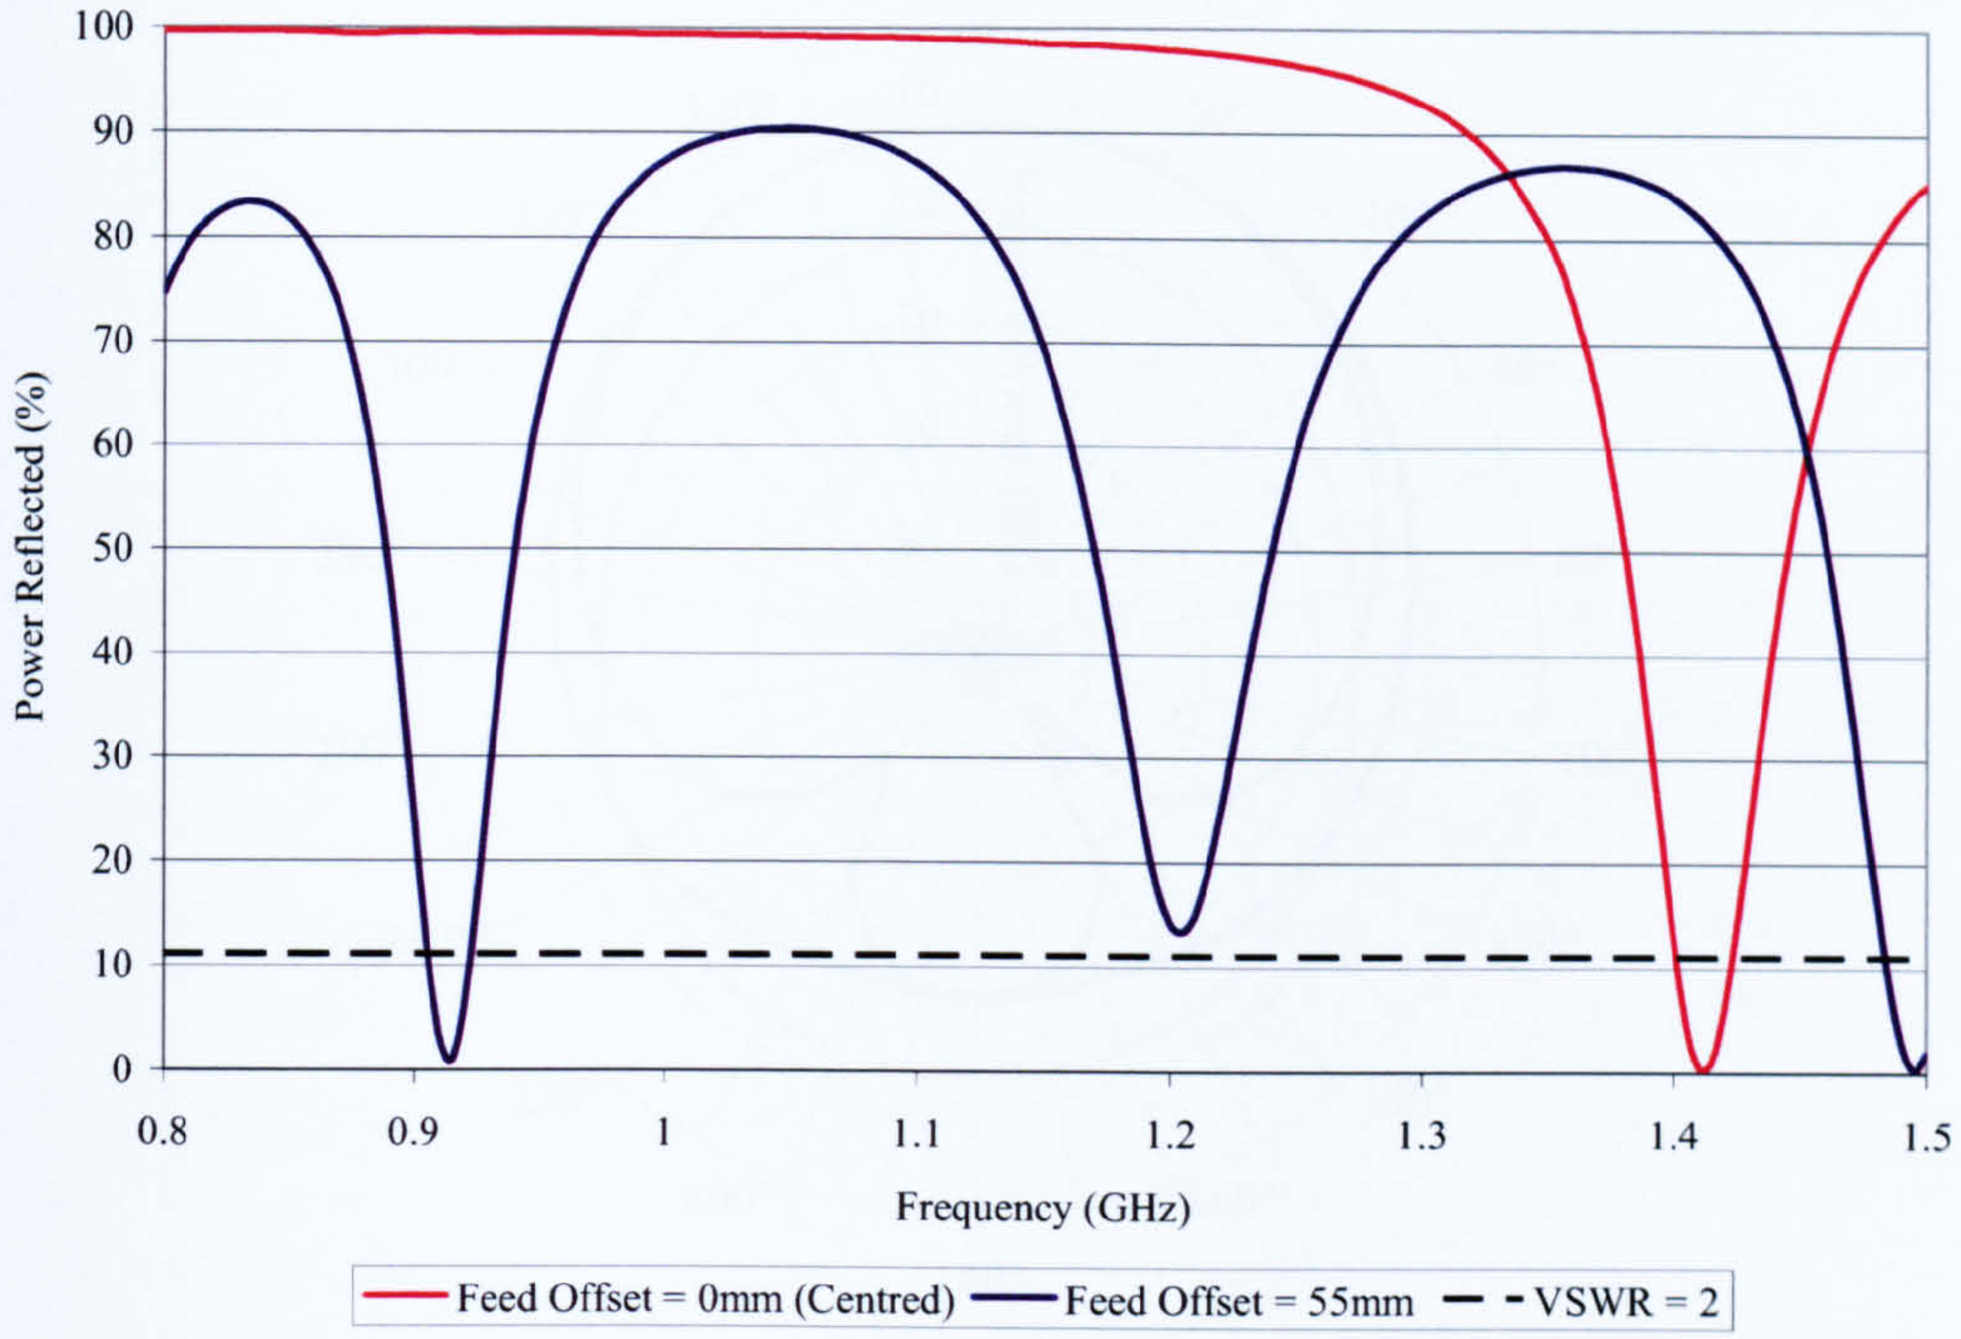

Figure 6.3 – Power reflected as a function of frequency for the patch antenna

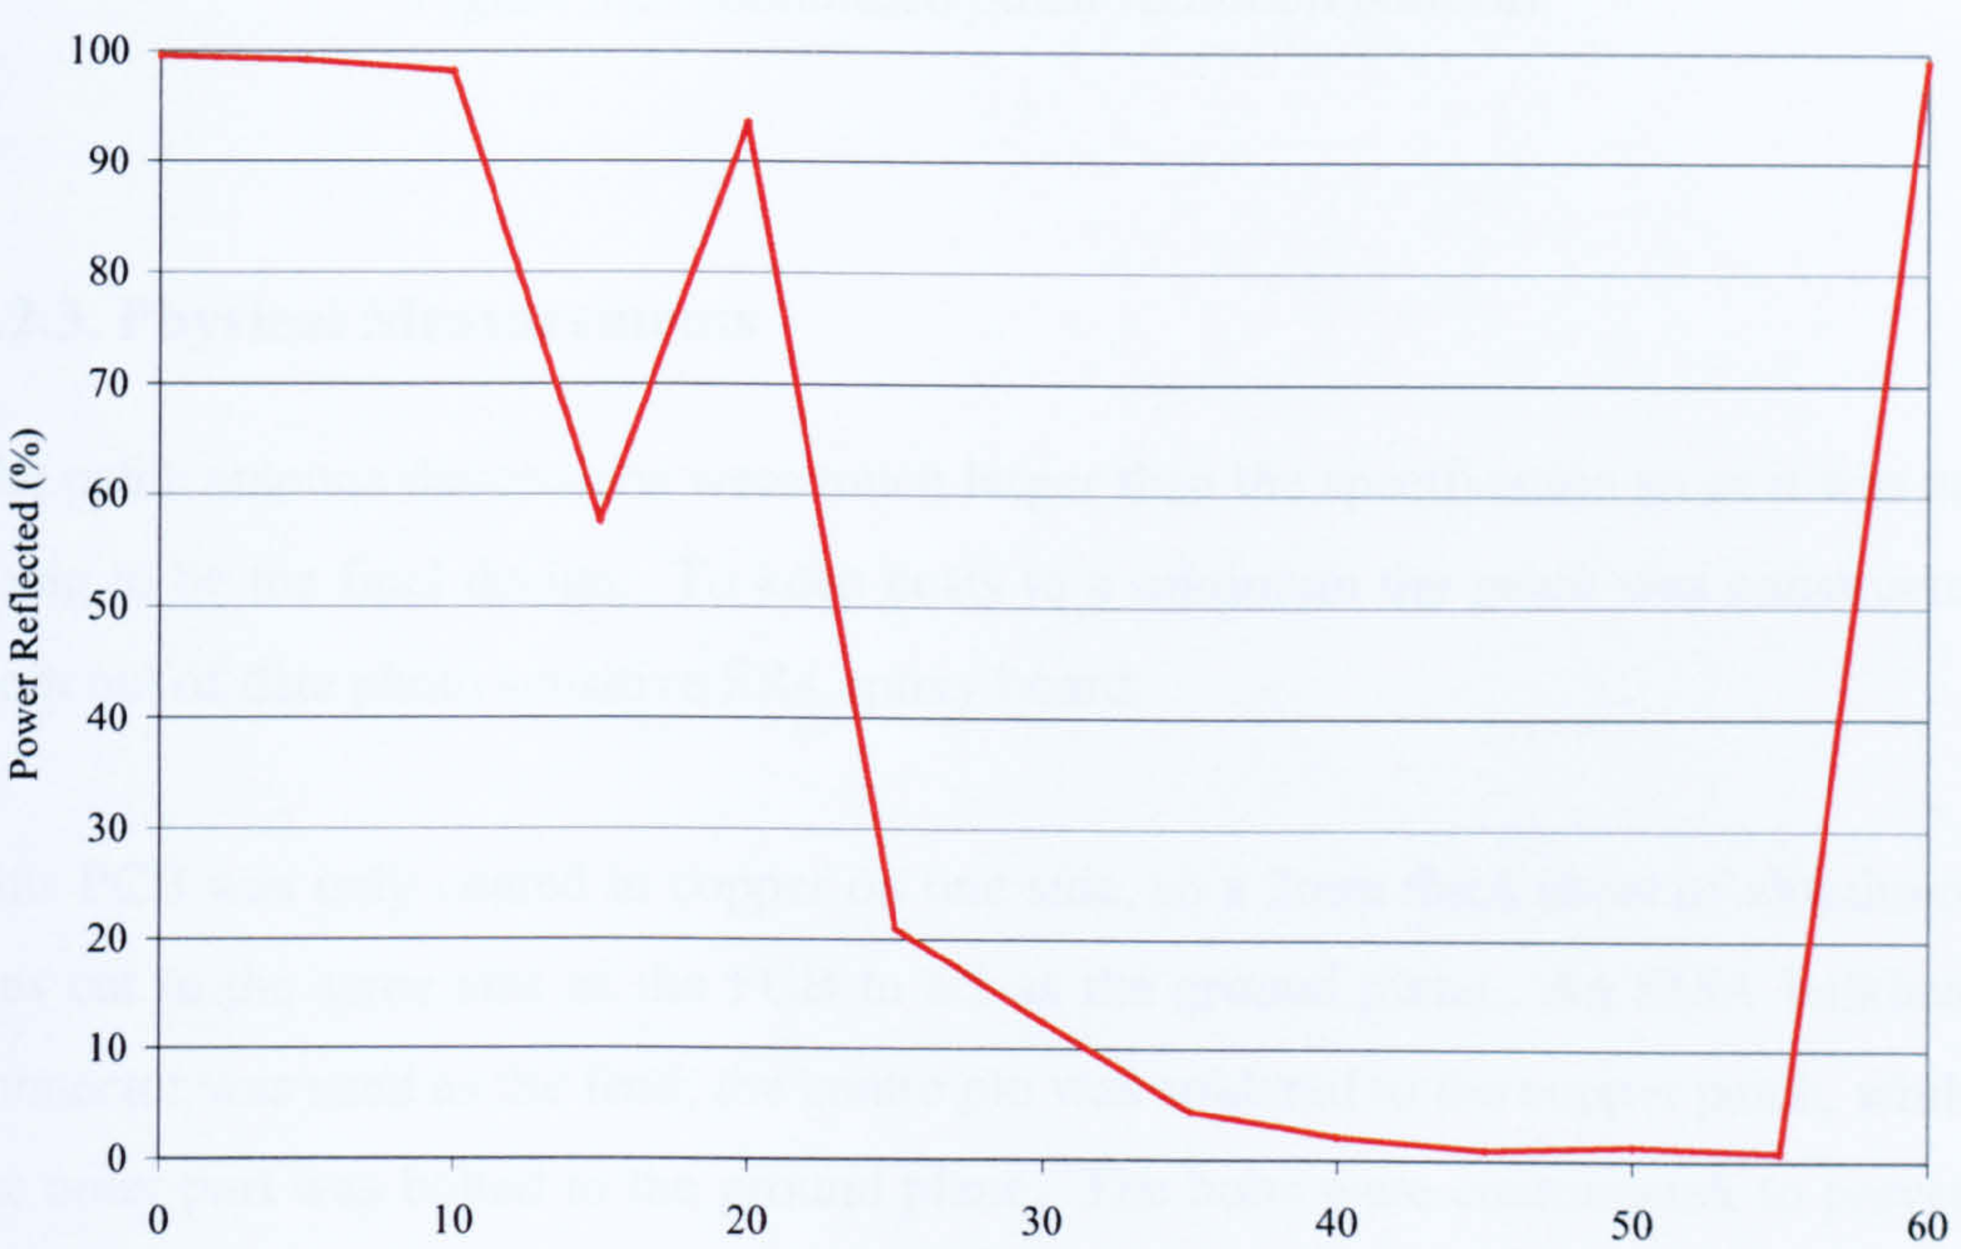

Figure 6.4 – Refinement of the patch antenna by moving the feed point

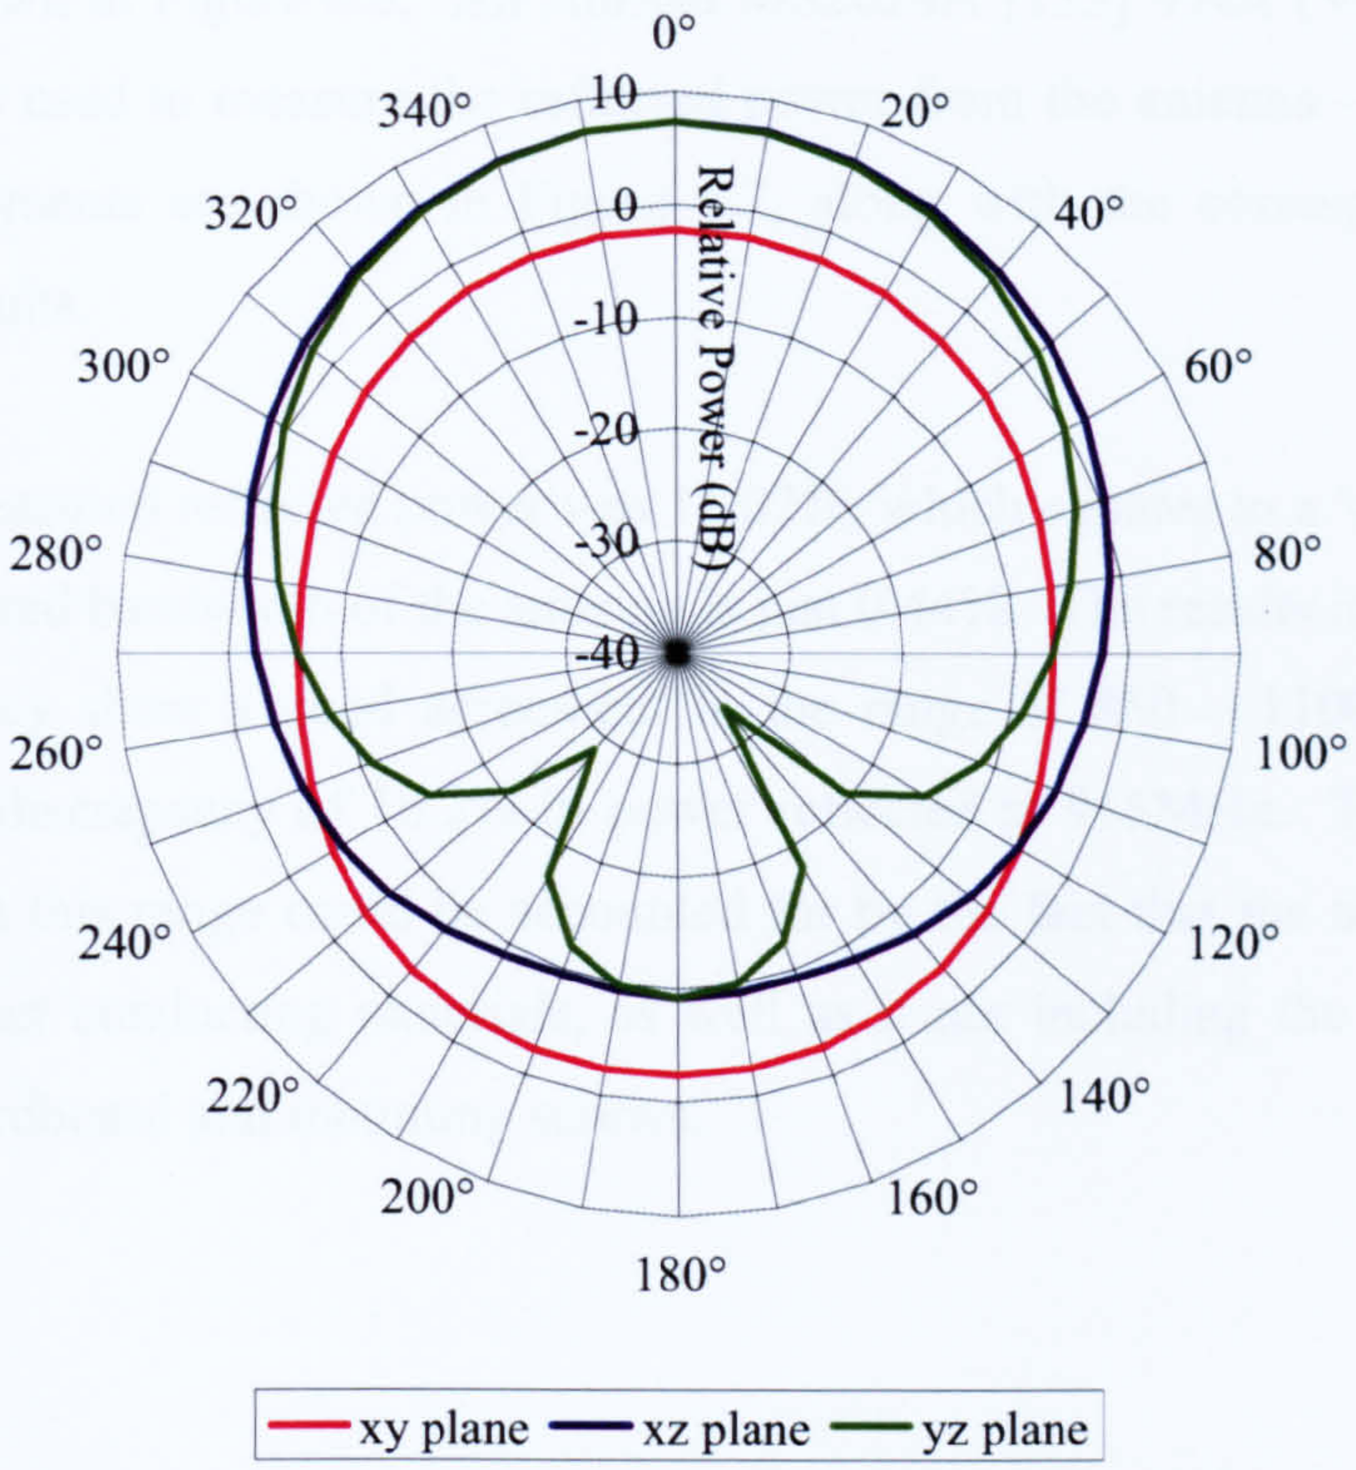

Figure 6.5 – Simulated patch radiation patterns

6.2.3. Physical Measurements

The patch antenna dimensions were much larger than the specification so as it was not going to be the final design. To keep costs to a minimum the patch was constructed from out of date photo sensitive FR4 epoxy board.

This PCB was only coated in copper on one side, so a 2mm thick sheet of aluminium was cut to the same size as the PCB to act as the ground plane. An SMA bulkhead connector was used as the feed; the centre pin was soldered to the copper patch, whilst the outer part was bolted to the ground plane. The bolts were countersunk to prevent an air gap appearing between the ground plane and PCB. The antenna was then mounted on hardboard using plastic spacers to provide a 28mm clearance between the ground plane and mounting board as shown in Figure 6.6; this allowed access to the SMA bulkhead connector and ensured the rigidity of the whole structure. The end

product is shown in Figure 6.6. An Anritsu MS2024A [155] VNA (Vector Network Analyser) was used to measure the reflected power from the antenna – the results of these measurements are shown in Figure 6.7, along with the corresponding HFSS simulation results.

The lowest measured reflected power was 11.02%, which equates to a VSWR of 1.99, and the measured bandwidth of the antenna is just 0.44%. The results in HFSS for the target frequency show a good agreement in the range of 850 – 1100MHz, despite there being a discrepancy of 10.2% in power reflected at 915MHz. The differences experienced in this range could be accounted for by the fact that the model in HFSS assumed perfect conducting materials, as well as it not including the extra material such as the hardboard and mounting screws.

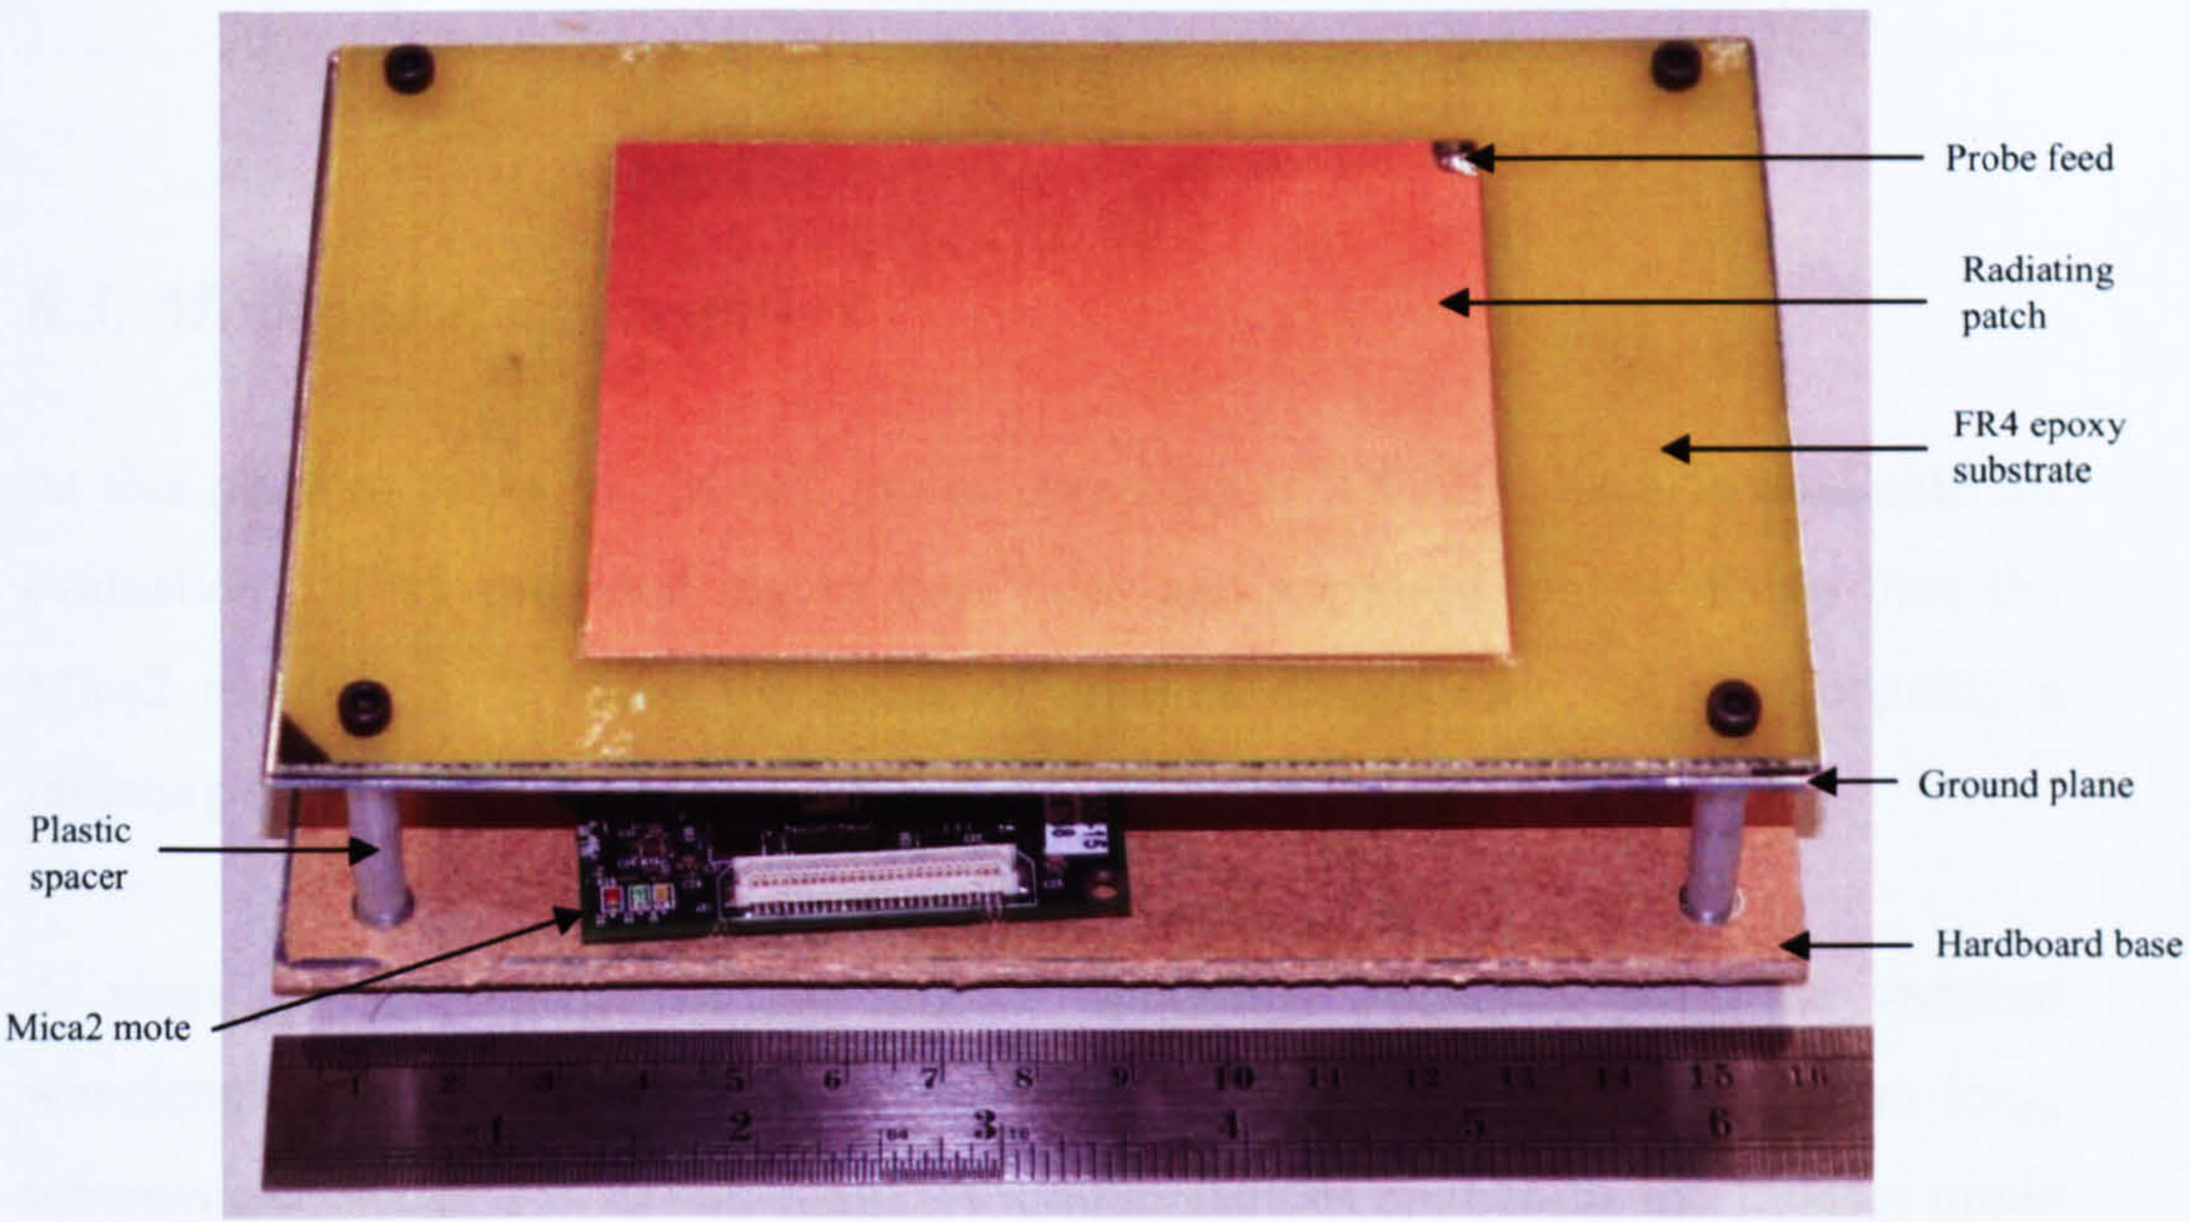

Figure 6.6 – Patch antenna

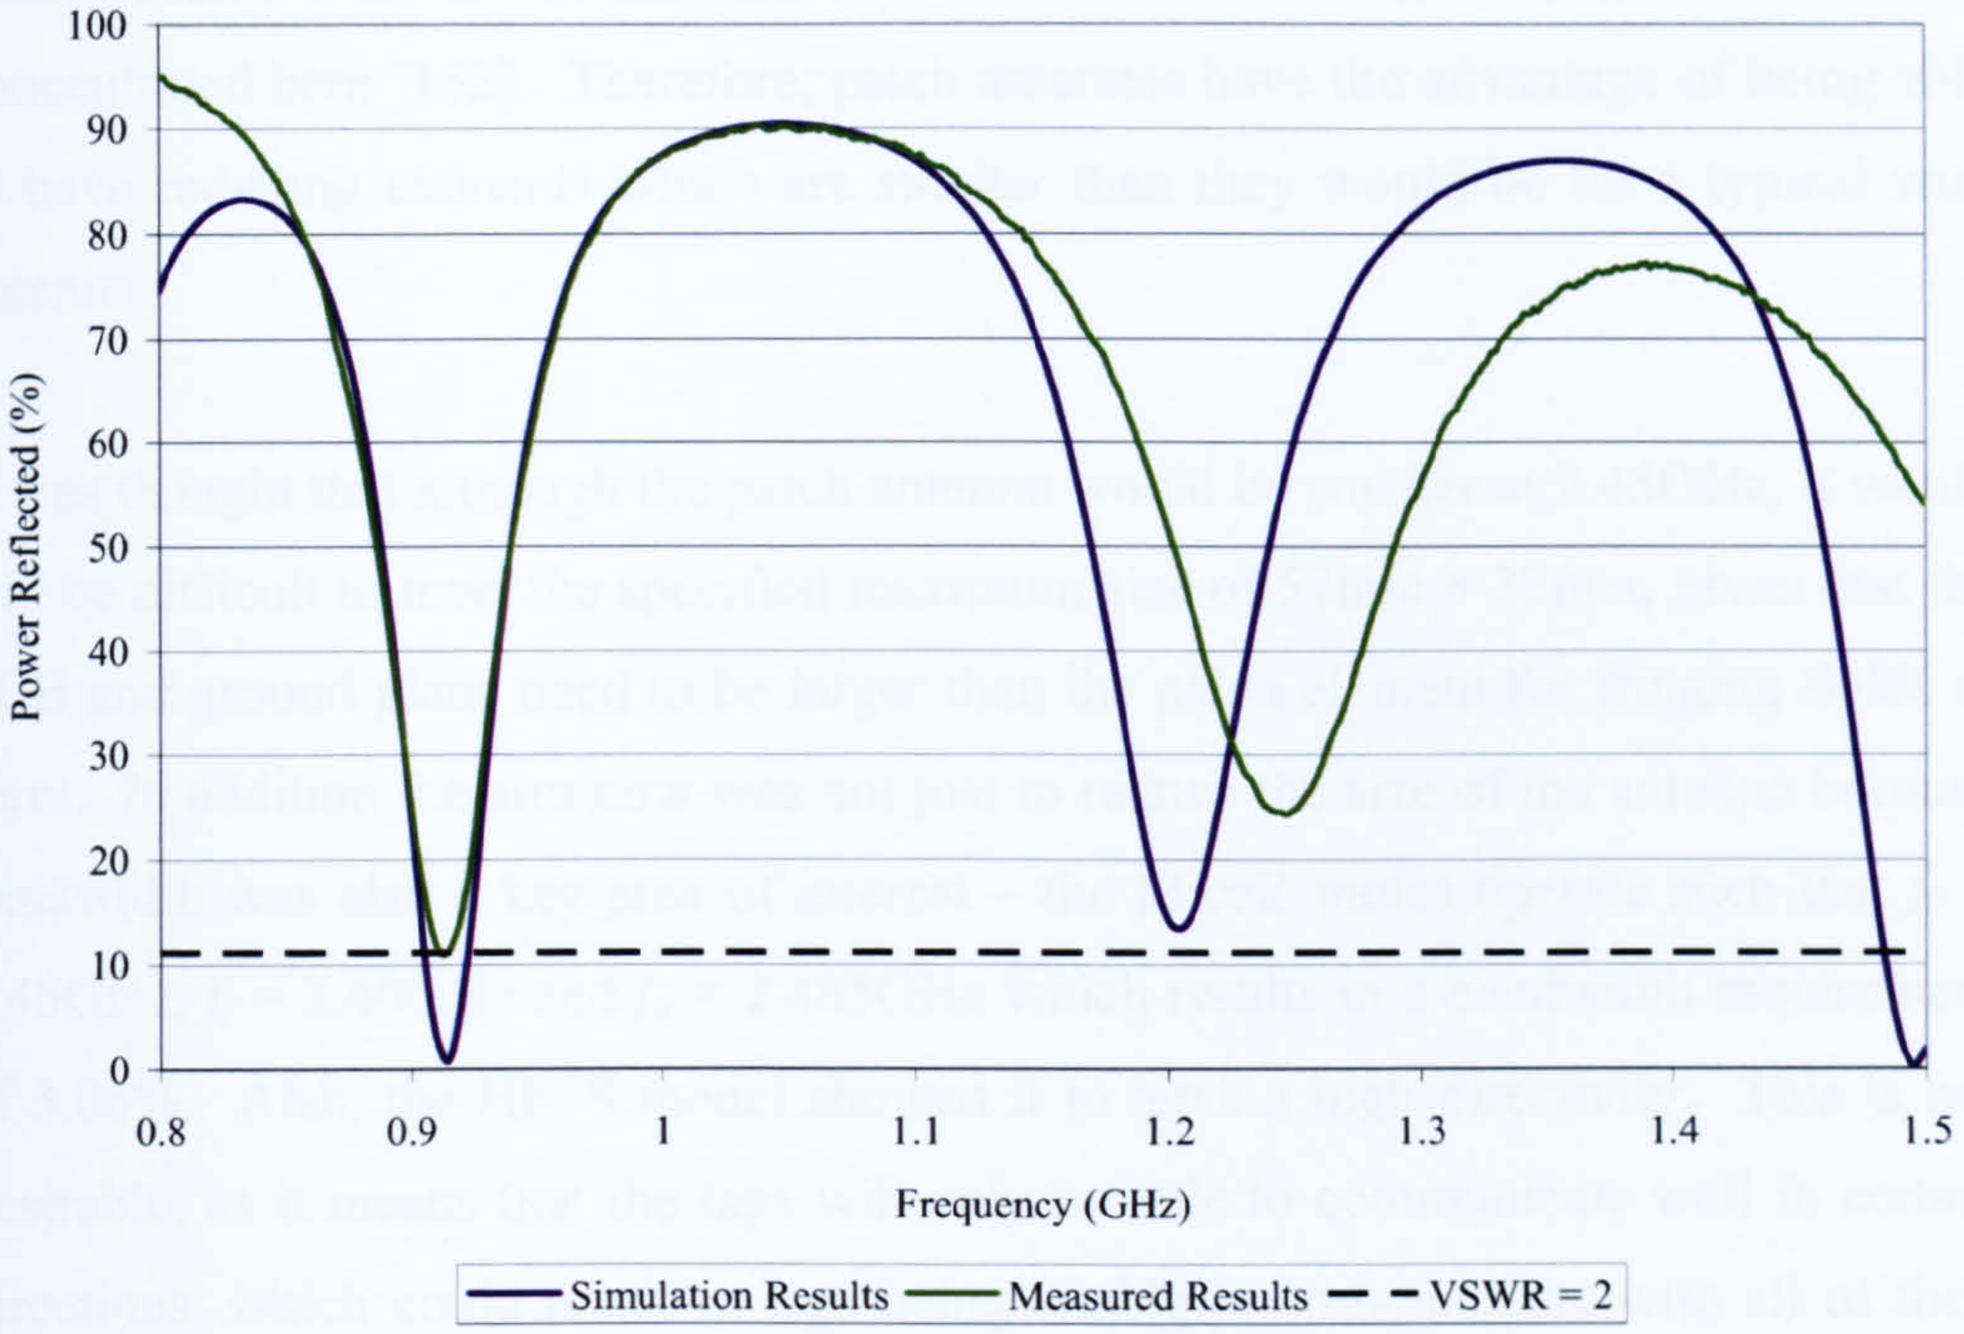

Figure 6.7 – Patch measured vs. simulation power reflection

6.3. Modified Requirements

At this point in the work, MicaZ motes operating at 2.45GHz became available for evaluation. They provided higher data rates and appeared more reliable than the Mica2 motes. As a result the focus of antenna design switched to producing a radiating structure for the MicaZ mote.

A significant advantage of the higher operating frequency was the reduced wavelength and resulting antenna size. The MicaZ monopole is just 31mm long, whereas the Mica2's is 82mm long. A similar feature applied to the Balanis guide since a patch antenna operating at 2.45GHz would have a radiating copper element of L=29.4mm × W=38.0mm. With PCB antennas however, the permittivity of the substrate effectively reduces the wavelength compared to that in free space. This is because the electric field lines are formed in both the air and the substrate, so the effective dielectric constant is somewhere in between that of air and the substrate. In addition, as the frequency increases, Balanis notes that the effective dielectric constant

becomes closer to that of the substrate since the electric field lines become more concentrated here [152]. Therefore, patch antennas have the advantage of being able to have radiating elements which are smaller than they would be for a typical wire antenna.

It was thought that although the patch antenna would be smaller at 2.45GHz, it would still be difficult to meet the specified maximum size of 57mm × 32mm, given that the PCB and ground plane need to be larger than the patch element for fringing fields to form. In addition the aim now was not just to reduce the size of the antenna because bandwidth was also a key area of interest – the MicaZ motes operate such that f_0 = 2.45GHz, f_l = 2.40GHz and f_u = 2.485GHz which results in a bandwidth requirement of 3.06%. Also, the HFSS model showed it to have a high directivity. This is not desirable, as it means that the tags will only be able to communicate well in certain directions, which could result in tags being unable to communicate with all of their surrounding neighbours. It seemed that a major cause of this was the ground plane – therefore it was thought that a PCB antenna without a ground plane on its underside might be a more promising solution.

6.4. Co-Planar Waveguide Antenna

6.4.1. Simulation Model

Nithisopa et al [156] have designed a broadband co-planar waveguide (CPW) fed slot antenna which, in simulations, has proved to be suitable for use over the range of approximately 2.1-2.9GHz, resulting in a bandwidth of 32.7%.

The term CPW itself refers to the way in which the antenna is fed; two parallel slots are cut into a copper surface to act as a transmission line feed to the radiating elements of the antenna itself. The radiating elements come in many different forms, although it appears that the slot type is quite popular. The copper surrounding the feed slots (see Figure 6.8) act as the ground plane, but since it is not directly behind the radiating part of the antenna it was thought that it may not result in directivity as high as the patch antenna.

Based upon the work conducted by Nithisopa, a HFSS model was created as shown in Figure 6.8. The dimensions used are shown in Table 6.2, where the substrate used is FR4. The model was set up by following strict guidelines [157], provided by the developer of HFSS, for the creation of CPW models. These guidelines give information particularly relating to the size and position of the waveport used to excite the antenna.

Substrate	ε_r	*pcbX*	*pcbY*	*h*	*W1*	*W2*	*H1*	*H2*	*L1*
Duroid	2.2	90mm	45mm	1.6mm	0.5mm	2.4mm	23mm	10.5mm	39mm
FR4	4.4	90mm	45mm	1.6mm	0.5mm	4mm	23mm	10.5mm	37mm

Table 6.2 – Comparison of parameters when using FR4 and Duroid substrates

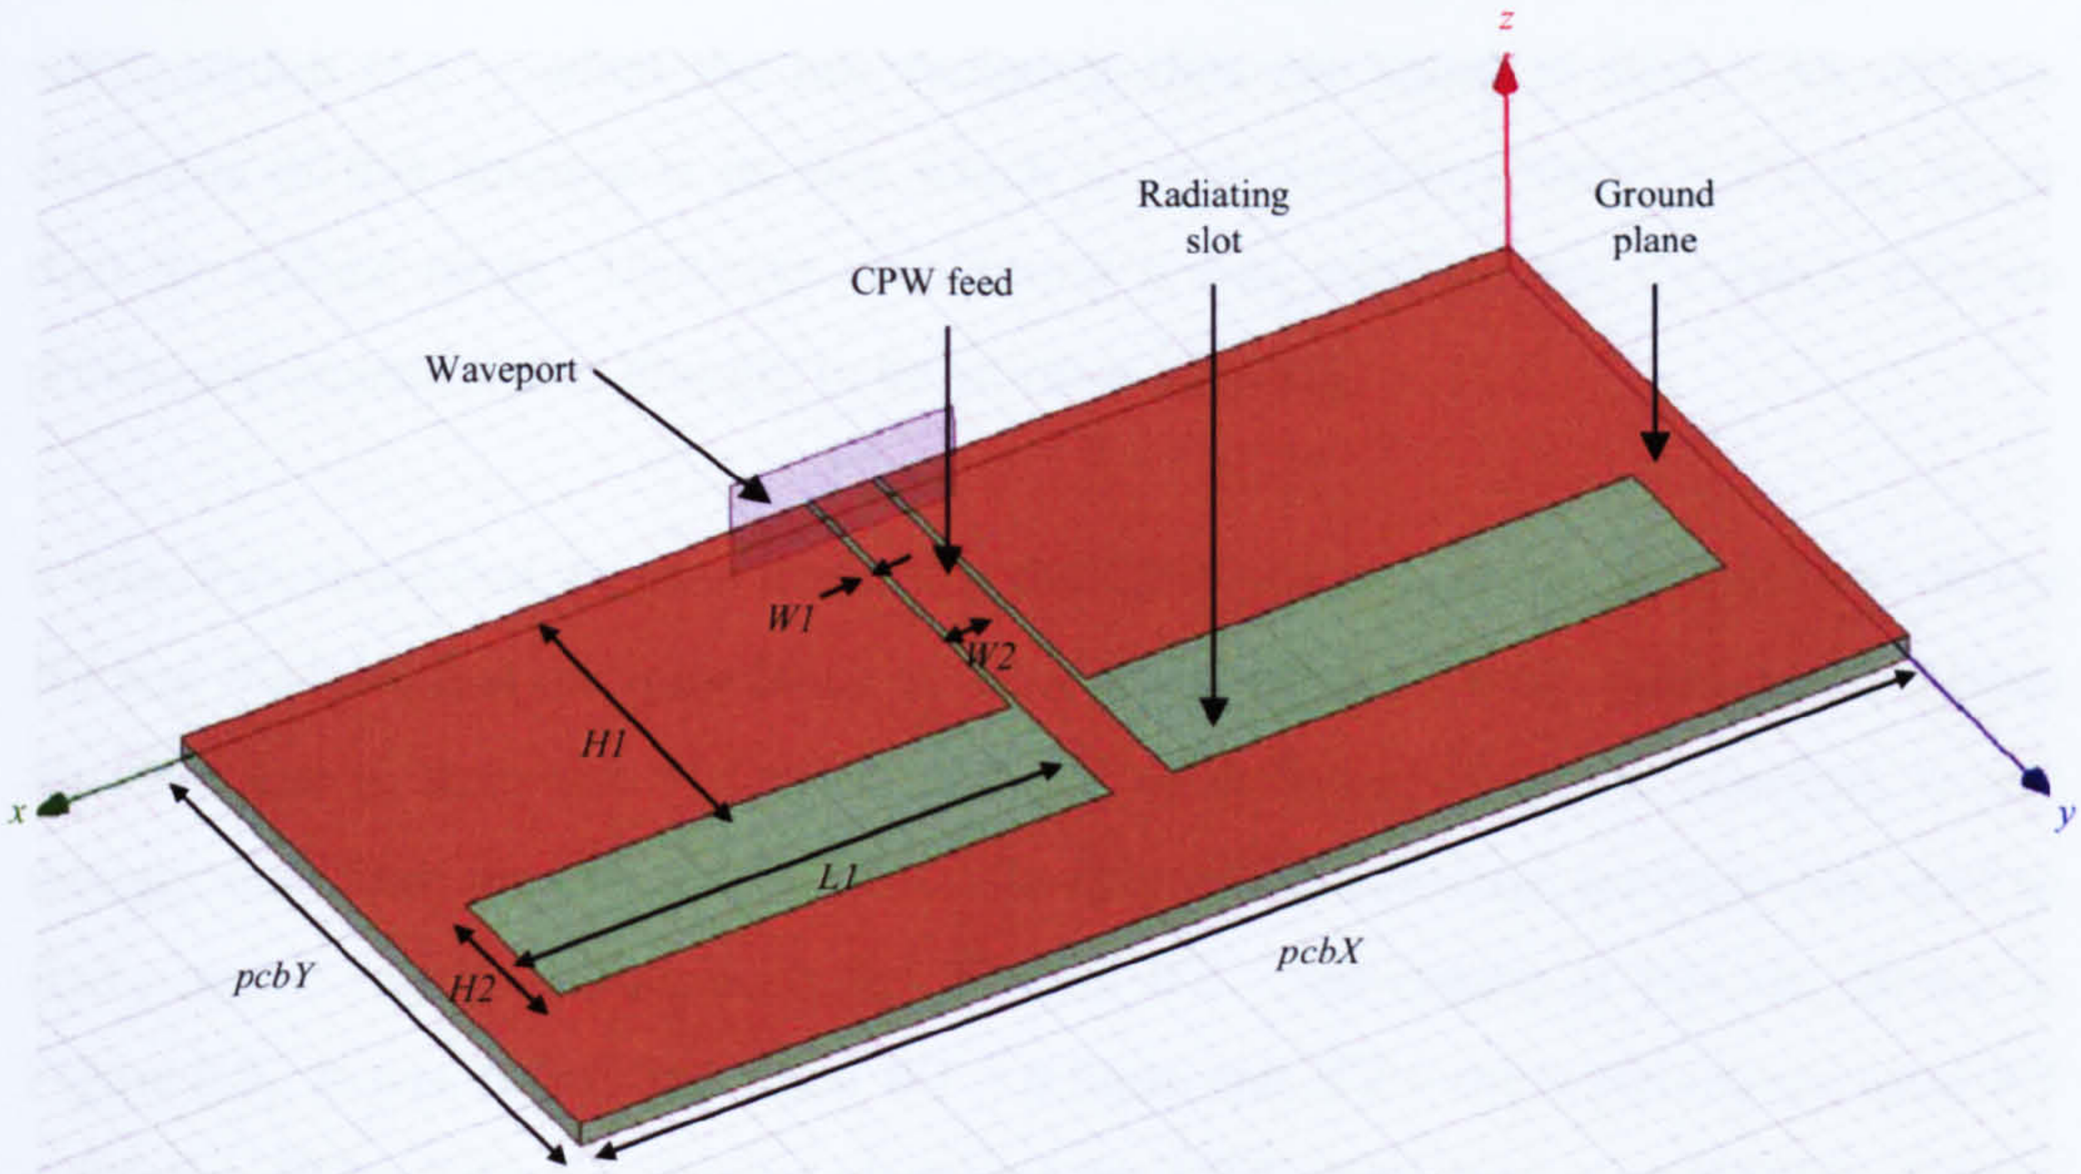

Figure 6.8 – CPW antenna simulation model

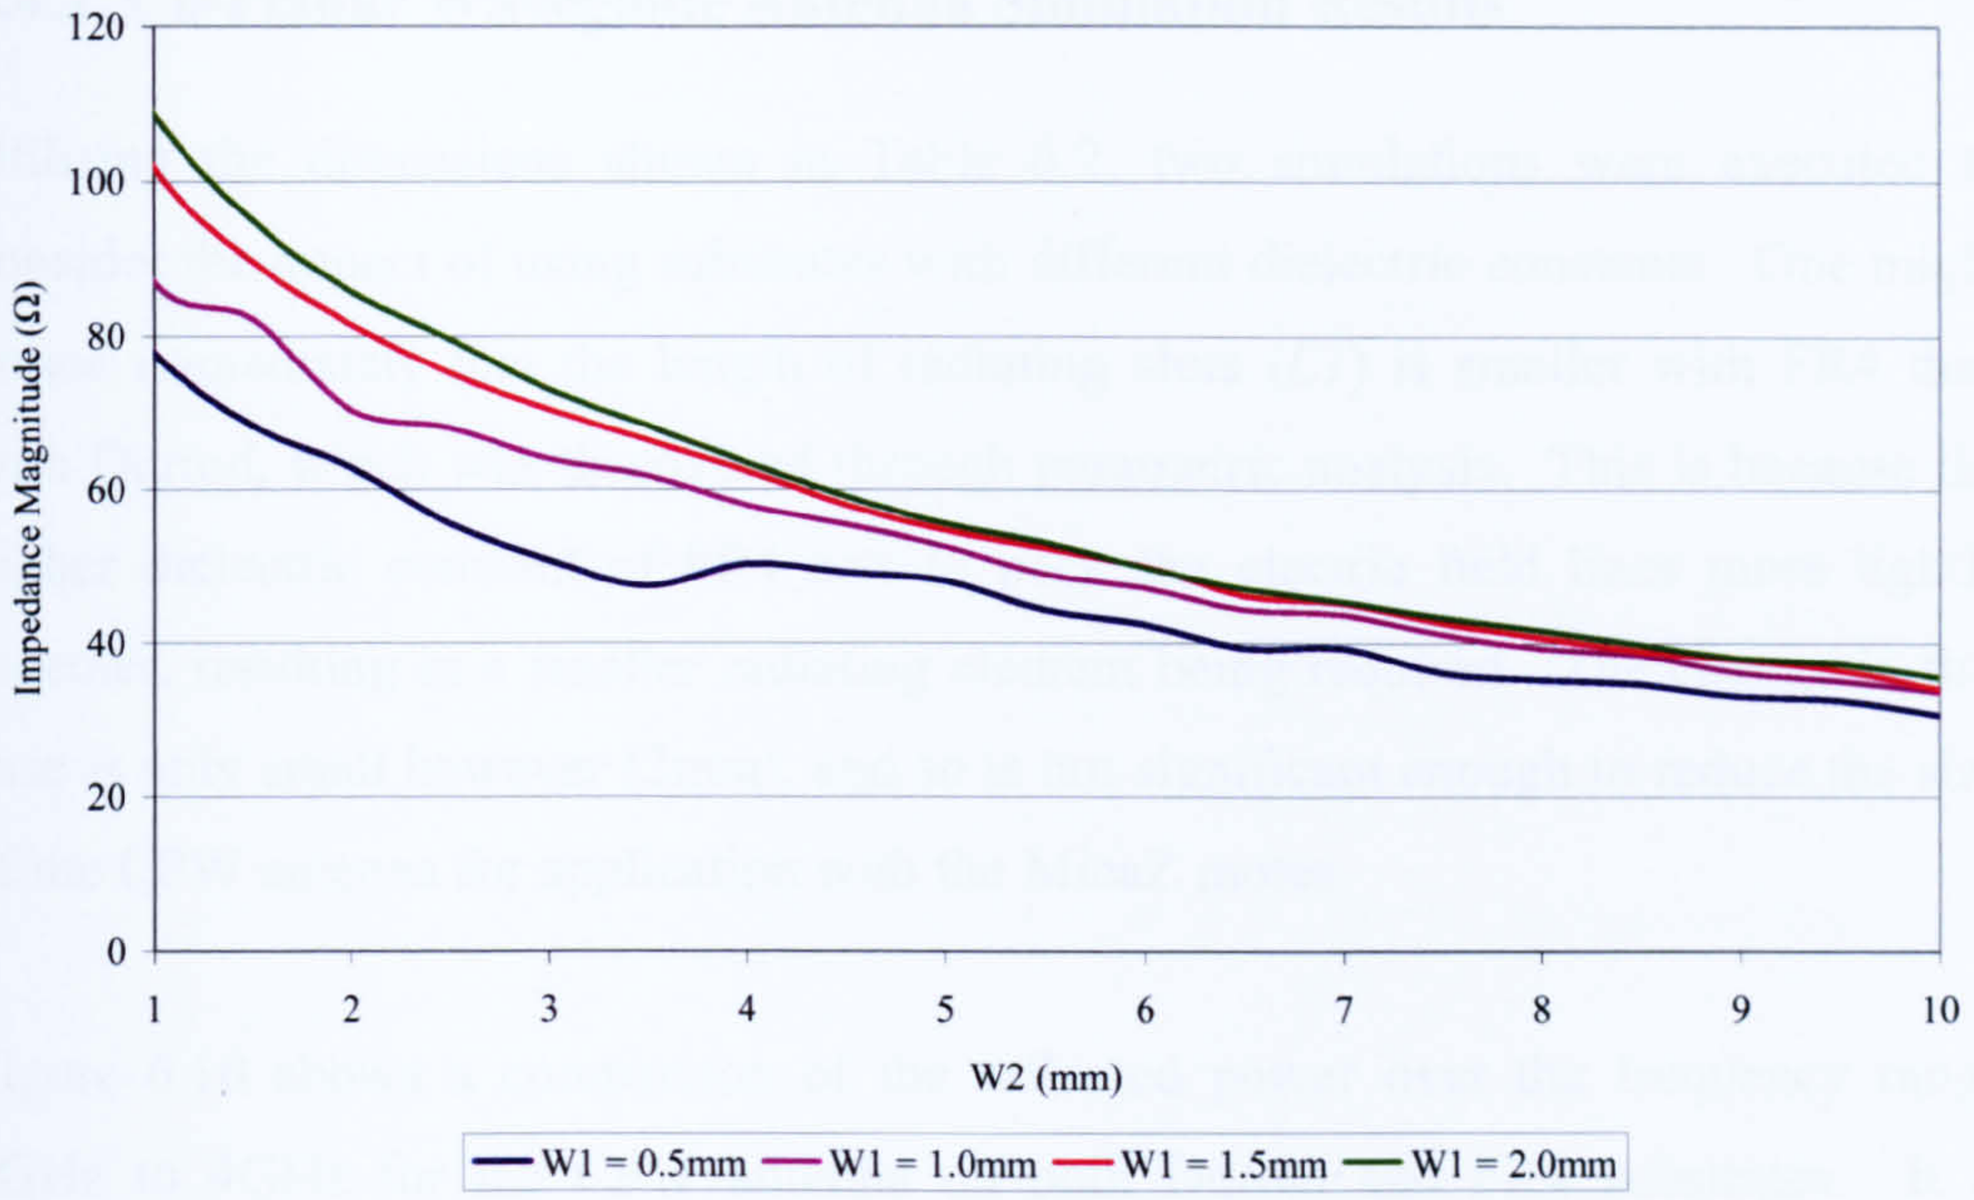

Figure 6.9 – Impedance as a function of W2 with incremental changes to W1

Nithisopa provides the dimensions and substrate type (Duroid) but gives no data relating to the radiation pattern of their antenna design when it is in a symmetrical configuration (i.e. – when the two radiating slots are equal in size). An obvious downside of this antenna is its size which is approximately 90mm × 45mm, larger than the MicaZ mote. However considering it was smaller than the patch antenna previously created it was thought to be worth evaluating.

It was discovered that, prior to modelling the CPW antenna, it was necessary to define *W1* and *W2* such that the impedance of the antenna matched 50Ω. Nithisopa had done this, but due to their use of Duriod material of differing dielectric constant, it was thought that there may be a difference. To this end, parametric analysis was used to find a set of suitable dimensions for *W1* and *W2* – the results of this are shown in Figure 6.9. It was found that when *W1* = 0.5mm and *W2* = 4mm (Nithisopa used *W1* = 0.5mm, *W2* = 2.4mm), the impedance of the CPW feed lines was 50.27Ω. Whilst other values for *W1* and *W2* also gave an impedance close to 50Ω it was decided to use the smallest values to minimise the antenna size.

6.4.2. Co-Planar Waveguide Antenna Simulation Results

Utilising the dimensions shown in Table 6.2, two simulations were executed to consider the impact of using substrates with different dielectric constants. One might notice immediately that the length of radiating slots (*L1*) is smaller with FR4 than with Duriod, which was discovered through parametric analysis. This is because the higher dielectric constant of FR4 acts to pack the electric field lines more tightly together, resulting in a smaller radiating element being required. The change in this case is only small however (2mm), and so is not significant enough to reduce the size of the CPW antenna for application with the MicaZ motes.

Figure 6.10 shows a comparison of the reflected power over the frequency range 2GHz to 4GHz for the CPW antenna on both Duroid and FR4 substrates. It is noticeable that the bandwidth of the antenna using Duriod is greater (25.2%) than the bandwidth when using FR4 (18.3%). So whilst an increase in dielectric constant allows a reduction in element size, it also reduces the bandwidth of the antenna [144].

The next consideration was the radiation pattern of the CPW antenna utilising the FR4 substrate. Since Nithisopa provided no data to which a reference could be made for Duroid, it was decided to focus solely on FR4. Figure 6.11 shows the simulated radiation pattern results obtained from HFSS. There are prominent peaks in power in the positive and negative *z* directions, which is clear from the plots in the *xz* and *yz* planes. Despite this HFSS reports the directivity of the antenna to be 4.65dBi; 2.65dB less than that of the previous patch antenna. This is most likely as a result of the fact that the CPW fed slot antenna can radiate in both the positive and negative *z* directions (i.e. perpendicular to the PCB) due to the absence of a ground plane on the underside of the PCB substrate.

In the *xz* and *yz* planes at 90° and 270° the radiated power is reduced, which means that there is little radiation in the plane of the PCB. These positions in the *xy* and *yz* planes represent the waveport and the edge of the PCB directly opposite the waveport. At these points there is little radiation because there is no radiating element along the line which runs perpendicular to *pcbX*. It is notable that the radiated power in the *xy*

plane is typically 10-15dB less than that in the *xz* and *yz* planes. A likely explanation for this is that the electric field lines formed between opposing charges converge here, but do not join, which causes a significant reduction in radiated power in the plane of the PCB. The fact that the *xy* plane pattern is asymmetric is probably because of the radiating slots not being central in the *y* direction.

At this stage there was much enthusiasm about this type of antenna. It provided a radiation pattern which exhibited less directivity than that of the original patch antenna, as well as providing a significant amount of bandwidth. In addition, its size was smaller than the patch so, although still too large to meet the design requirements, it was at least a step in the right direction.

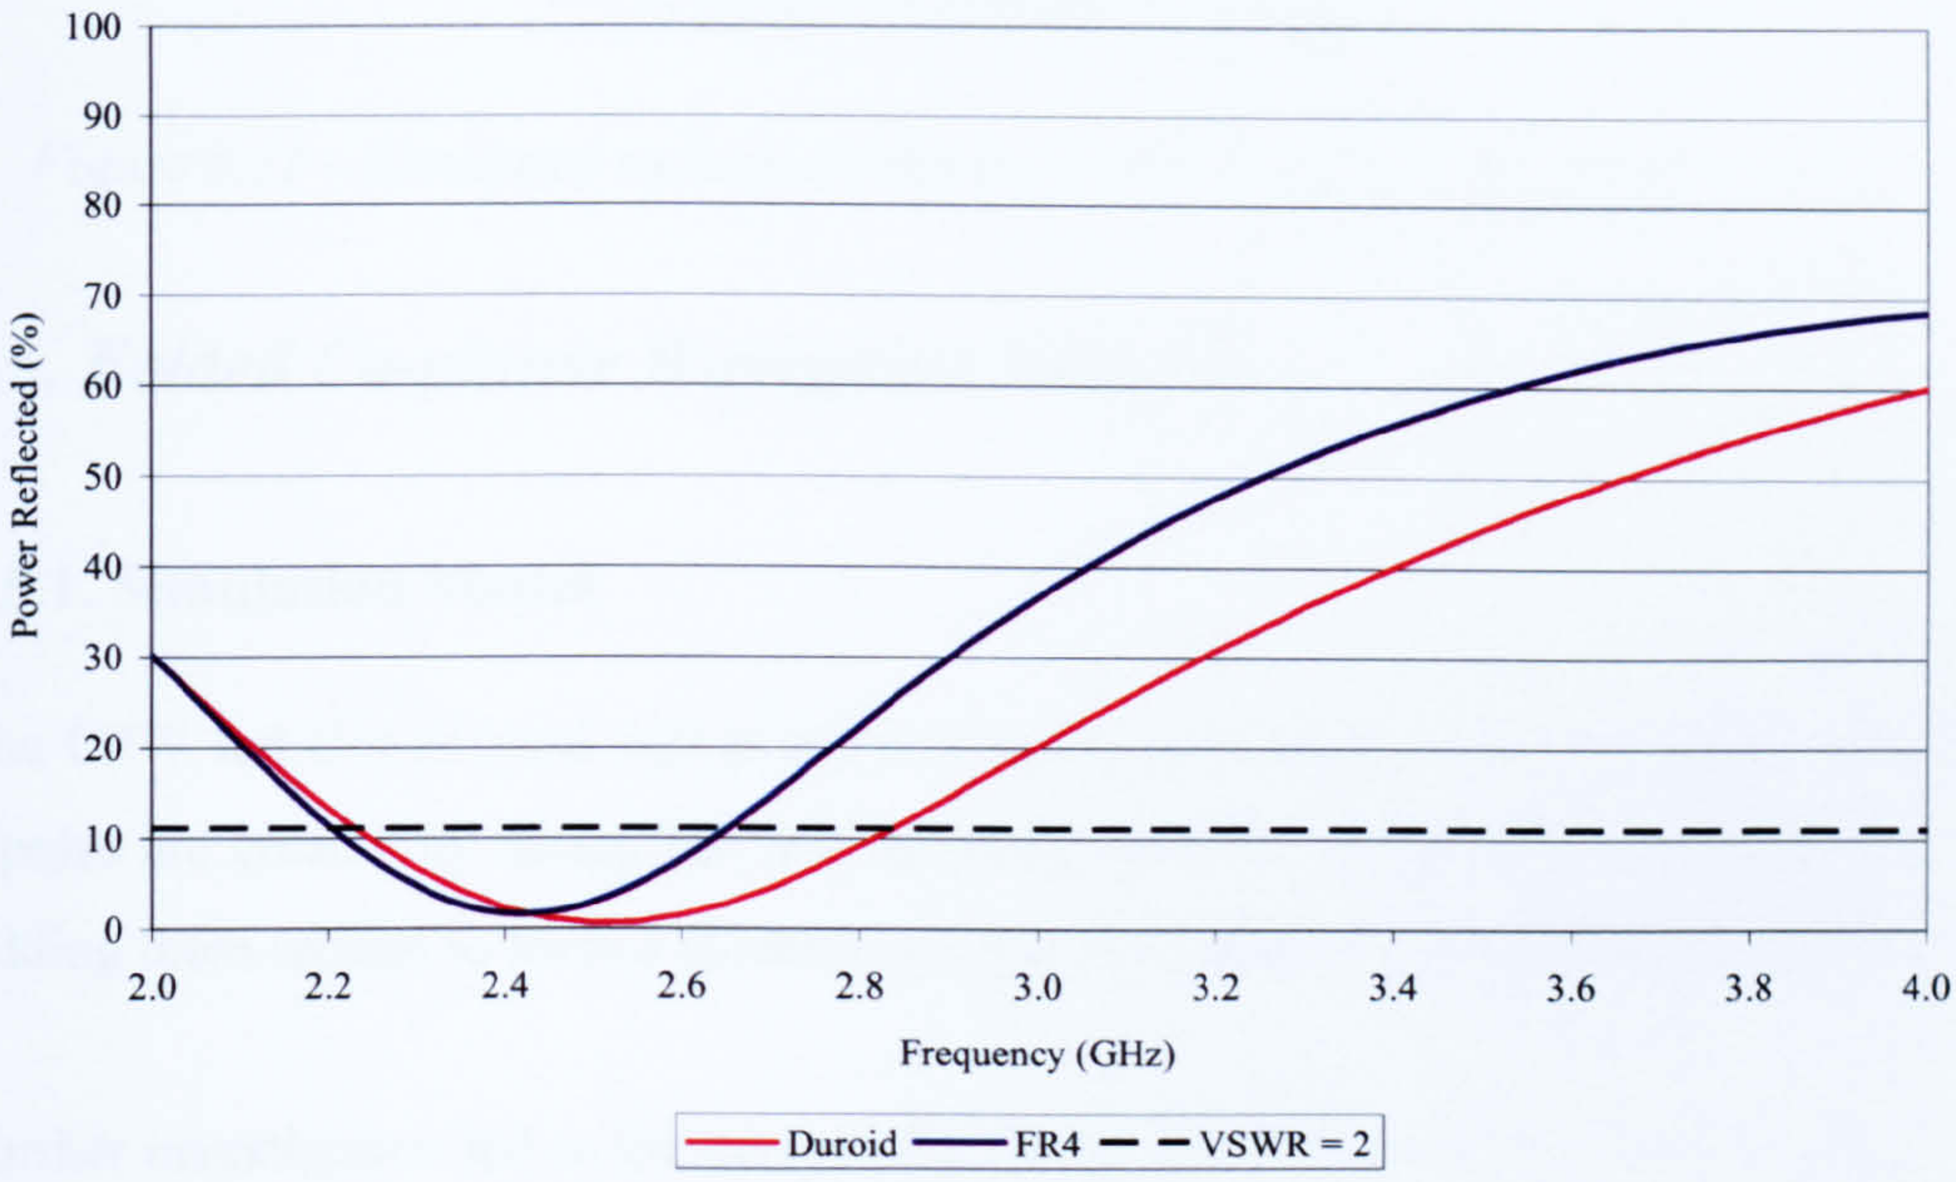

Figure 6.10 – Reflected power as a function of frequency on FR4 and Duroid substrates for the CPW antenna

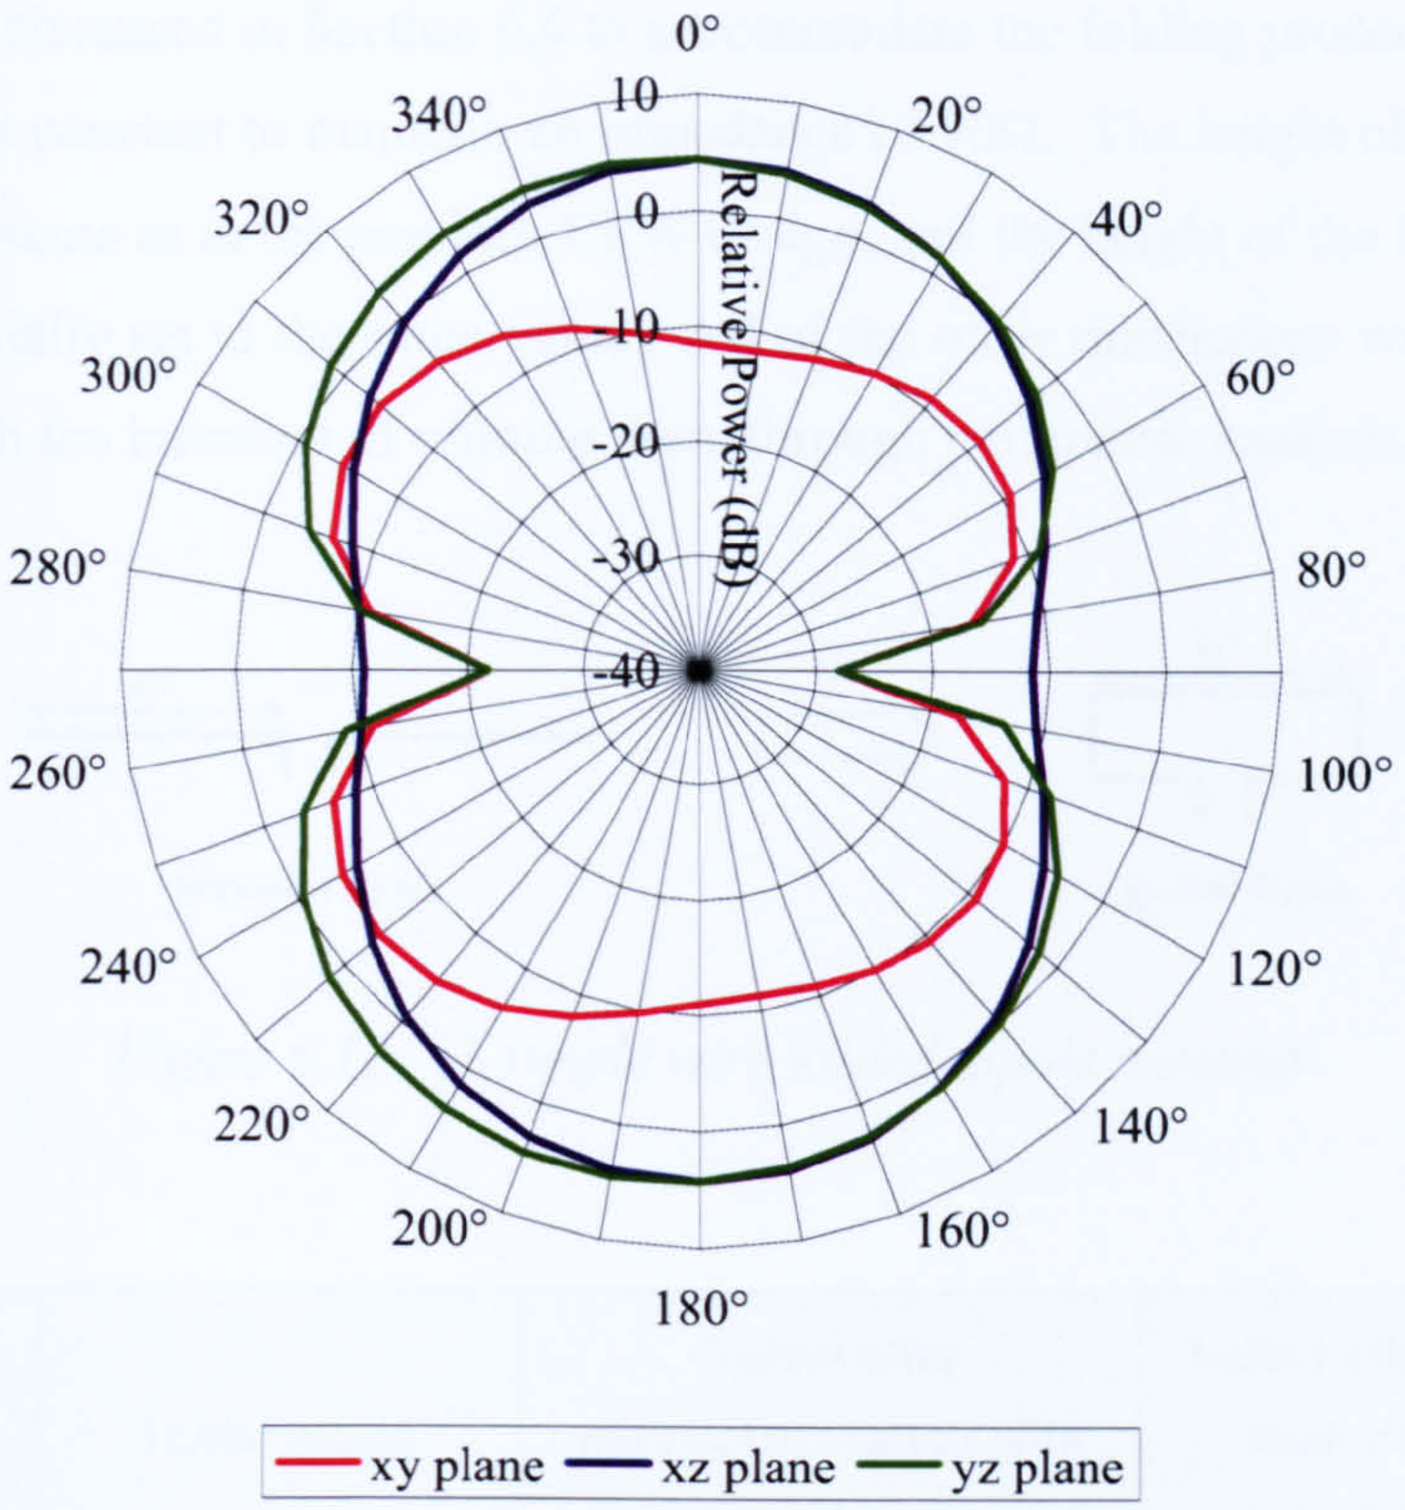

Figure 6.11 – Simulated radiation pattern for the co-planar waveguide antenna

6.5. Folded Co-planar Waveguide Antenna

6.5.1. Simulation Model

The CPW fed slot antenna was promising but improvements were required. Folded dipoles are created by taking the two radiating elements of a standard $\lambda/2$ dipole and folding them around to form a closed loop; this transition is shown in Figure 6.18.

Further investigation led to the design of a folded CPW (FCPW) antenna. In order to test this design, a HFSS model was created which included three additional parameters (*H3*, *H4* and *W3*). This model is shown in Figure 6.13 and Table 6.3 (column two) shows the initial values used in this model.

The initial size of the PCB (FR4 substrate) was chosen to match that of the mote PCB (i.e. 57mm × 32mm). *L1* was chosen to be approximately half the value used in the

CPW design discussed in Section 6.4 to accommodate the folding procedure. *W1* and *W2* were kept constant to maintain an impedance of 50Ω. The height of the slot (*H2*) was kept the same as in the previous CPW design, and the height of the folded section (*H4*) was initially set to the same value. All of the other dimensions were arbitrarily assigned, with the intention of refining them through parametric analysis.

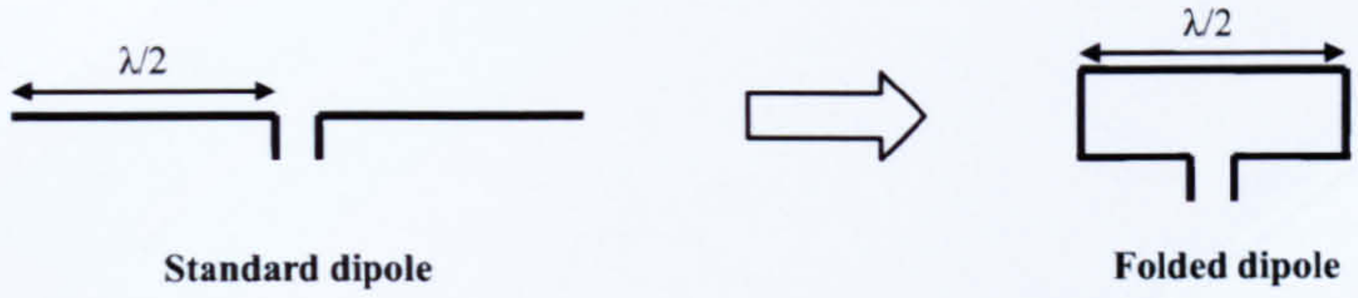

Figure 6.12 – A simple wire folded dipole antenna

Parameter	Initial values	Values after parametricanalysis with default mesh	Values after parametric analysis with high resolution mesh
pcbX	57mm	38mm	40mm
pcbY	32mm	28mm	27.5mm
h	1.6mm	1.6mm	1.6mm
W1	0.5mm	0.5mm	0.5mm
W2	4mm	4mm	4mm
W3	2mm	1.5mm	1.5mm
H1	2mm	2mm	2mm
H2	4mm	4mm	4mm
H3	1mm	0.5mm	0.5mm
H4	4mm	8.5mm	8.5mm
L1	18mm	15.5mm	14.3mm

Table 6.3 – Parameters used during modelling of the FCPW fed slot antenna

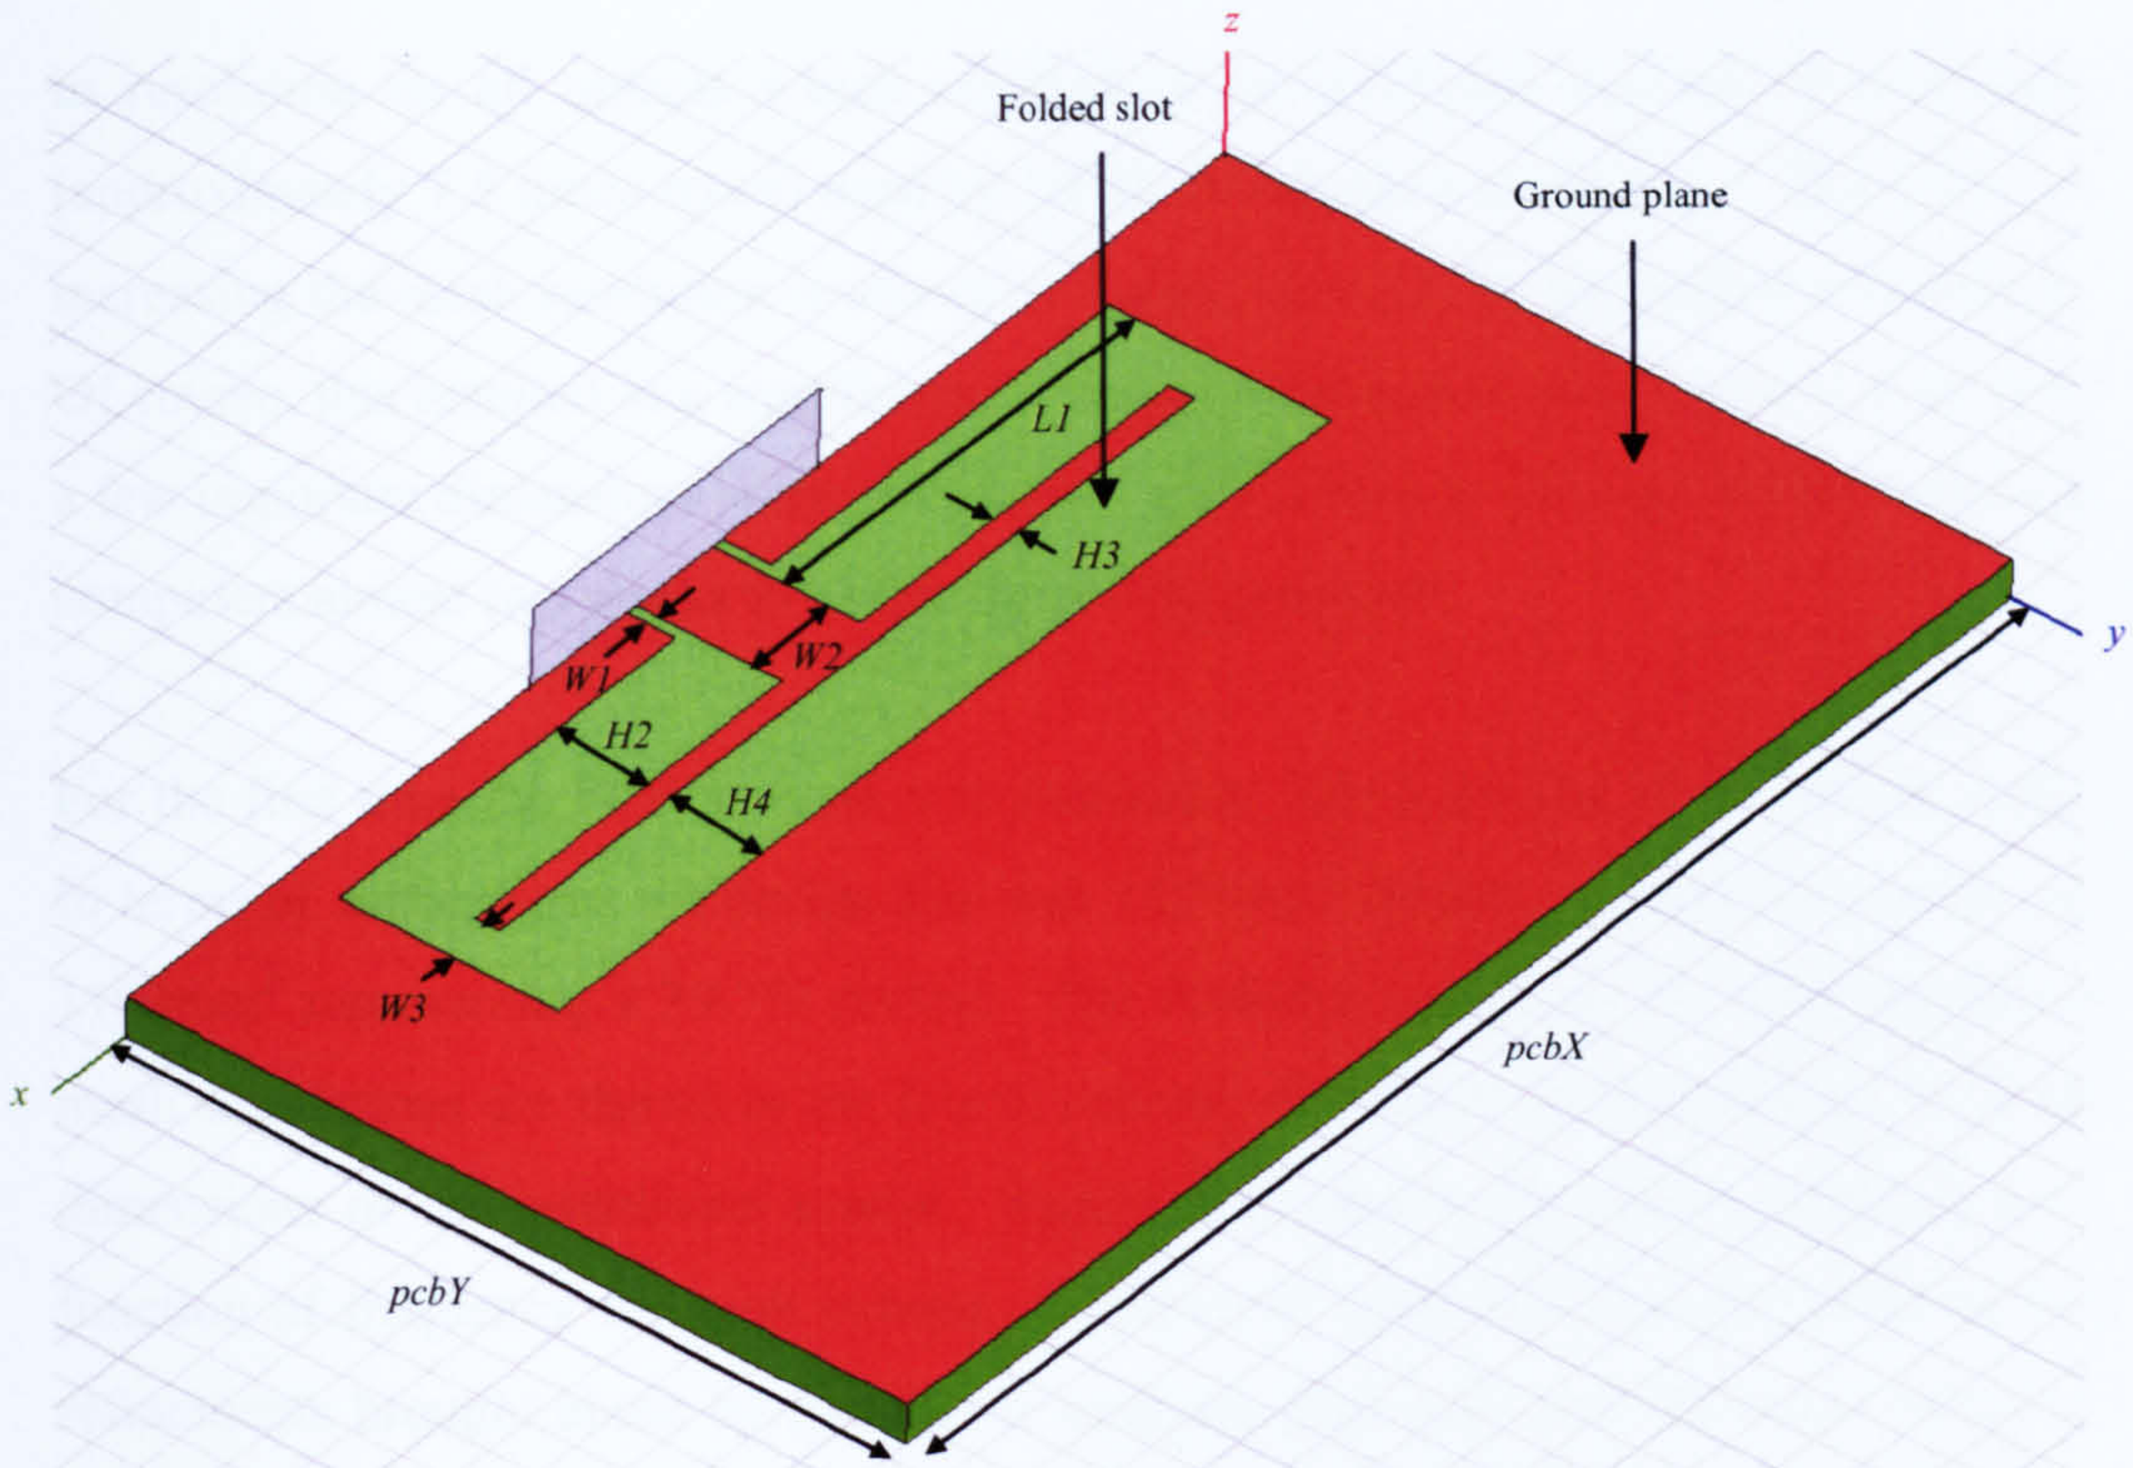

Figure 6.13 – HFSS model of the FCPW antenna

6.5.2. Refining the Model

At first all the dimensions were subject to parametric analysis and the optimal value for each dimension was determined to be the one which offered the lowest reflected power; these values are shown in the third column of Table 6.3. These values were obtained using the default mesh resolution assigned by HFSS since a large number of simulations had to be carried out – therefore time was an important factor. Each dimension was altered in increments of 0.5mm over an appropriate range. In the case of *L1*, a number of possible lengths resulted in very little reflection, however, the smallest of these lengths was utilised to minimise the size of *pcbX*, and therefore the antenna.

Following these first set of parametric simulations, a second set were run. These focused on the *L1*, *pcbX* and *pcbY* dimensions and looked to optimise them to the nearest 0.1mm. For increased accuracy in these simulations, the mesh was increased

in resolution on all objects which resulted in an increased number of tetrahedra (approximately 7.5 times more) used by HFSS to calculate a solution. The effect of increasing the mesh resolution is shown in Figure 6.14 for the air radiation boundary. Of course, this resulted in a significant increase in simulation time which is why only a few important dimensions were considered. This still took four reasonably specified computers almost two weeks to obtain the results presented.

For the first time the PCB antenna was reduced to less than that of the MicaZ mote. In terms of surface area the design aim was 1824mm^2, whilst the achieved result was 1100mm^2 representing a 39.7% saving. The complete results of the high resolution mesh simulations are shown in the fourth column of Table 6.3 and represent the final dimensions of the simulation model. Figure 6.15 shows the power reflected as a function of frequency for each column of values in Table 6.3, and the effect of the refinements brought about by parametric analysis. One may note that even with the initial parameters the amount of reflected power near the target frequency is very small. However, with the later refinements it is possible to see that the FCPW structure has been optimised for minimal reflection at 2.45GHz.

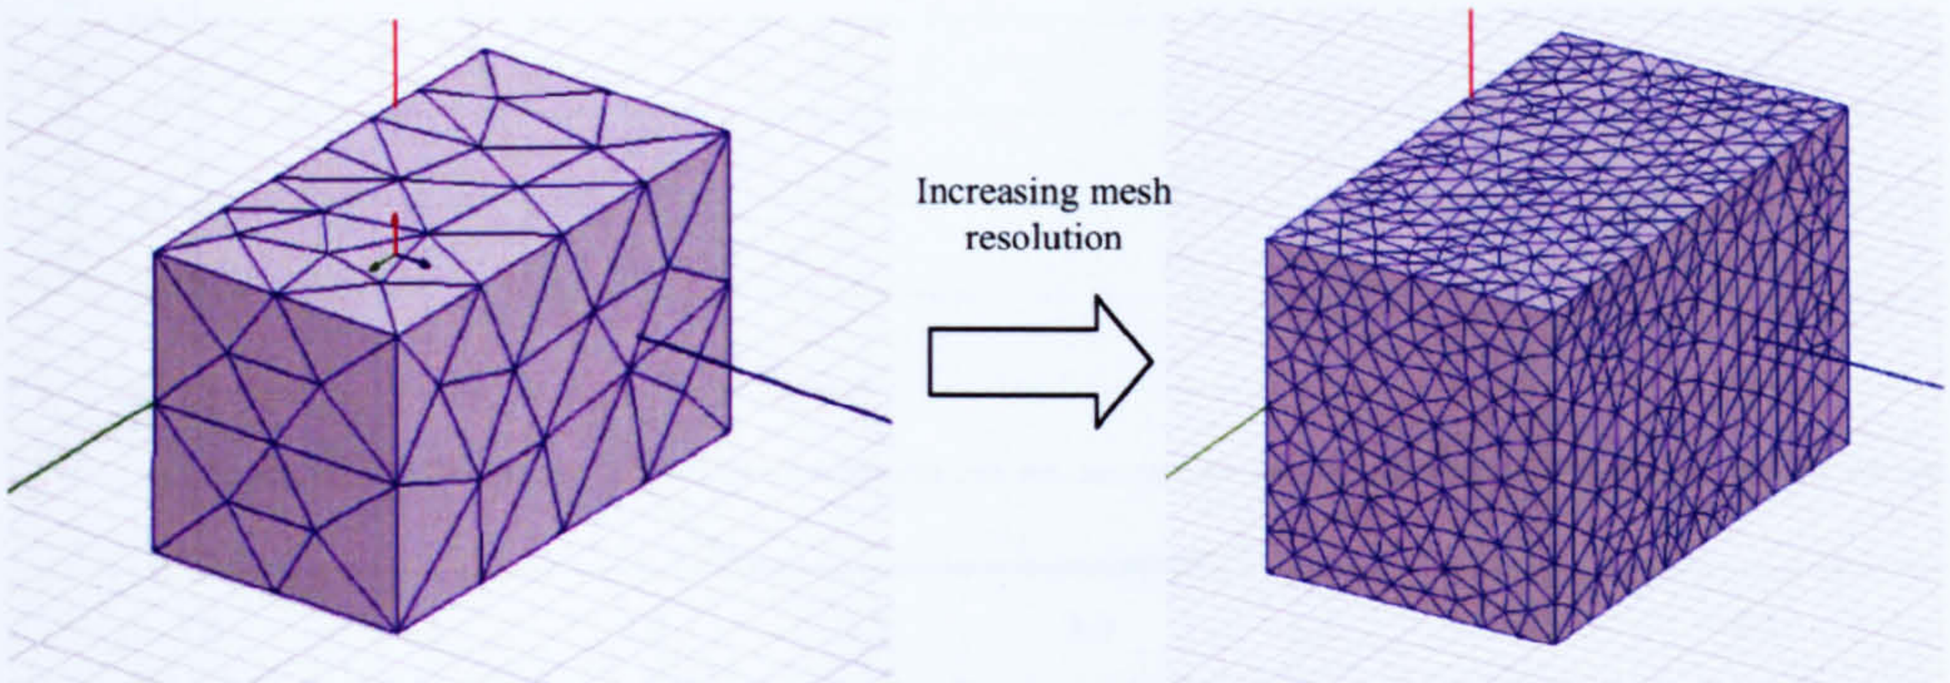

Figure 6.14 – Changing the mesh resolution

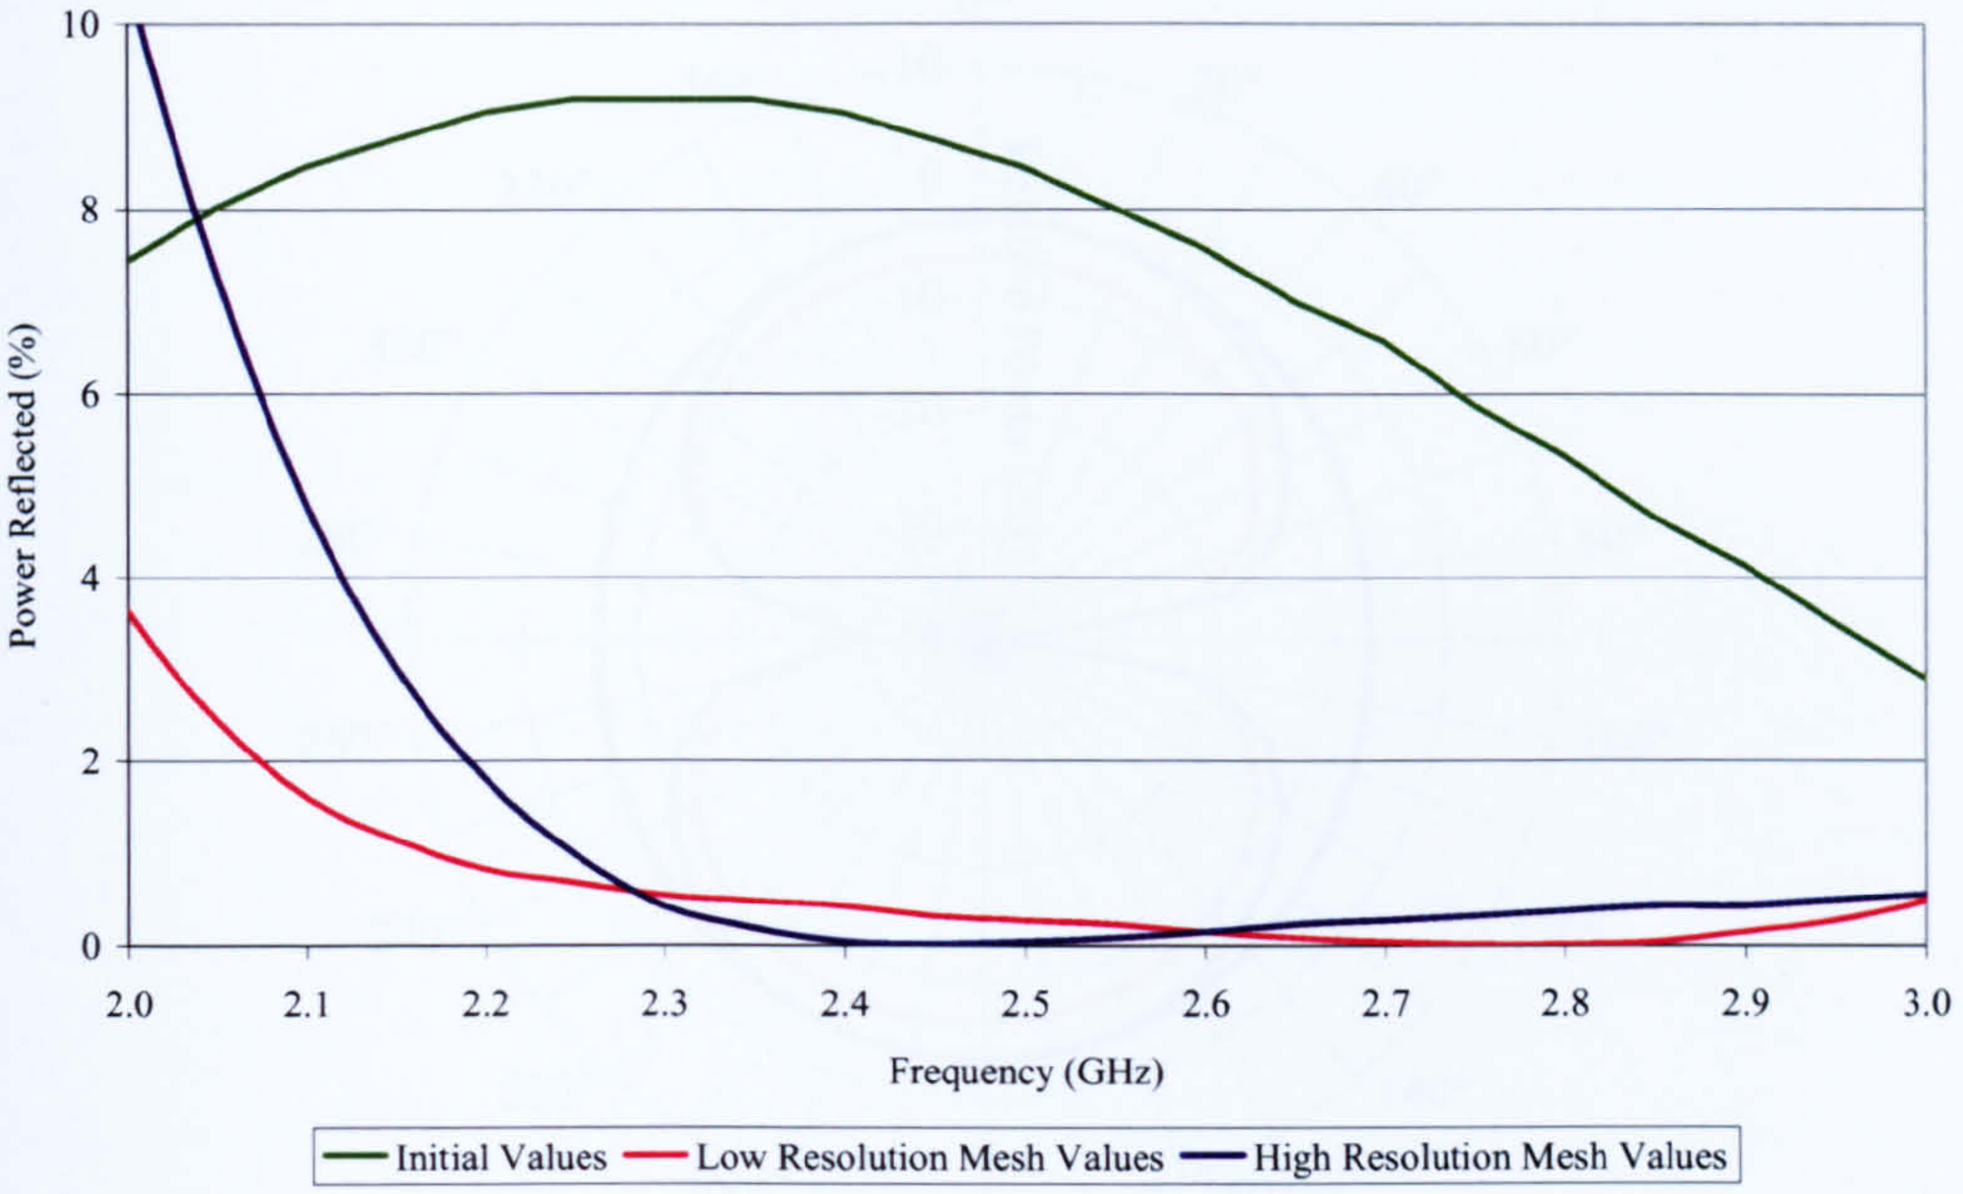

Figure 6.15 – Power reflected as a function of frequency for each set of values in Table 6.3

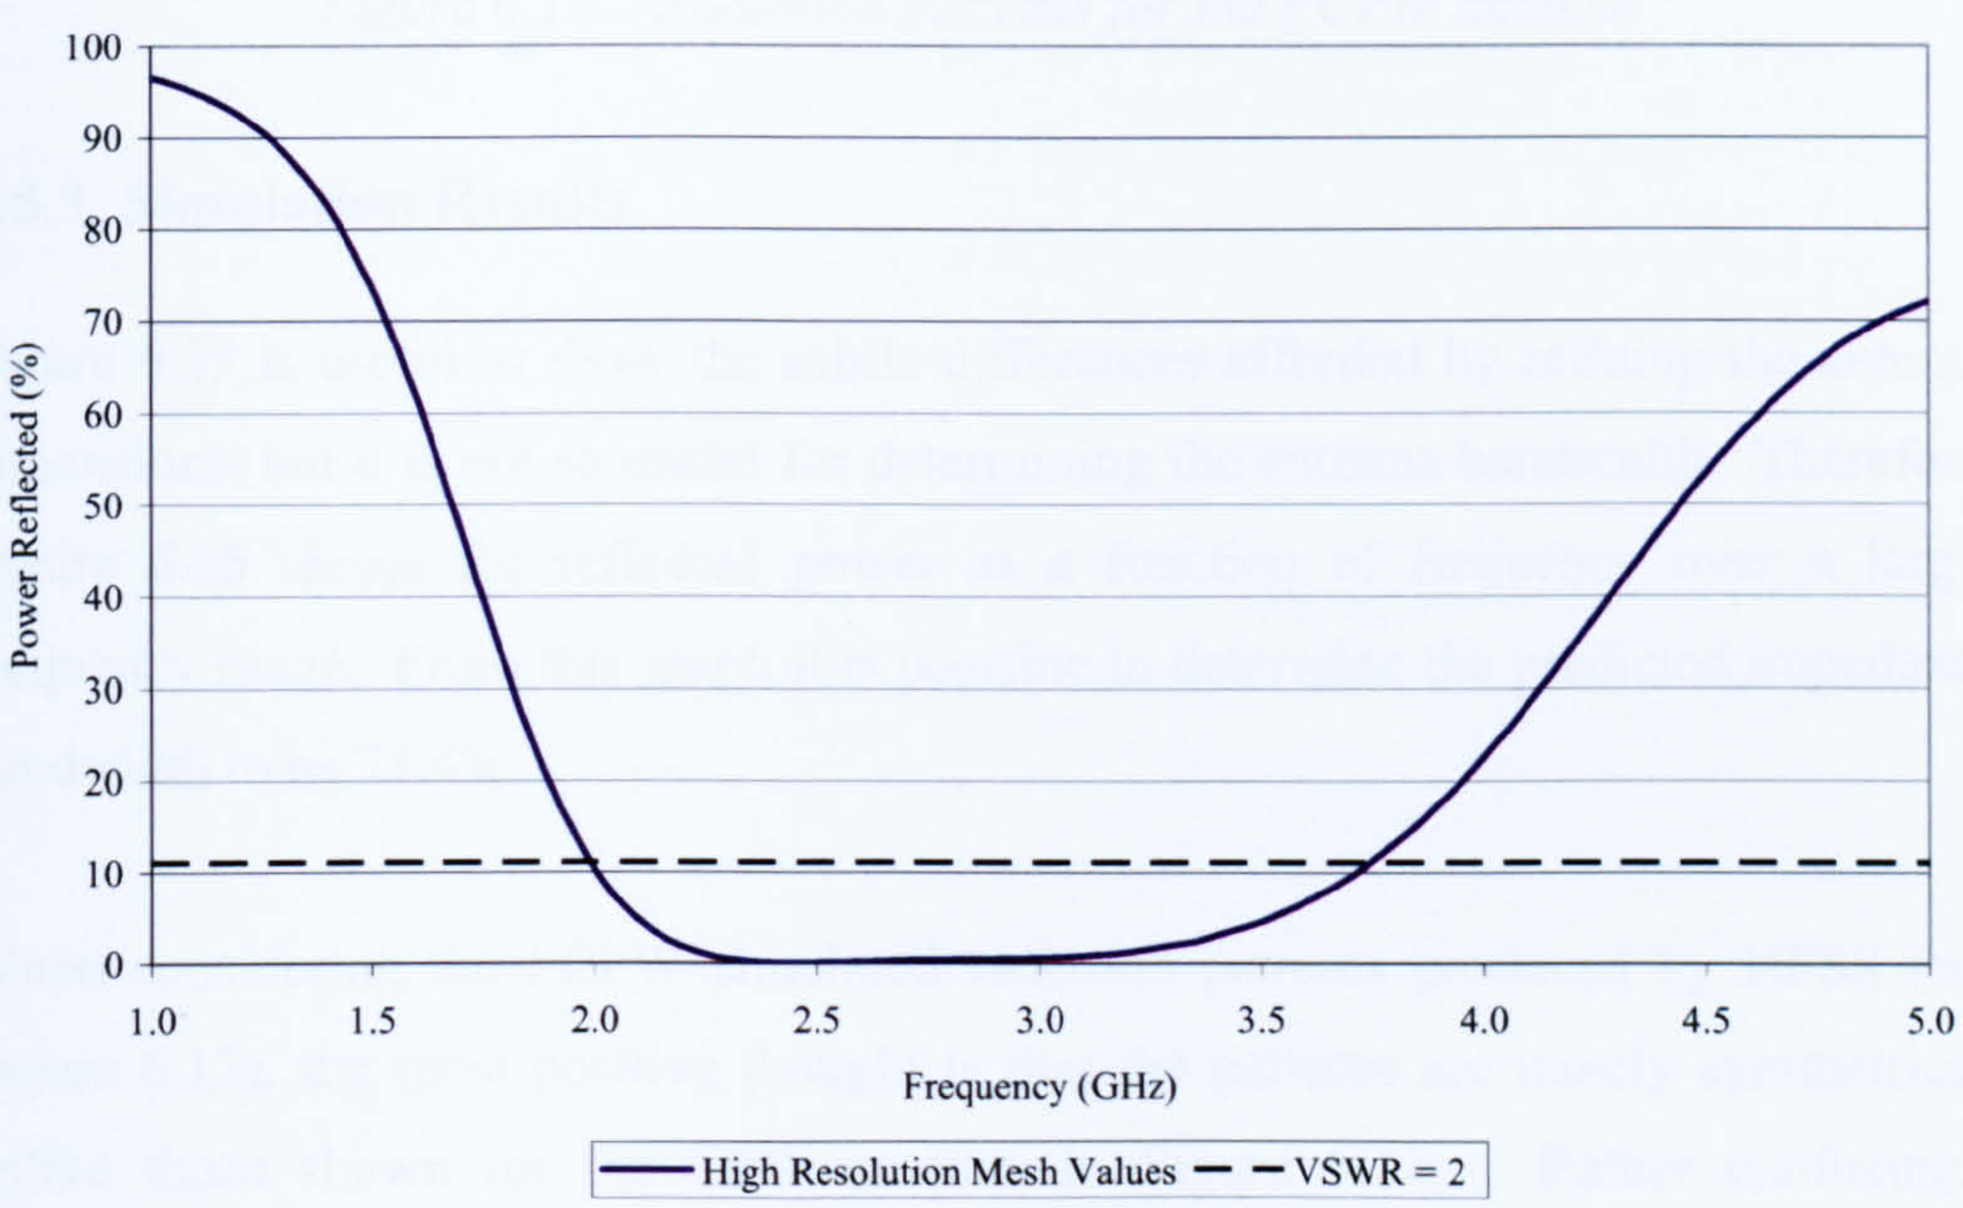

Figure 6.16 – Power reflected as a function of frequency for the final design parameters of the FCPW antenna

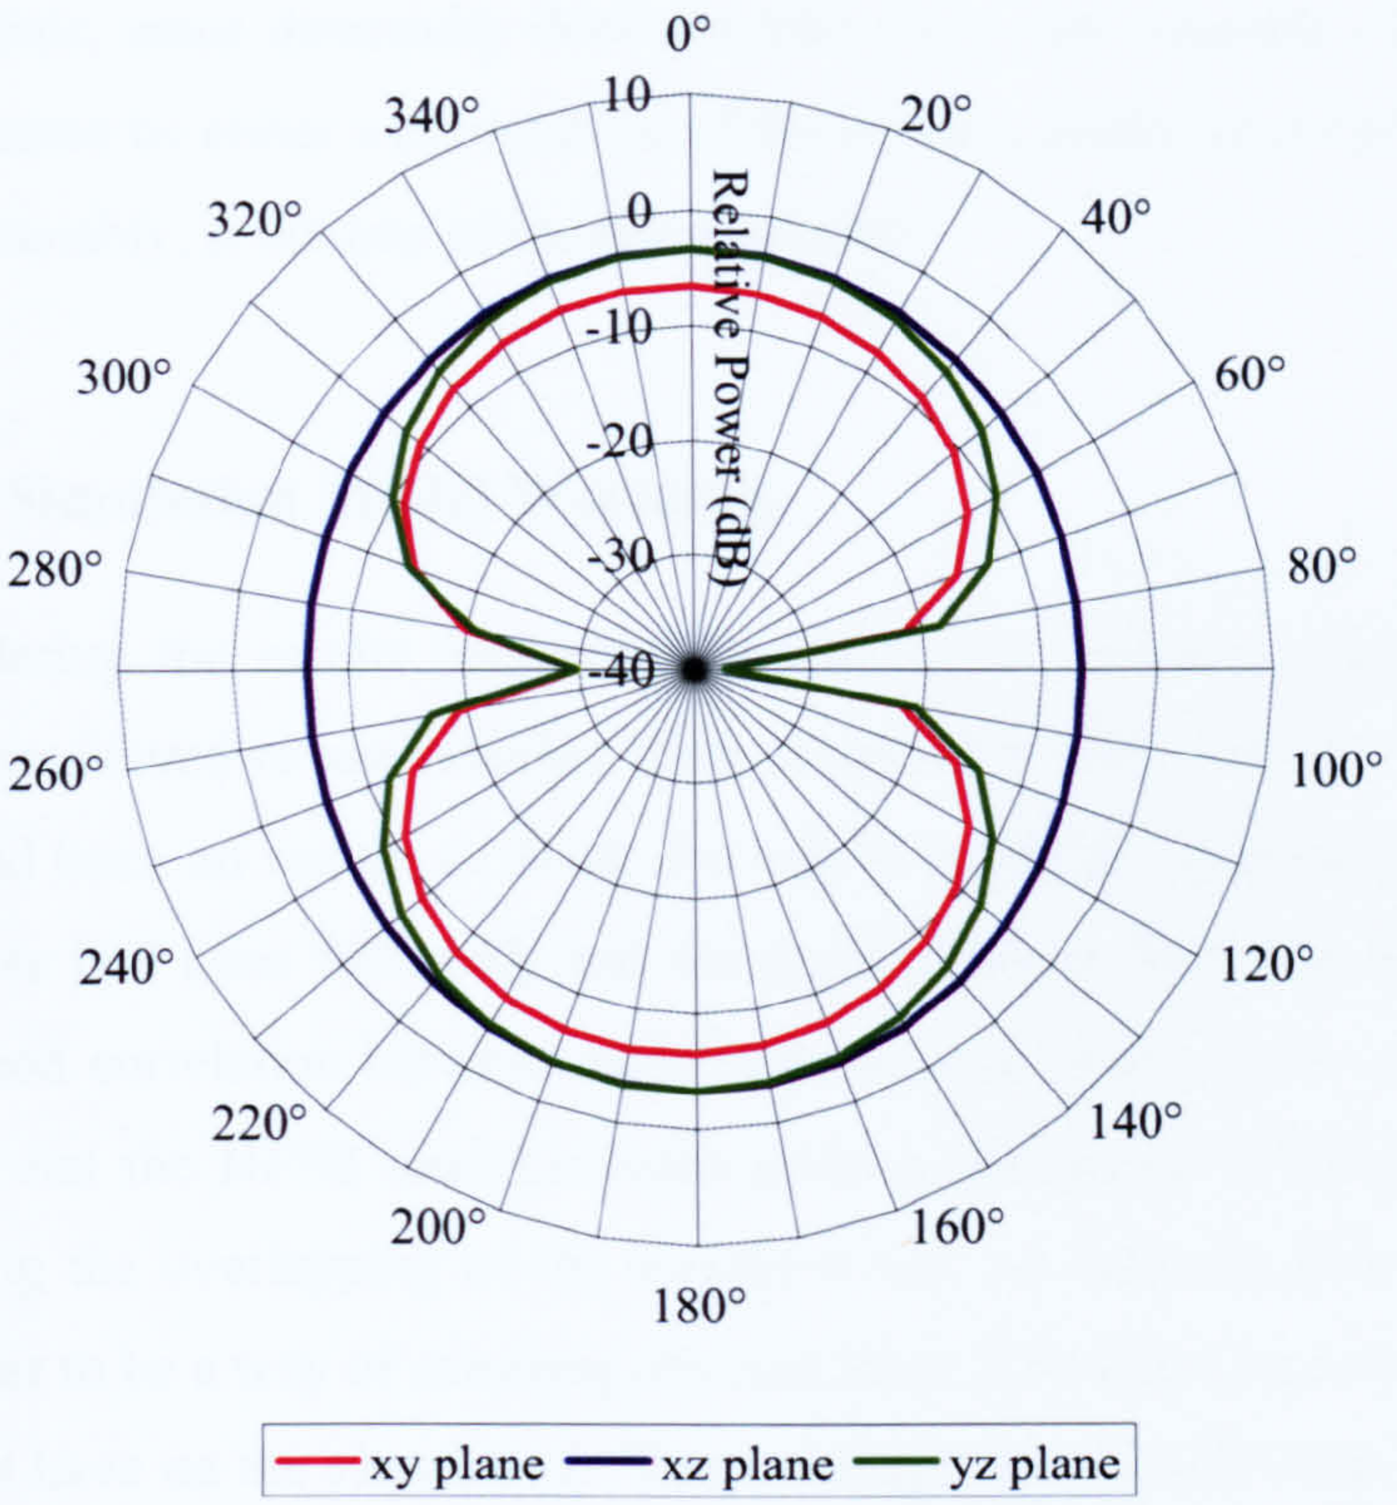

Figure 6.17 – Radiation patterns for the FCPW antenna

6.5.3. Simulation Results

Figure 6.15 is useful to show the subtle differences afforded by refining the antenna dimensions, but it is not so useful for determining the antenna bandwidth. Therefore, Figure 6.16 shows the reflected power as a function of frequency over a larger frequency range. From this graph it is possible to determine the predicted impedance bandwidth to be 71.4%.

When considering the FCPW simulated radiation patterns produced by HFSS (see Figure 6.17), the most positive thought is that the patterns are mostly symmetrical, unlike those shown for the CPW antenna in Figure 6.17. Rather confusingly however, HFSS calculates the directivity of this antenna to be -3.31dBi; even though an antenna cannot have a directivity less than 0dBi [159]. One possible assumption is that the antenna is not particularly effective at radiating energy. Considering this observation, it is possible to see from Figure 6.16 that none of the input power to the antenna is reflected. It is also known that the loss in power cannot be due to the

antenna materials, since directivity does not take these into account. Therefore the loss in power must be either a consquence of the antenna geometry, or an error within HFSS which, notably, is not one of the latest releases.

6.5.4. HFSS Simulation Model Warnings

Whilst considering the results presented for the FCPW antenna, the HFSS model created was considered to see whether there could be a fault with it. Up until this point there had been no reason to doubt the results of HFSS; guidelines for creating CPW structures had been followed, and the patch antenna discussed in Section 6.2 had shown good correlation between simulated and measured results. However it is worth noting that the HFSS model created gives a simulation warning (see Figure 6.18) regarding the overlapping of the waveport and air radiation boundary. There does not appear to be a way of avoiding this and there is no clear indication as to what effect it might have on the simulation. The warning occured with both the CPW and FCPW models; it was initially ignored in the CPW case because the results appeared feasible. However, in the case of the FCPW the results were disappointing particularly since a large amount of time had been spent running simulations to refine the model. Therefore, it was decided that the antenna should be constructured and tested physically to see just how well it operated in the real world.

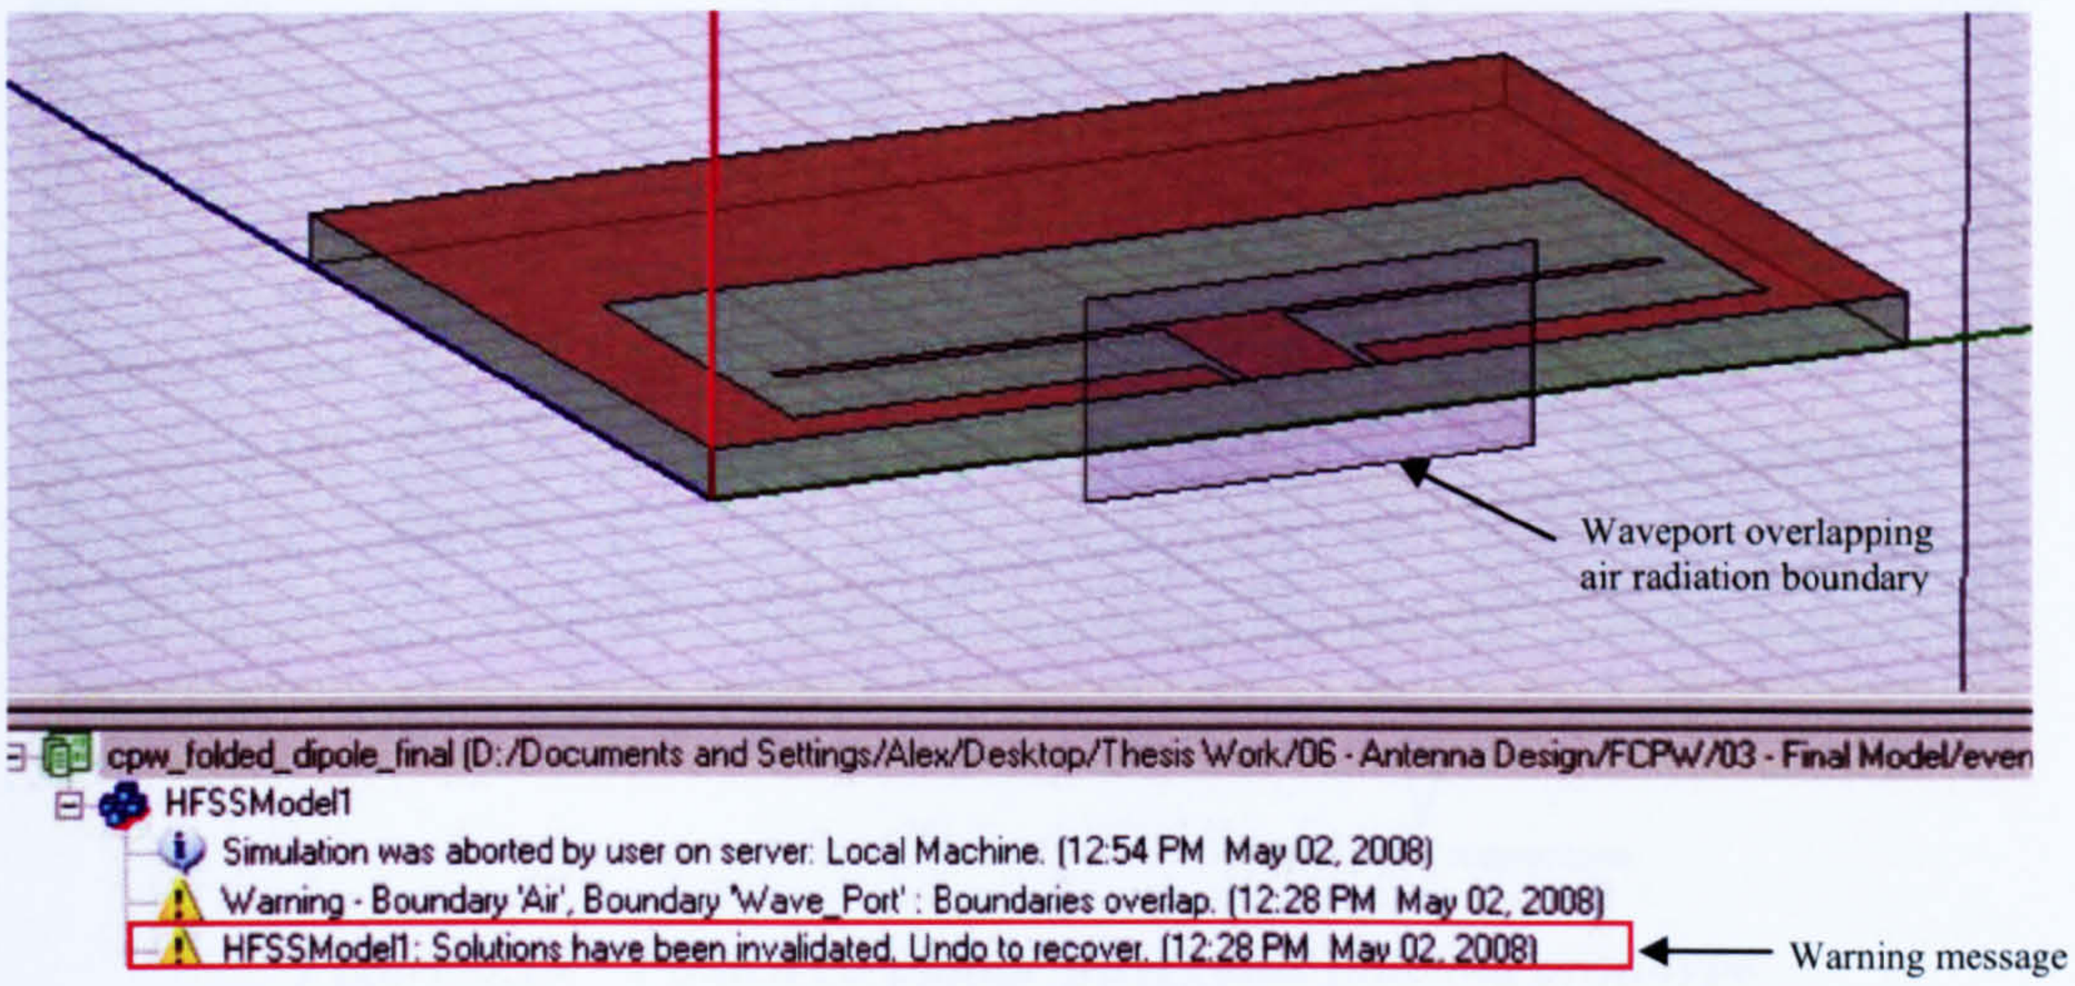

Figure 6.18 – HFSS boundary warning

6.6. Physical Construction

Using the Eagle PCB CAD software [115] a mask of the FCPW antenna was made so that photosensitive FR4 epoxy PCB could be exposed to ultraviolet light, developed and then etched. This was preferable, for the sake of accuracy, to the manual method used for creating the patch antenna previously discussed.

The finished prototype FCPW antenna is shown in Figure 6.19. It has a bulkhead type SMA connector attached so as to allow connection to various devices. This has been convenient for experimentation purposes, the results of which are presented in the next section. It is imagined that this connector would not be necessary if the antenna were used in practise – instead the antenna could be connected directly to an appropriate radio transceiver. The centre conductor of the SMA connector is soldered to the centre copper strip, whilst the outer conductor is connected to the antenna ground plane on either side of the centre strip.

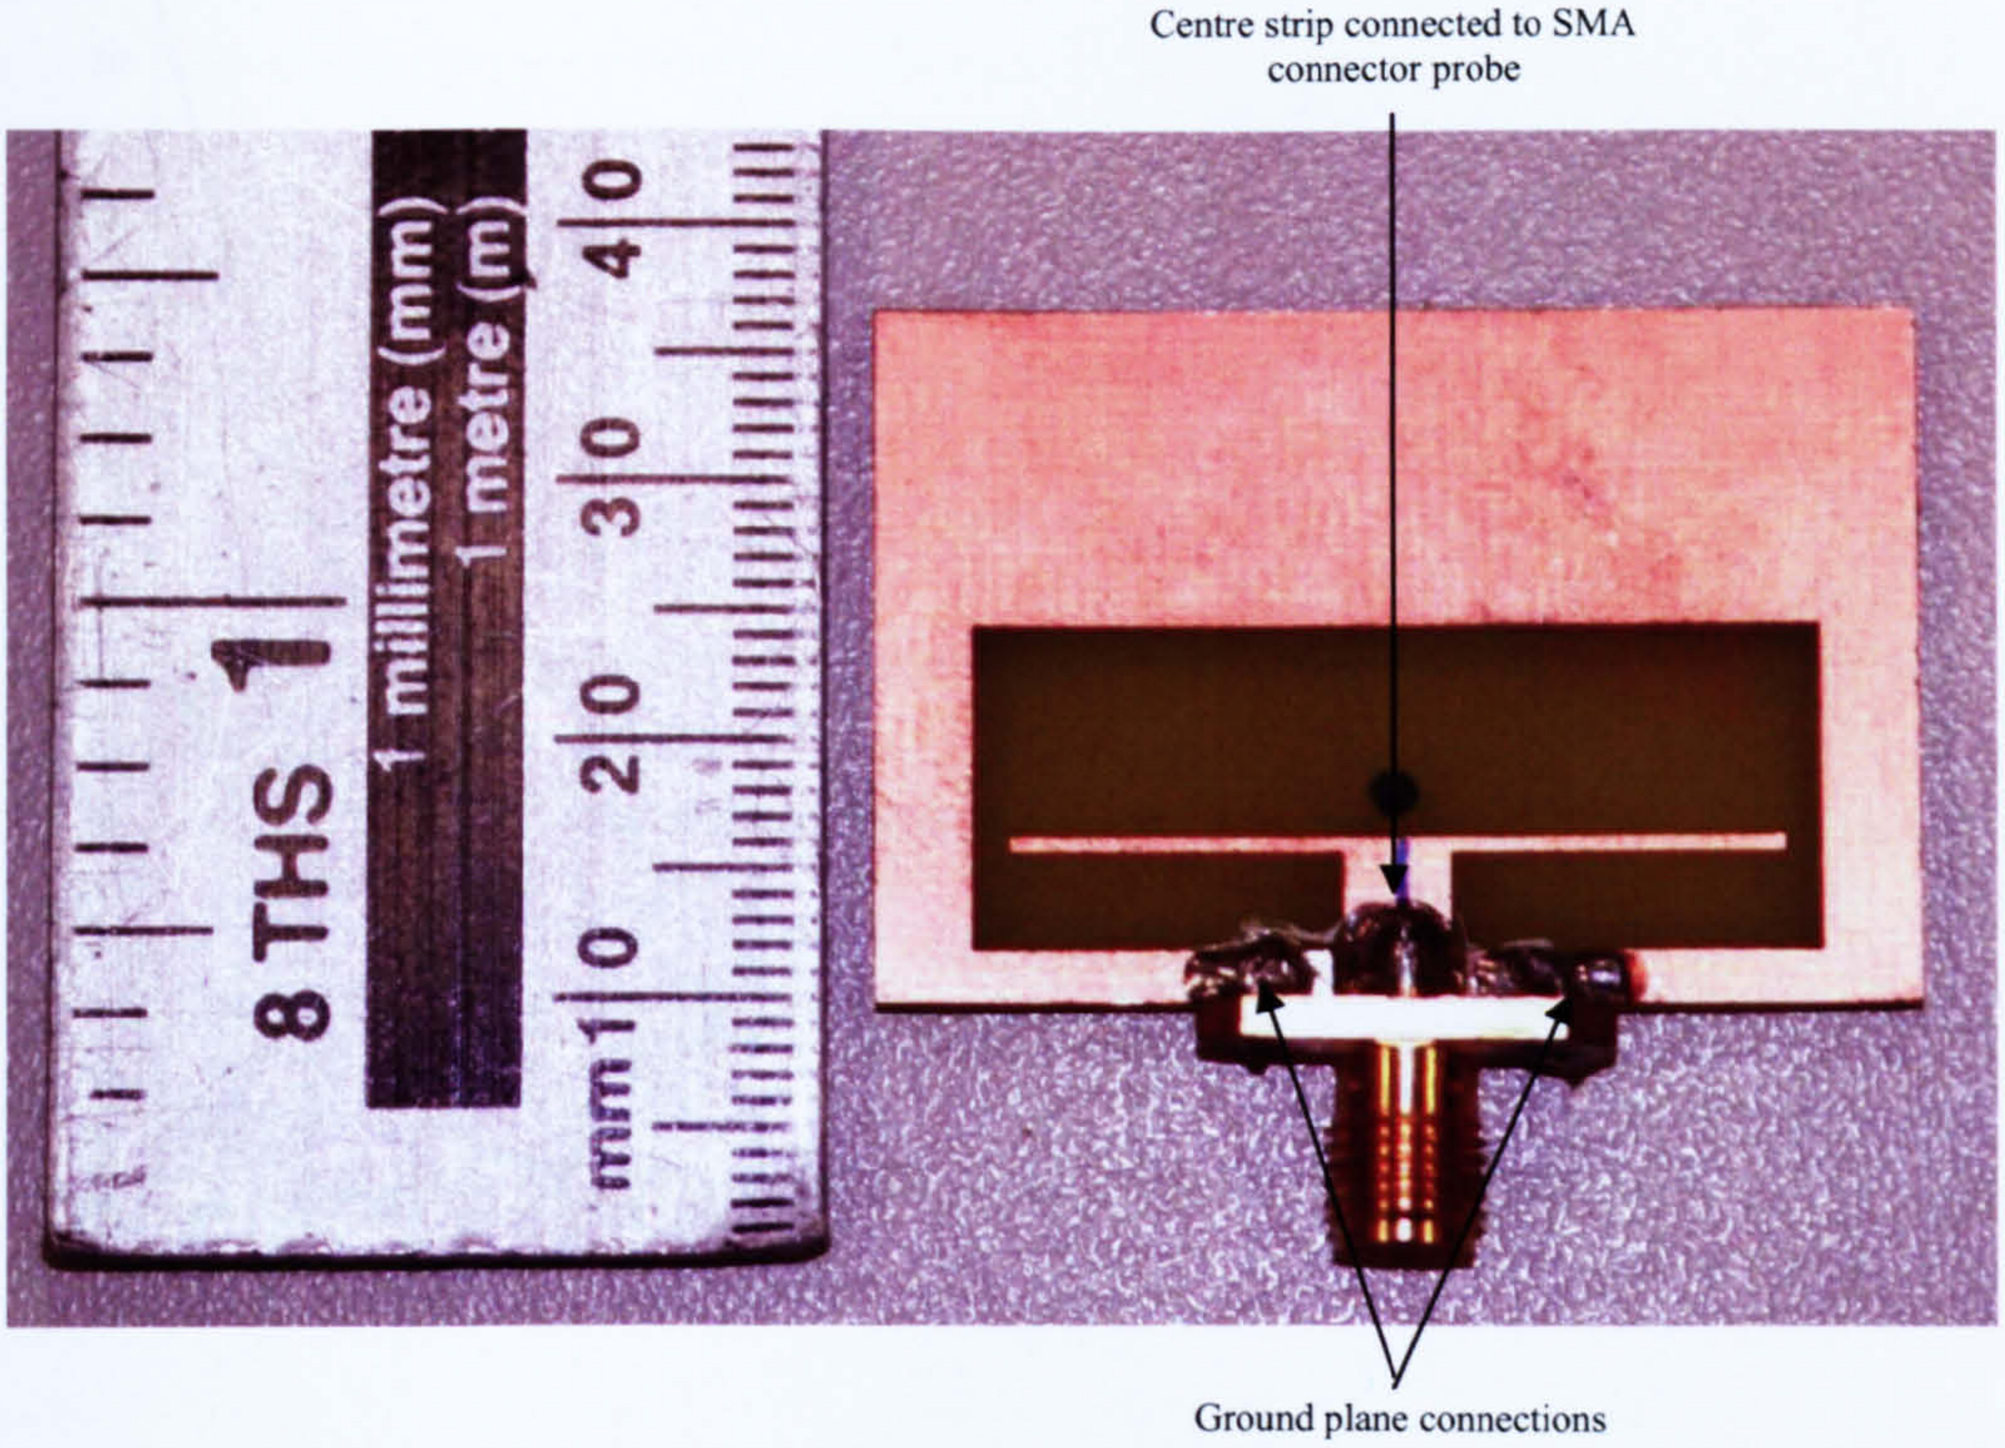

Figure 6.19 – Folded co-planar waveguide antenna prototype

The simulated results for power reflection and those measured using the Anritsu VNA show a reasonable agreement (see Figure 6.20), although the measured results suggest that slightly less power is reflected at higher frequencies than expected (e.g. 3.2GHz). At 2.45GHz the reflected power was just 2.6% - this means that 97.4% of the power incident to the antenna should be radiated. The antenna bandwidth in the simulation was found to be 71.4%, whilst the measured bandwidth was 65.3%. This is more than adequate for the successful operation of the MicaZ mote over all of its selectable frequencies. These results restored some confidence in HFSS at least with regard to the matching of the CPW transmission lines, but further experimentation was required to better determine the real world performance of the FCPW antenna.

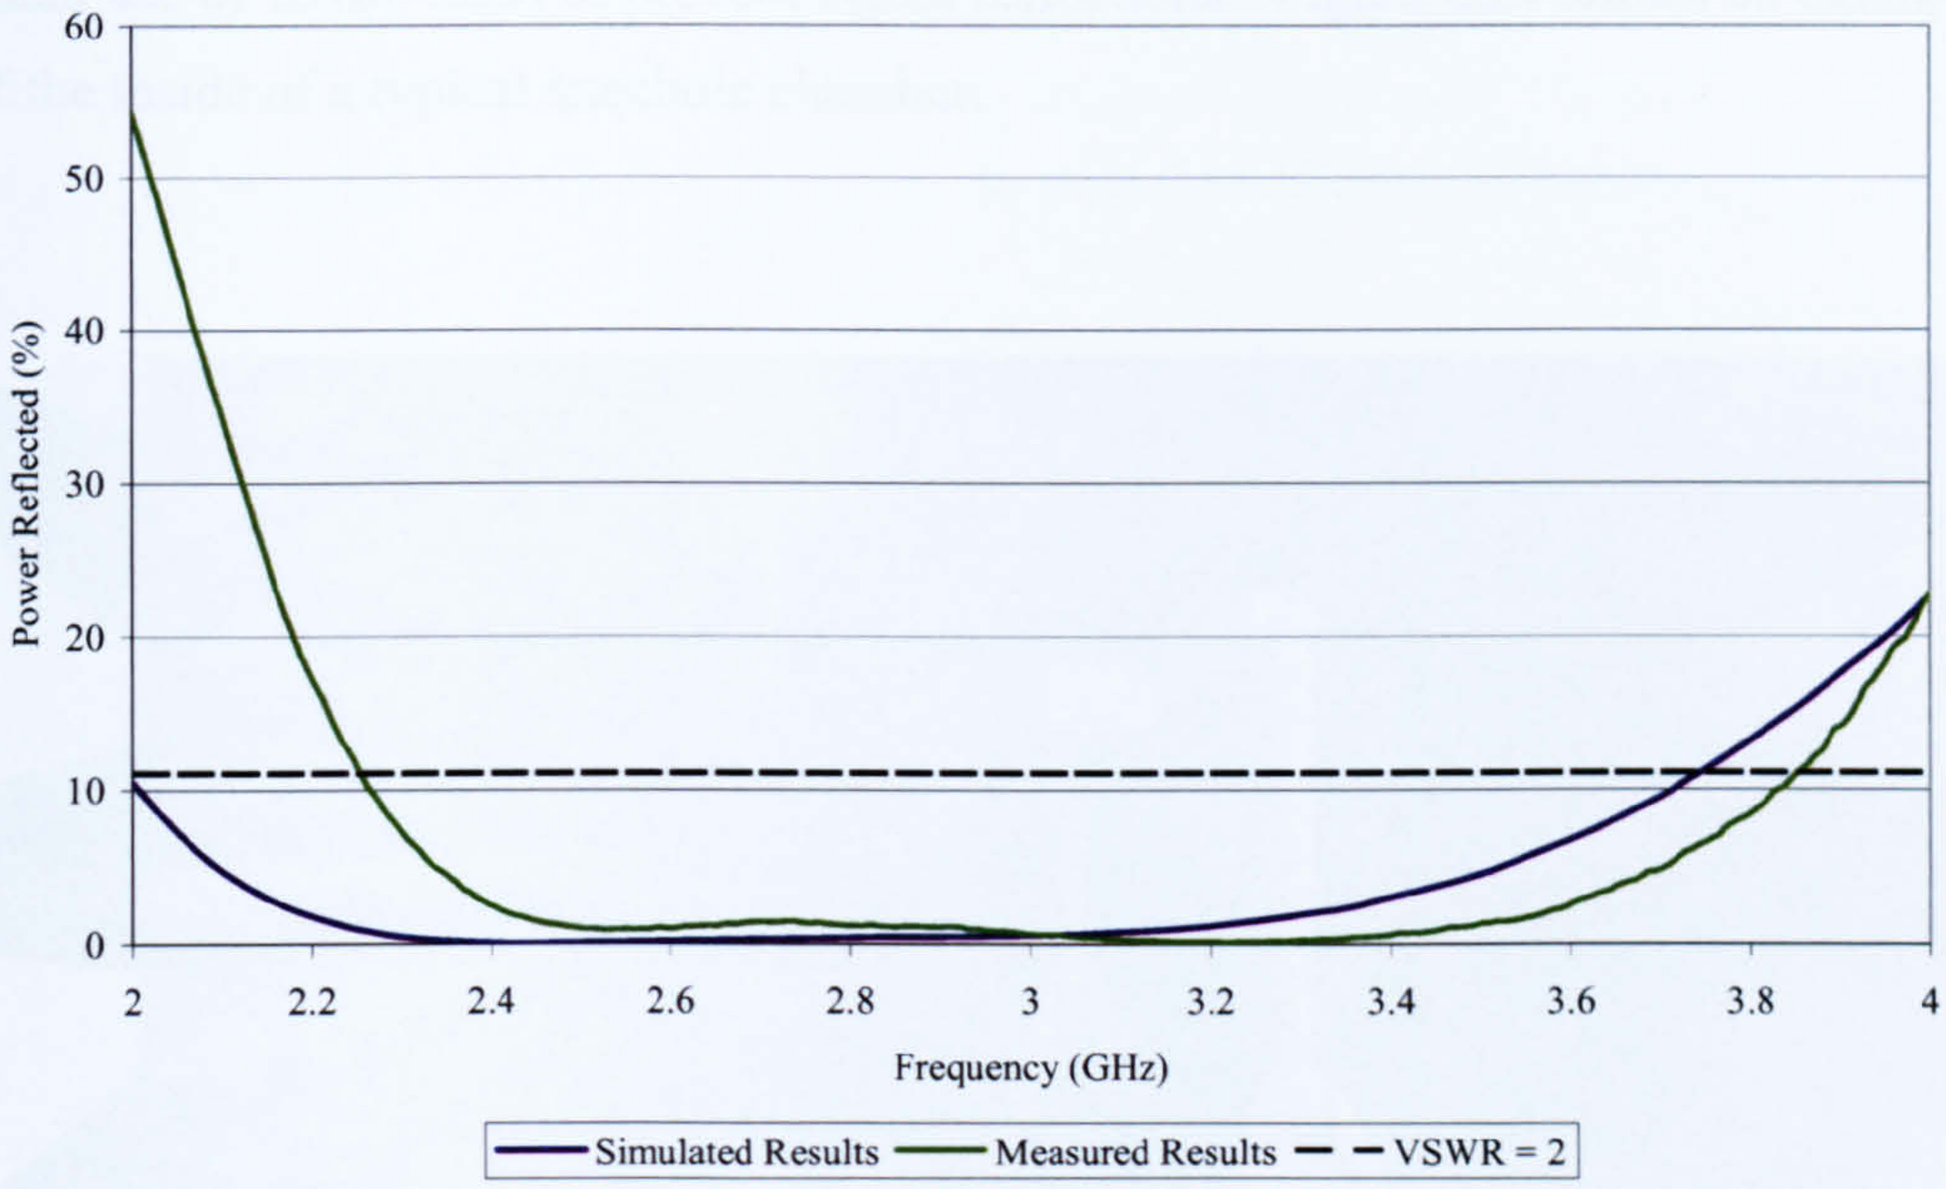

6.20 – FCPW measured vs. simulated power reflection

6.7. Radiation Pattern Measurement

6.7.1. Experimental Setup

In order to measure the radiation pattern of the FCPW antenna, a number of steps had to be taken to ensure that measurements were not affected by interference from other devices. Since the intended operating frequency was 2.45GHz, the main concern was avoiding interference from existing WLAN infrastructures.

Small antennas are typically tested in almost ideal conditions inside a structure known as an anechoic chamber [160]. Such chambers are usually shielded with a thick conducting metal, thus utilising the skin effect to prevent signals from the outside world entering the chamber. Inside the chamber, the walls are lined with absorbing material (typically pyramidal in shape, although lower frequency applications can make use of ferrite tiles) to prevent signal reflections. Figure 6.21 shows an example of the inside of a typical anechoic chamber.

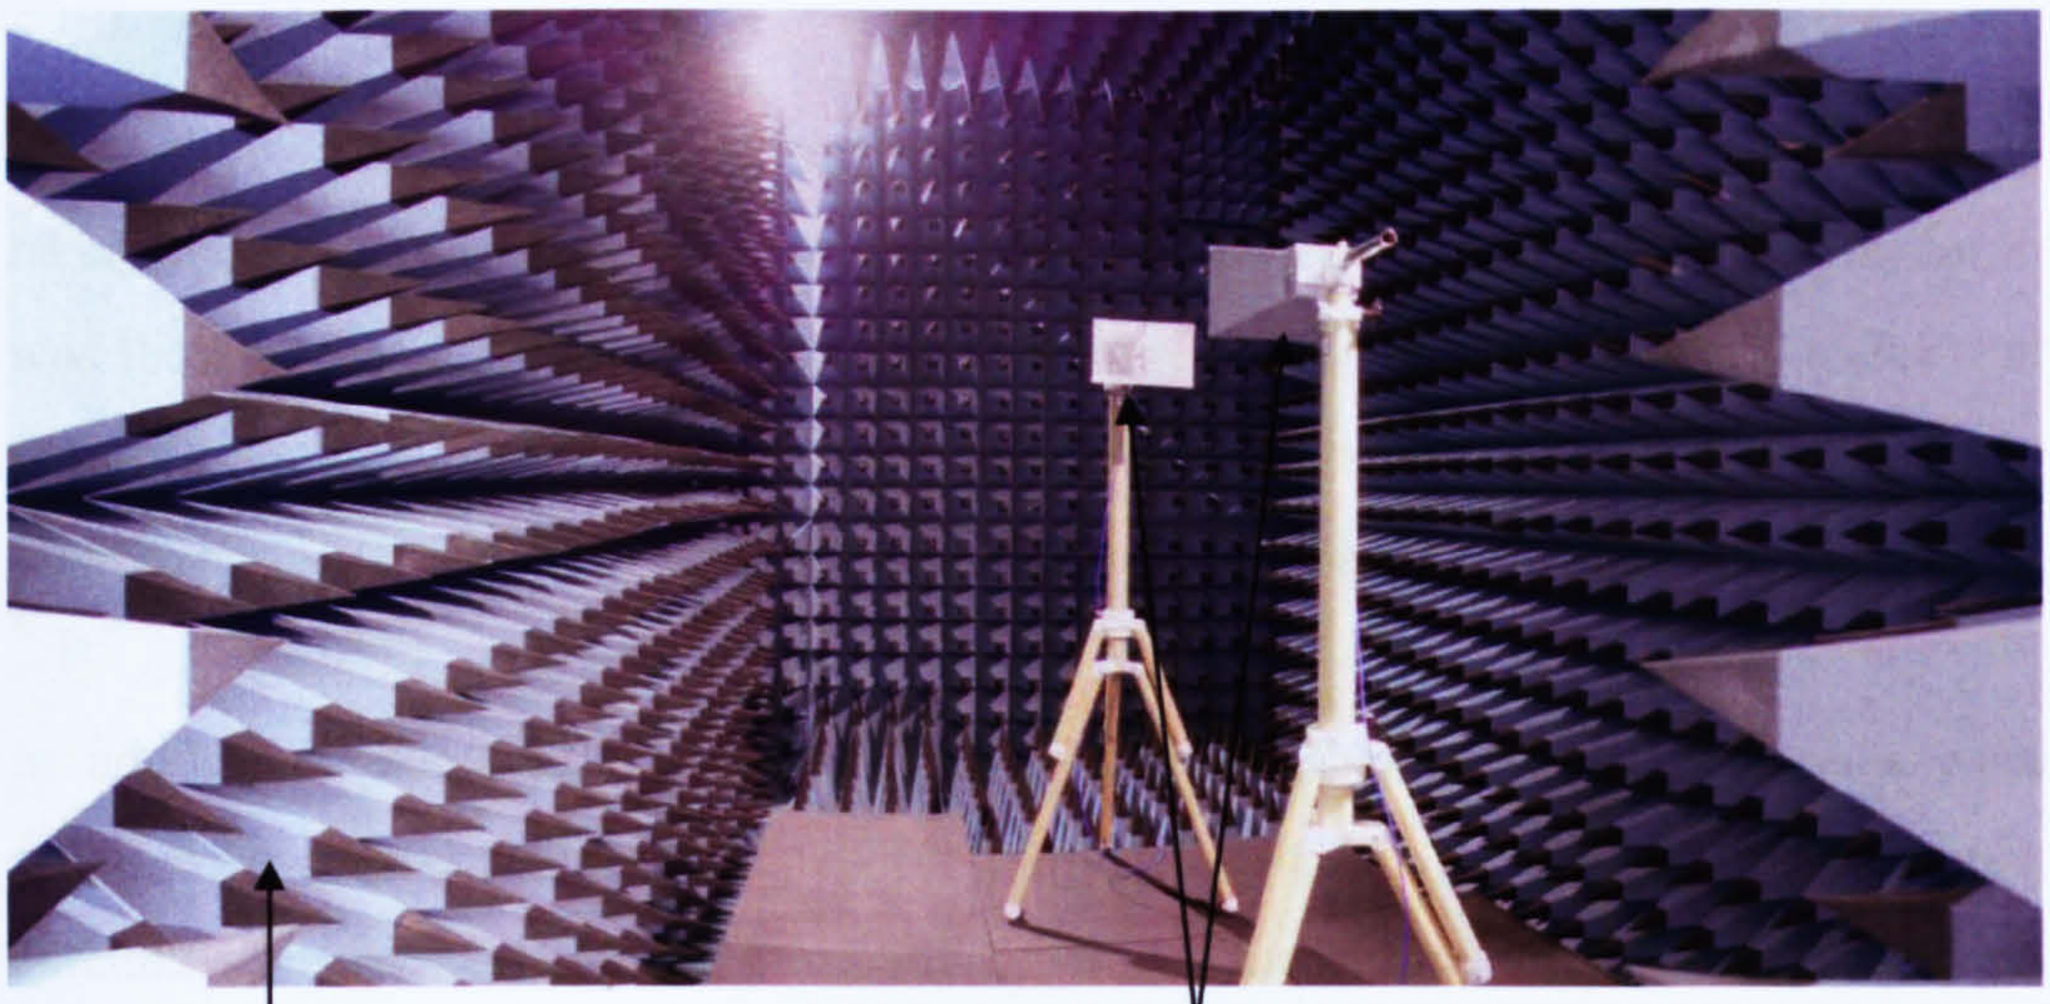

Figure 6.21 – A typical anechoic chamber

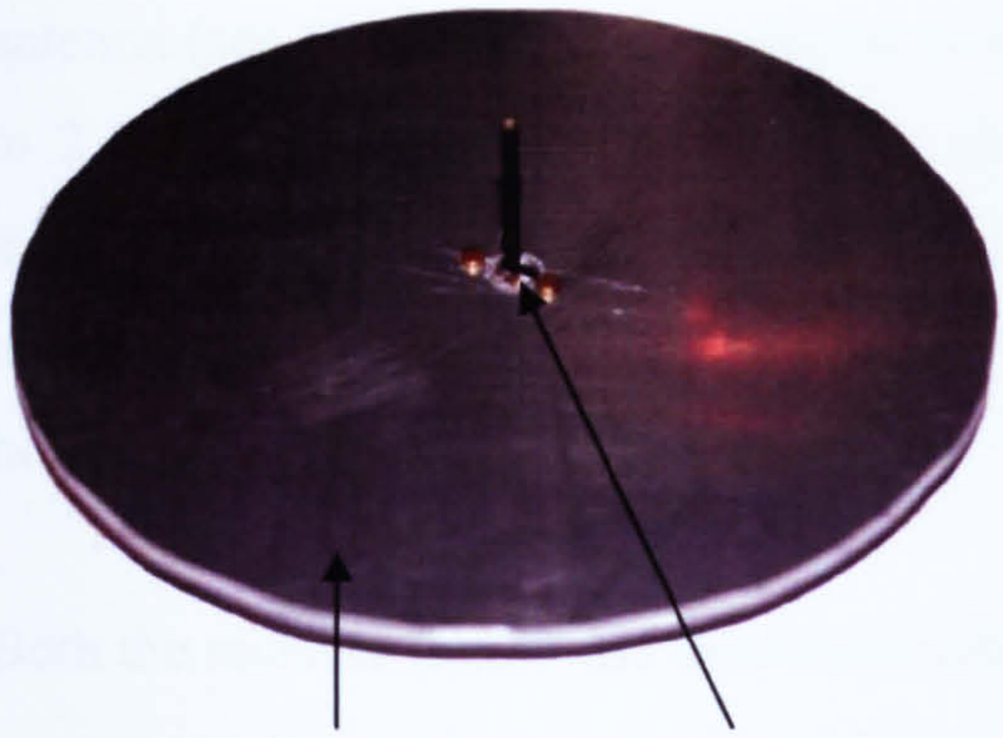

Figure 6.22 – The reference monopole antenna

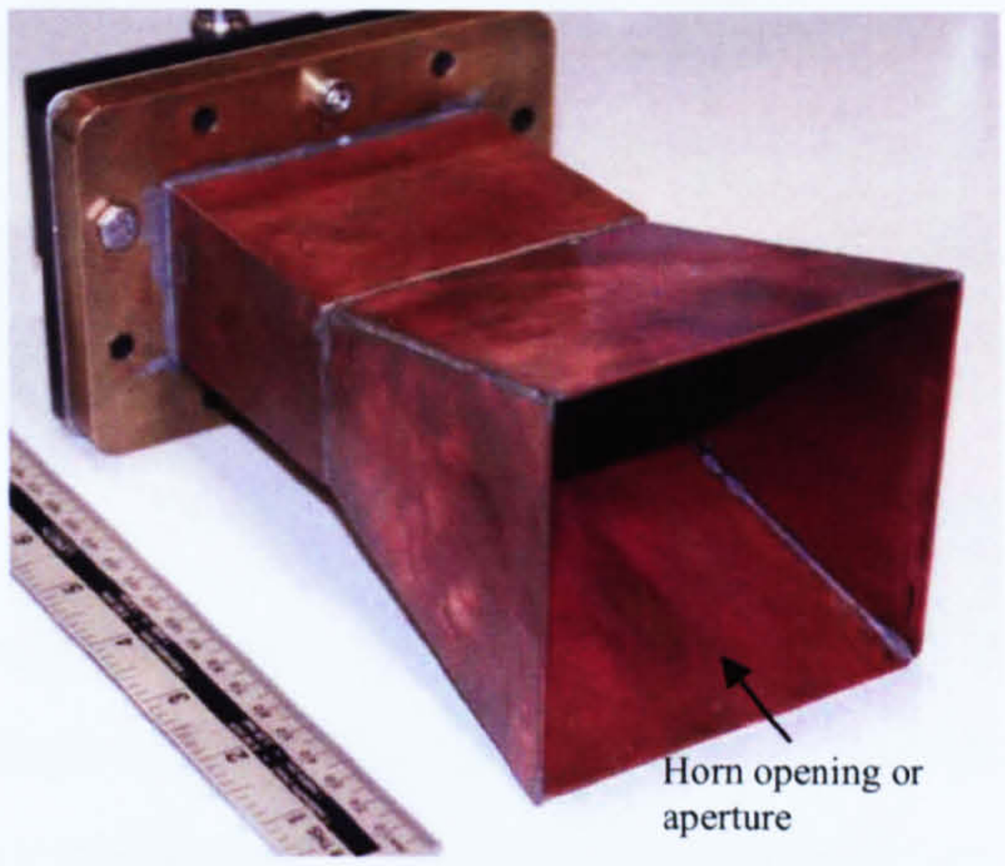

Figure 6.23 – The receiving horn antenna

Unfortunately, an anechoic chamber was not available for this project since they are costly to rent, and even more costly to build. The absorbing materials are expensive and one has to also factor in the cost of the space required to set up such a chamber. The measurements were still required so an alternative testing solution was put into operation.

It was thought that since WLAN infrastructure is localised to urban areas, tests could be carried out in a rural scenario where interference would be reduced. In addition, it was thought that a rural location, such as a large open field, would serve to prevent the occurance of signals being reflected since the the lack of obstacles would allow them to simply travel away from the antenna and into space.

In addition to the FCPW antenna, it was decided that two other antennas were required; a reference (for comparison purposes), and a receiver. The reference antenna constructed for the purposes of these measurements was a $\lambda/4$ length monopole with a ground plane (radiating element = 36mm, ground plane diameter = 130mm - see Figure 6.22) because the expected radiation pattern for this antenna is described in literature; it is also included in the Appendix C.2 as a reference. Another reason for choosing this type of antenna is that it is similar to the existing MicaZ antenna, and therefore, it provides a useful comparison with the FCPW. A horn

antenna (see Figure 6.23) was selected to be the receiver; this was prebuilt and suited to 2.45GHz. The reason for choosing the horn antenna as the receiver is that, as mentioned in Chapter 5, horn antennas have a high directivity and are therefore unlikely to be affected by signals which do not originate from directly in front of the horn aperture.

Both the receiver (horn) and transmitter (monopole or FCPW) needed to be suspended away from the earth in order to prevent reflections from the floor causing multipath interference. In order to do this, tripods were employed. The receiving horn antenna was bolted to the top of a large surveyor's tripod. Since it was assumed that the ground in a rural setting would not be particularly level, a base was made for this tripod with a large bolt at each corner allowing the base to be levelled – a plumb line hanging from the centre of the tripod was used for levelling. This setup is shown in Figure 6.24.

A standard camera tripod with full height adjustment was used to mount the transmitter. Much of the tripod was made out of plastic rather than metal which was considered to be useful in terms of reducing its impact on the experimental results. An attachment was made to fit the standard tripod ¼ inch thread, which allowed for the rotation of the transmitting antenna. Integrating a protractor and pointer into the attachment allowed radiation patterns to be measured with ease. Although the tripod was made mostly of plastic, a further precaution was taken by fixing a PVC pole of approximately 30cm in length to the top of the tripod attachment – the antenna could then be fixed at the highest point away from the tripod. The pole was hollow which allowed a coaxial cable to run down its centre and through the rotating part of the attachment. By positioning the cable in this way it was prevented from becoming deformed during antenna rotation, a possible source of error in the results. In order to change between the *xy* and *xz* / *yz* planes an SMA elbow joint was employed. The attachment for the top of the standard camera tripod is shown in Figure 6.25, whilst the experimental setup with both tripods is shown in Figure 6.26.

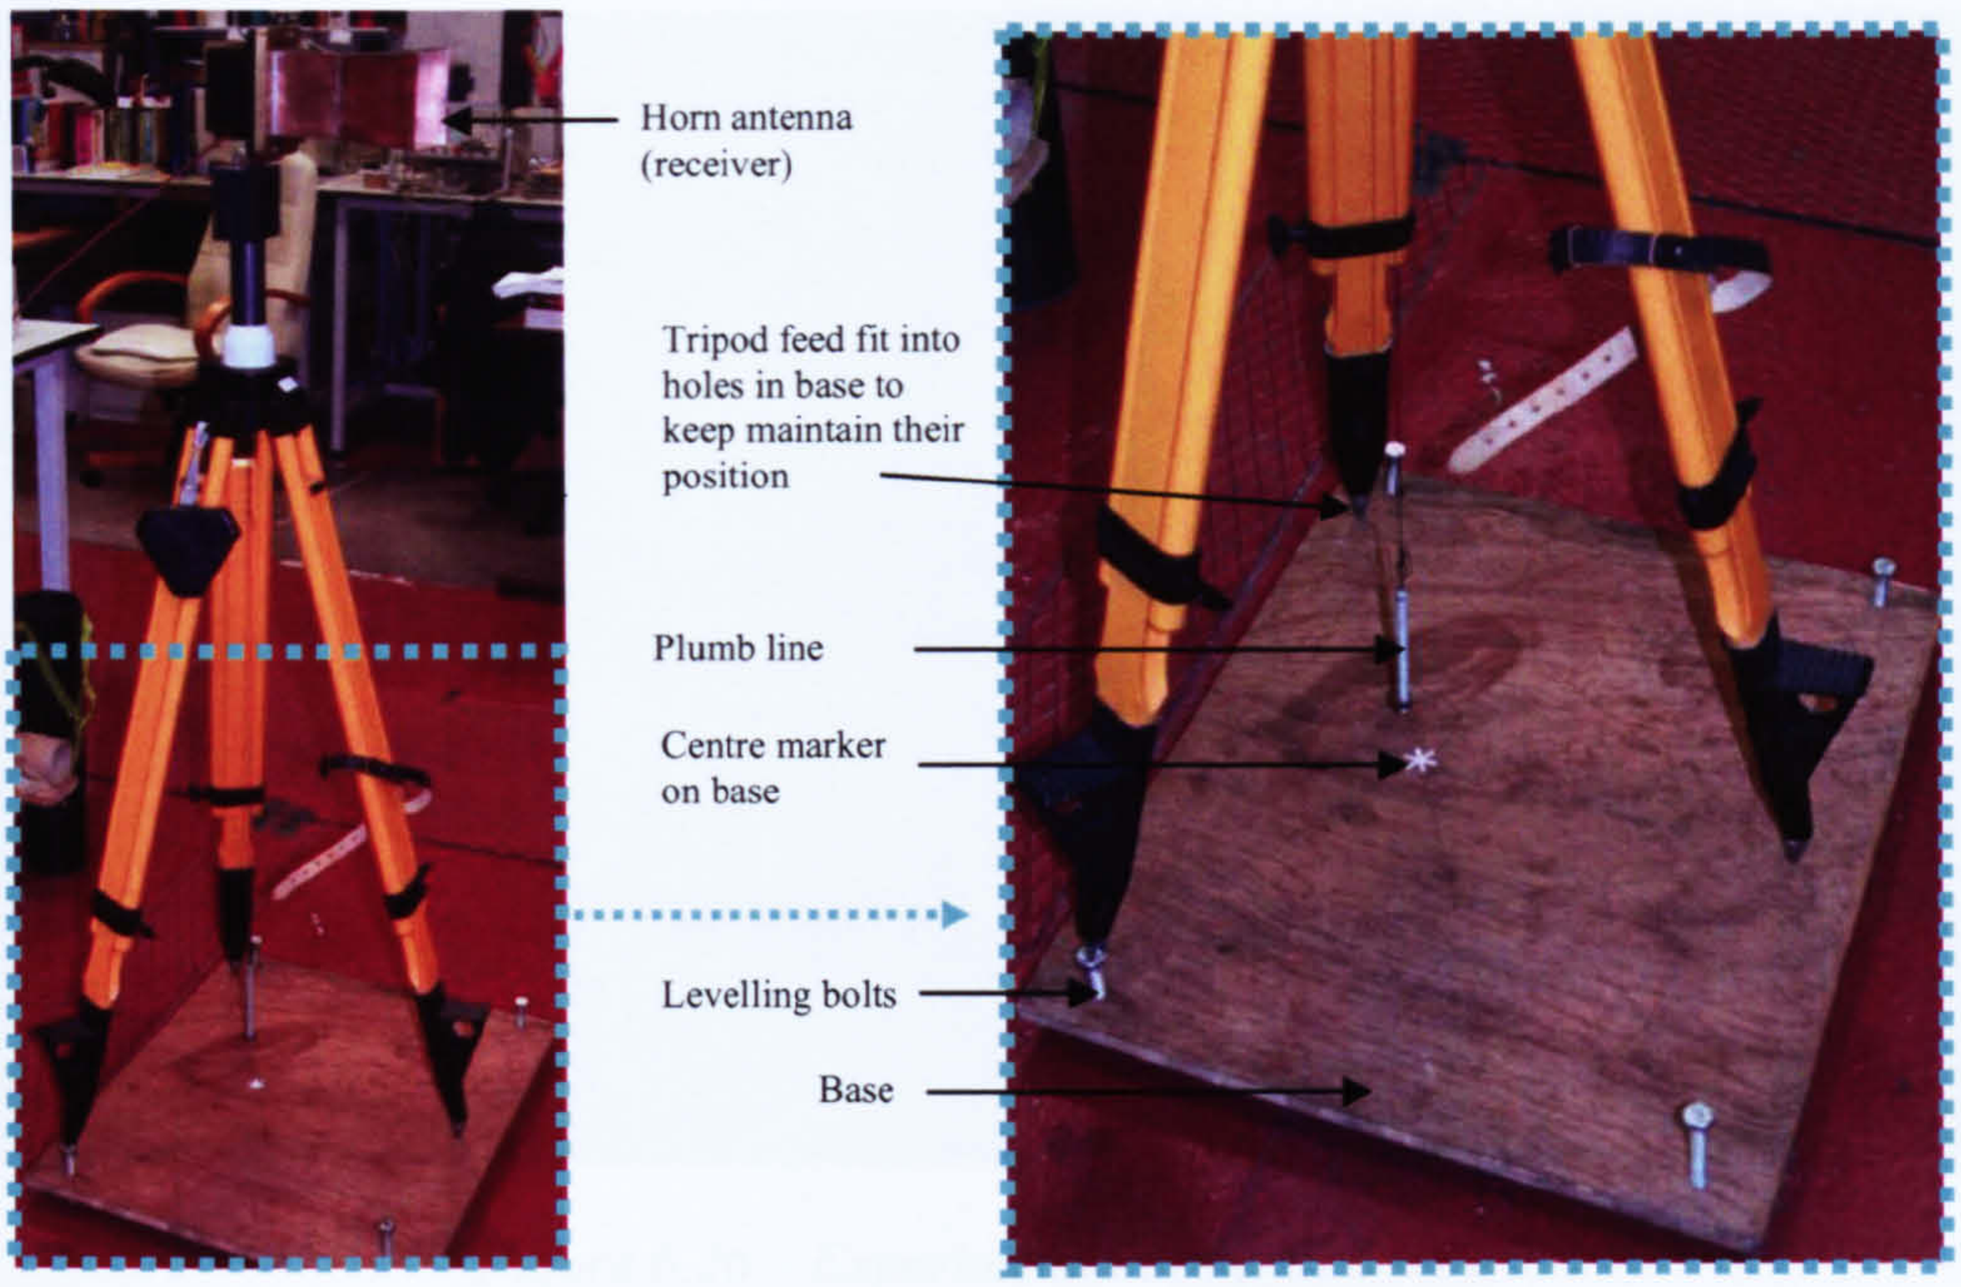

Figure 6.24 – Horn set up on surveyor tripod

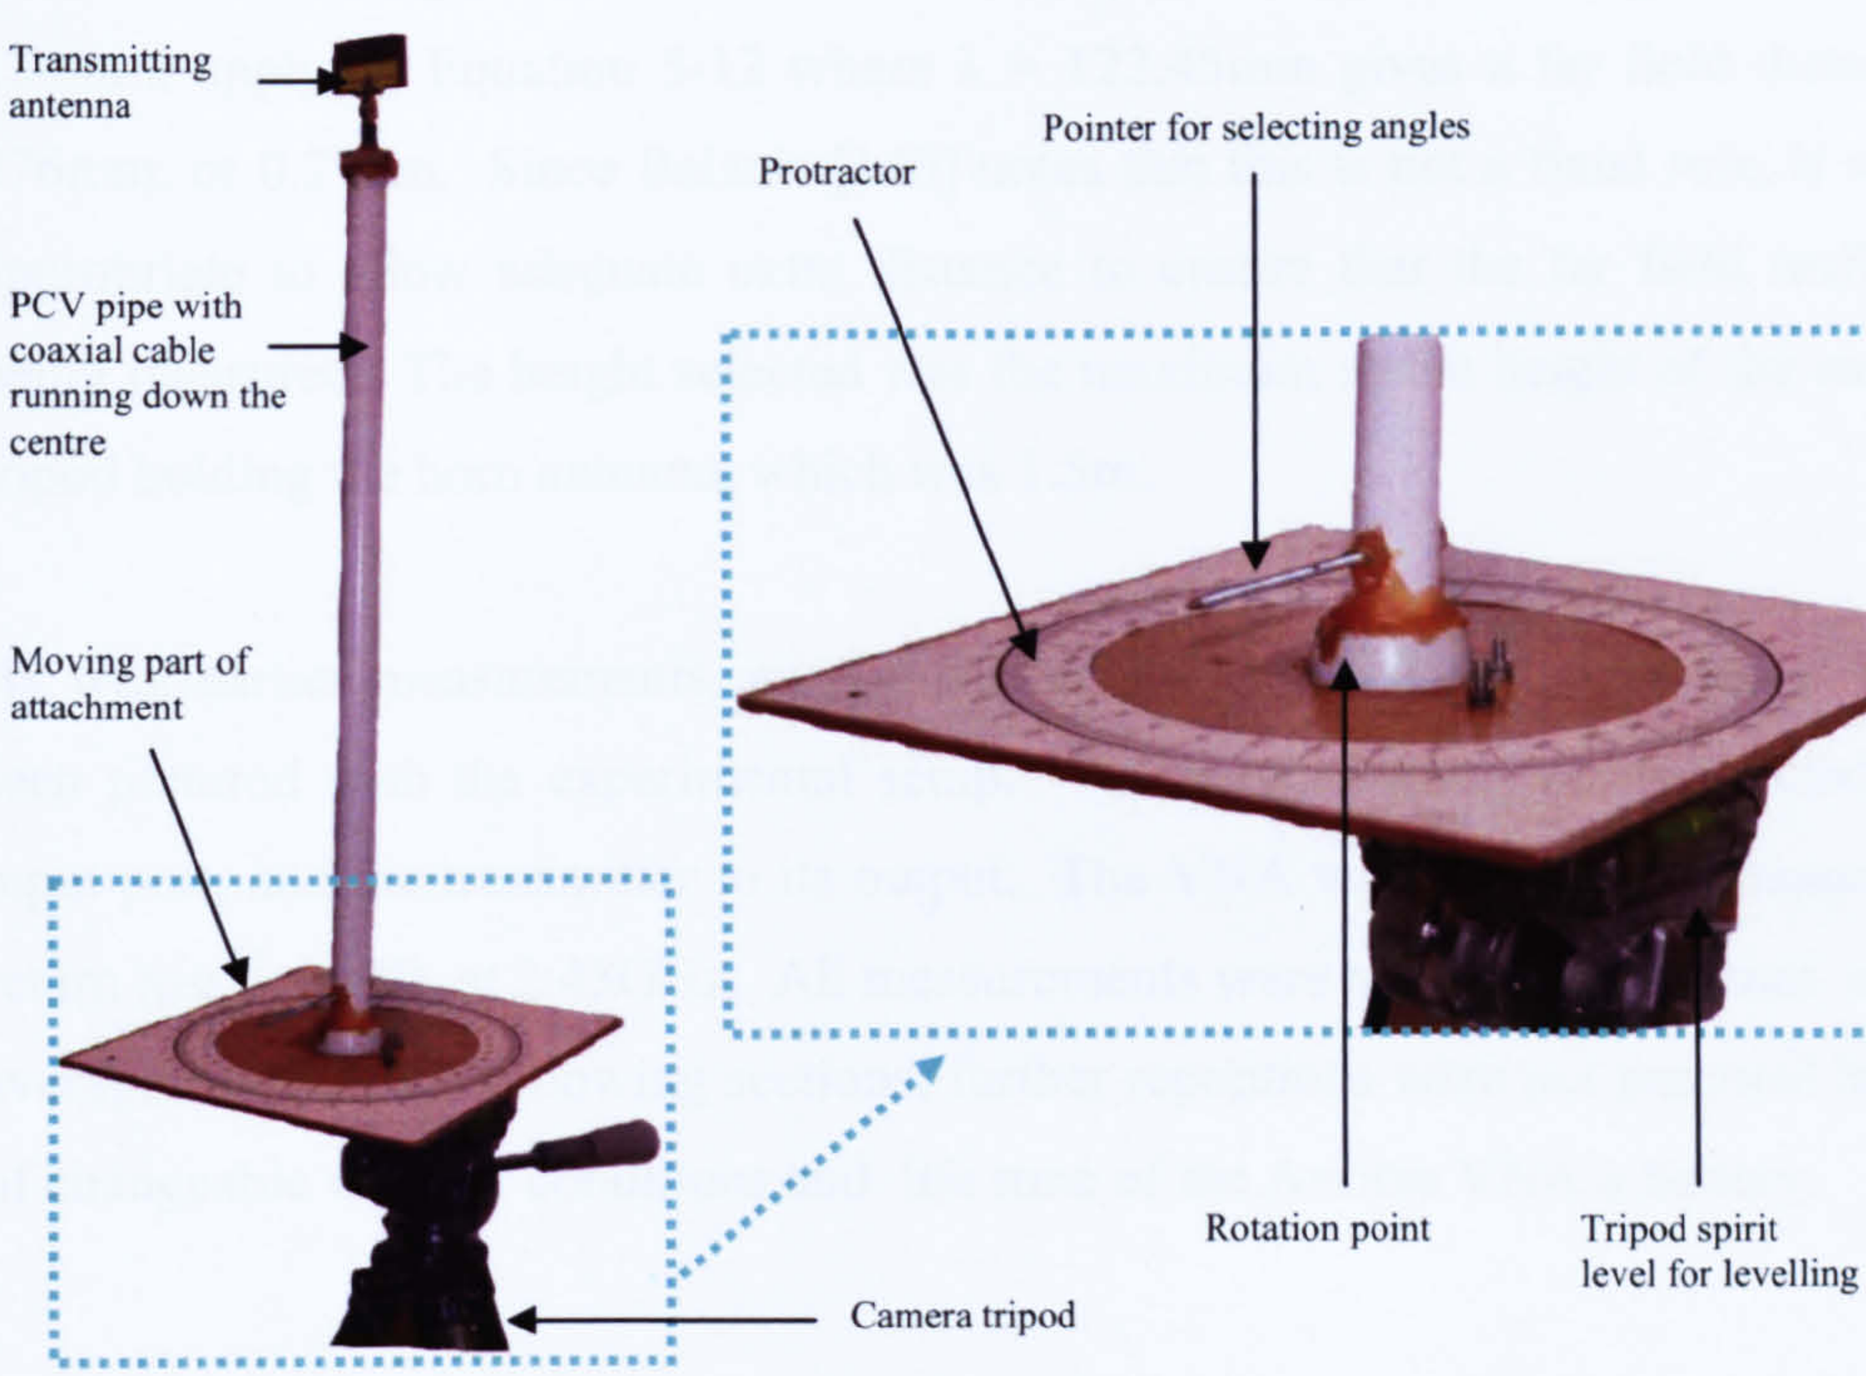

Figure 6.25 – Standand camera tripod attachment

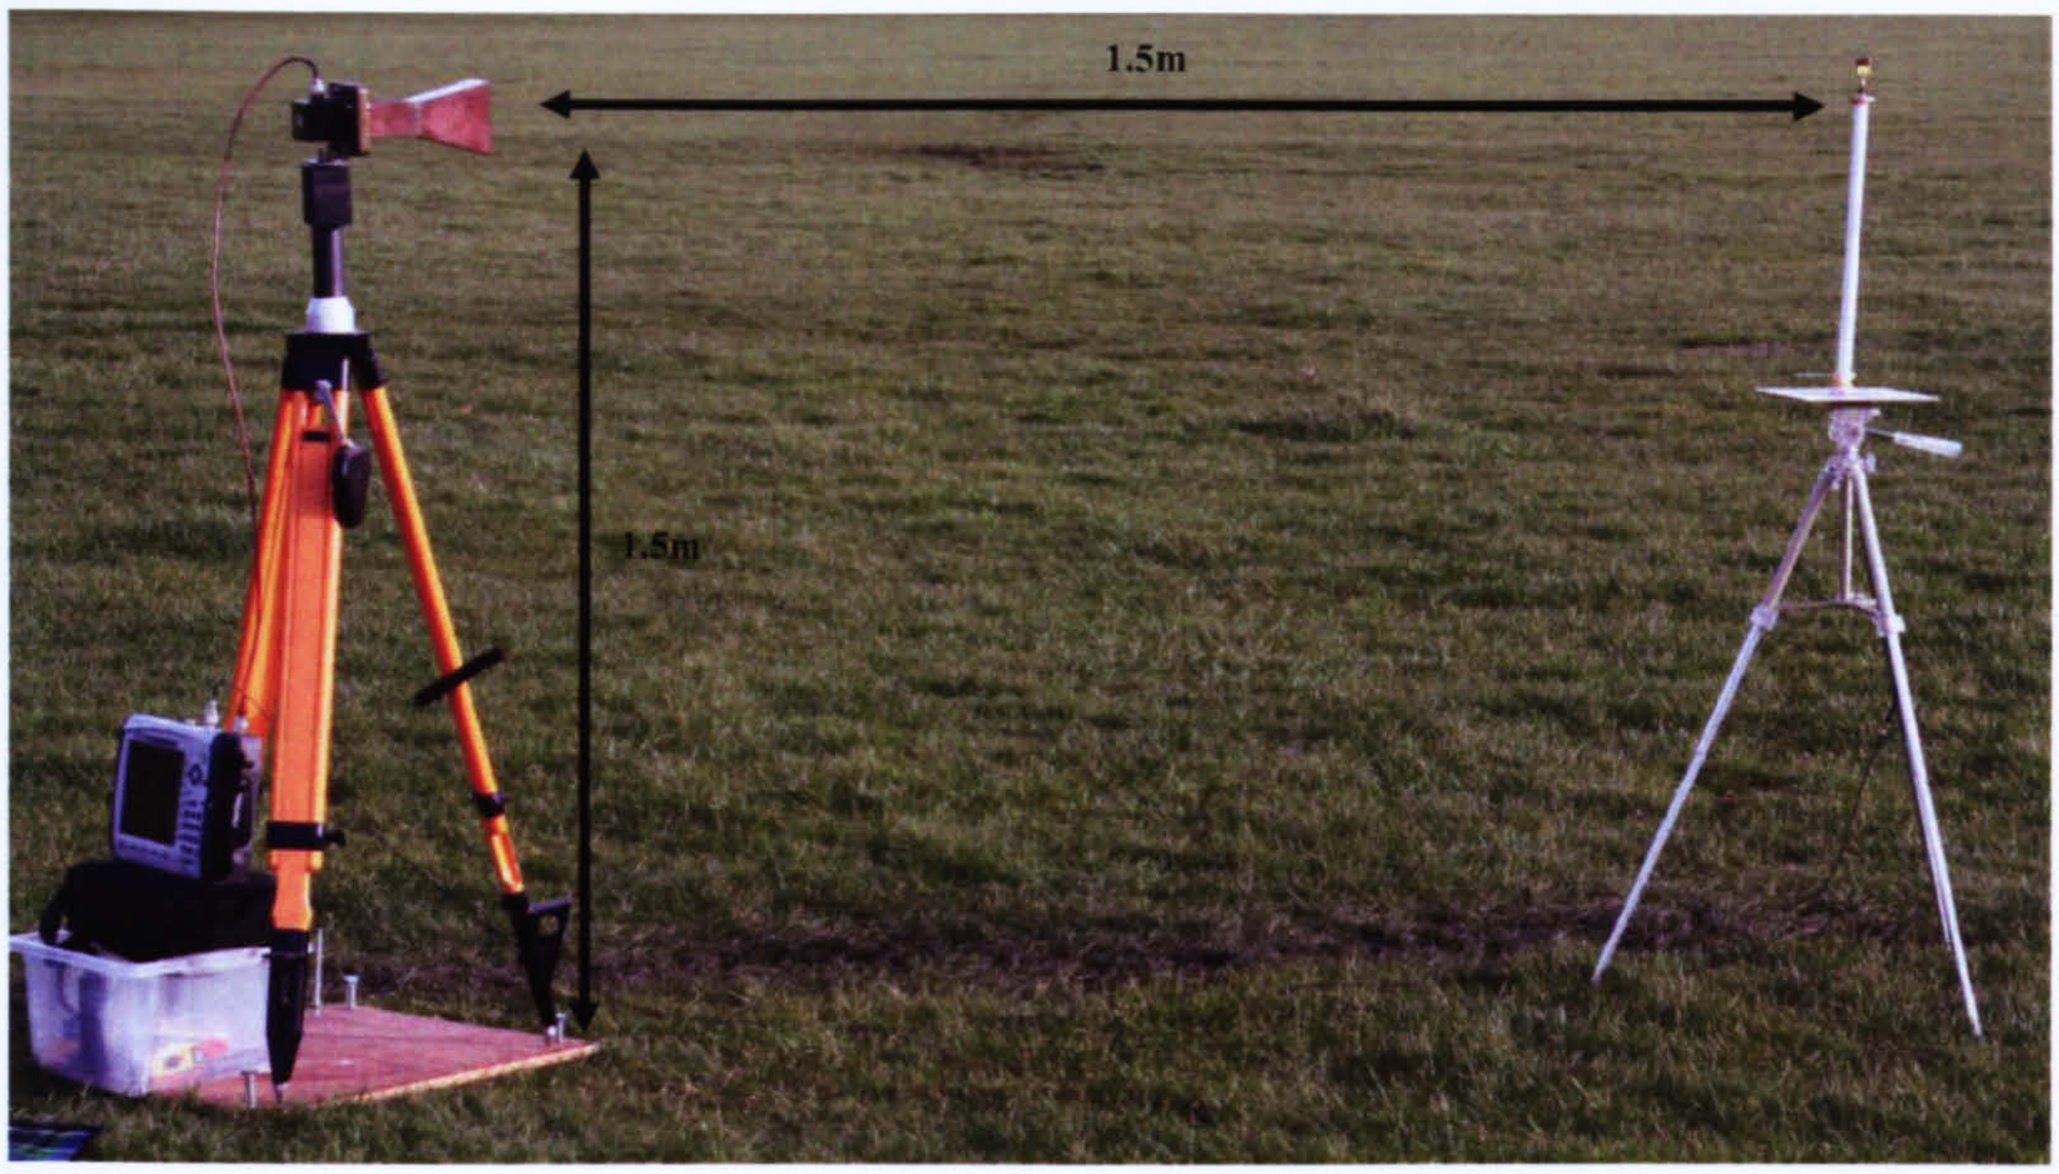

Figure 6.26 – Experimental setup in a field

As shown in Figure 6.26 the transmitter and receiver were placed 1.5m apart. This distance was selected to ensure that the measurements were for the far field of the antennas. The largest dimension of the monopole antenna was its ground plane of 130mm; applying Equation 5-12 where λ = 122.45mm gives a far field distance of 276mm, or 0.276m. Since Balanis [140] notes that this is not a fixed rule, it seemed appropriate to allow adequate extra distance to ensure that the far field really was being measured. The height selected was the maximum stable height of the surveyor tripod holding the horn antenna, which was 1.5m.

As with earlier measurements, an Anritsu MS2024A VNA was used and it can be seen pictured with the experimental setup. The horn antenna was connected to its input port, and the transmitter to its output. The VNA was then used to measure the return loss[17], in dB, at 2.45GHz. All measurements were repeated three times, and the averages used in the following sections; further repetitions were not practical because of changeable weather conditions and life time of the Anritsu VNA's battery.

[17] The ratio of received power (input) to transmitted power (output).

6.7.2. Radiation Patterns

The results of these experiments represent the measured return loss arbitrarily normalised to -50dB. The measurements shown are relative to the antenna orientations shown in Figures 6.27 (monopole) and 6.28 (FCPW). For the monopole antenna the electric field is polarised in the *z* direction, whilst for the FCPW it is taken to be in the *y* direction. Figure 6.30 shows the measured radiation pattern for the monopole antenna, whilst Figure 6.31 shows those for the FCPW antenna.

Looking first at the results for the monopole antenna, these show a good correlation with the literature. In the *xy* plane, the antenna acts isotropically, whilst in the *xz* and *yz* planes significant attenuation occurs in the *z* axis. As previously discussed, finite ground planes do not have the effect of completely blocking signals but there is definite attenuation (in the order of 5-10dB) due to its presence in the positions between 90° and 270° (not inclusive). The maximum measured power occurs in the *yz* plane at 310°, and is 8.07dB. The patterns are also generally quite symmetrical in the *z* axis, again as one would expect, so a similar peak can be found at 50°.

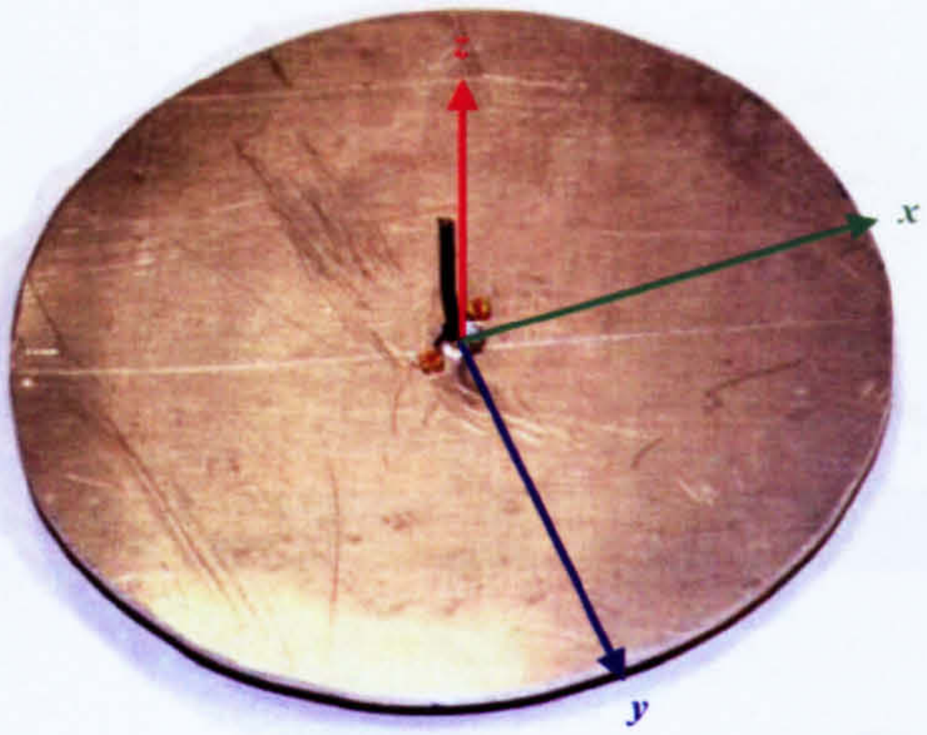

Figure 6.27 – Monopole orientation

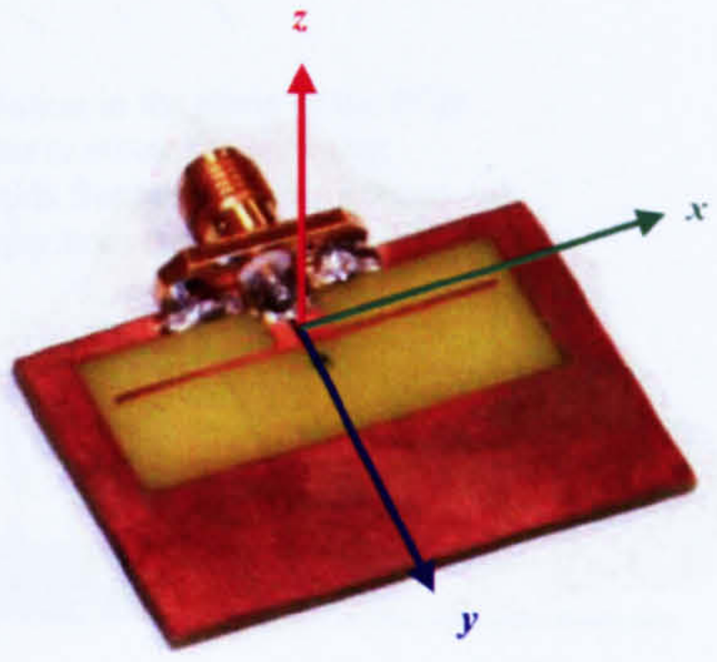

Figure 6.28 – FCPW orientation

The maximum measured power of the FCPW antenna also occurs in the *yz* plane and is 2.5dB, 3.6 times less power than that of the monopole. As with the monopole patterns there is a reasonable amount of symmetry, although the *xz* plane is an exception with it radiating 3dB more at 270° than at 90°. There does not appear to be any obvious answer as to why this occurs, although one possible explanation is that the copper track has been removed a little on one side of the PCB where the ground plane is soldered to the SMA bulkhead connector. Whilst there was no obvious sign of this happening, a number of previous prototypes have suffered from this fault when not handled with the upmost care. This would not be a problem in the future however, since the SMA connector is only for testing purposes in this project.

The radiation is weakest in the plane of the PCB, as expected and explained earlier; a HFSS electric field plot is shown in Figure 6.29 to further highlight why this is the case. A drop in the radiated power of approximately 10dB is present at the 90° and 270° degree positions, as expected, due to the lack radiating elements covering these positions. Whilst a 10dB loss is significant, it is not nearly as significant as the loss that HFSS reported (shown as a 35dB loss in Figure 6.17). Another noteworthy feature is that the FCPW antenna, whilst not outputting peak power near that of the monopole antenna, does not drop much below -10dB, whilst the monopole antenna drops to almost -20dB along the *z* axis.

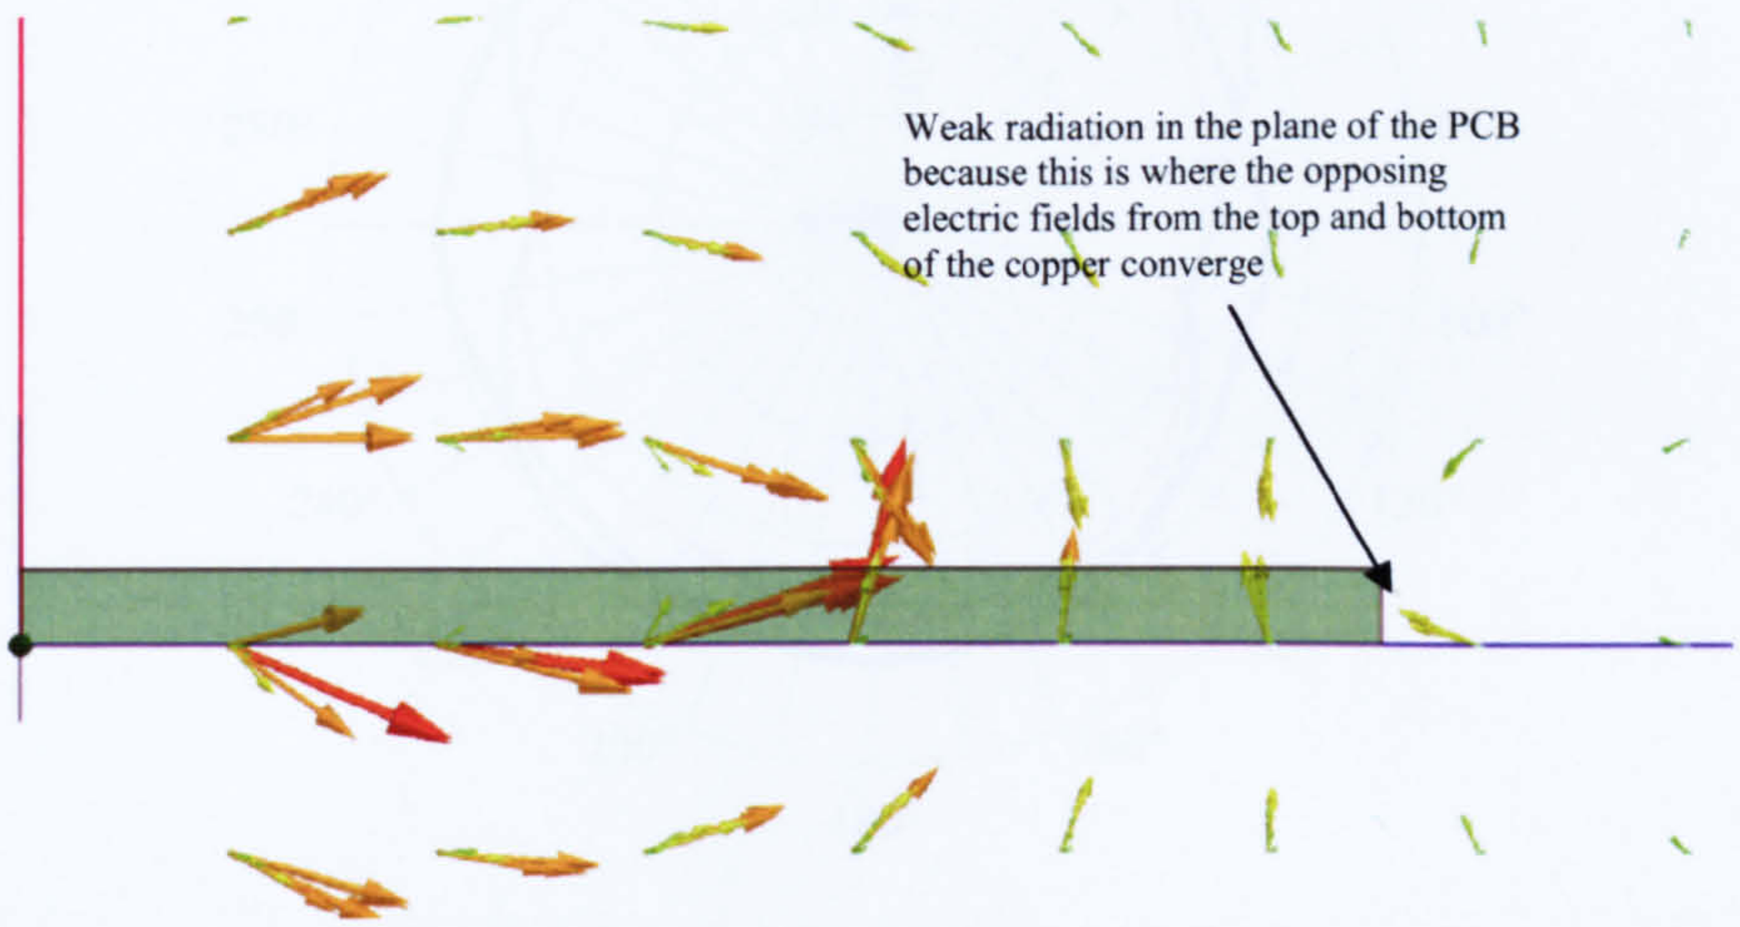

Figure 6.29 – FCPW simulated electric field vector plot showing the formation of fields at the edges of the PCB

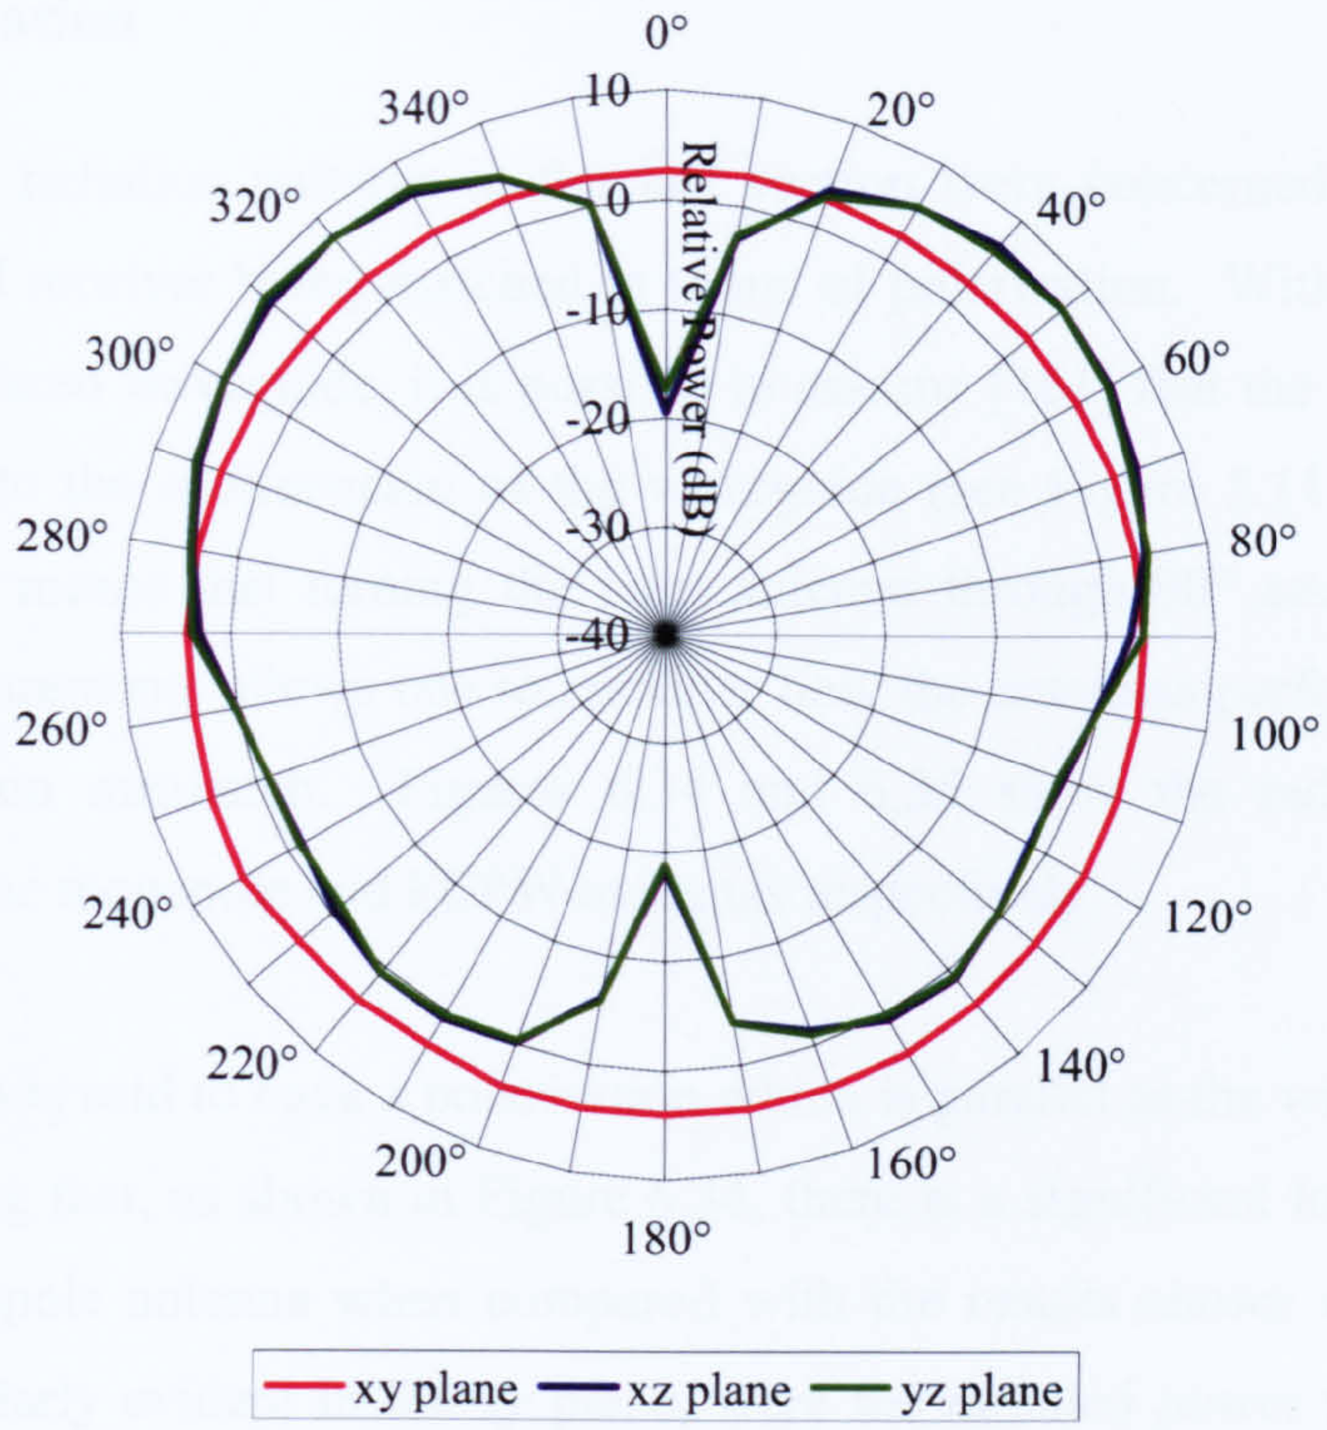

Figure 6.30 – Monopole antenna measured radiation patterns

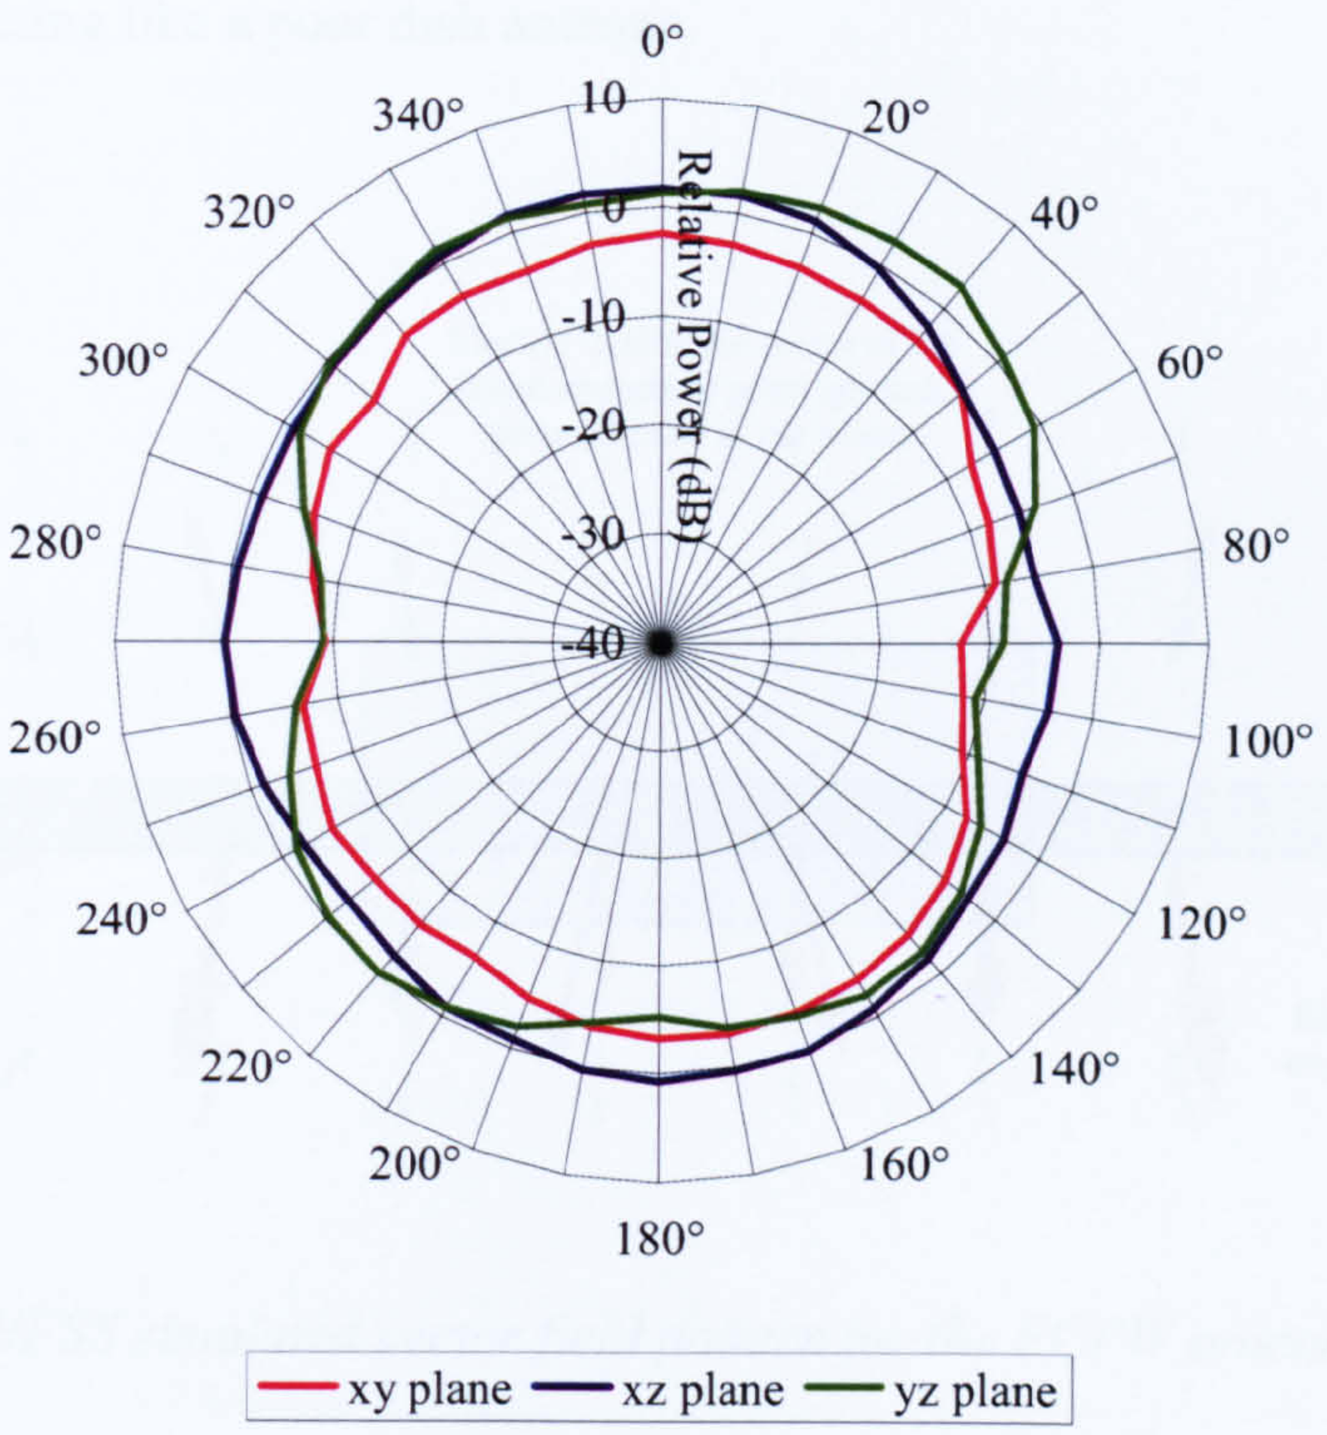

Figure 6.31 – FCPW antenna measured radiation patterns

6.7.3. Polarisation

The measured radiation patterns in the last section were concerned with both the transmitter and receiver being matched in terms of polarisation. With the pyramidal horn being a flared waveguide, it is possible to assume [161] that the electric field is perpendicular to the *a* dimension of the waveguide (see Figure 5.11 for dimension details). This means that turning the horn antenna through 90° and repeating the previous measurements allows one to consider how the antennas perform when there is a polarisation mismatch. Figures 6.34 and 6.35 show the radiation patterns measured for the monopole and FCPW antennas respectively.

A wire antenna is said to have a polarisation which is parallel to the wire [162]. So it is not surprising that, as shown in Figure 6.34, there is a significant loss experienced with the monopole antenna when compared with the results shown in Figure 6.30. This is particularly evident in the *xy* plane, were the radiated power reaches -40dB. There is some improvement in the *xz* and *yz* planes, although it is suspected that this is a result of reflections incident to the ground plane; it is possible that the ground plane is effectively acting like a poor dish antenna.

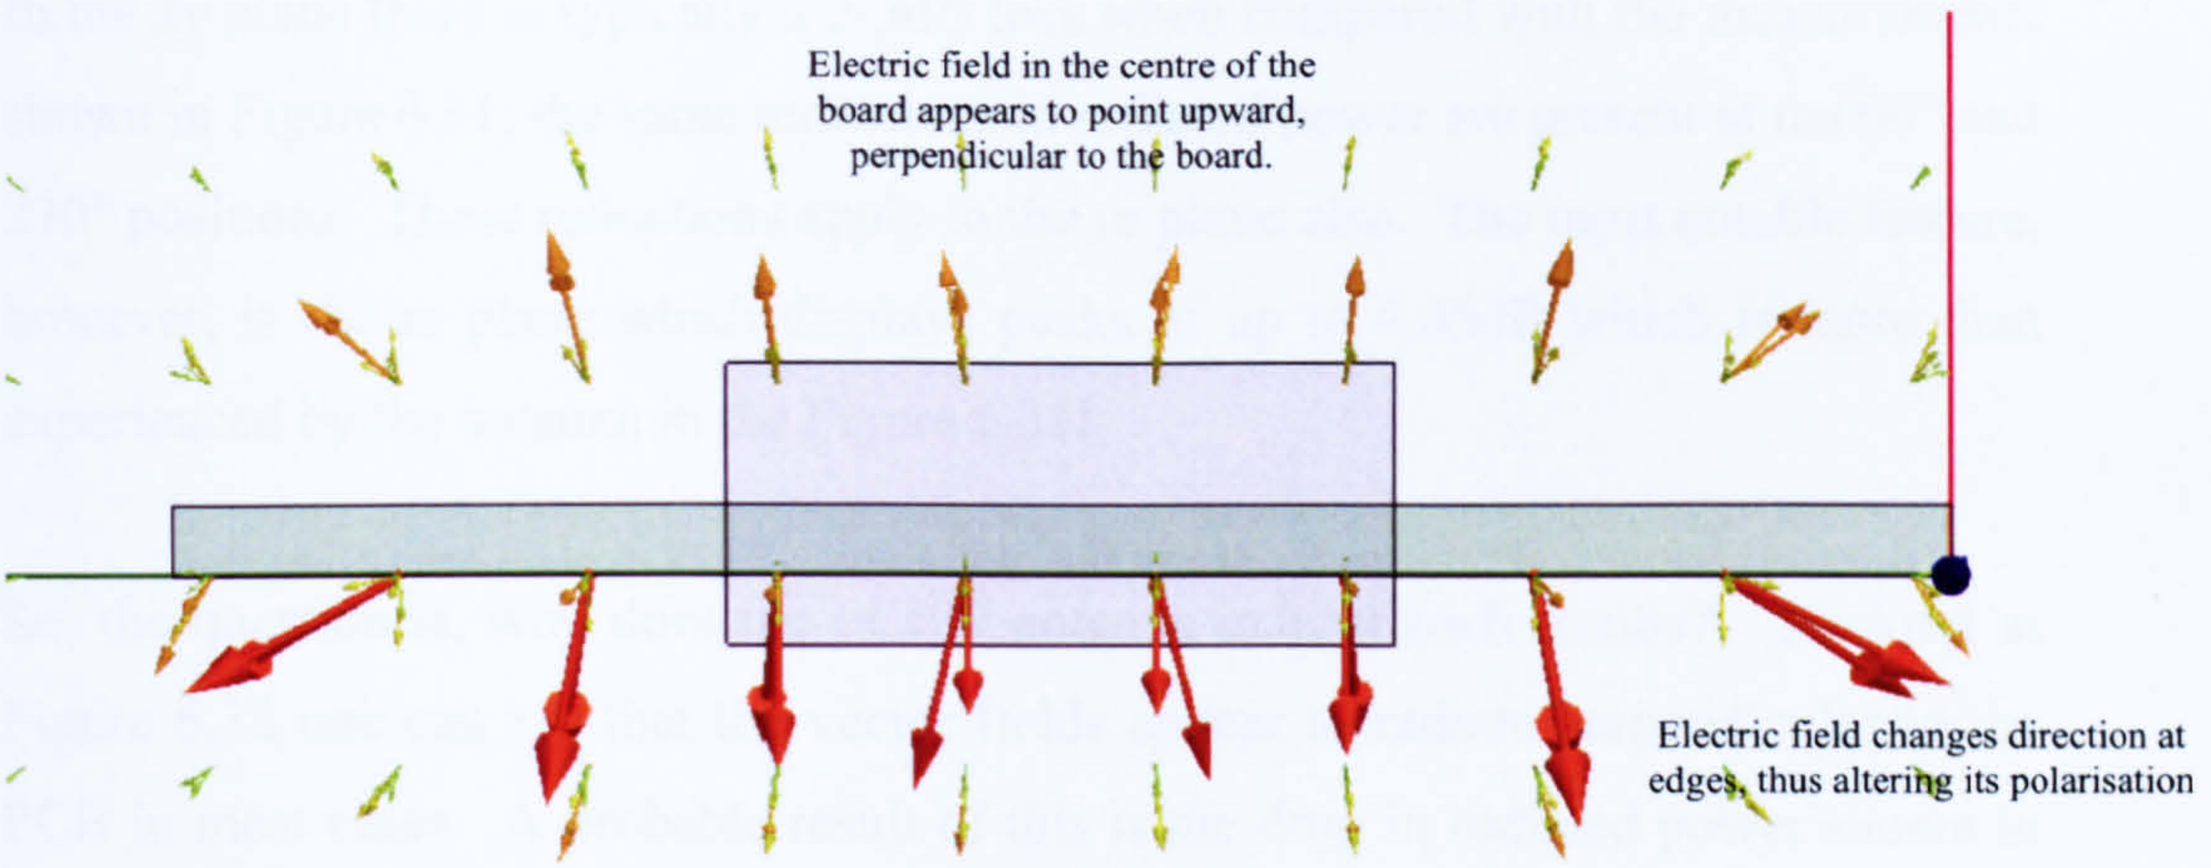

Figure 6.32 – HFSS simulated vector field pattern for the FCPW antenna showing the formation of multiple polarisations (side view, xz plane)

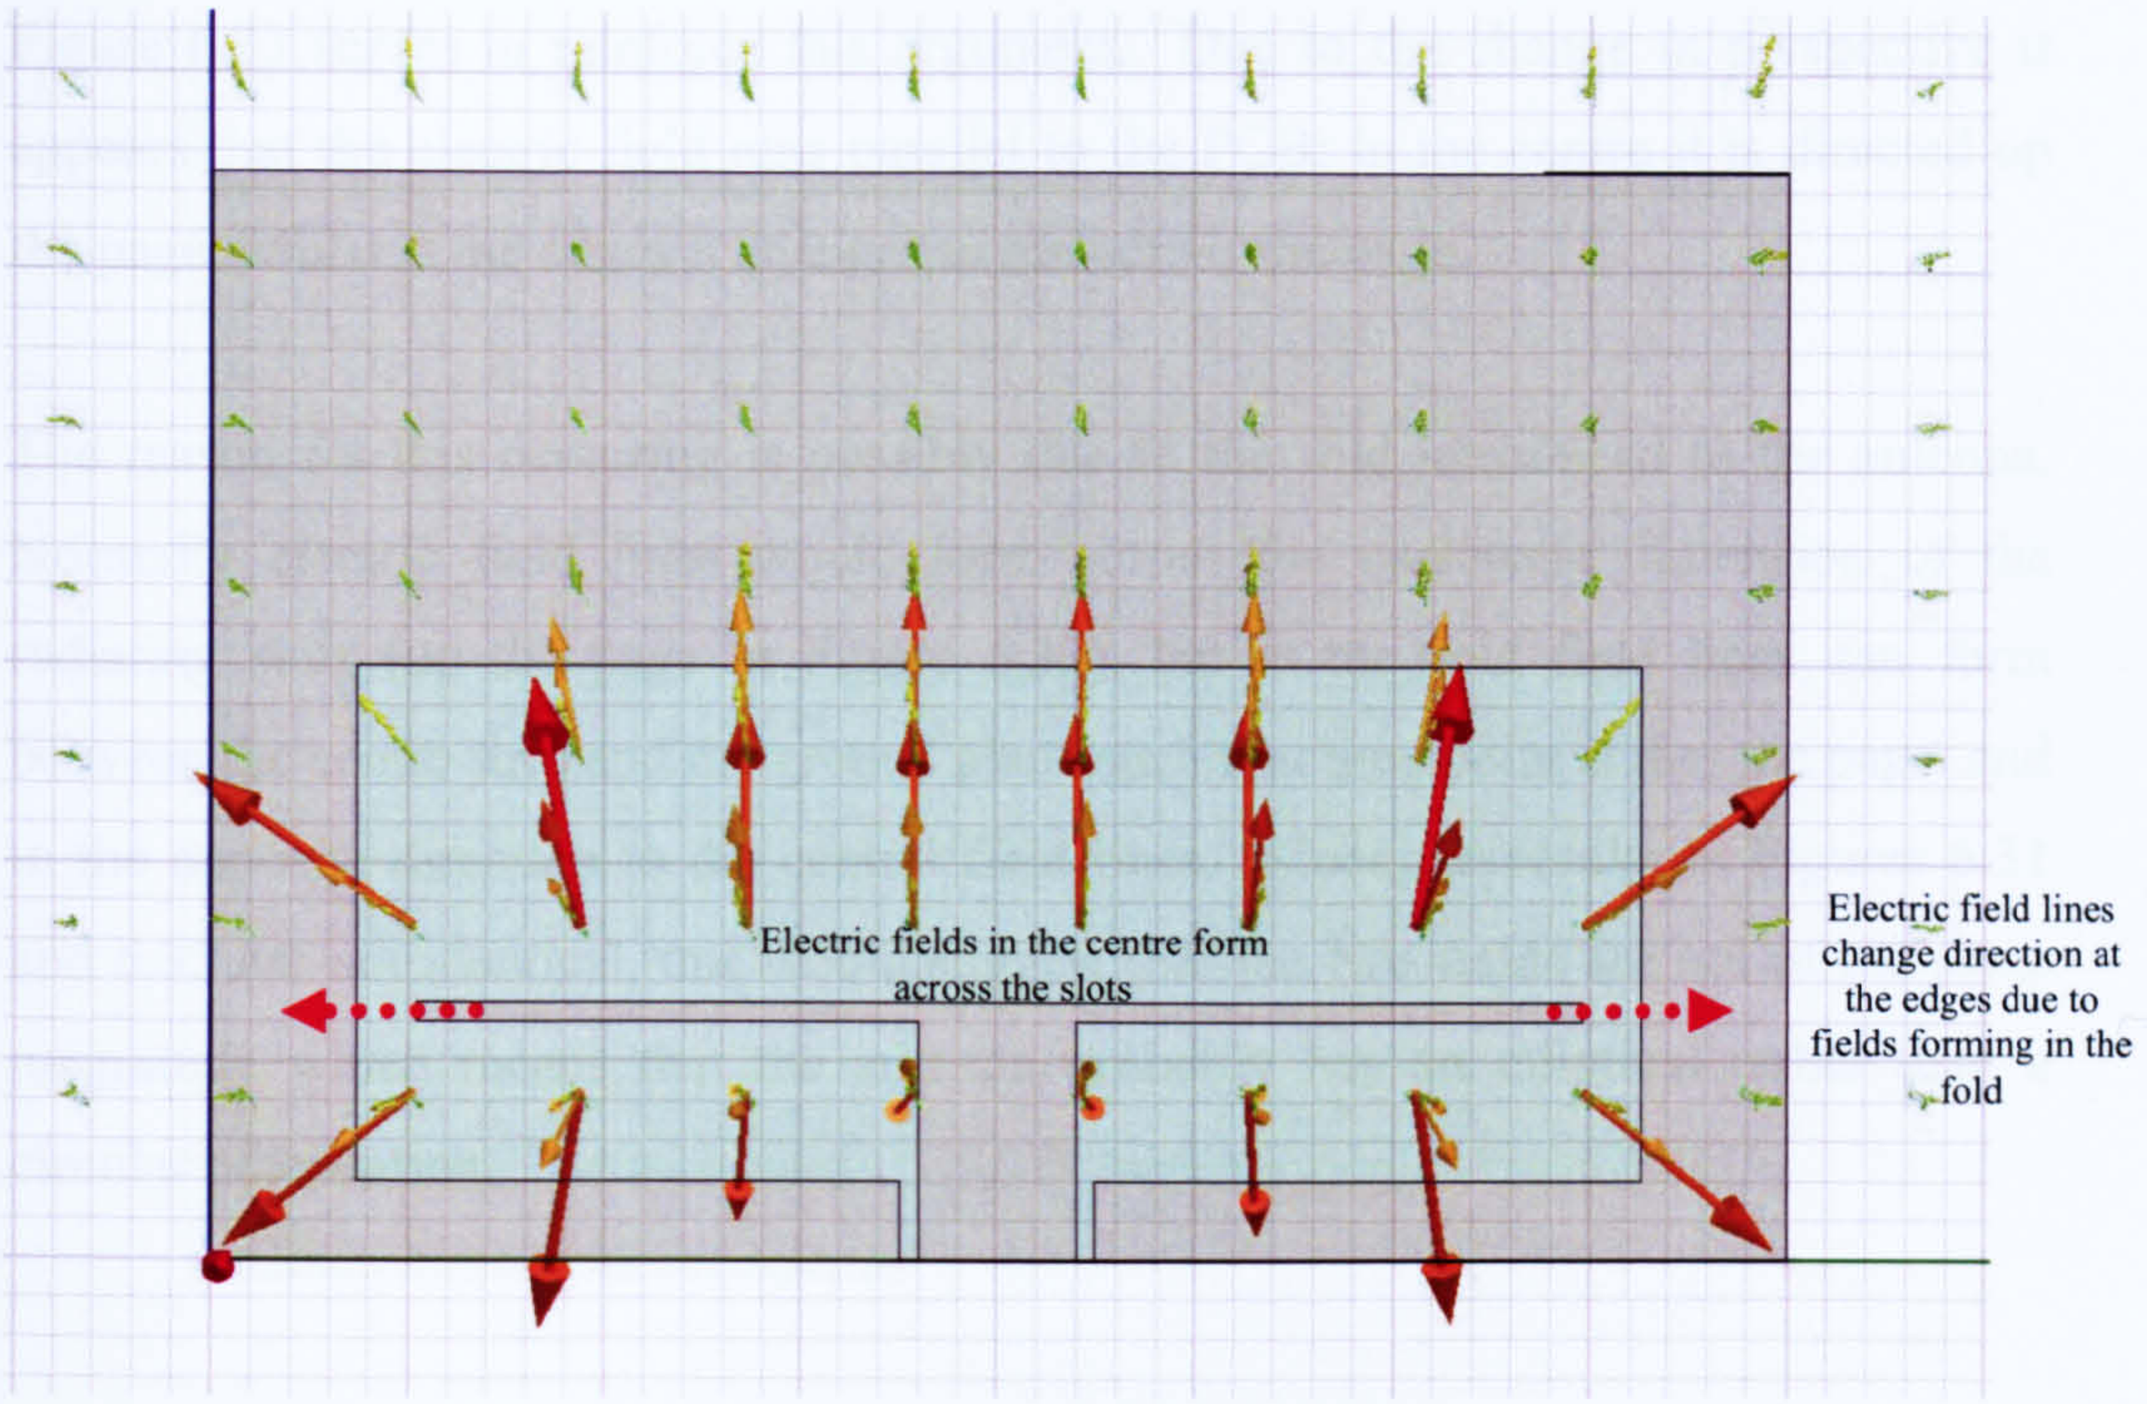

Figure 6.33 – HFSS simulated vector field pattern for the FCPW antenna showing the formation of multiple polarisations (top view, xy plane)

Considering now the results from the FCPW antenna (Figure 6.35), it appears that this antenna is much more tolerant of polarisation mismatch than the monopole antenna. In the *xy* plane there is typically a 2-3dB loss when compared with the measurements shown in Figure 6.31; the same reductions in radiated power are present at the 90° and 270° positions. These reductions apply to the *yz* plane also. The most notable feature, however, is the *xz* plane which displays peaks of up to 4.05dB which is more than experienced by the antenna in the Figure 6.31!

So, the question is, why does the FCPW antenna exhibit such results? Looking at Figure 6.32 one can see that the vector fields appear to radiate perpendicular to the PCB in most cases. A probable result of this is the drop in radiated power shown in Figure 6.35 in the *xz* plane at 0°, because directly above the antenna there is a polarisation mismatch, whilst there is an increase shown nearer the edges of the board. At the edges of the PCB the electric field appears to change direction so that it is nearly orthogonal to the field in the centre. This indicates that the polarisation of the electric field is not strictly linear, as is the case with the monopole antenna.

Figure 6.33 serves to reinforce this argument. Due to the change in perspective it appears that the electric field runs parallel to the PCB; in the centre it is directed up the page, whilst at the edges it changes to run across the page.

The reason for this occurring is possibly due to the fold introduced to the antenna. Normally electric field lines would form across the narrowest dimension of the radiating slots (up the page in Figure 6.33), but at the fold field lines can form between the centre strip and the ground plane such that they point across the page, and in the opposite direction to the central field lines. Since the results in Figures 6.31 and 6.35 are not identical, this would suggest that the two fields are not of the same magnitude which means that the antenna probably has an elliptical rather than a circular polarisation.

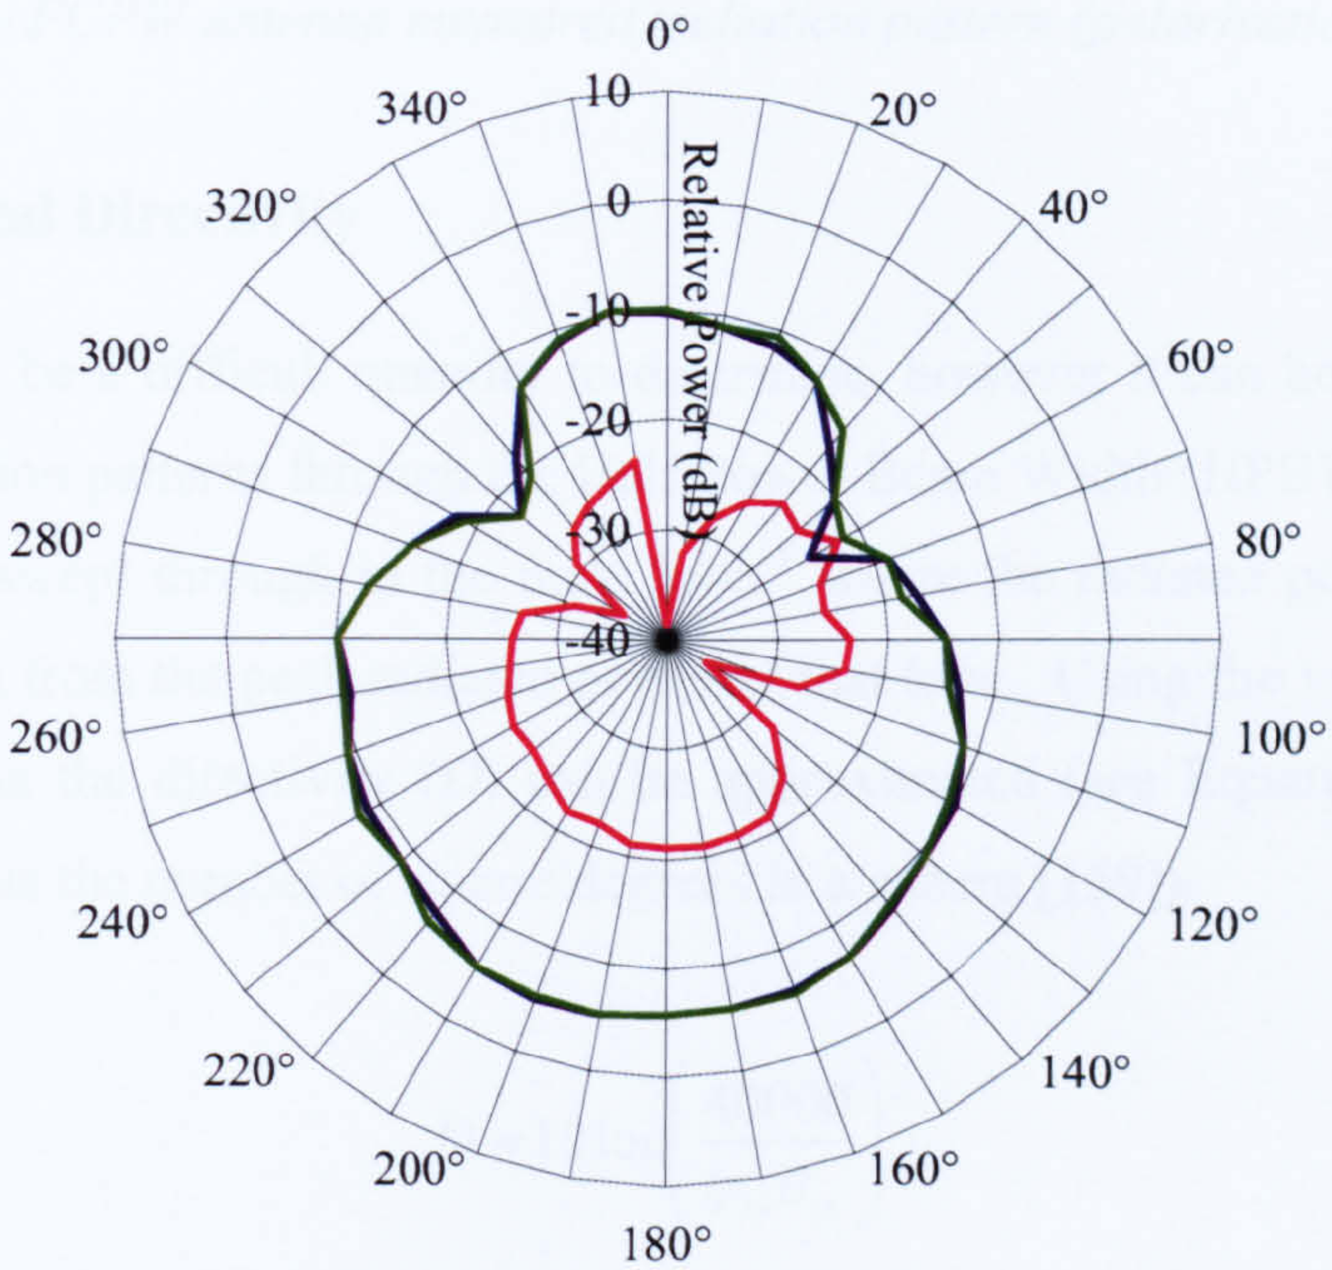

Figure 6.34 – Monopole antenna measured radiation pattern (polarisation mismatch)

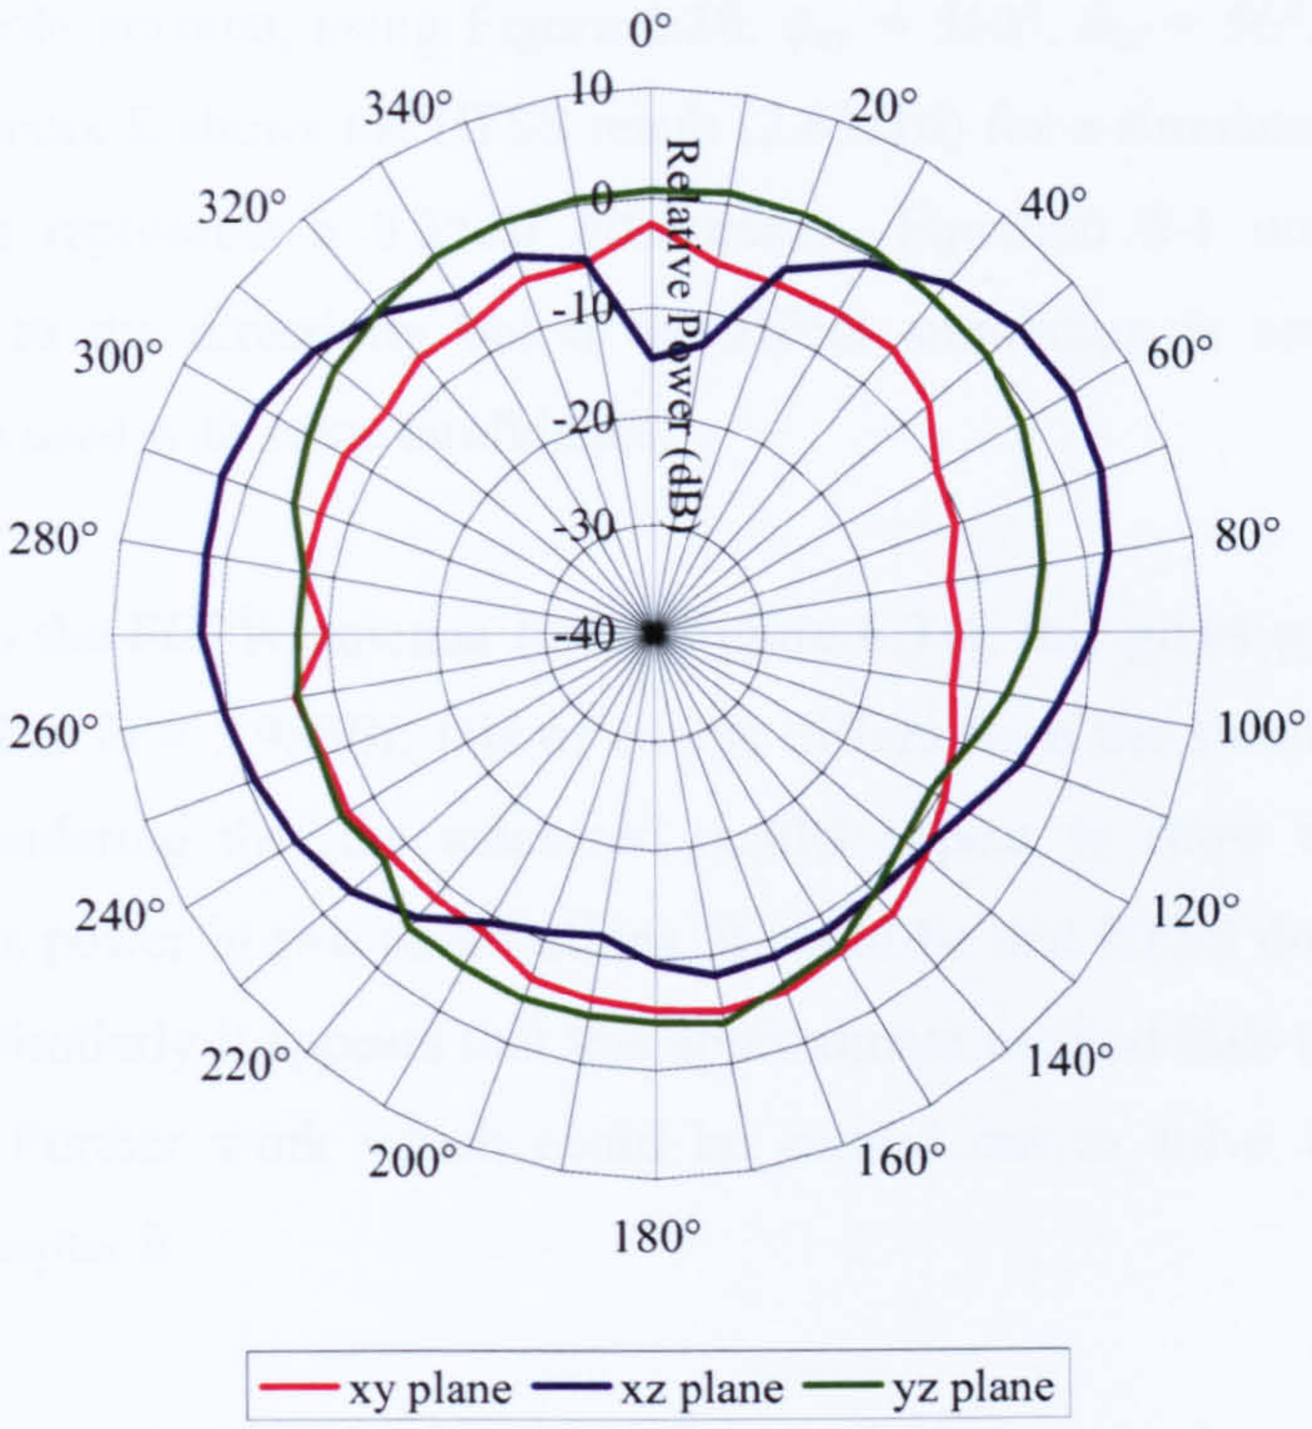

Figure 6.35 – FCPW antenna measured radiation pattern (polarisation mismatch)

6.7.4. Gain and Directivity

Directivity can be a difficult quantity to determine, however it can be approximated from the radiation patterns through the Half Power Beam Width (HPBW) [159] – that is – the angle swept through in the main lobe[18] where the radiated power is no less than 3dB down from the peak radiated power in that lobe. Using the *xy* and *zx* planes of each antenna the directivity (D) can be approximated (see Equation 6-1, where 40000 represents the number of square degrees in a sphere [159]).

$$D \approx 10\log\left(\frac{40000}{\varphi_{xy}\theta_{xz}}\right) \qquad \text{(Equation 6-1)}$$

[18] Lobes identifiable segments of a radiation pattern, where the main lobe is that were the most power is radiated.

For the monopole antenna, using Figure 6.30, $\varphi_{xy} = 360°$, $\theta_{xz} = 56°$, therefore $D \approx$ 2.98dBi. Appendix E shows the HFSS result (2.63dBi) for a simulated model of this antenna, which represents a 0.35dB difference. Equation 6-1 only provides an approximation of the directivity but if the HFSS simulation is accurate then the equation can be used with some confidence.

Moving now to the FCPW antenna (using Figure 6.31), this gives $\varphi_{xy} = 102°$, $\theta_{xz} =$ 99°, and therefore $D \approx 5.98$dBi; this of course differs from the simulated result of -3.31dBi. Considering that the measured results appear to show that the FCPW antenna radiates power in two polarisations, it could be that HFSS does not take this into account. Similarly it appears that this approximate method fails to consider both polarisations. Further work which could be carried out to solve this problem is discussed in Chapter 8.

In terms of gain, which takes into account ohmic losses in the materials of the antenna, this can be calculated using Equation 5-13. With the original model shown in Section 6.5.1, where the copper ground plane is assumed to be perfect and thin, HFSS calculated the radiation efficiency to be 0.9203. A further model was created with a 35μm layer of copper, and the radiation efficiency was reduced to 0.9164. Using this later value, the gain of the FCPW antenna can be calculated to be 5.60dBi.

6.8. Mote Antenna vs. Folded Co-Planar Waveguide Antenna

6.8.1. Experimental Setup

Whilst a measurement of the antenna's radiation patterns was useful to understand how power was distributed around the FCPW antenna, the calculated values for gain and directivity did not match the simulation so there was a need to consider a different method of evaluating the performance. Since the objective is to replace the MicaZ antenna, it was thought that this could involve a direct comparison of the existing antenna and the FCPW antenna. Therefore experimental results were obtained which compared the two antennas in operation on the MicaZ.

A PC application was written which utilised the RSSI feature of the motes to calculate a ten second average for signal strength, along with the number of data packets received (out of a maximum of 100). It was decided to record both of these items of data to ensure not only a high signal strength, but also that data transmission was taking place. Figure 6.36 shows a screenshot of this application. Results from this application also appear in Chapter 7 – it is noteworthy that it accommodates both the Mica2 and MicaZ motes since they have different methods of calculating the RSSI.

Tests were conducted outdoors, and set up as shown in Figure 6.37. The base station and a moveable MicaZ node were placed 1.5m above ground level, and the distance between the two was then increased in 1m increments, starting at 10cm and ending at 25m[19]. The transmission power level of both the base station and the moveable node were set to 0dBm. Various orientations of both the standard monopole and FCPW antenna were investigated, including what happened when a polarisation mismatch occurred.

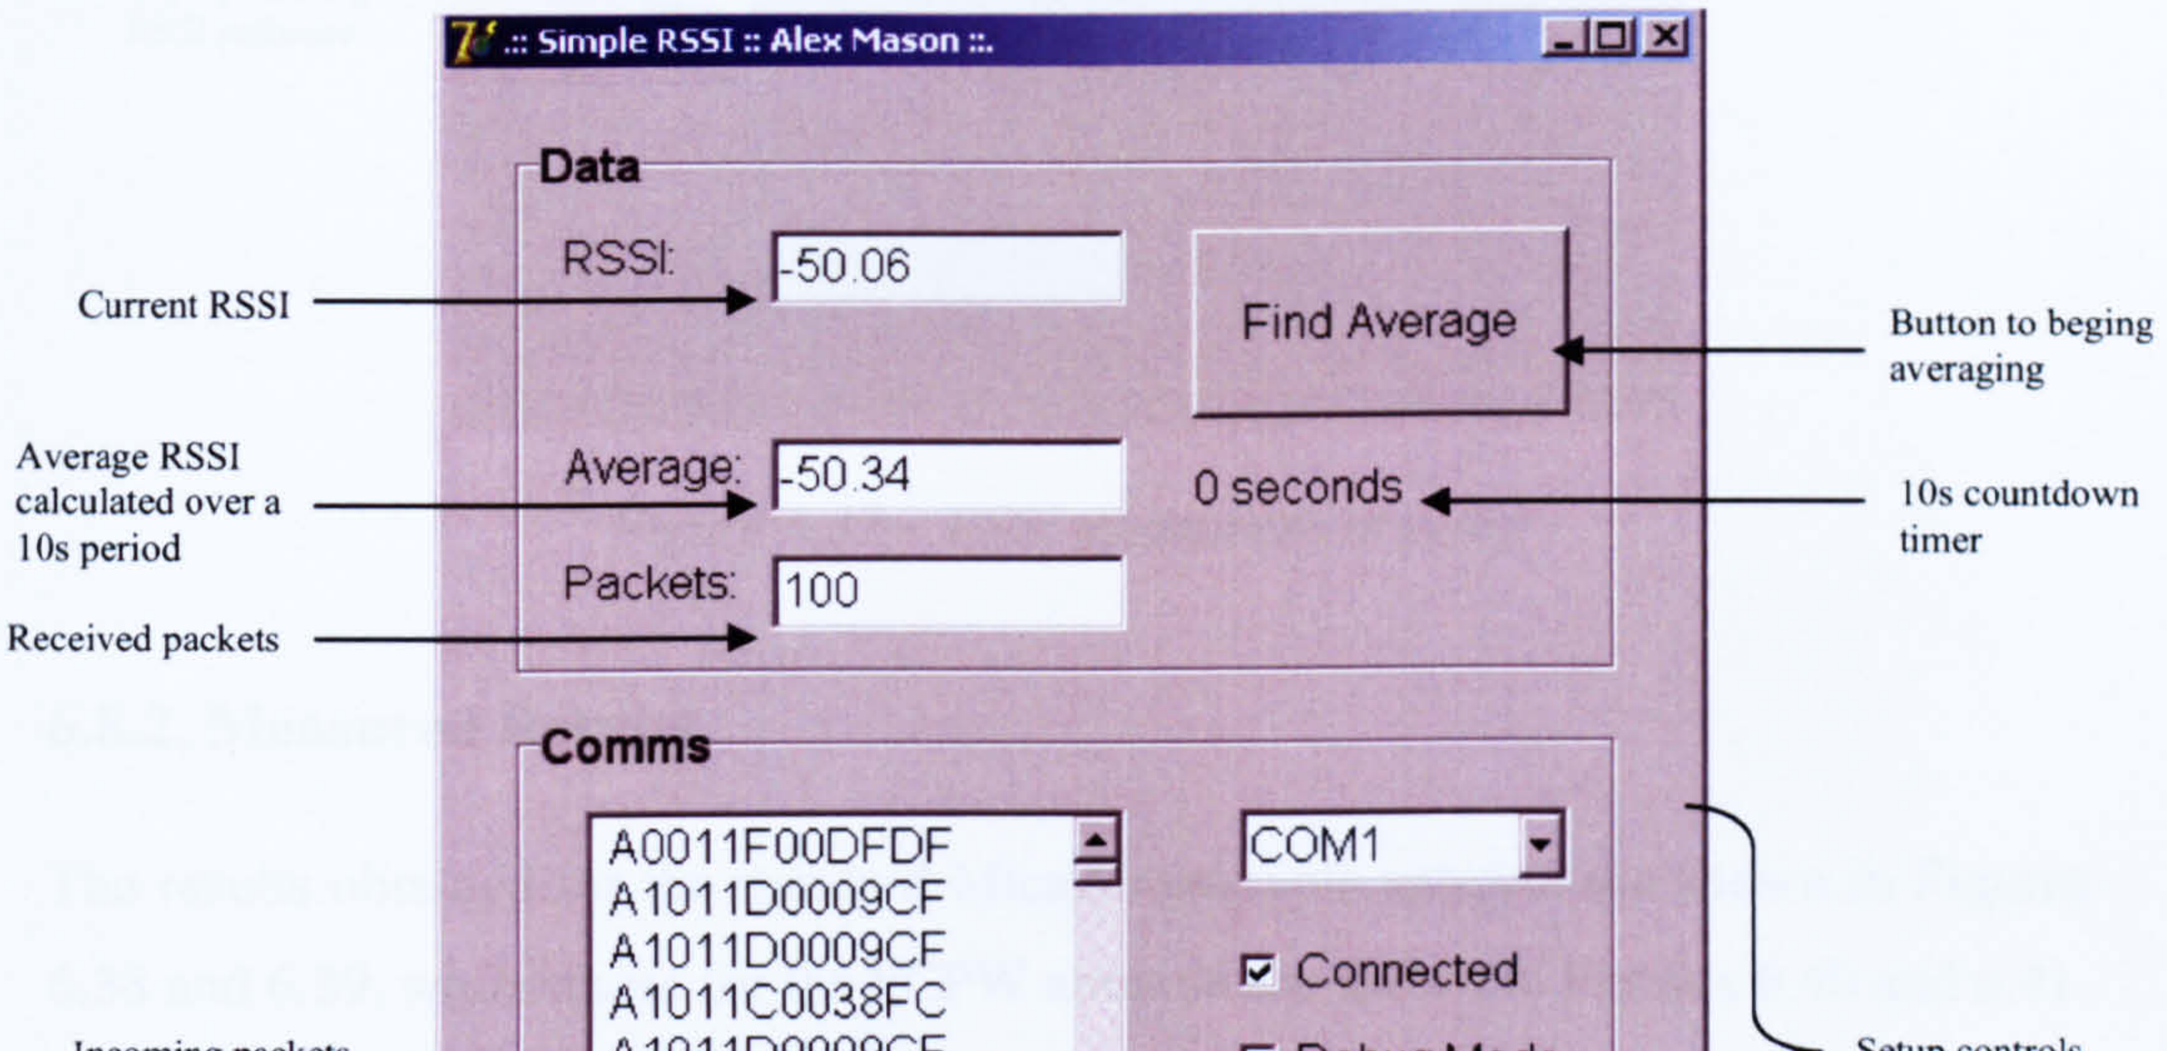

Figure 6.36 – Application for recording mote data

[19] Note – the first interval was 0.9m in order to accommodate the 0.1m initial spacing.

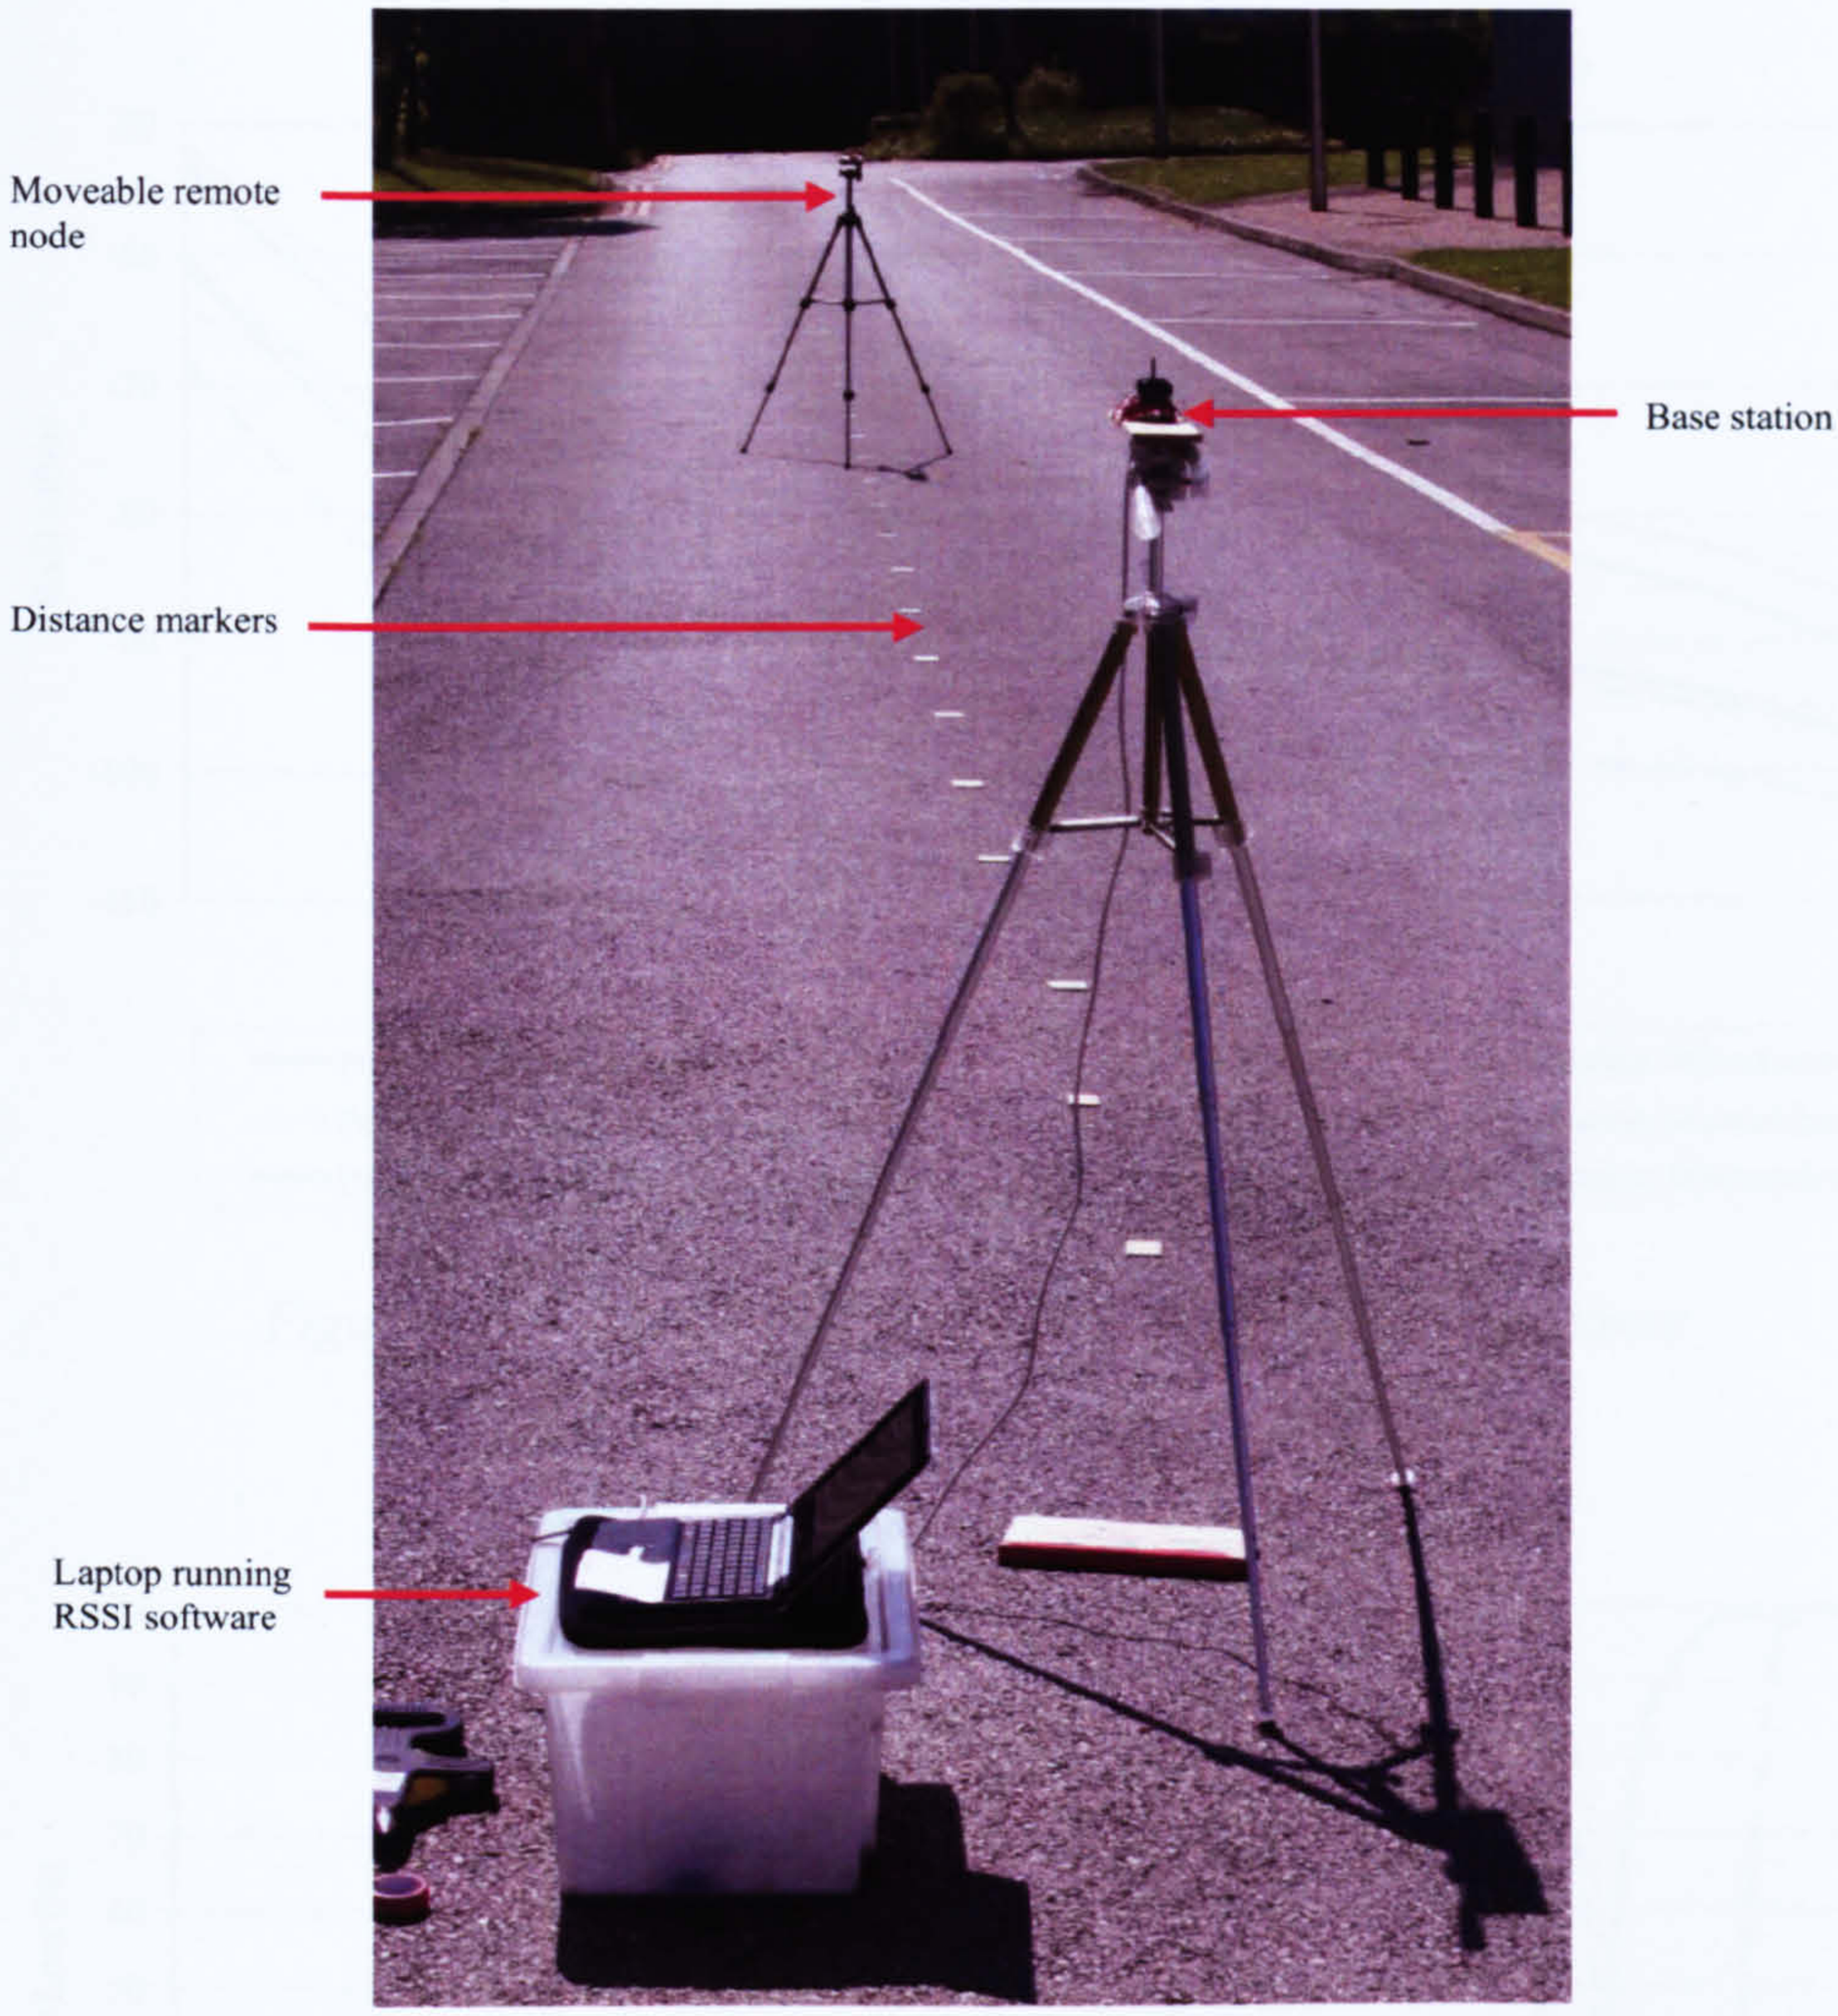

Figure 6.37 – RSSI measurement setup

6.8.2. Measured Results

The results obtained for the standard MicaZ monopole antenna are shown in Figures 6.38 and 6.39, whilst those for the FCPW antenna are shown in Figures 6.40 and 6.41. Each figure includes six sets of results. The associated orientation of each set of results is given in terms of the plane which is parallel to the base station monopole antenna. The orientations are relative to those shown in Figures 6.27 (monopole) and 6.28 (FCPW). For a polarisation mismatch the base station antenna was moved from being vertical to being horizontal with respect to the Earth.

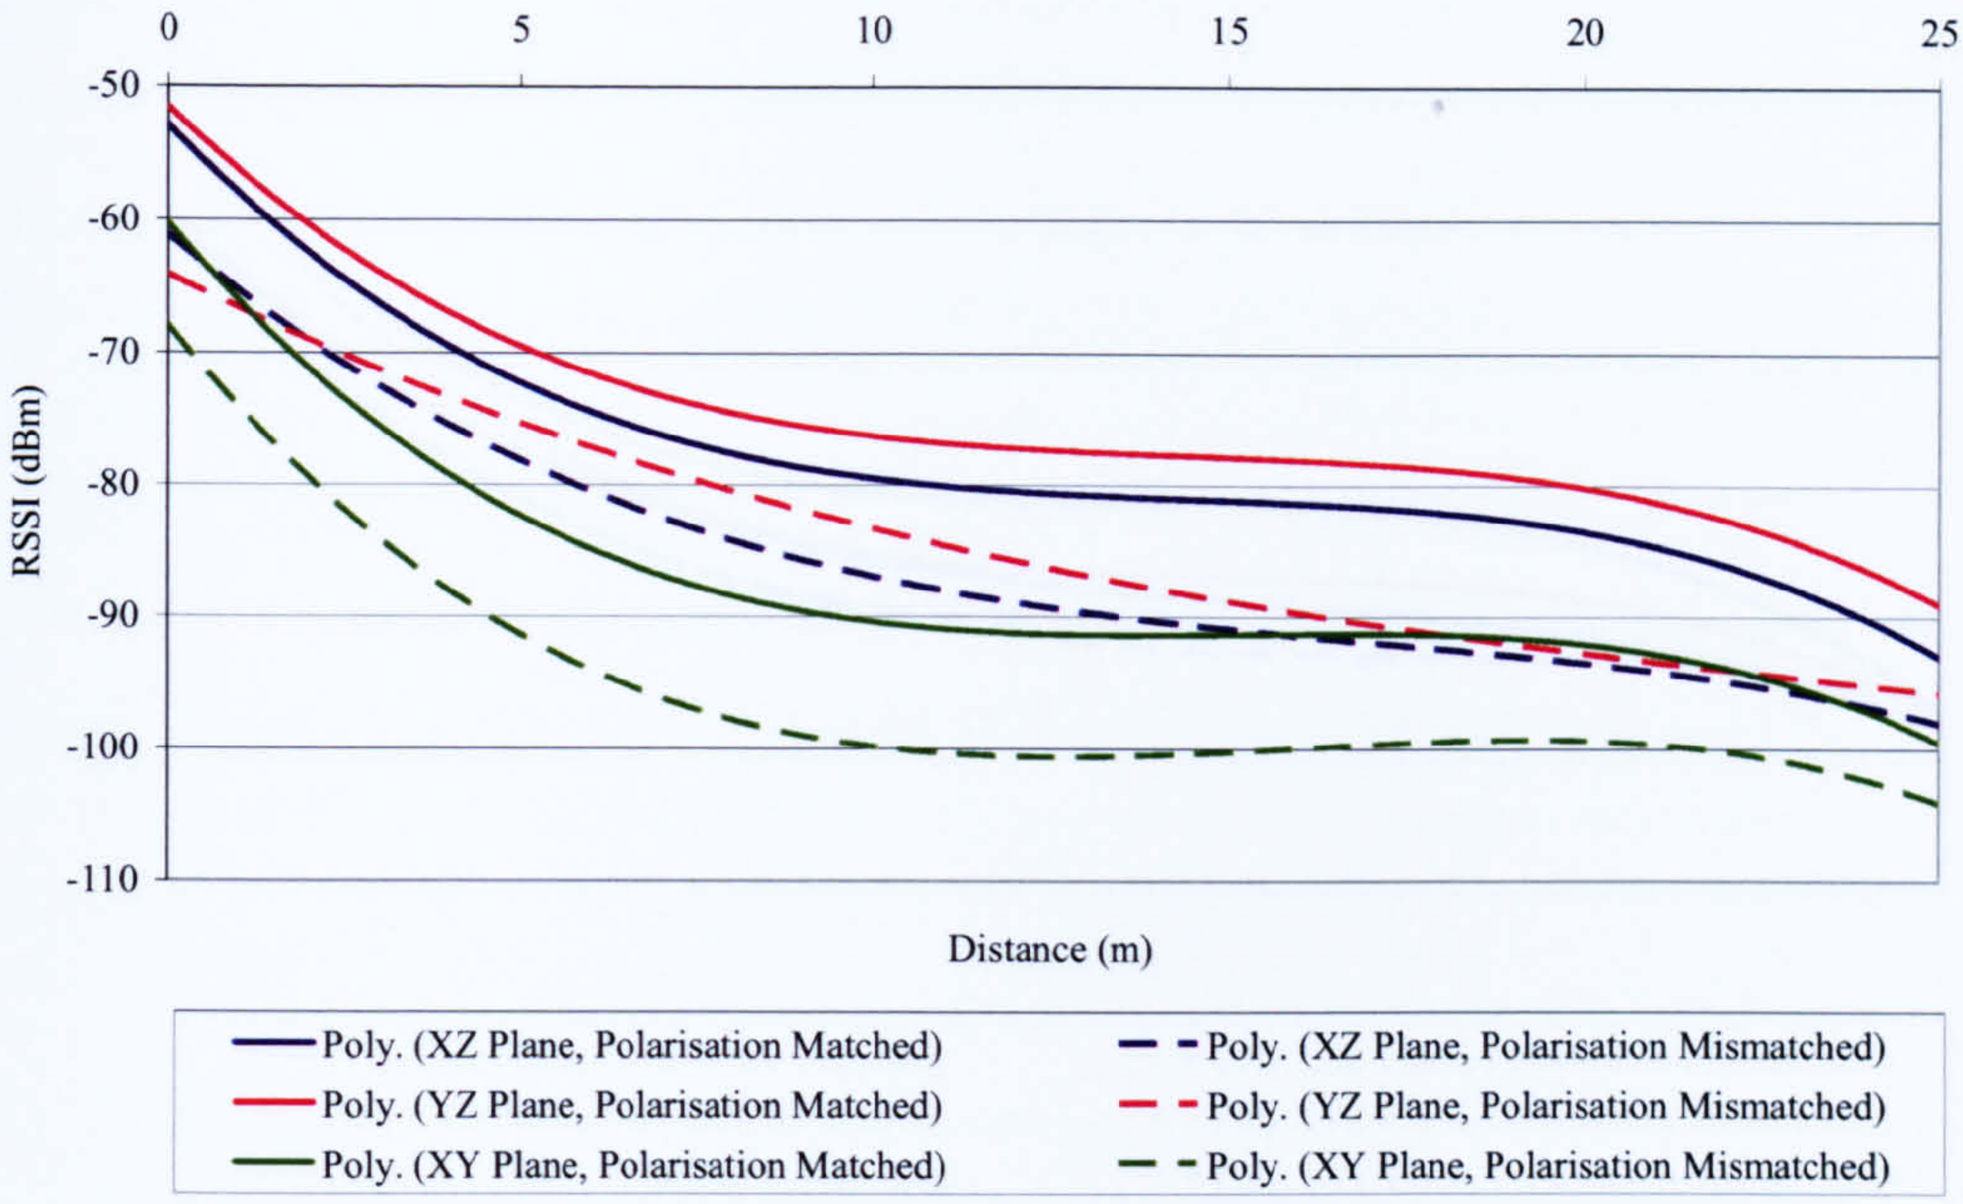

Figure 6.38 – MicaZ monopole RSSI as a function of distance

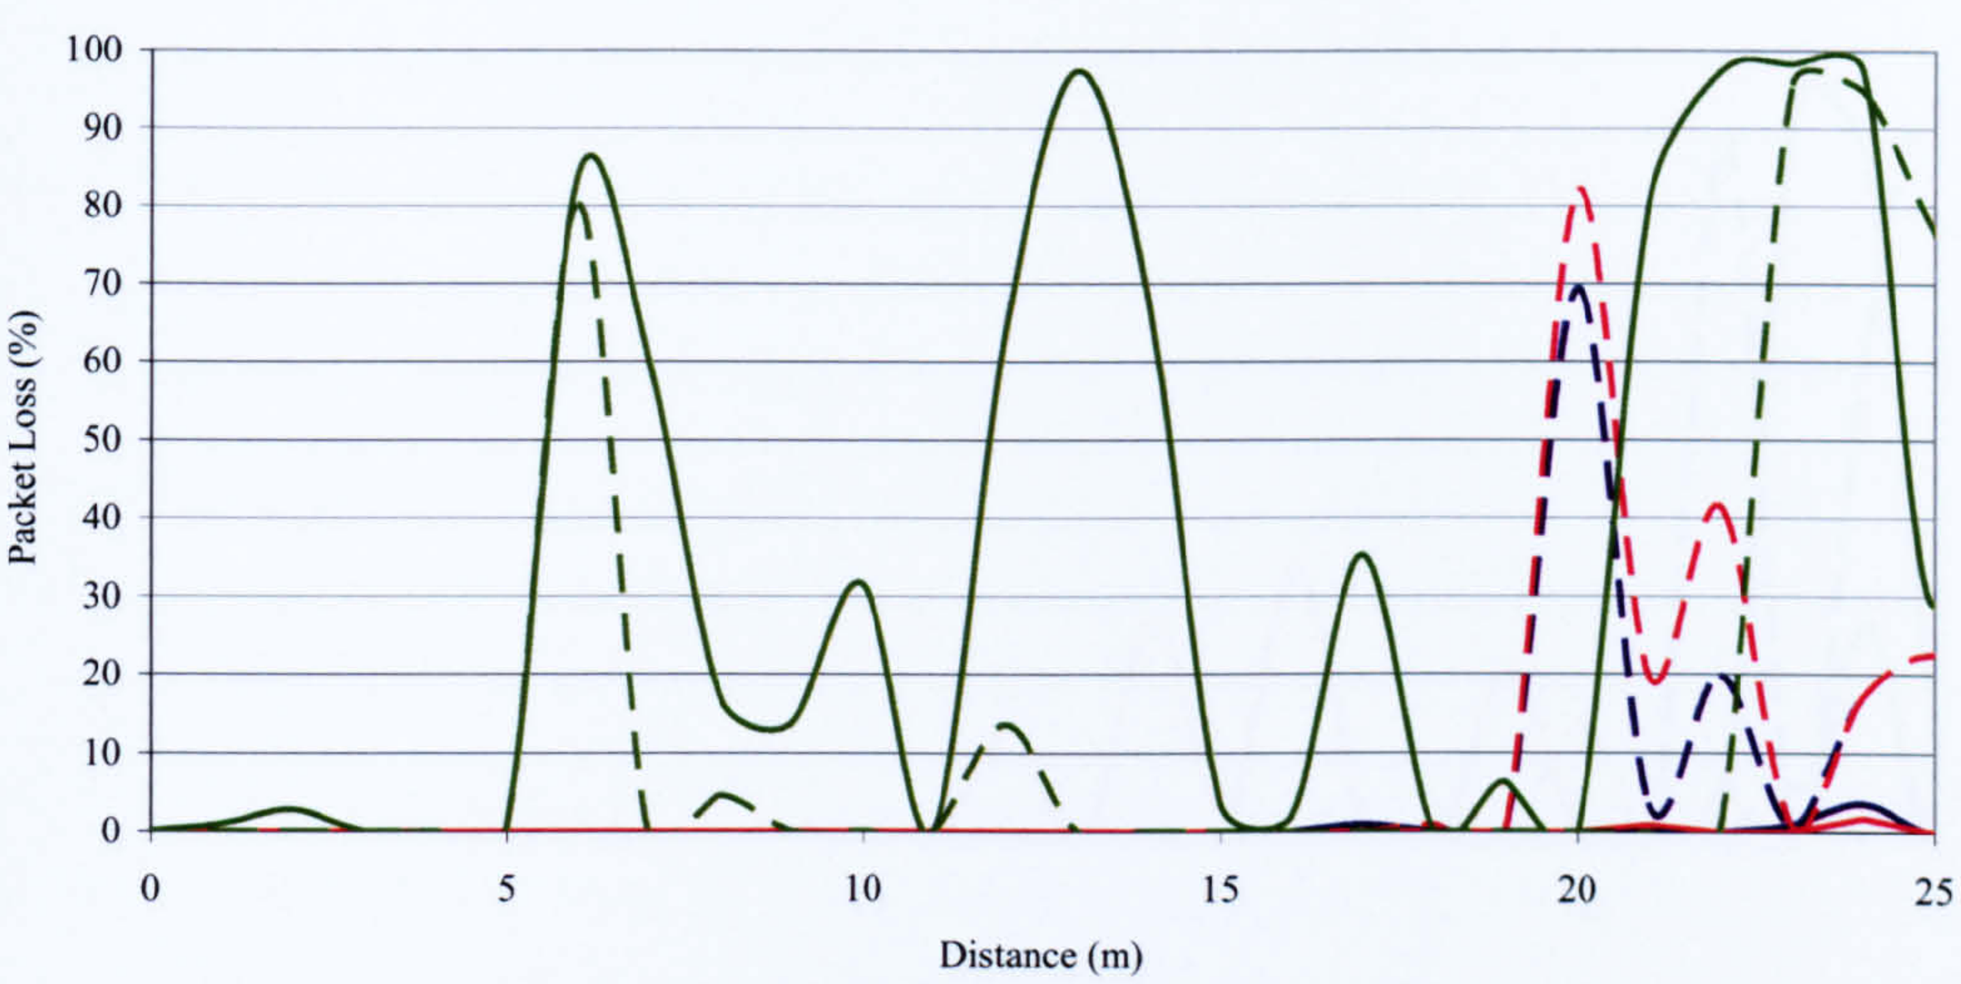

Figure 6.39 – MicaZ monopole packet loss as a function of distance

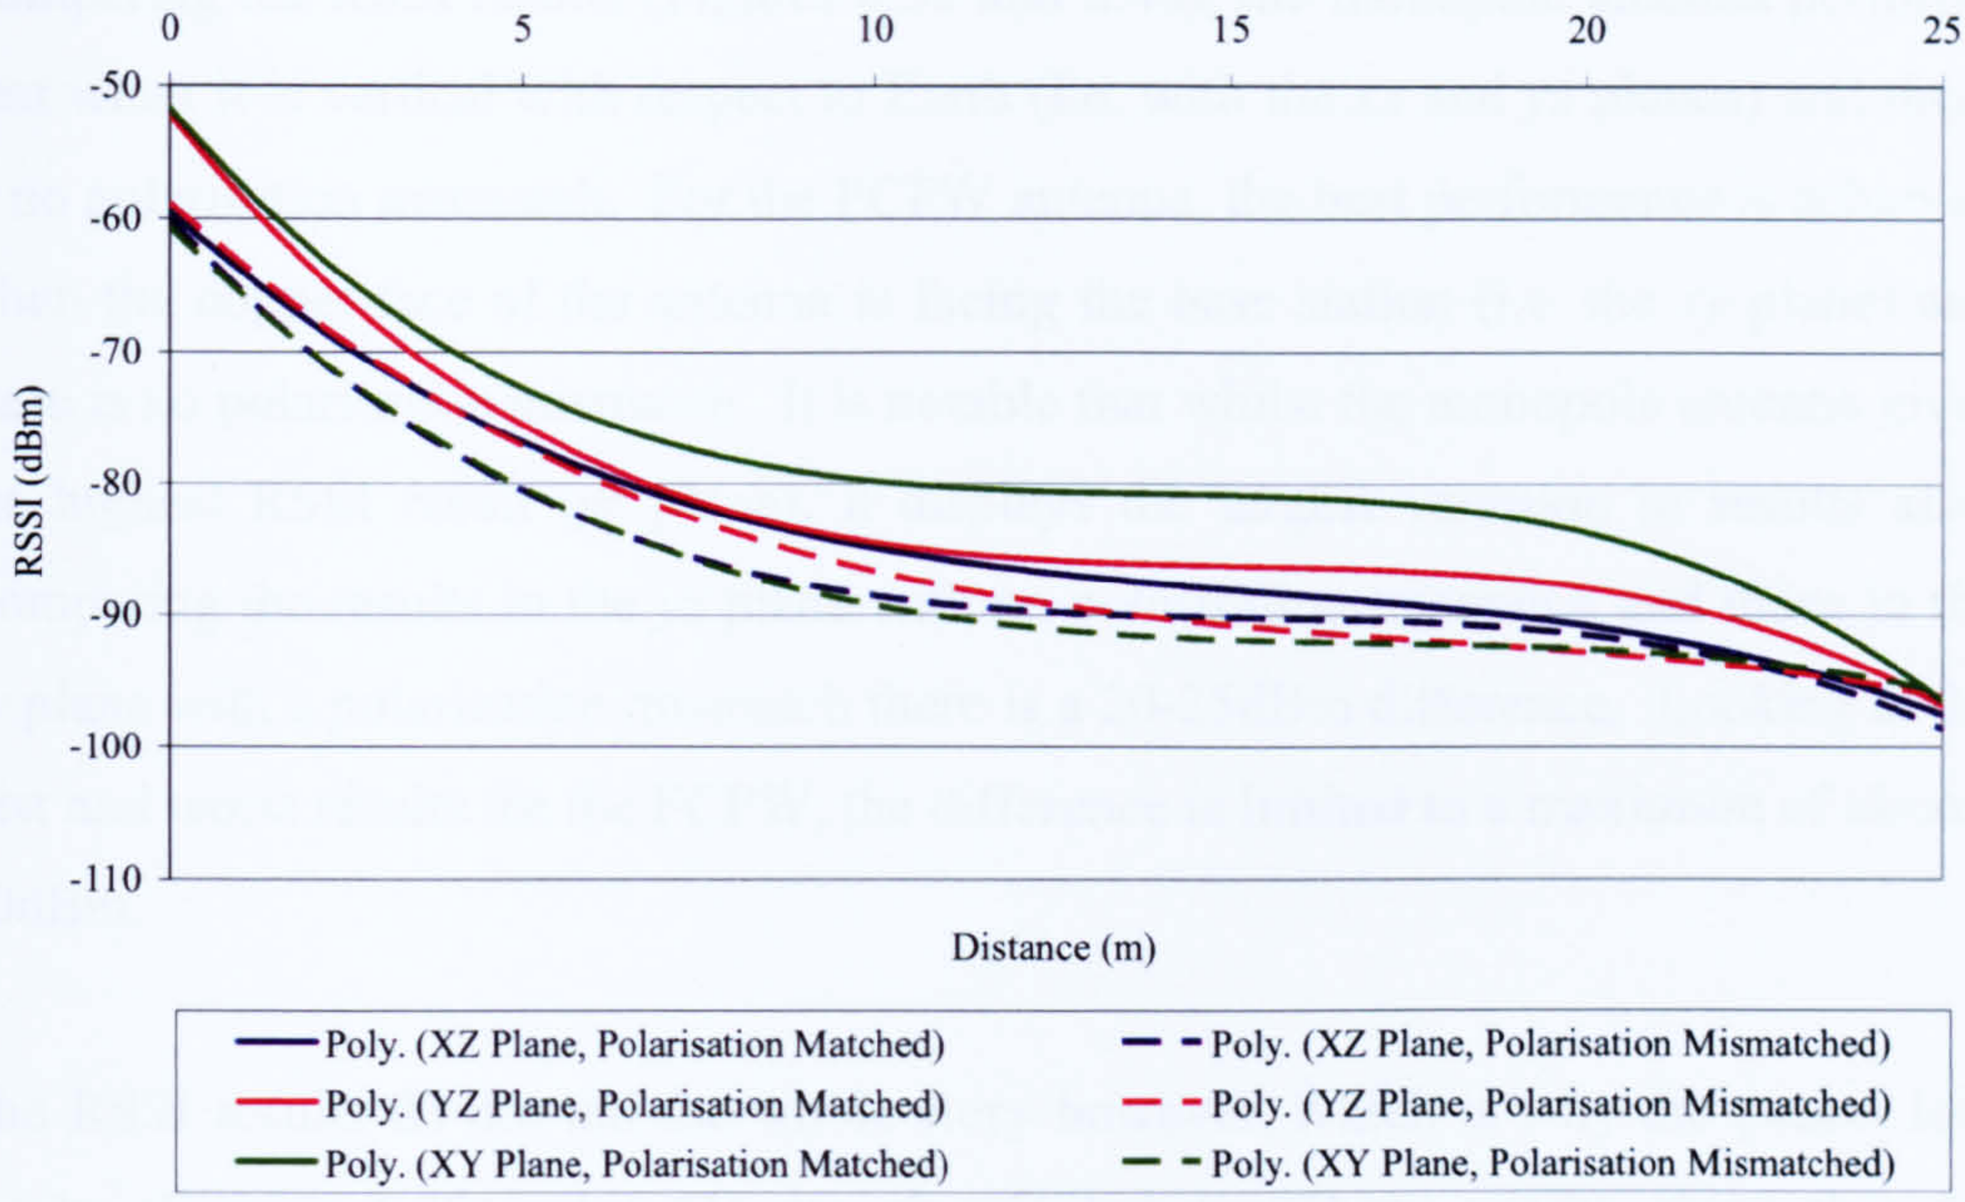

Figure 6.40 – MicaZ FCPW RSSI as a function of distance

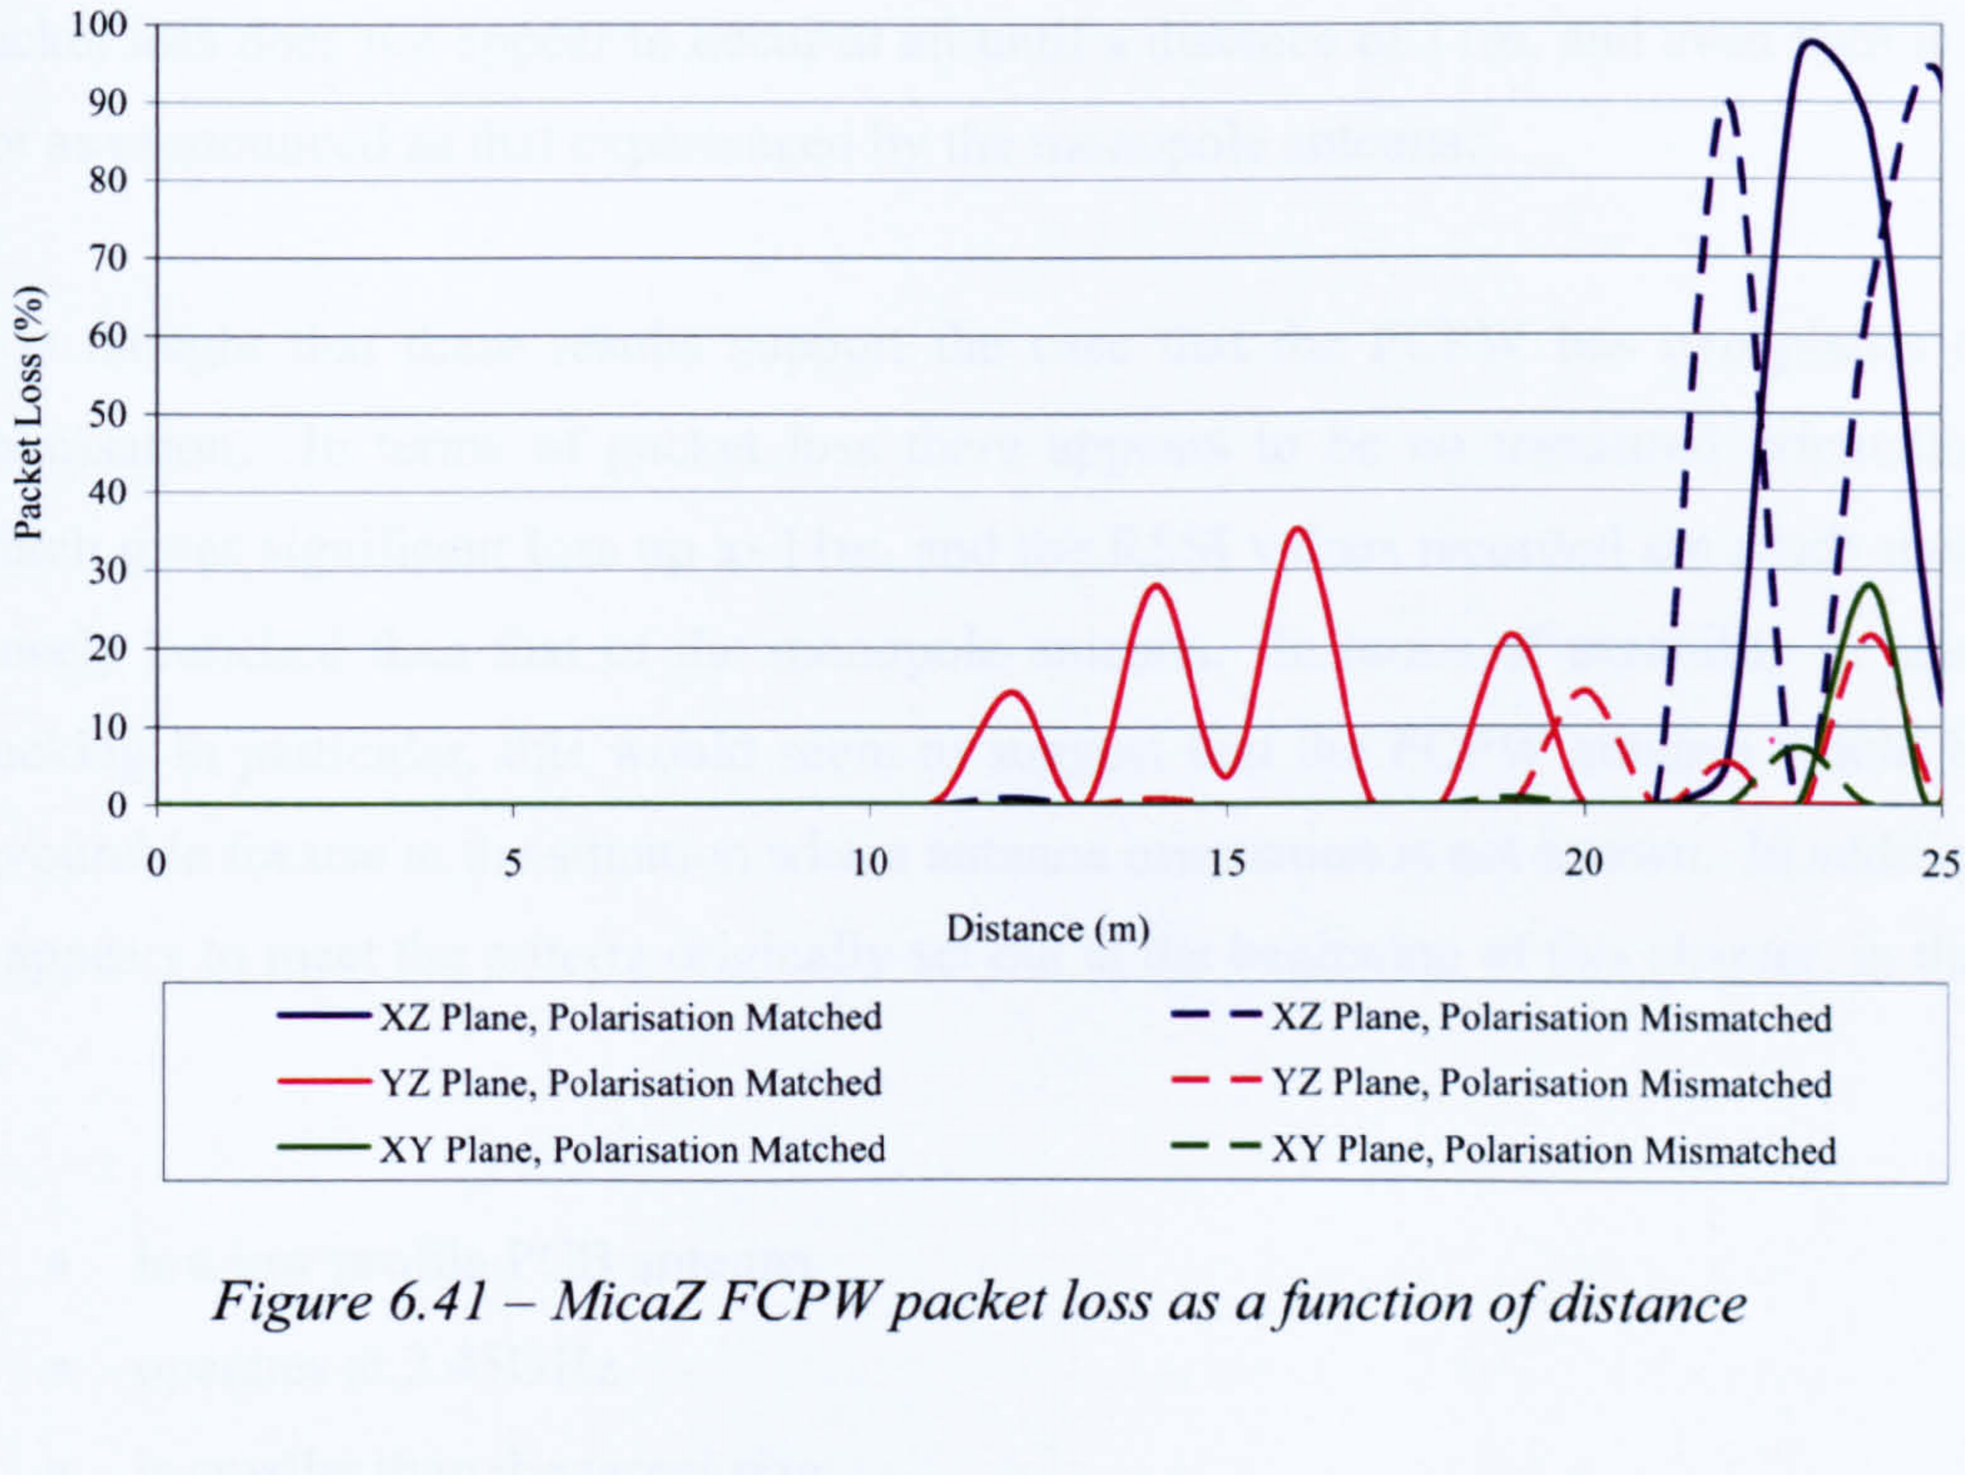

Figure 6.41 – MicaZ FCPW packet loss as a function of distance

Comparing the RSSI results (Figures 6.38 and 6.40), the monopole antenna performs best when it is vertical with respect to Earth (i.e. with the *xz* and *yz* planes) and there is no polarisation mismatch. For the FCPW antenna, the best performance is achieved when the copper face of the antenna is facing the base station (i.e. the *xy* plane) and there is no polarisation mismatch. It is notable that whilst the monopole antenna gives the highest RSSI result (*yz* plane), it displays the largest variation in results also. Comparing the results in the *yz* plane with no polarisation mismatch and those in the *xy* plane with a polarisation mismatch there is a 20-25dBm difference. Looking at the best and worst results for the FCPW, the difference is limited to a maximum of almost 10dBm.

The RSSI results do not tell the whole story however, which is why the packet loss results (Figures 6.39 and 6.41) are also included. These results show that the monopole antenna experiences heavy packet loss when it is orientated in the *xy* plane, even at a distance of just 2m. As the distance increases this packet loss varies greatly, and on numerous occasions peaks at over 50% loss. For the FCPW antenna however, packet loss does not appear to occur at all until a distance of 11m, and even then it is not as pronounced as that experienced by the monopole antenna.

It is thought that these results support the case that the FCPW has two planes of polarisation. In terms of packet loss there appears to be no measured orientation which gives significant loss up to 11m, and the RSSI values recorded are much more closely bunched than that of the monopole antenna. In terms of useability in asset tracking in particular, this would seem to suggest that the FCPW antenna would be favourable for use in the situation where antenna orientation is not known. In addition it appears to meet the criteria originally set out at the beginning of this chapter, in that it:

- is a low profile PCB antenna
- operates at 2.45GHz
- is smaller than the target size
- exceeds its bandwidth requirements.

7. Industrial Applications

7.1. Overview

The purpose of this chapter is to bring together some of the work discussed so far and show how it has been demonstrated to be, or could be, applicable to industry. Each application has its own requirements based upon various factors – in Chapter 2 the different asset life cycles were outlined in both the airline and packaged gas industry. As a result of this, Chapter 4 sought to discuss the creation of a generic prototype system based upon the fact that each application has a number of fundamental similarities. This applies in the case of applicability for industry too, since key factors such as operational life time and operating environment remain similar regardless of application. Therefore, the first part of this chapter considers these topics in detail before moving on to look at the specific industrial applications.

7.2. General Applicability

7.2.1. Battery Life

The battery life of the motes is paramount when they are acting as tags for inventory management purposes. In the automotive industry, this is less of an issue as it is imagined that the devices will be able to source power from wherever they are mounted. That said however, there may be future applications in the automotive industry where power does become an issue, and therefore this data could become applicable.

In order to test the battery life of the motes, a TinyOS application was created which forced the motes to perform a hundred arbitrary calculations, sample the ADC to obtain a battery measurement and then transmit a message. This occurred once every 10 seconds, and the motes actively performed these tasks for approximately 100ms thus representing a duty cycle of 1%. A computer application, shown in Figure 7.1, was then used to record the transmitted data, which included the battery reading.

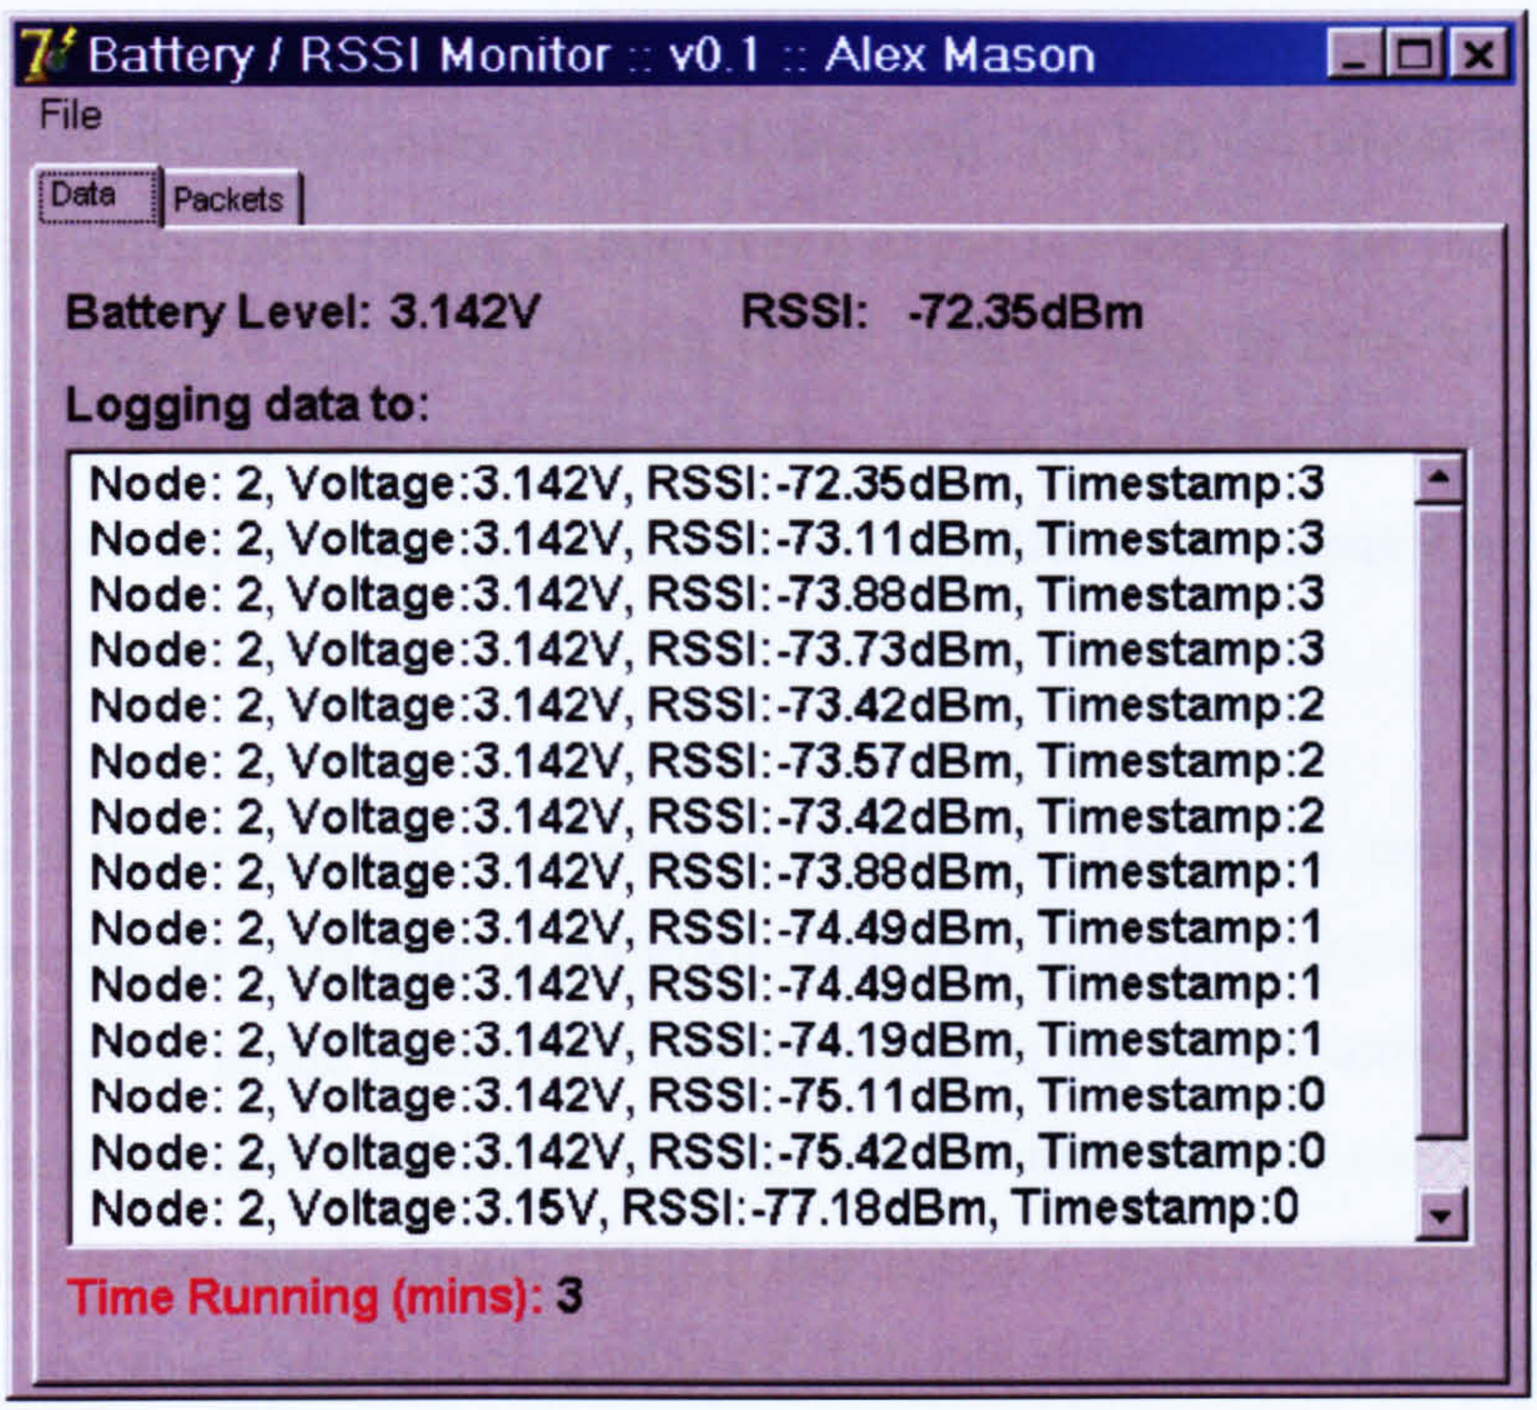

Figure 7.1 – Logging software for battery data

Figure 7.2 – Results showing the batteries running out

In Chapter 4 power management was discussed extensively because it could significantly reduce the consumption of the motes' finite energy resources. To investigate this two motes were employed, but only one had the power management enabled. This experiment ran for a little over 6 days (148 hours) – the time it took for the battery voltage of the mote without power management to drop to 2.5V. It is reported that the mote will operated to 2.1V; the datasheets for the microcontroller and transceivers support this report, however the ADC is not capable of accurately reading voltages this low.

The results of the experiment are shown in Figure 7.2. The motes used two type AA alkaline batteries for each mote (1500mAh capacity). Studying Figure 7.2, there is an obvious difference in the gradient of the two lines; in the time it took the mote with the power management disabled to drop 0.74V, the mote with it enabled drops just 0.09V. This initial result would indicate that the later mote would operate 8 times longer due to power saving being enabled, but this does not take into account the faster decay rate shown between approximately 3.2V and 3.0V when the power management was disabled. Therefore, if the graph is analysed in a linear way and split into two parts (i.e. 3.2-3.0V and 3.0-2.5V), it can be shown that the voltage decay rate is approximately five times greater in the first part than the second. Using this result it is possible to estimate that power saving enables a mote to operate for 254 days from a voltage of 3.2-2.5V. Using the same linear relationship and the assumption that the motes will operate to a voltage of 2.1V it is possible to estimate a maximum lifespan for the motes of 1.24 years which is comparable with the manufacturer's claim of 1.5 years considering the batteries used in these experiments were relatively low capacity. Batteries rated in excess of 2000mAh are easily available, and would increase the mote lifespan significantly; as an example, 2500mAh batteries should be capable of providing over 2 years of operation based on the experimental results obtained in this work.

The temperature of the experiment was not recorded but it is likely that the blips shown in Figure 7.2 at regular intervals are a result of temperature change in the laboratory between day and night. It is apparent from literature [163] that changes in temperature can have a significant effect on battery voltage; higher temperatures make

the energy within a battery more readily accessible resulting in higher voltage, whilst lower temperatures have the opposite effect. By not having the motes in a temperature stable environment, the results obtained here give a realistic rather than ideal figure for mote battery life.

7.2.2. The Effects of Different Operating Environments

Aside from power, the other main concern of this work is how the motes fare in terms of communication in different environments. In industry the assets, and therefore the motes, may be stored either indoors or outdoors. In addition, the storage areas may well be made of metal (e.g. metal storage containers).

To reflect these different storage scenarios a number of tests were conducted indoors and outdoors – close quarter indoor conditions were considered with the use of a 1m wide tunnel and a metallic storage container which is introduced later in this chapter (see Figure 7.14). These tests, as with those considered in Sections 7.2.3 and 7.2.4, are conducted using a set up similar to that shown in Section 6.8.1. The standard mote antenna is used unless otherwise stated. Results are typically taken over a range of 10m in 0.5m increments; the exception to this rule is the metallic chamber which had a maximum internal dimension of a little over 4.5m. LOS conditions are used. The results for the Mica2 and MicaZ motes are shown in Figures 7.3 and 7.4 respectively. Packet loss results are not shown here as the results show no more information than can be obtained from the RSSI data. As in Chapter 6, best fit curves are presented as this makes the graphical information much more legible.

As reported in Chapter 6, the RSSI results in open environments appear to decay with a reasonable amount of predictability such that the general trend is that increasing distance results in decreasing signal strength. Furthermore, the Mica2 and MicaZ appear to display very similarly shaped decay curves.

For the indoor case, the MicaZ mote results show an approximate average of -5dBm difference from those taken outdoors, and the general trend of the RSSI decay with distance is maintained. The Mica2 mote initially follows a similar pattern to this,

however a notable increase in RSSI is shown in the 4-8m range. It is thought that this is a result of reflection and diffraction since the indoor environment contained desks, chairs and other typical laboratory equipment. Since these items are all quite large, it is thought they are more comparable to the Mica2 wavelength than the MicaZ. Although in this case a positive effect is recorded, in other situations the interference resulting might be destructive.

Interestingly the MicaZ fairs better in the tunnel (or enclosed space) than in the typical indoor environments, whilst the Mica2 does not. It is thought that this is again a result of wavelength differences since the tunnel is only 1m wide. The motes were placed in the centre of the tunnel and so it is thought that signals from the Mica2 will be much more affected by the proximity of the tunnel walls than those of the MicaZ.

In the metal chamber a large loss was noted at the 4.5m mark for both motes. This is shown by a sudden drop in the RSSI and appears to be a result of the last reading being taken when the remote node is place almost flush with the metal wall of the chamber. This drop appears to be much more prominent with the Mica2 mote than the MicaZ and gave an early indication that the MicaZ might be favourable in metallic environments.

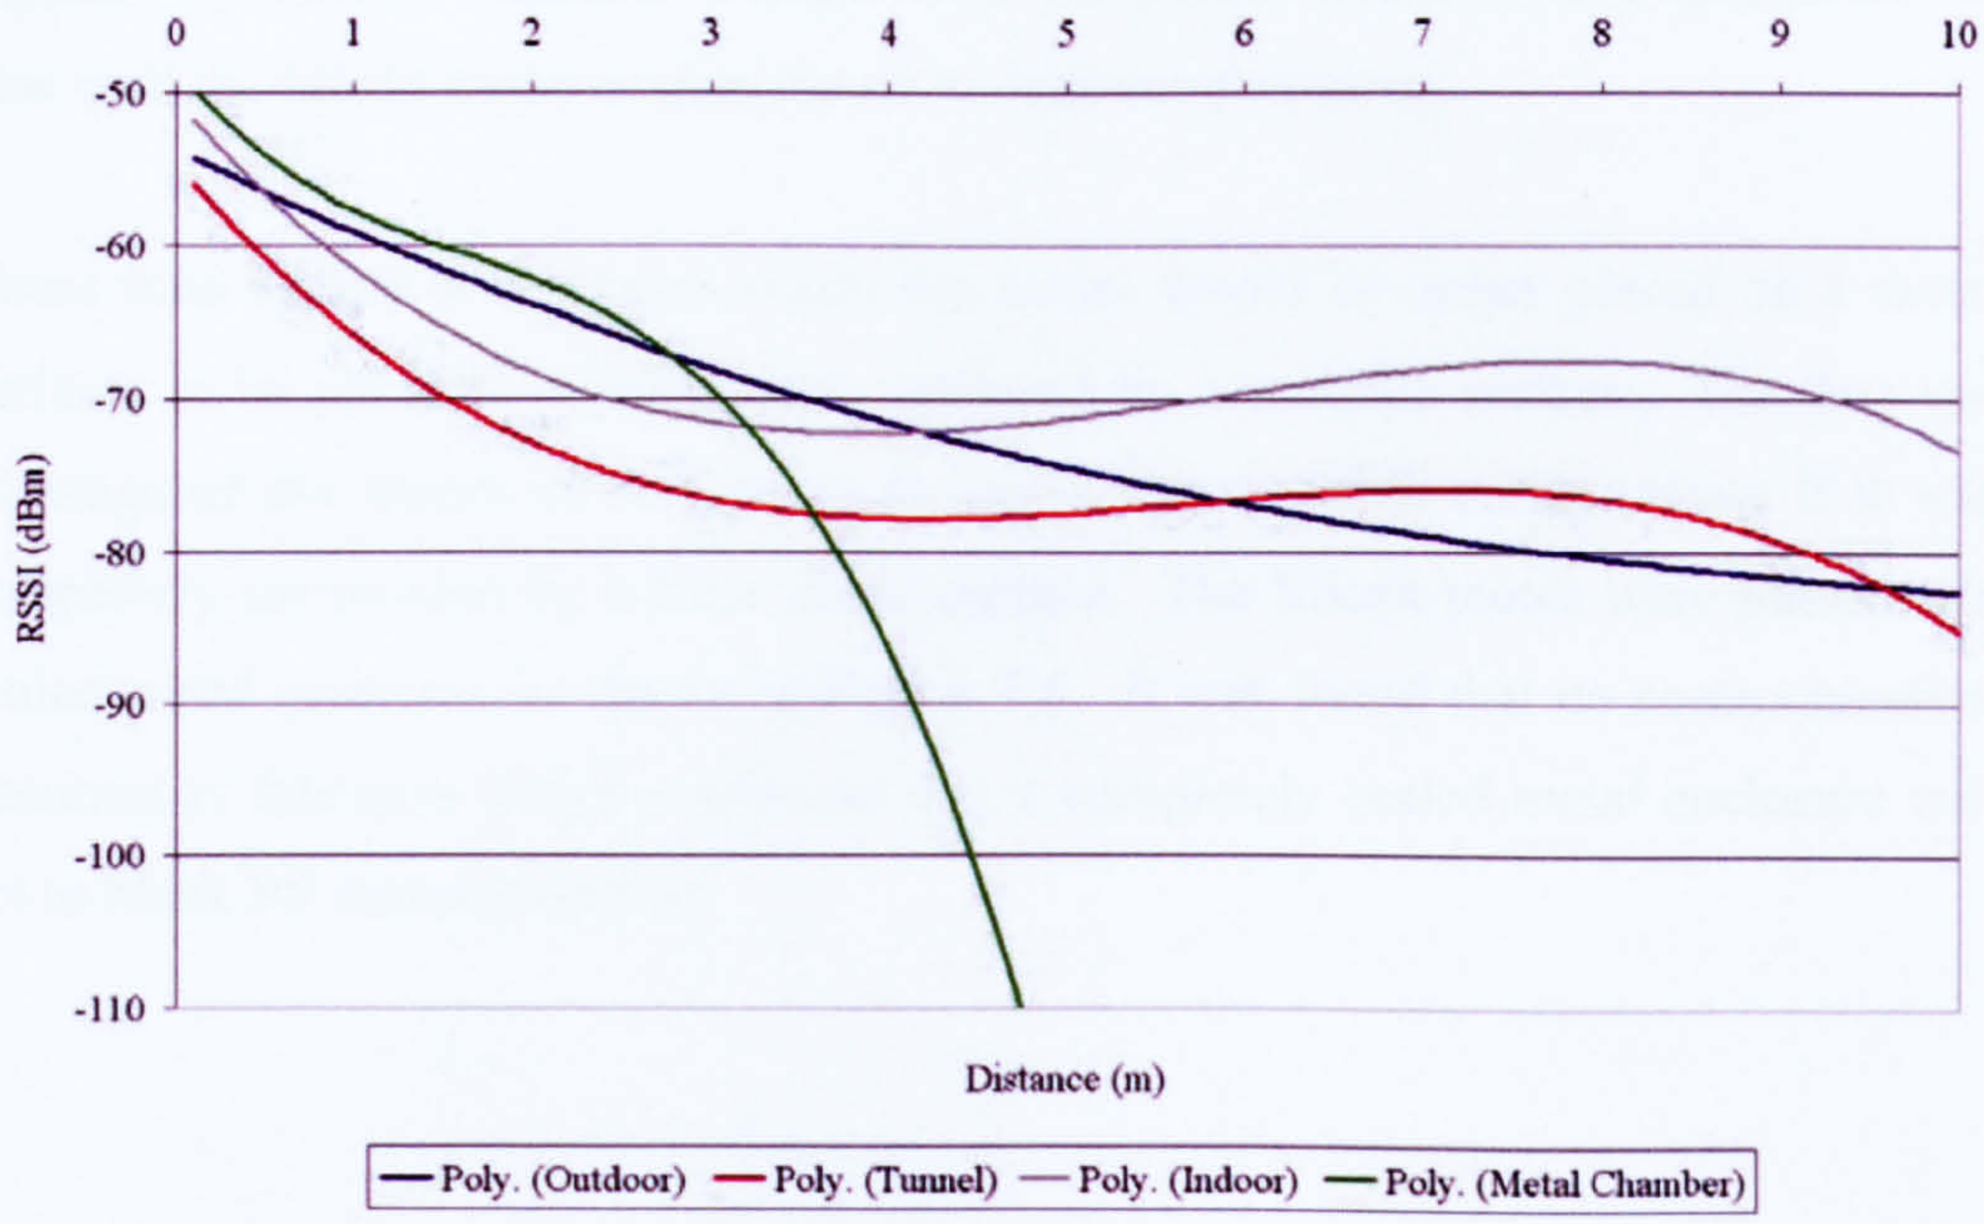

Figure 7.3 – Mica2 RSSI results in different operating environments

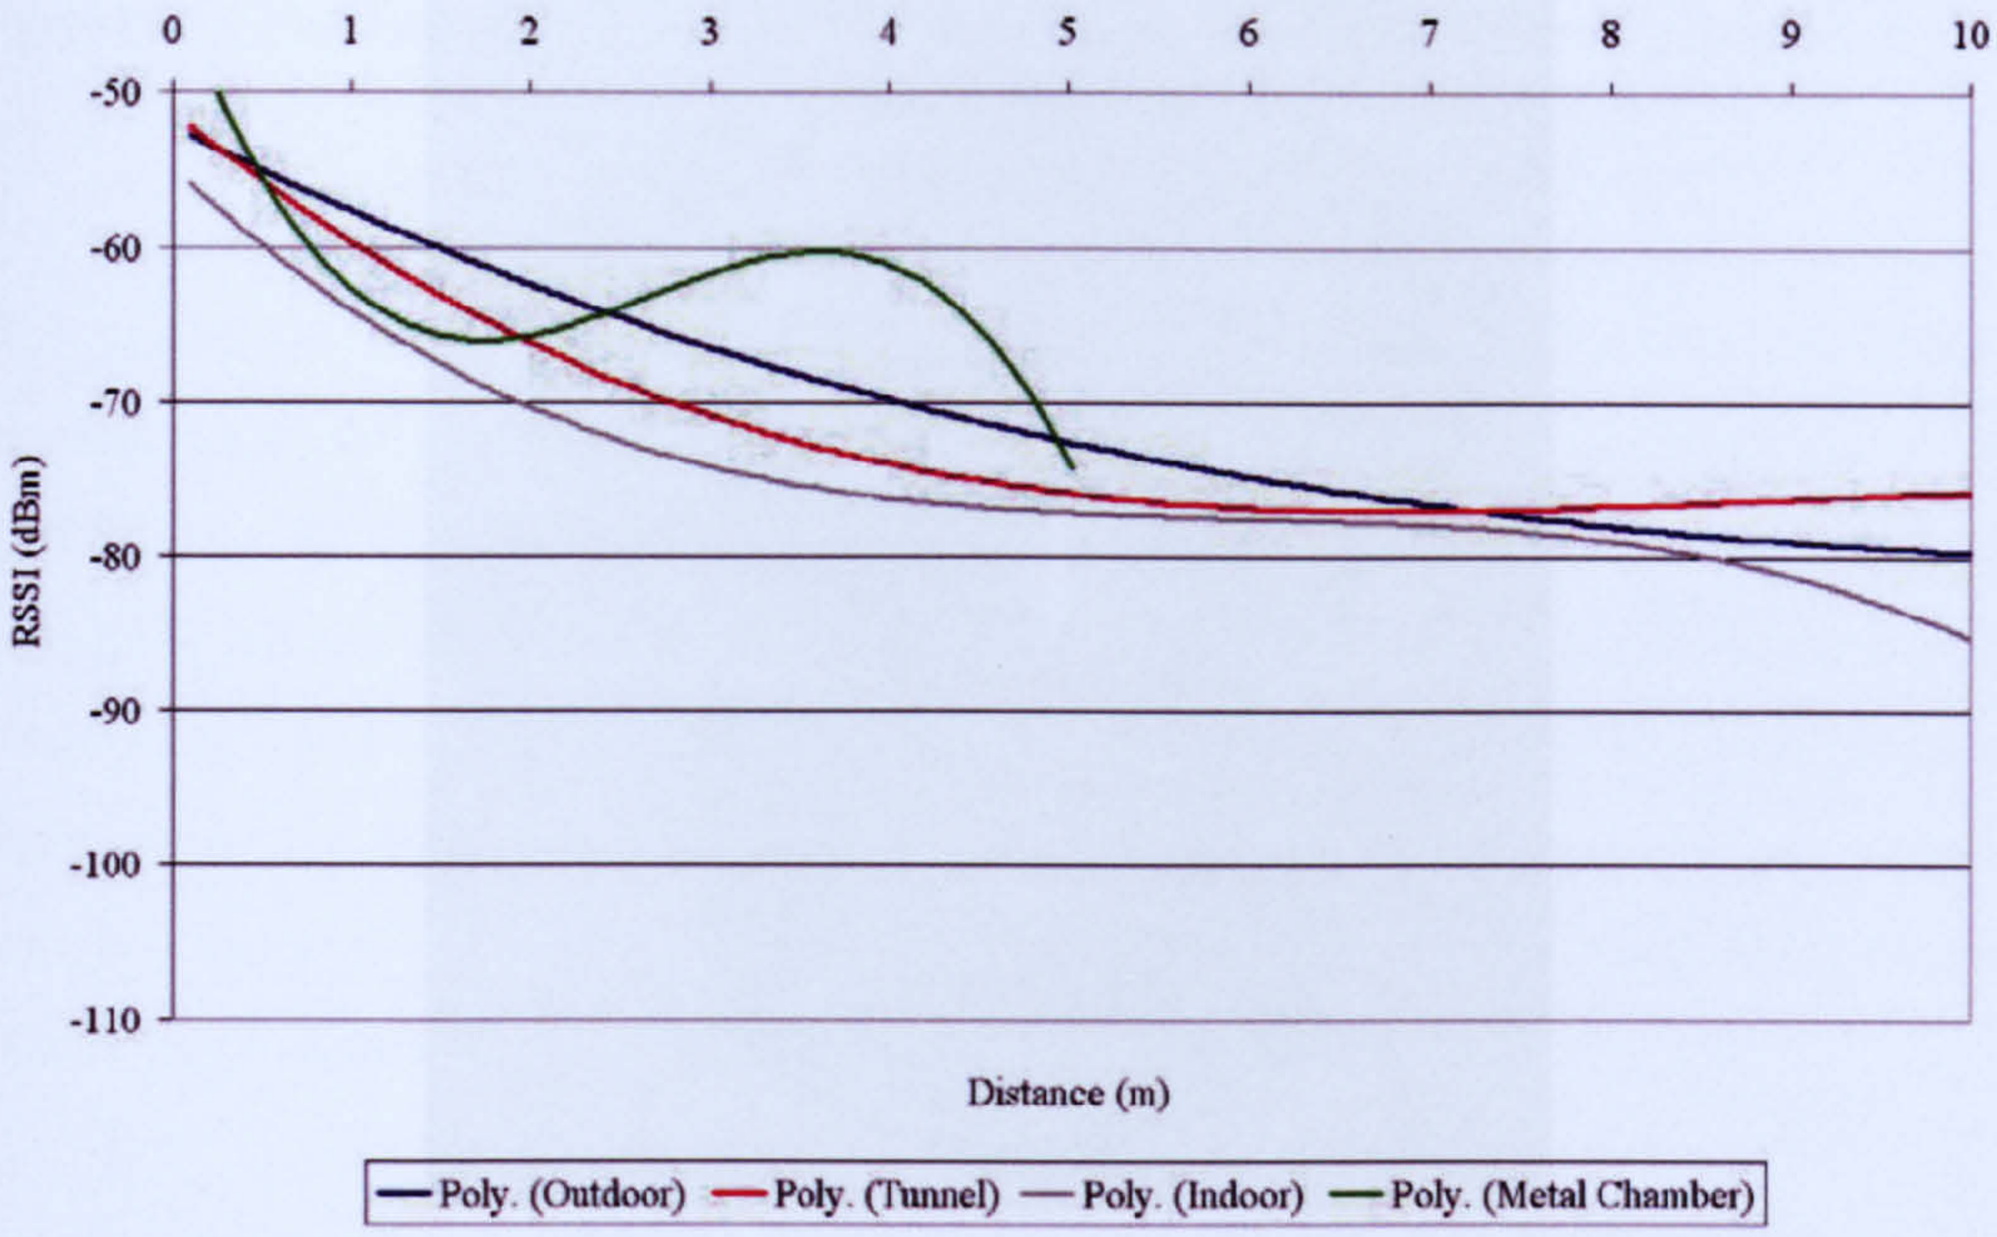

Figure 7.4 – MicaZ RSSI results in different operating environments

7.2.3. The Effect of Metallic Objects

One of the main concerns of this project is in environments which contain metallic objects. Therefore a number of tests were conducted which looked specifically at how well the Micax motes communicate in such environments.

These tests looked at the cases where the motes would be either placed on a metal surface, or be partially or completely enclosed by a metallic surface. The first test investigated the theory of skin depth to see if a mote could communicate if it was completely surrounded by a thick metal surface. The Micax motes were placed in a sealed metal container as shown in Figure 7.5. It was found that no communication occurred in this case which confirmed that a completely sealed metal enclosure will act to block RF communication.

Figure 7.5 – Sealed metal chamber not allowing RF communication

Next, the case of motes being placed on metal surfaces was considered. This is an important scenario since the motes acting as tags are likely to be placed on metallic objects such as luggage and gas cylinders. Also in the automotive industry the motes are likely to be mounted on machinery which is largely constructed of metal. There is no guarantee of how a mote might be arranged in these situations and so conditions where the mote antenna was touching the metal were considered as well as when it was not. Figures 7.6 and 7.7 show results obtained for the Mica2 and MicaZ motes respectively. A reference set of results from the previous section is included which shows how the RSSI of each mote varies under the same conditions when no metal was present. Unless otherwise stated, all the results represent matched polarisation between the antennas.

When the antenna is not touching the metal, it appears that the results for both motes follow a trend similar to that of air; in some cases it appears that the metal sheet has the positive effect of boosting signal strength (by 2-3dBm in some instances). Surprisingly the MicaZ antenna does not behave similarly to the Mica2 when it is touching the metal sheet – the Mica2 suffers in the order 20dBm power loss but with the MicaZ this varies from a 2dBm gain, to a loss of 5dBm. This would appear to

suggest that the higher frequency device (i.e. – the MicaZ) has the upper hand in this set of tests.

One may note from Figure 7.7 that there are some results pertaining to the FCPW antenna discussed in Chapter 6. When the FCPW antenna is not touching the metal its RSSI results are down by 1-2dBm when compared with the results given by the monopole MicaZ antenna. When the FCPW antenna is touching the metal sheet (i.e. – the copper face is touching the metal sheet) an initial 15dBm loss is experienced at short range, but as the range increases the RSSI values do not decay at the normal rate; at a range of 10m there appears to be no difference in performance whether the FCPW antenna touches the metal sheet or not.

Finally, Figures 7.6 and 7.7 display results taken from the MicaZ mote when it was placed in a metal tin (see Appendix F). It was thought that the motes in an asset tracking scenario will not just be placed on metal, they will be surrounded by it so how does this affect communication? To answer this question results were taken with the lid on and off. With the lid on, the Mica2 did not communicate at all. For the MicaZ however communication was possible over a range of 2.5m, albeit with the RSSI being significantly reduced. It is suspected that communication is possible in the case of the MicaZ simple because of the shorter wavelength. The tin lid was not sealed and so there will be small gaps out of which some EM waves could radiate.

When the tin lid was removed, both motes could communicate over the full 10m range, although the Mica2 suffers a 15dBm loss when compared with results in air whilst the MicaZ suffers a 10dBm loss. In addition, the difference between the air and metal tin results appear to diminish with distance in the case of the MicaZ – this is not the case with the Mica2.

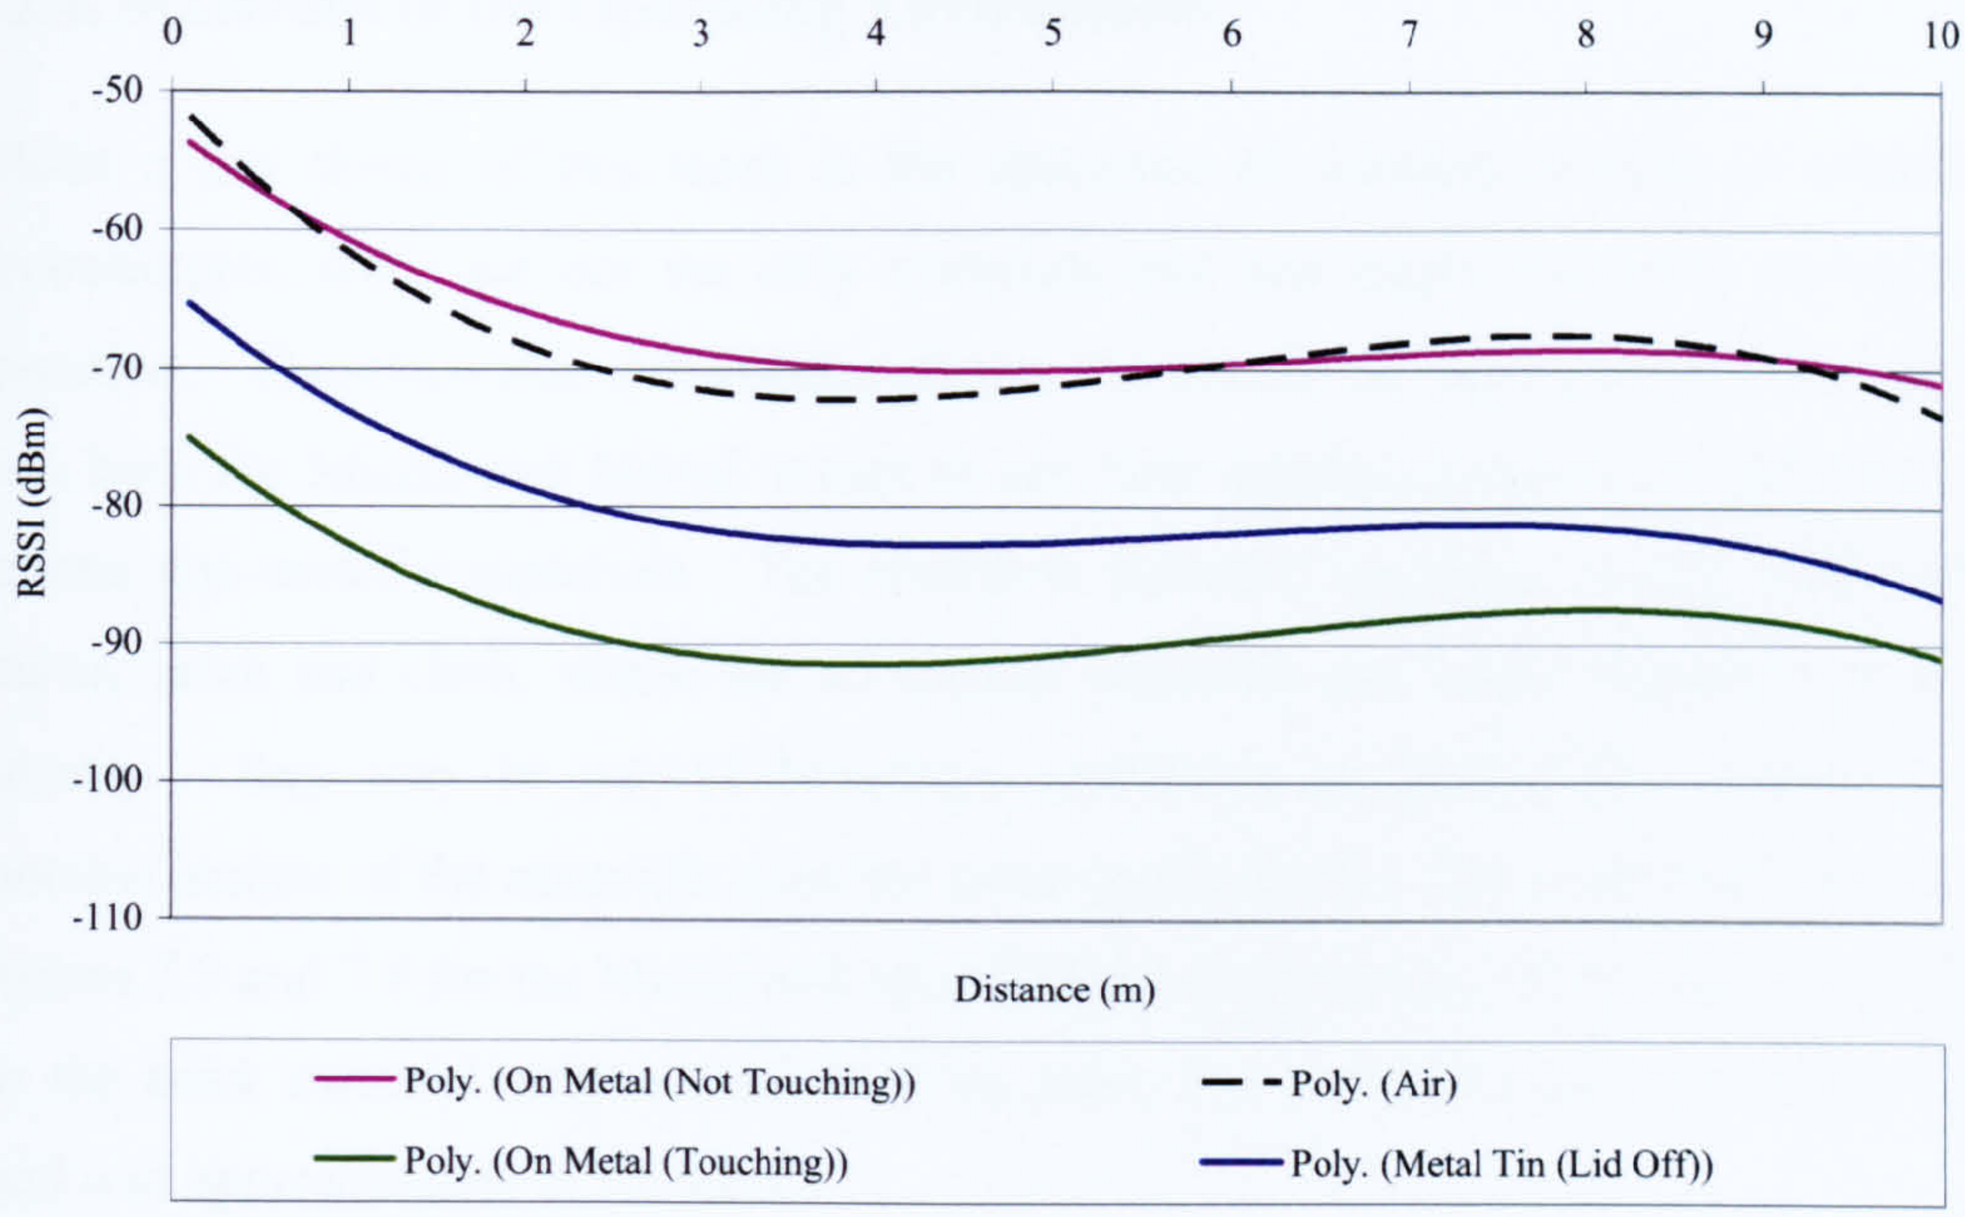

Figure 7.6 – Mica2 RSSI results in metallic scenarios

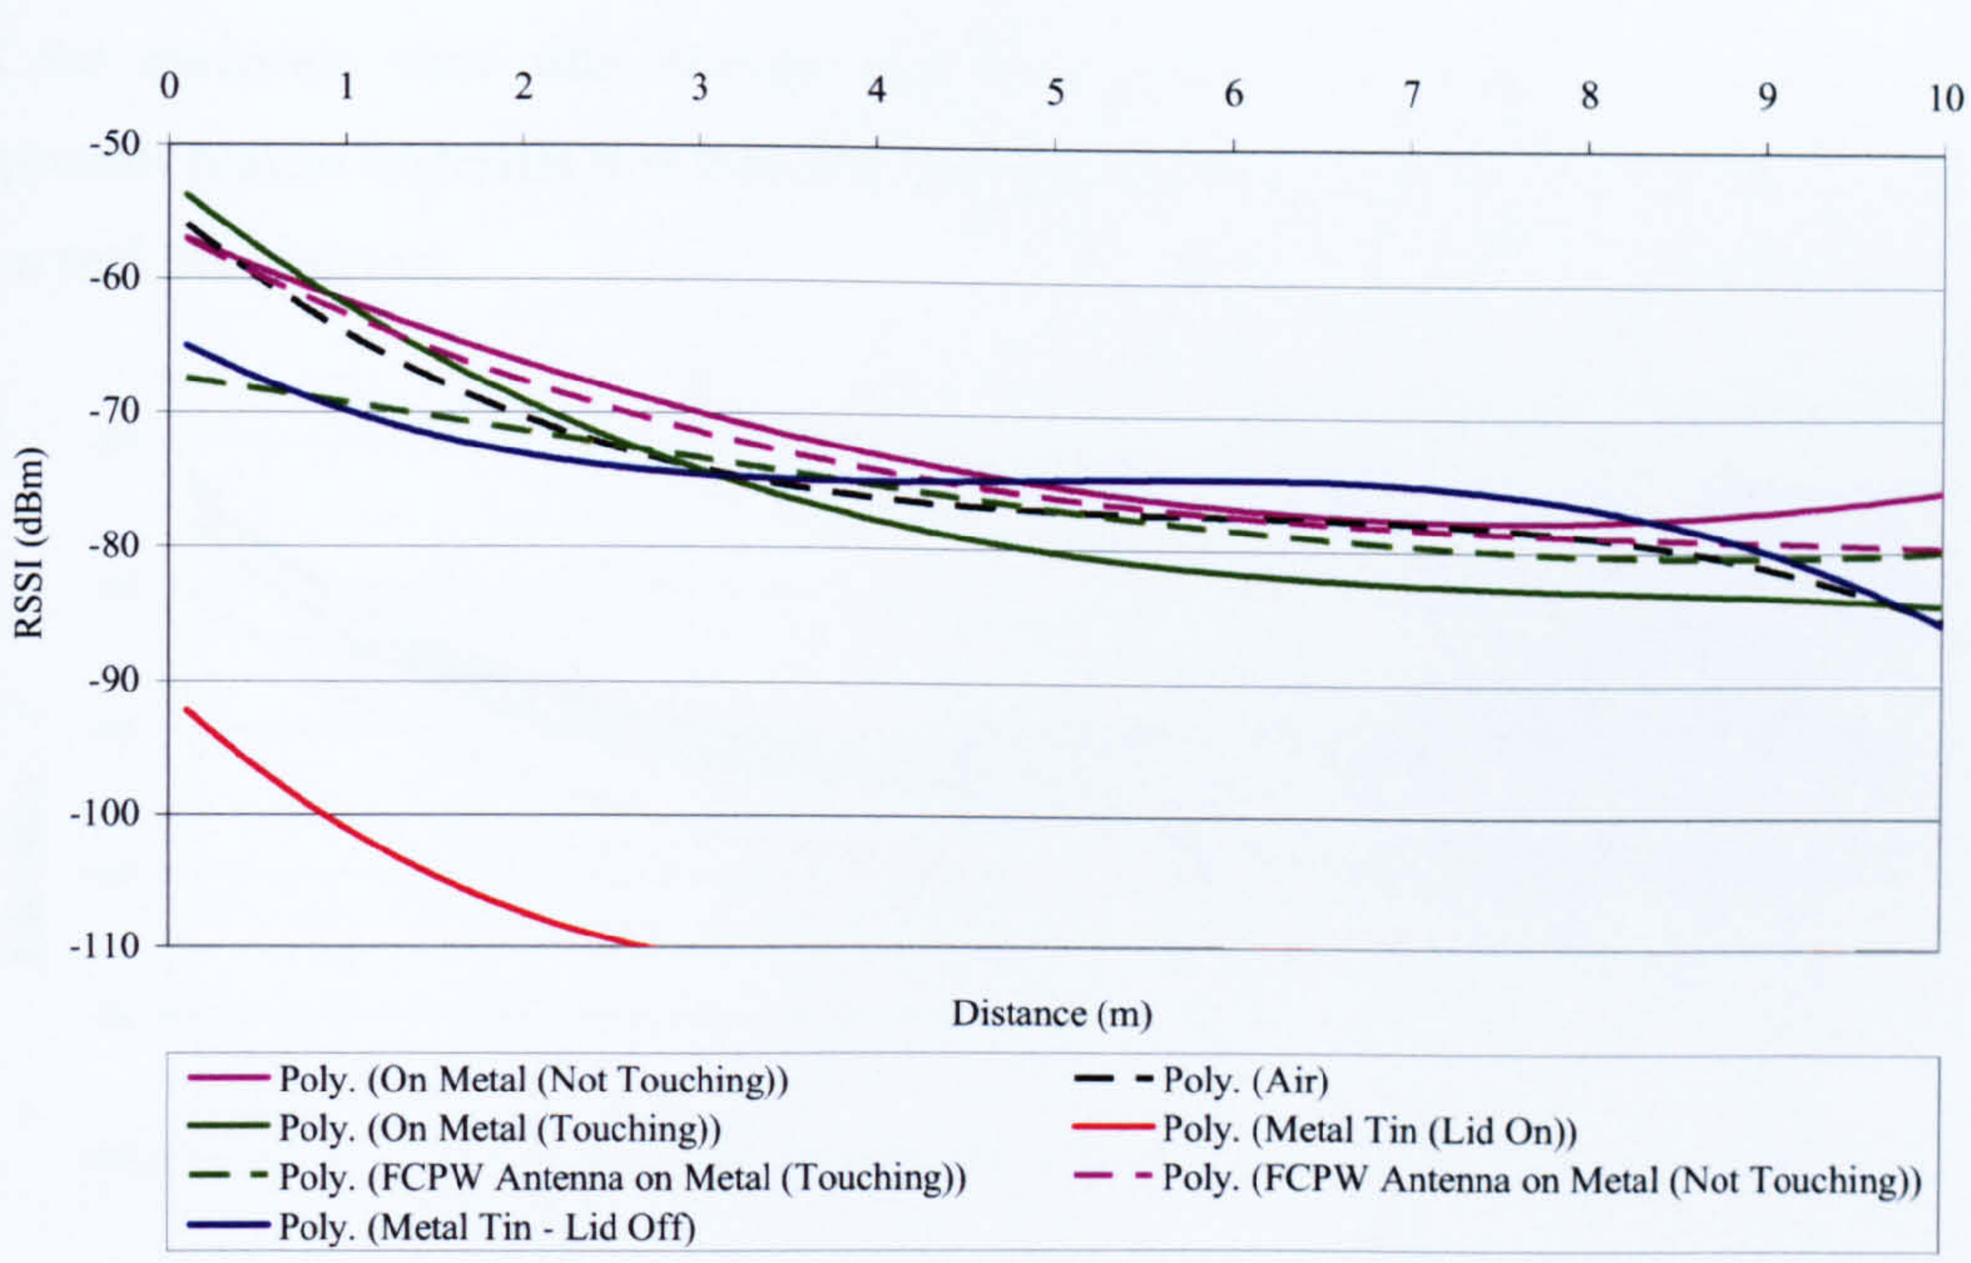

Figure 7.7 – MicaZ RSSI results in metallic scenarios

7.2.4. Materials in the Operating Environment

Whilst a key theme of this work is the operation of wireless devices in metallic environments, these are not the only materials that one might encounter in typical operation. Therefore this sub-section shows the results of experiments undertaken with both the Mica2 and MicaZ motes to see how communications are affected by various non-metallic materials. The materials included are glass, wood, cardboard, plastic, brick and cloth, which are all typical materials one might expect to find in industry – they may be part of buildings, machinery or packaging. Appendix F contains images of the materials used for these experiments. The results are shown in Figures 7.8 and 7.9 for the Mica2 and MicaZ motes respectively. Note that the results for the brick material were started at 0.5m rather than 0.1m because the brick wall used was approximately 0.3m thick.

It is apparent that the air medium does not give the optimum results in terms of RSSI, however it is likely this is due to variations in the sampled signal strength. As a result it is difficult to comment on minor differences experienced. It is possible that many of the materials were thin enough that their effect was negligible – since they represent typical materials it is possible to assume that their effect would be negligible in a real life situation.

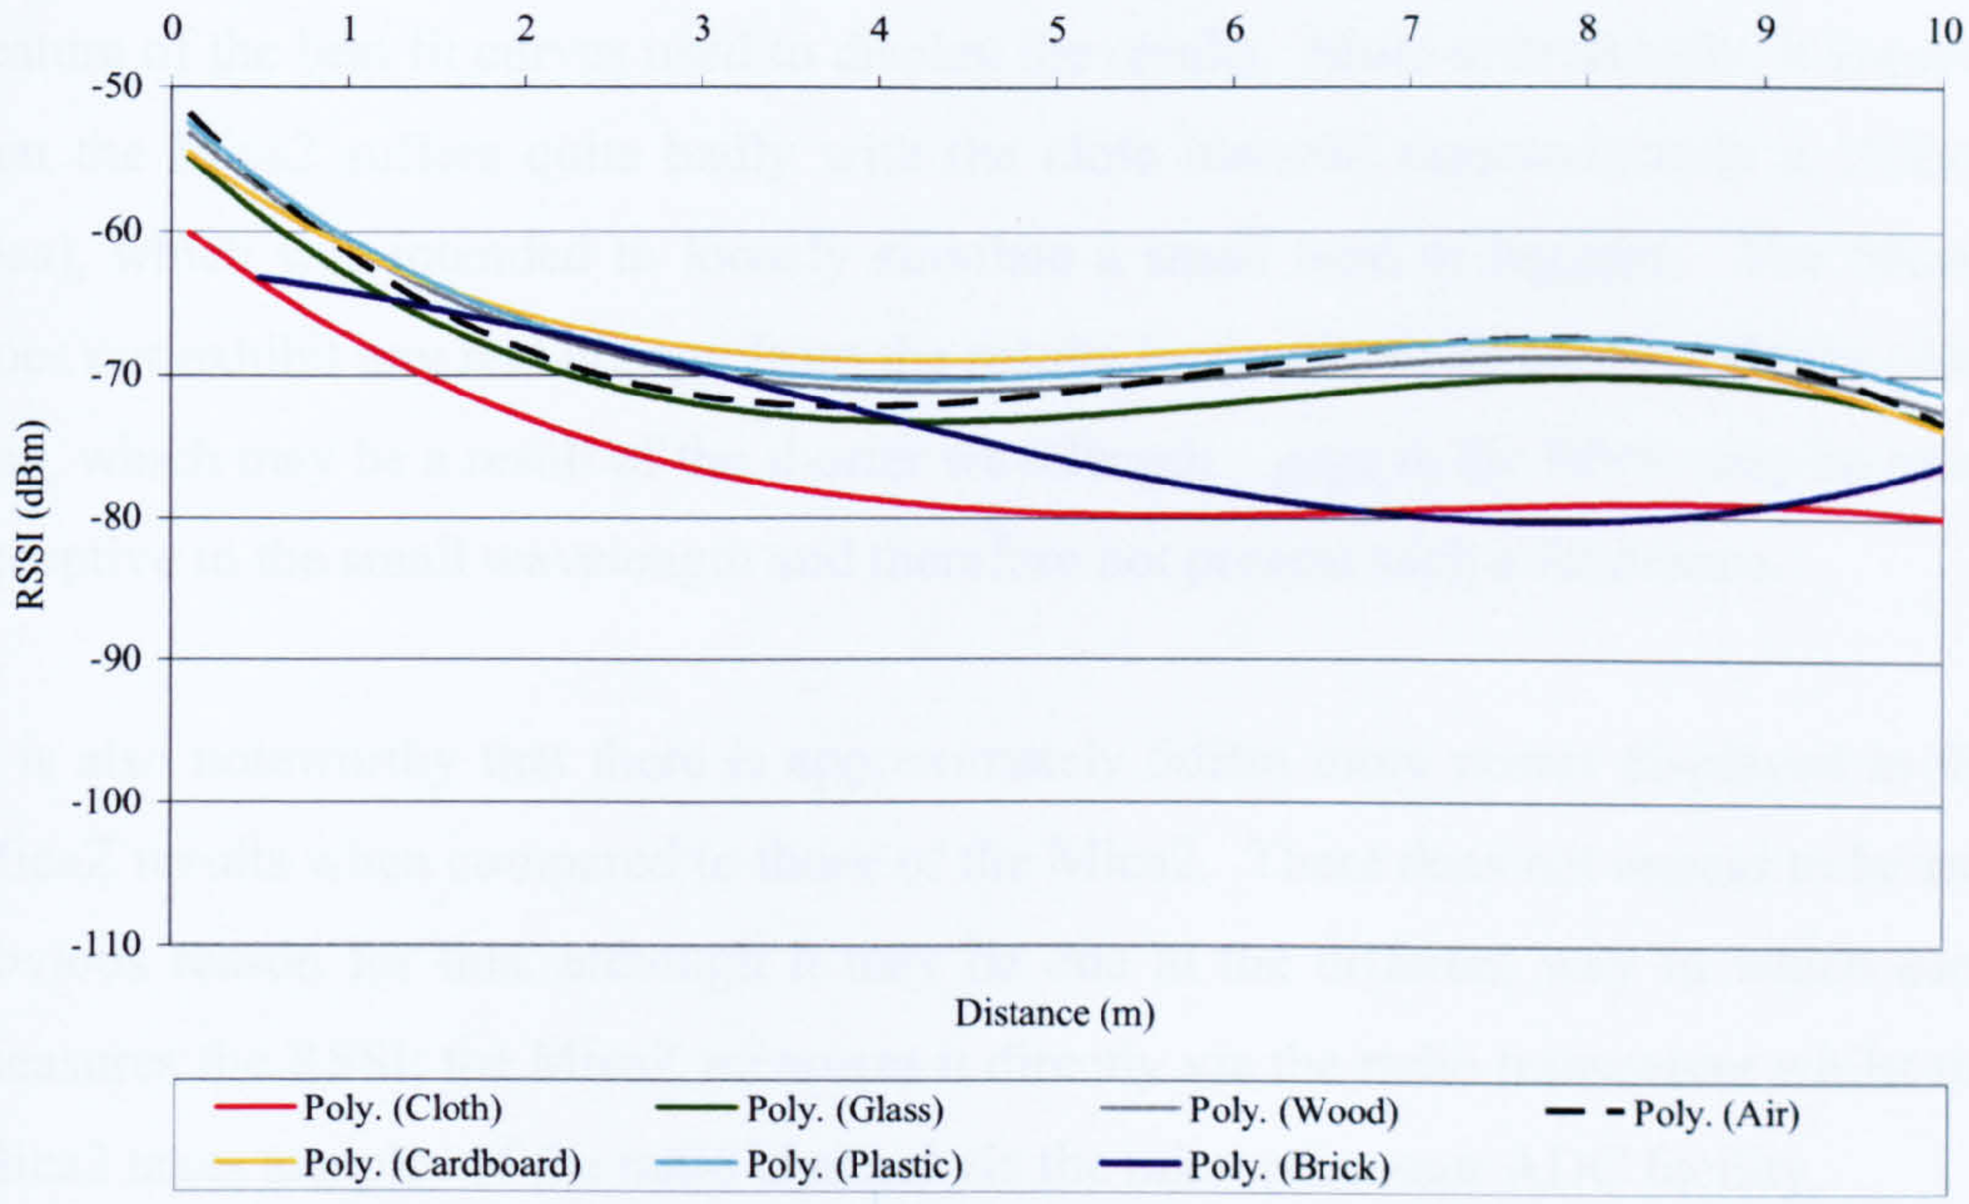

Figure 7.8 – Mica2 RSSI results with non-metallic materials

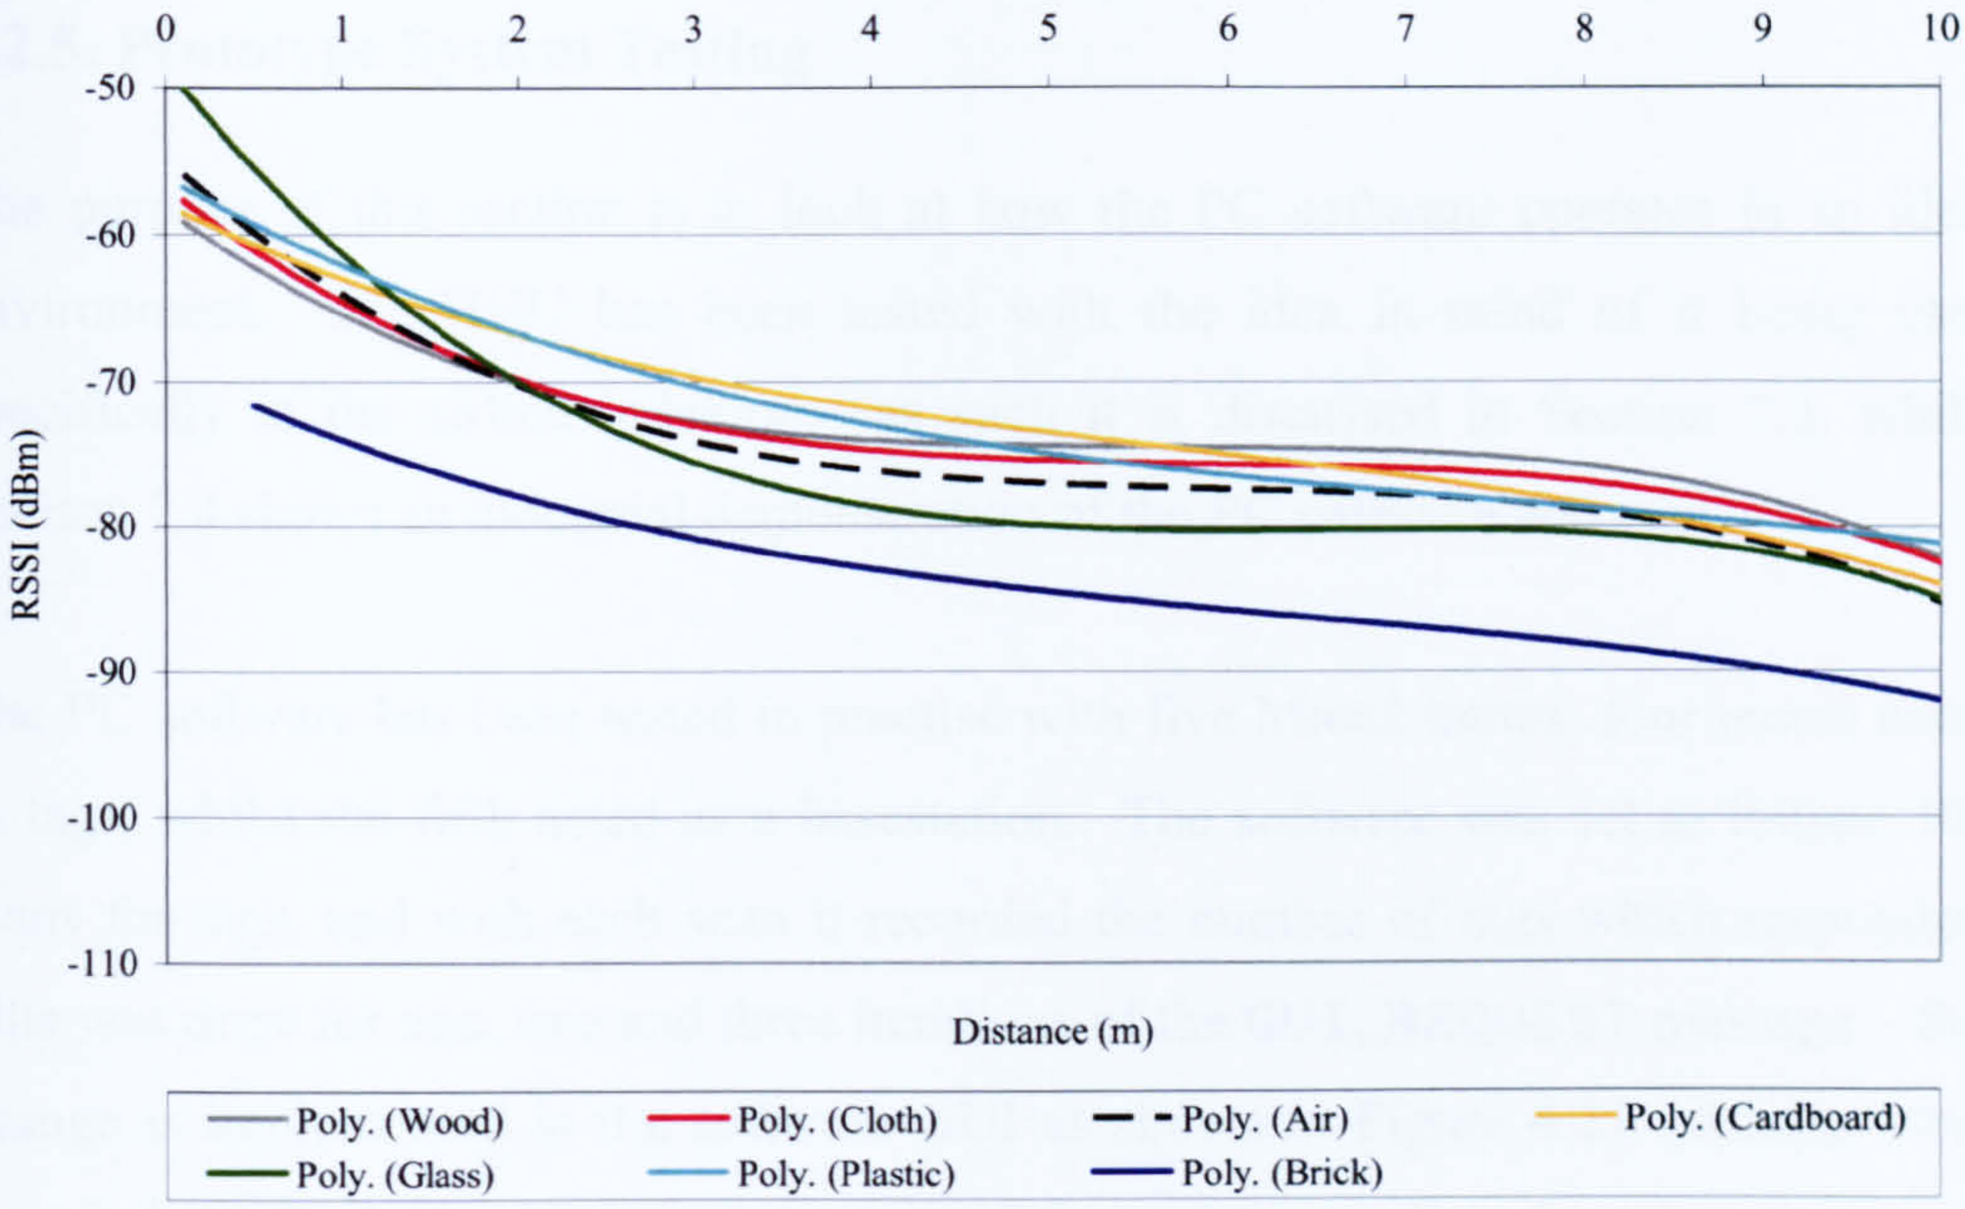

Figure 7.9 – MicaZ RSSI results with non-metallic materials

For the brick wall both the Mica2 and MicaZ showed a drop in RSSI which was relatively equal. It appears that the loss experienced by the MicaZ is approximately 5dBm more than that experienced by the Mica2. It is thought that this is due to absorption (see Section 5.4.2) having more of an effect as frequency increases – Equations 5-4 and 5-5 support this idea. The Mica2 appears to display an increase in signal strength between 8m and 10m with the brick wall which is thought to be a feature of the best fit curves used to display the results. More interestingly, it appears that the Mica2 suffers quite badly with the cloth material (approximately a 10dBm loss), which was intended to loosely simulate a small item of luggage. The MicaZ does not exhibit any real change from the results in air when enclosed inside the cloth bag, which may be a result of the shorter wavelength – gaps in the fabric may be more receptive to the small wavelength and therefore not present such a hindrance.

It is also noteworthy that there is approximately 5dBm more power displayed in the MicaZ results when compared to those of the Mica2. There does not appear to be any obvious reason for this, although it may be due to the different way in which each measures the RSSI; the MicaZ measures it directly via the radio transceiver whilst the Mica2 takes samples of the radio channel via the microprocessor ADC facility.

7.2.5. Prototype System Testing

The purpose of this section is to look at how the PC software operates in an ideal environment. The HHU has been tested with the idea in mind of it being used specifically in the airline industry – as such it is discussed in Section 7.3, whilst Section 7.4 shows an industrial demonstration of the PC software and tags.

The PC software has been tested in practise with five Mica2 motes; four motes acted as tags, whilst the fifth acted as a basestation. The software was set to initiate 100 scans for tags, and with each scan it recorded the number of tags which responded. This was done for one, two and three iterations of the `GUI_REQUEST` message – this change is implemented in the software GUI as shown in Figure 4.27. Results were recorded for the case were the motes were 0.5m apart as well as for when they were extremely close to one another to show what effect tag proximity might have – the later setup is shown in Figure 7.10.

As can be seen from Table 7.1, having the software only interrogate the tags once resulted in 2.75% of the tags not being read when they were spaced by 0.5m – this decreased a small amount to 1.75% when they were in close proximity. These results were calculated by assuming that 100 scans could result in a maximum of 400 successful reads; therefore only 200 successful reads would have resulted in 50% of tags not being read. The losses experienced are due to collisions occurring between the motes – whilst this is reduced by the prototype code having a random offset for sending messages (and utilising TinyOS MAC) it is not completely eliminated. Increasing the number of times the tags were interrogated increased the reliability of the system to 100%.

	Successfully discovered tags (%)	
Iterations	**0.5m spacing**	**Close proximity**
1	97.25%	98.25%
2	100%	100%
3	100%	100%

Table 7.1 – Results of testing the prototype system in an ideal scenario

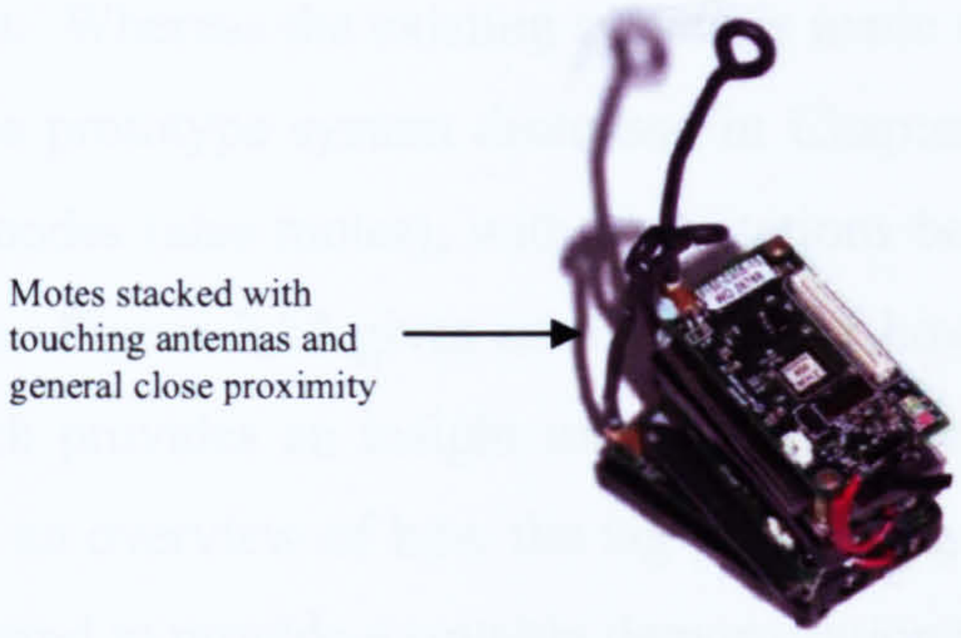

Figure 7.10 – Tags stacked in extremely close proximity

The implications of these results are that a reliable system (i.e. – one which identifies all tags on each scan) would need to have some level of data redundancy. Using the data presented in Figure 4.6, a five node single hop network would require 21 messages to be sent if it was using broadcasting. Using reverse routing and a single iteration, 9 messages would be sent; using two iterations means that this doubles to 18 but is still a saving in terms of energy over the broadcast method even on a network of just five nodes. One should consider that the 18 messages are spread over two separate time periods when using the reverse routing mechanism, thus reducing the load on the network at any one time.

An area of future work for this project is to look further at the operation of the software to see if and where it requires improvement in order to operate on a large scale. This could be done through simulation or industrial trials, a matter which will be discussed in Chapter 8.

7.3. The Airline Industry

7.3.1. Applying the Prototype Tags and Software

This sub-section deals with how the prototype system software and tags could be applied in the airline industry. There is further discussion regarding the HHU, as this was created originally as an airline specific extension of the prototype system. In terms of implementing the prototype system it is thought that it will have much in common with the existing system but add further features and seek to provide further

automation. Whereas the existing system is made up of bar codes and many bar code readers, the prototype system discussed in Chapter 4 is made up of RF tags (motes) and relay nodes (also motes), with base stations being used to determine the position of the tags. Figure 2.14 gives an overview of how luggage flows around an airport, and as such provides an insight into how base stations might be positioned; Figure 7.11 gives an overview of how the tag positioning window of the prototype software may be altered to provide a suitable demonstration to the airline industry.

A mote/tag could be attached to an item of luggage at check in, much in the same way that bar codes are now. The tag could be encased in a hard resin container to protect it from damage. The luggage could then traverse the conveyor and the DCV system; base stations could be placed at key locations in order to interrogate the tags as they are moved. The positioning of the base stations could be similar to that of the current bar code readers, since this arrangement appears to provide adequate resolution for tracking the luggage. The key benefit would be that the orientation of bags would not matter, so LOS between a tag and the base station would not be an issue, although it is noted that any RF technology would provide this advantage. Of course, what other technologies discussed in Chapter 2 cannot provide is the ability to enable data delivery in the case of heavily laden conveyors, where multiple items of luggage may become stacked on top of each other – the stacking may cause severe signal attenuation which may well be overcome due to the multihop routing implemented in the prototype system. In terms of comparison with passive RFID (see Table 2.2) the motes have a significant range advantage with all types of material. This includes metal, which is significant when one considers that much of the luggage handling system is constructed from this.

Unlike the current setup in airports, relay nodes could be positioned at the departure and arrival gates removing the need for checking by luggage handlers since the software could automatically create an inventory and alert the luggage handlers to missing luggage via intelligent database querying. The prototype system, whilst not including complex querying features, implements a basic database structure (see Figures 7.12 and 7.13 for an airline specific adaptation) which could form the basis of this functionality. This may even allow for luggage lost within the DCV system to be

recovered before a flight leaves the ground since automation of the inventory process will allow luggage handlers to prioritise their time. There is also the likelihood that reduced LOS issues will result in fewer losses within the conveyor and DCV systems, and so a higher percentage of luggage will make it to the correct departure gate.

Luggage then has to be transported from the departure gate to the aircraft. Rather than tracking luggage as it moves across the tarmac, it is envisaged at this stage that it would be inventoried again as it is loaded onto the aircraft. Further base stations on the conveyor leading to the cargo hold (or on the cargo hold door itself) could allow this. This inventory could be cross matched with the inventory created at the departure gate and any discrepancies could then be rectified – whilst this would be a largely manual process, it is feasible since the luggage loading process itself requires manpower. Once in the cargo hold the tags will enter their sleep state and should typically stay like this until the aircraft lands at the destination airport.

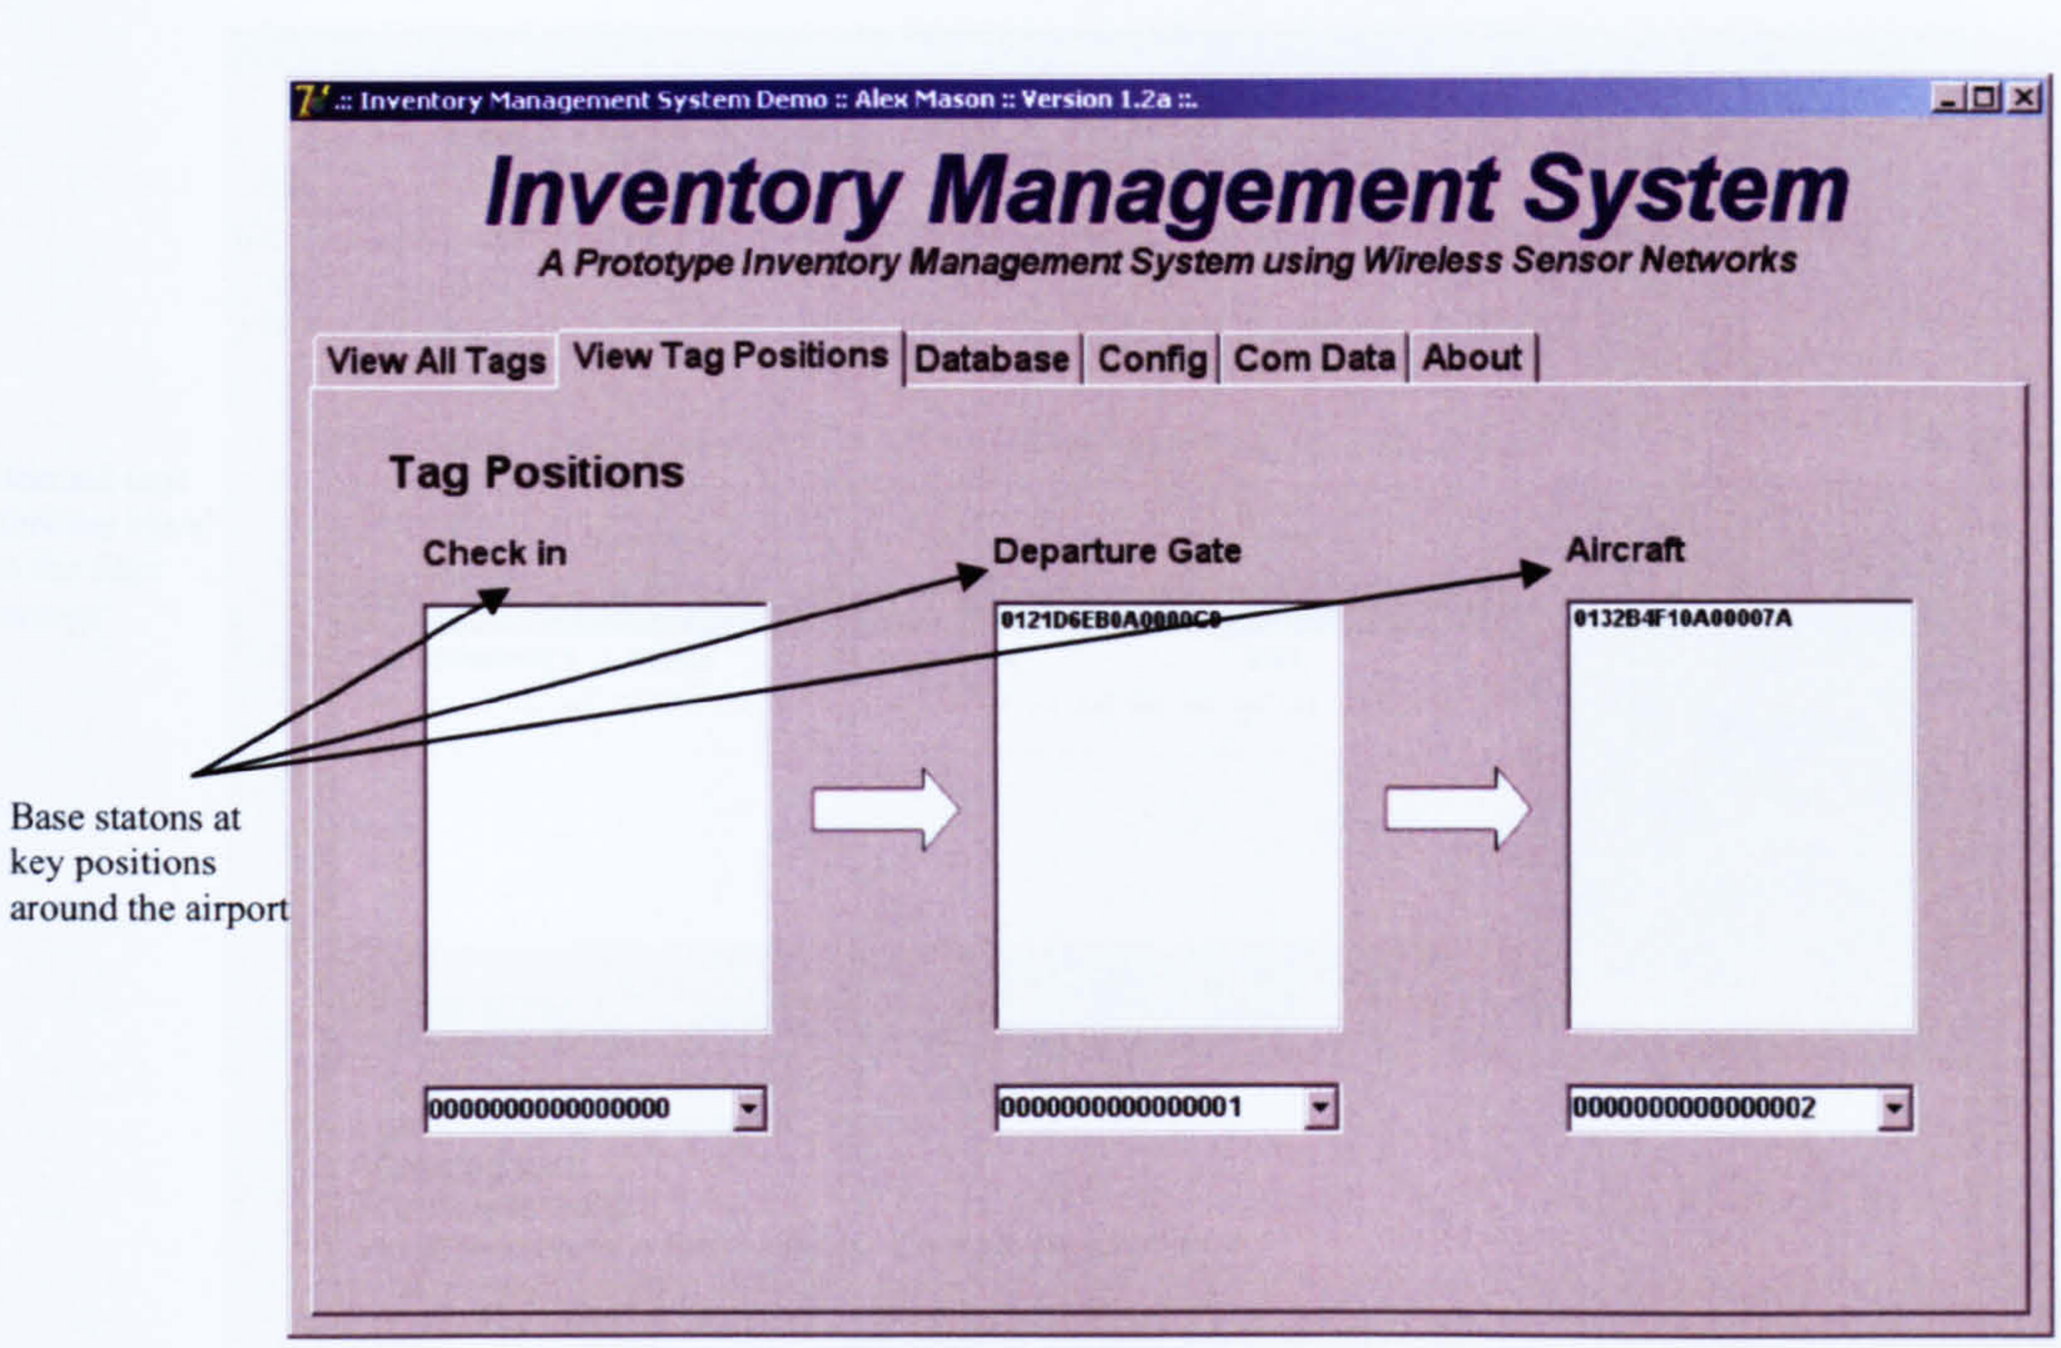

Figure 7.11 – Airline departing luggage position example

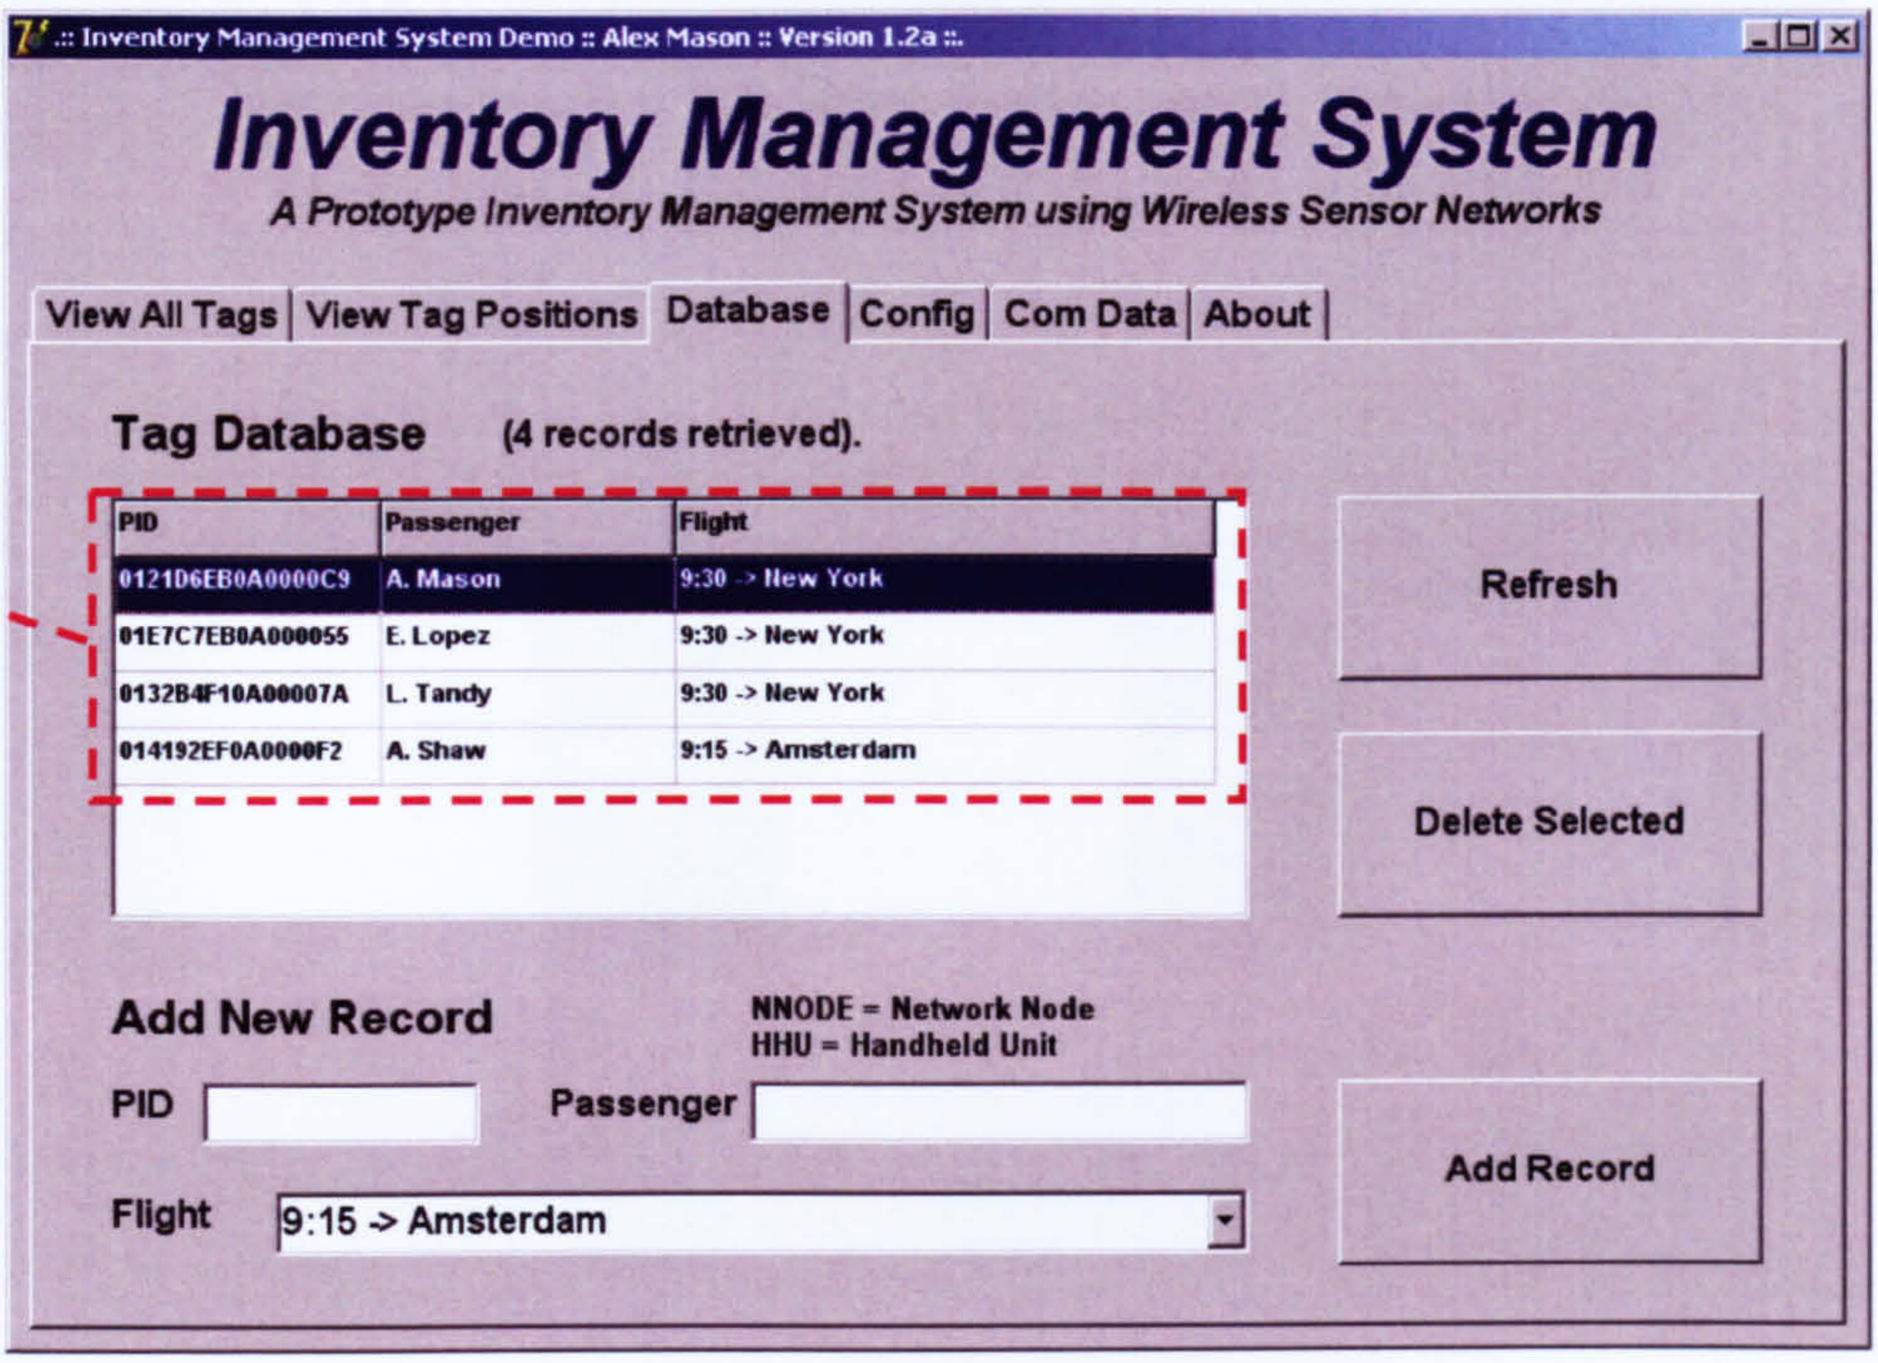

Database structure showing different passengers and their flights

Figure 7.12 – Airline specific database structure

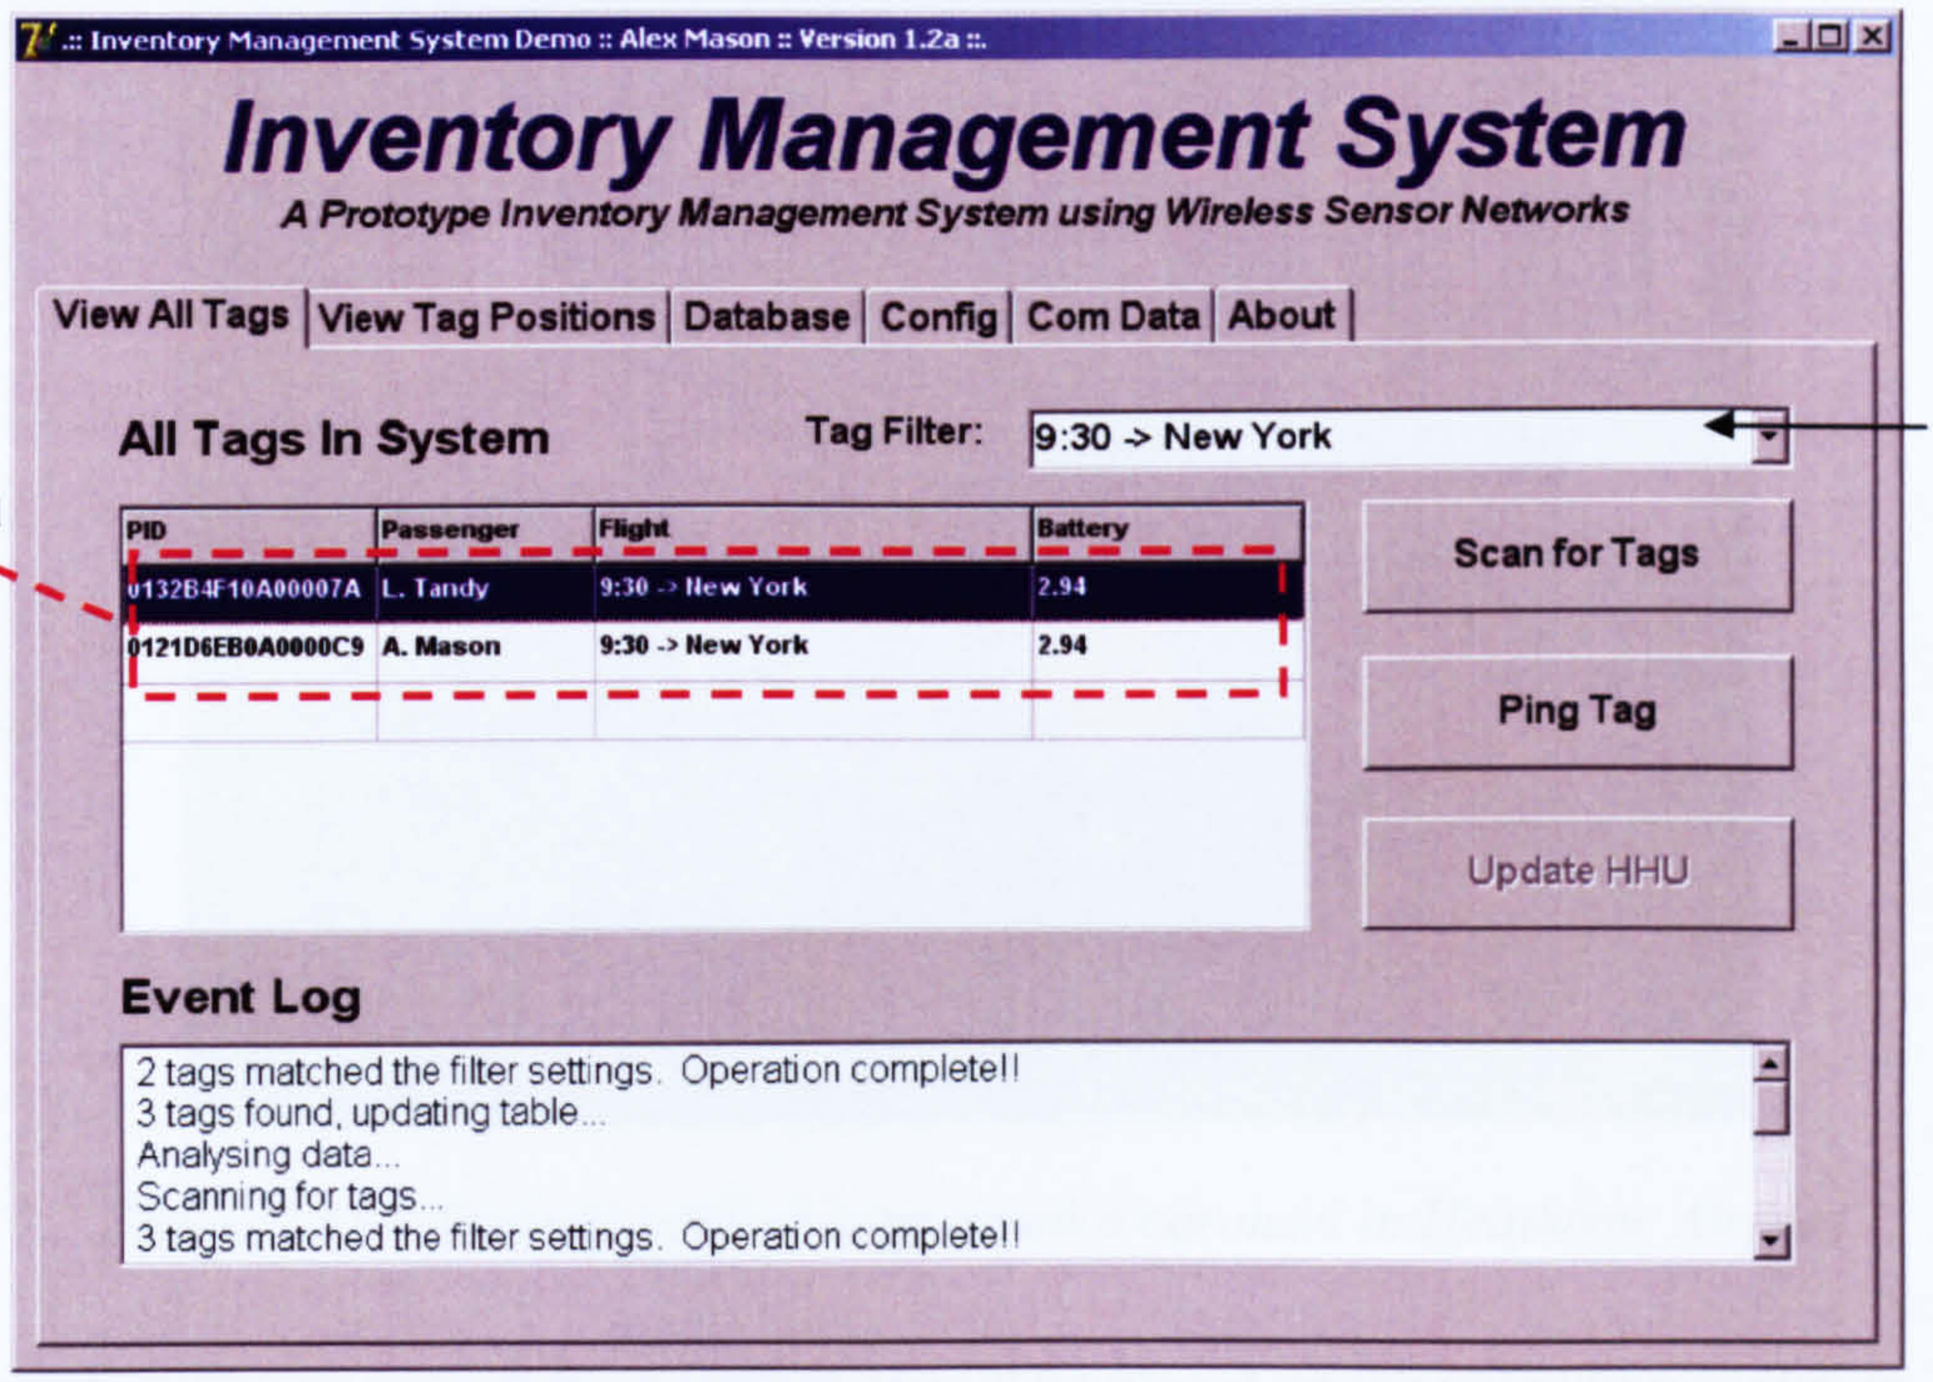

Scanned tags showing based on the filter settings

Filter setting

Figure 7.13 – Airline specific tag window

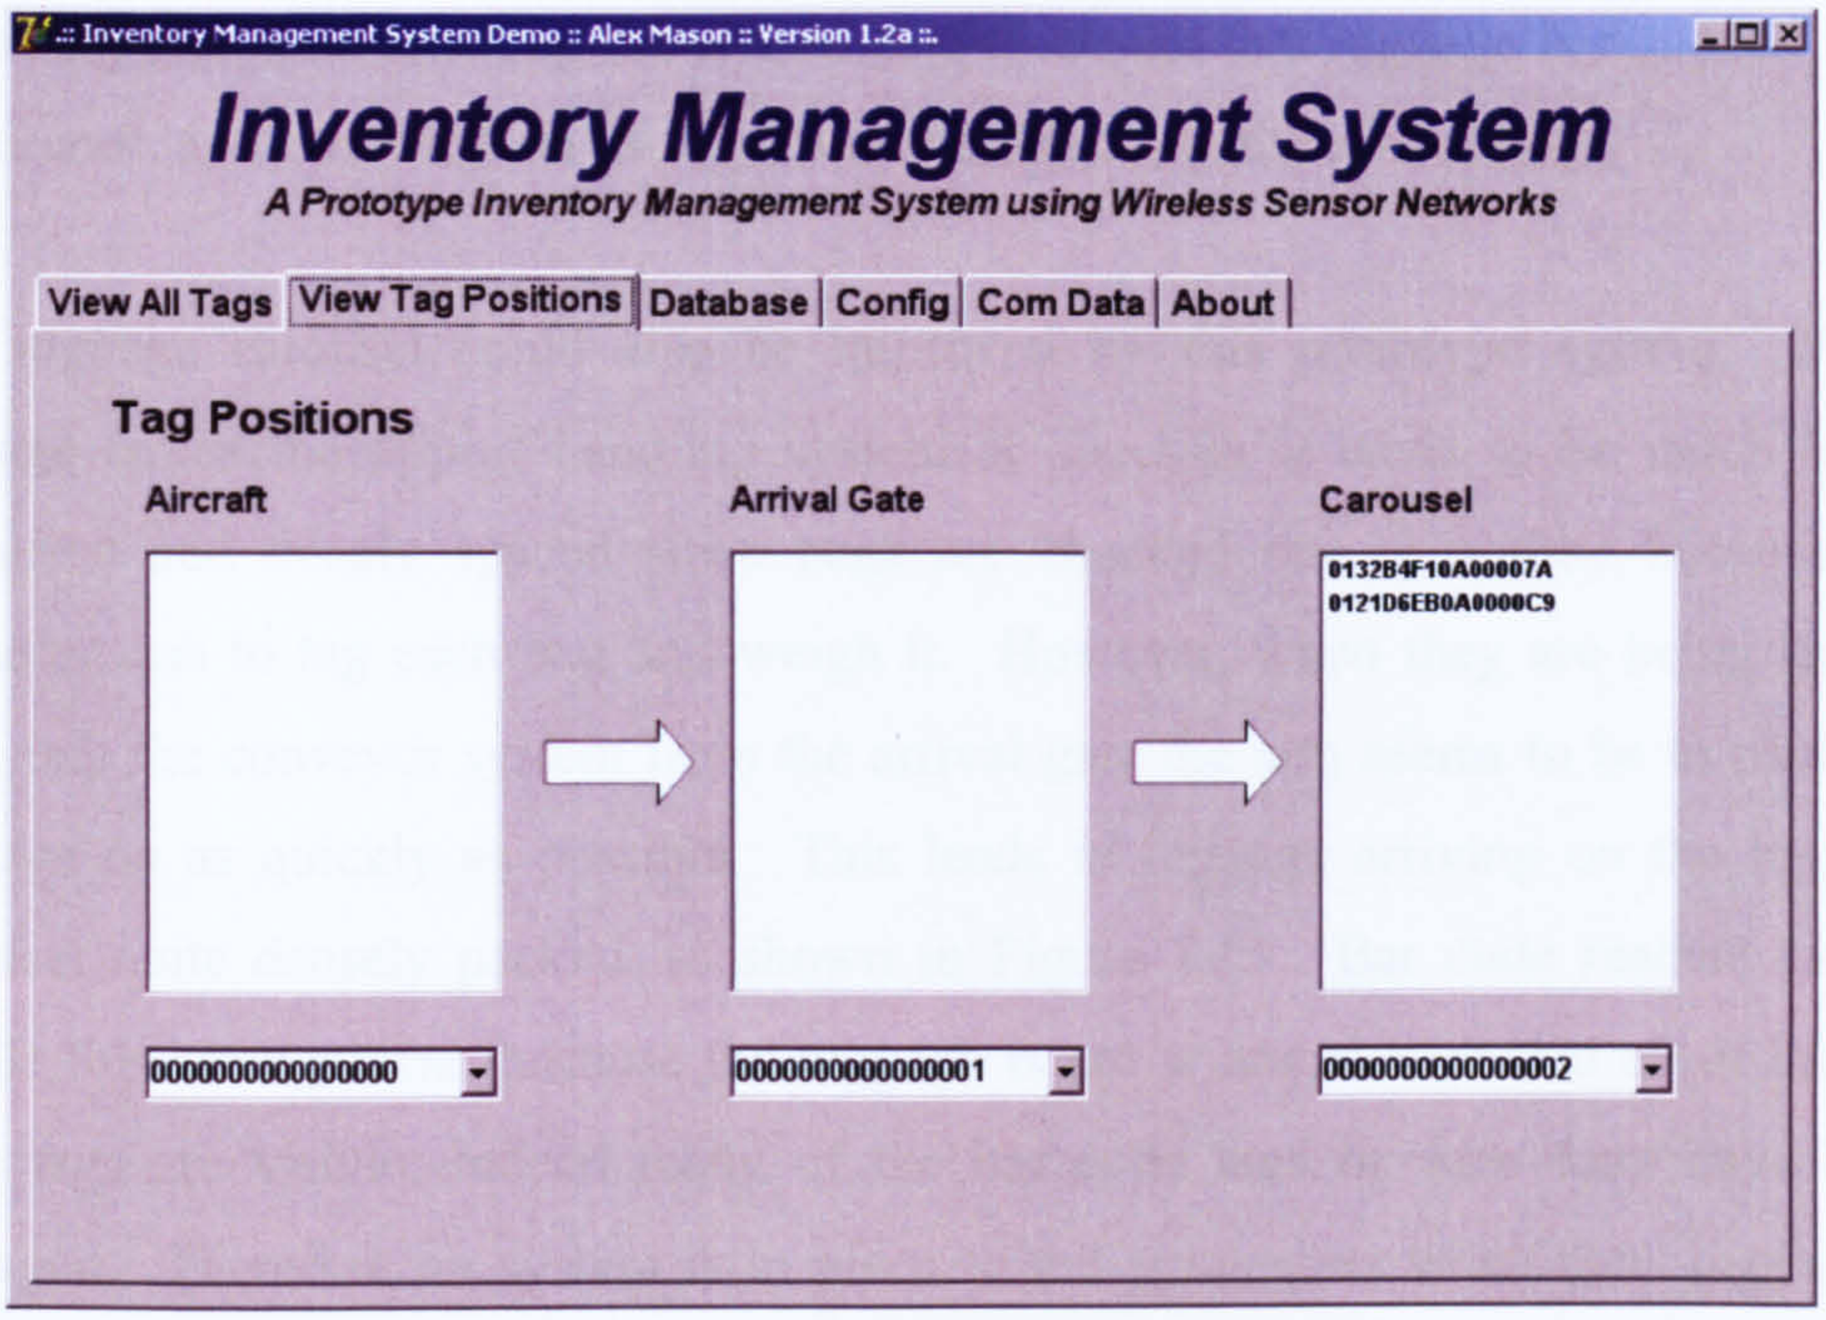

Figure 7.14 – Airline incoming luggage position example

Luggage densely packed on a busy carousel

Figure 7.15 – Densely packed luggage on a carousel in Heathrow Airport

For unloading an aircraft base stations would be placed on the unloading conveyor, as well as at the arrival gate and finally at the luggage carousel, as shown in Figure 7.14. The loading inventory list can be cross referenced with inventories created whilst

unloading and upon reaching the arrival gate to ensure that luggage is removed from the aircraft, and is not lost on its way from the aircraft to the arrival gate.

The luggage carousel could also be improved by this prototype system. When luggage enters the airport handling system at check-in it tends to be much better organised and evenly spaced since bags are checked one at a time because the attendant has to tag each bag and weigh it. However, when they are being loaded back into the conveyor system from the arrival gate the aim seems to be to pack the luggage on as quickly as possible. This leads to luggage arriving on the luggage carousel quite densely packed, as shown in Figure 7.15. Bar code readers cannot handle this arrangement, because the luggage is not arranged such that all of the bar code tags are visible; indeed many of the bar code tags by now may have been damaged. Therefore no system is in place to tell passengers when their luggage is available for collection, and instead they are required to try and find their luggage through their knowledge of what it looks like. Naturally this leads to many problems – pushing and shoving can occur, people can pick up the wrong luggage, and some times it may not arrive.

It was noted earlier however that the prototype system can overcome LOS issues due to its use of both RF communications, as well as multihop routing which overcomes the attenuation which is inevitiable in such a cluttered scenario. Therefore, the prototype system could easily be adapted so that when it senses the arrival of luggage at the carousel (see Figure 7.14) it could notify the appropriate passengers. This means that whilst some passengers may still have to wait longer than others, the collection process could be more orderly. Notifications could be via a simple VDU, or alternatively there could be a system in place which notifies passengers via mobile phone – such data could easily be held in the passenger database.

Upon collecting their bags, some mechanism through which the tag is returned to the airline would be necessary – since the mote based tag would cost significantly more than a bar code or even a simple passive RFID tag (as demonstrated in Section 3.9) it is thought that they should be reusable. It is also thought that this may add an element of security to luggage being claimed because if the passenger has to return the tag

then it could be cross referenced against either their boarding card or passport, and therefore seek to minimise luggage theft, or simple passenger mistakes.

7.3.2. Using the Hand Held Unit for Luggage Discovery

One of the initial problems identified in this project, and discussed in Chapter 1, was that luggage inside an aircraft cargo hold is difficult for luggage handlers to find. Added to this, they are required to find what they are looking for quickly which often leads to damaged luggage. Although it was mentioned in Chapter 2 that airlines in New Zealand utilise hand-held bar code readers to authenticate luggage being loaded onto an aircraft, there is no indication that similar technology is being used to assist its discovery once loaded.

Therefore the HHU (see Chapter 4) was created with this particular problem in mind; it utilises the Mica2 RSSI feature in order to attempt to discover the distance between the HHU and the target tag. An experiment to test the HHU was set up inside a metal shipping container (5m long, 2.5m wide, 3m tall) which was chosen in order to emulate the conditions inside a metallic aircraft cargo hold. Since sufficient luggage was not available for use, 24 plastic crates (stacked in four columns, each six crates high) were used – one of the crates contained a mote, which was moved at random. This setup is shown in Figure 7.16.

To perform the tests, firstly the HHU was moved along the length of the container in 0.5m increments from the back wall to the door. At the point were the lowest RSSI value was found the HHU was then moved toward the stacked crates in 0.2m increments, maintaining a fixed height of 1m above the ground. Finally the HHU was moved vertically from the floor to the top of the crates (0.3m increments from 0-1.5m) – the aim of these sweeps was to see if the crate with the mote in could be discovered reliably through a procedural method.

Due to time limitations imposed upon utilising the container, ten iterations of this test were recorded and it was found that 80% of these tests resulted in the correct crate being discovered. It was found that moving along the length of the container proved

far more useful in terms of a varying RSSI with distance than moving width ways; this is most likely because the container was longer than it was wide. This meant that there was much more room for definite variation in the RSSI results over the greater distance – experience with the RSSI during this project has shown that it can vary rather unpredictably over short distances. In the cases when this procedure failed to accurately locate the correct crate, the maximum error recorded was 1.4m; this is large error when one considers that the room is only 5m in length!

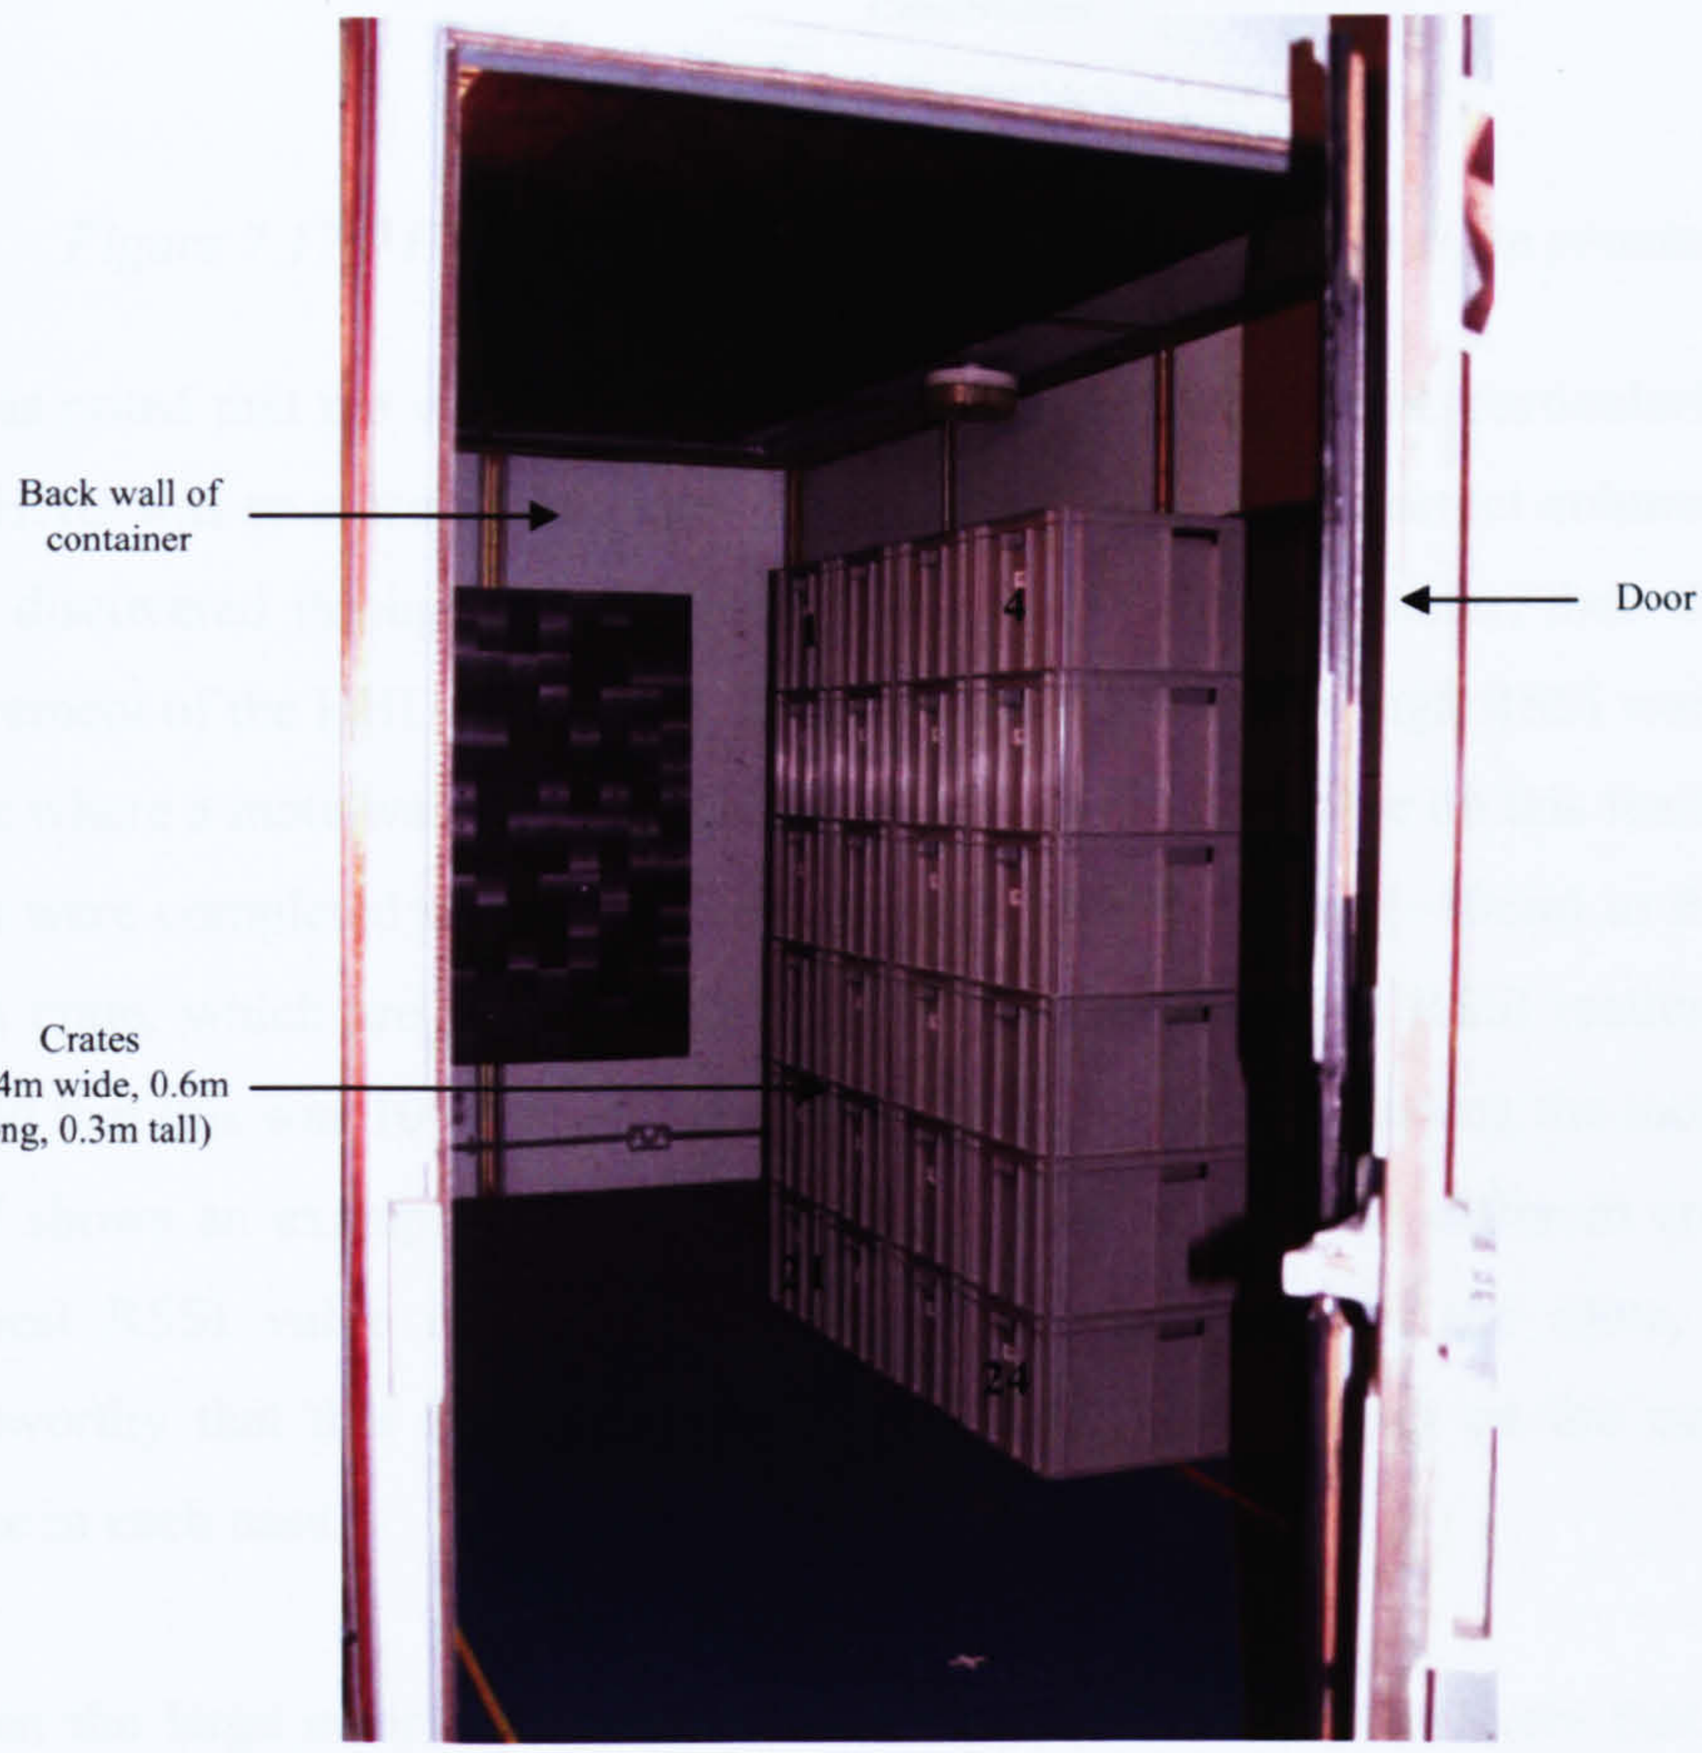

Figure 7.16 – Shipping container containing stacked plastic crates

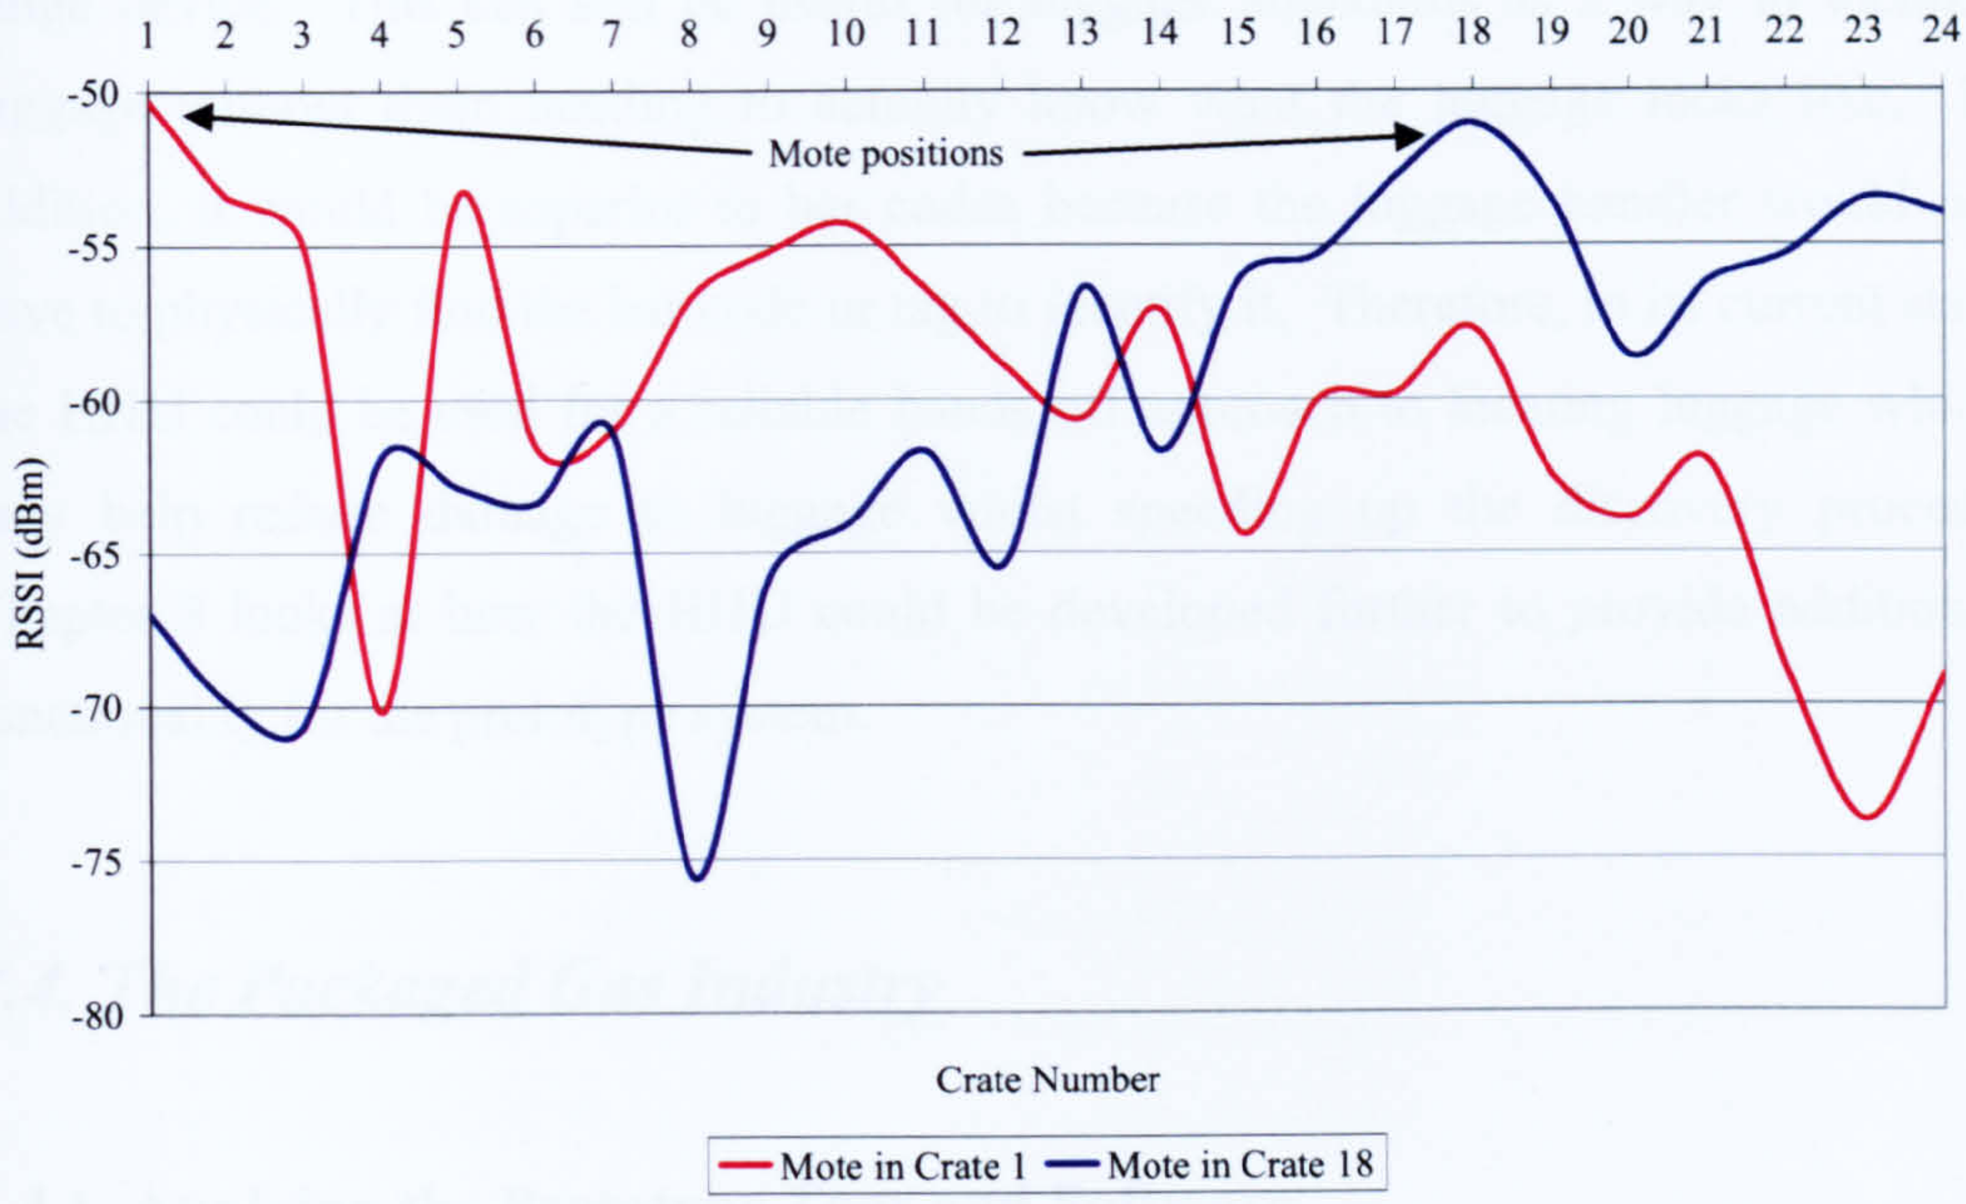

Figure 7.17 – HHU results when sweeping each crate in close proximity

It was noted that the vertical movement of the HHU was useful, particularly because the HHU was so close to the crates. It was found that if the correct column of crates was discovered through moving along the length of the container, then the vertical movement of the HHU next to the crates always resulted in a high RSSI reading at the crate where a mote was located. Therefore in order to capitalise on this feature further tests were completed which involved holding the HHU close (~10cm) to the front of each crate, which are numbered in Figure 7.16, and taking a RSSI reading. It was found that this was 100% reliable in identifying the crate containing the mote. Figure 7.17 shows an example of two[20] tests where the mote was in different crates. The highest RSSI value in each test represents the position of the mote, and it is noteworthy that this RSSI value is approximately 3-4dBm up on the next highest value in each case.

Given the large error experienced in the first set of tests, it is likely that the HHU would be more suitable for use in close proximity to luggage, rather than as a long

[20] Note that displaying all of the data graphically, whilst possible, is not particularly useful in terms of readability and so a few results have been selected for clarity.

range device. This can still be useful for luggage attendants as a way to identify luggage without them needing to actually know what the luggage looks like. In addition, it would be superior to bar codes because the luggage handler would not have to physically find the bar code or tag to identify it. Therefore, in its current state the HHU could be used for a reliable hands off approach to locating luggage which may help reduce damage to luggage whilst speeding up the discovery process. Chapter 8 looks at how the HHU could be developed further to provide additional functionality for the prototype system.

7.4. The Packaged Gas Industry

7.4.1. Applying the Prototype Tags and Software

Unlike the airline industry, there is not a widespread automated inventory management system in place for managing packaged gases, and as mentioned in Chapter 2 those large organisations which do have bar code systems have only recently been implementing them.

The prototype software could be applied at this level in order to track a cylinder as it is filled, stored, delivered, checked and then refilled – this would simply require the prototype software to be given a greater resolution in terms of the number of base stations it tracked.

The prototype in its current form was aimed more specifically at inventory management within individual gas cylinder storage compounds, with the aim of creating the potential for automated gas cylinder storage facilities – work relating to this has been published recently [164]. Such compounds may store hundreds or even thousands of gas cylinders, and the idea is that cylinders can be returned to and taken from the compound without the need for personnel to manually scan each cylinder with a handheld bar code reader.

When cylinders enter a facility, base stations can be used to interrogate the tags so that a record is kept of the cylinders which enter the facility. Figure 2.19 shows a

typical lorry load of cylinders which might be used to carry cylinders into and out of a compound. The cylinders may not be in direct LOS with a base station but it was proven in Section 7.2.2 that the motes are capable of communicating unless enclosed completely by metal, so LOS is not required. Whilst signal attenuation will occur, the gas cylinders are densely packed such that motes surrounded by other cylinders only have to be able to communicate over a short distance to a neighbour. These neighbours can then pass messages to the base stations, ensuring that all the cylinders are identified.

Once in the compound, cylinders can be appropriately stored. The example setup of the prototype system in Figure 7.18 shows a general storage area – further resolution could readily be added if, for example, the compound were split into distinct areas (e.g. – there may be areas allocated for different gases and empty cylinders). Stored cylinders can be periodically scanned in order to check that they are still within the compound – this scanning should be quite infrequent though. If the tags are powered up for 10 seconds each time they are scanned, a 1% duty cycle would entail each tag being powered up every 1000 seconds (or 16.7 minutes). Depending on the availability of security it may be possible to reduce the duty cycle such that it is much less than 1%. This would result in a significant battery life extension, which would be important considering that the gas cylinders may be stored for months before being sent out to customers. The inventory collected upon each scanning operation can be utilised for security purposes, and the prototype software could be modified to send an alert.

Cylinders leaving the facility could also be interrogated, much in the same way as when they enter. It is likely that those picking up products will have some sort of identification card (maybe based around RFID), which could be used to link the cylinders leaving the compound to a particular customer account for automated billing.

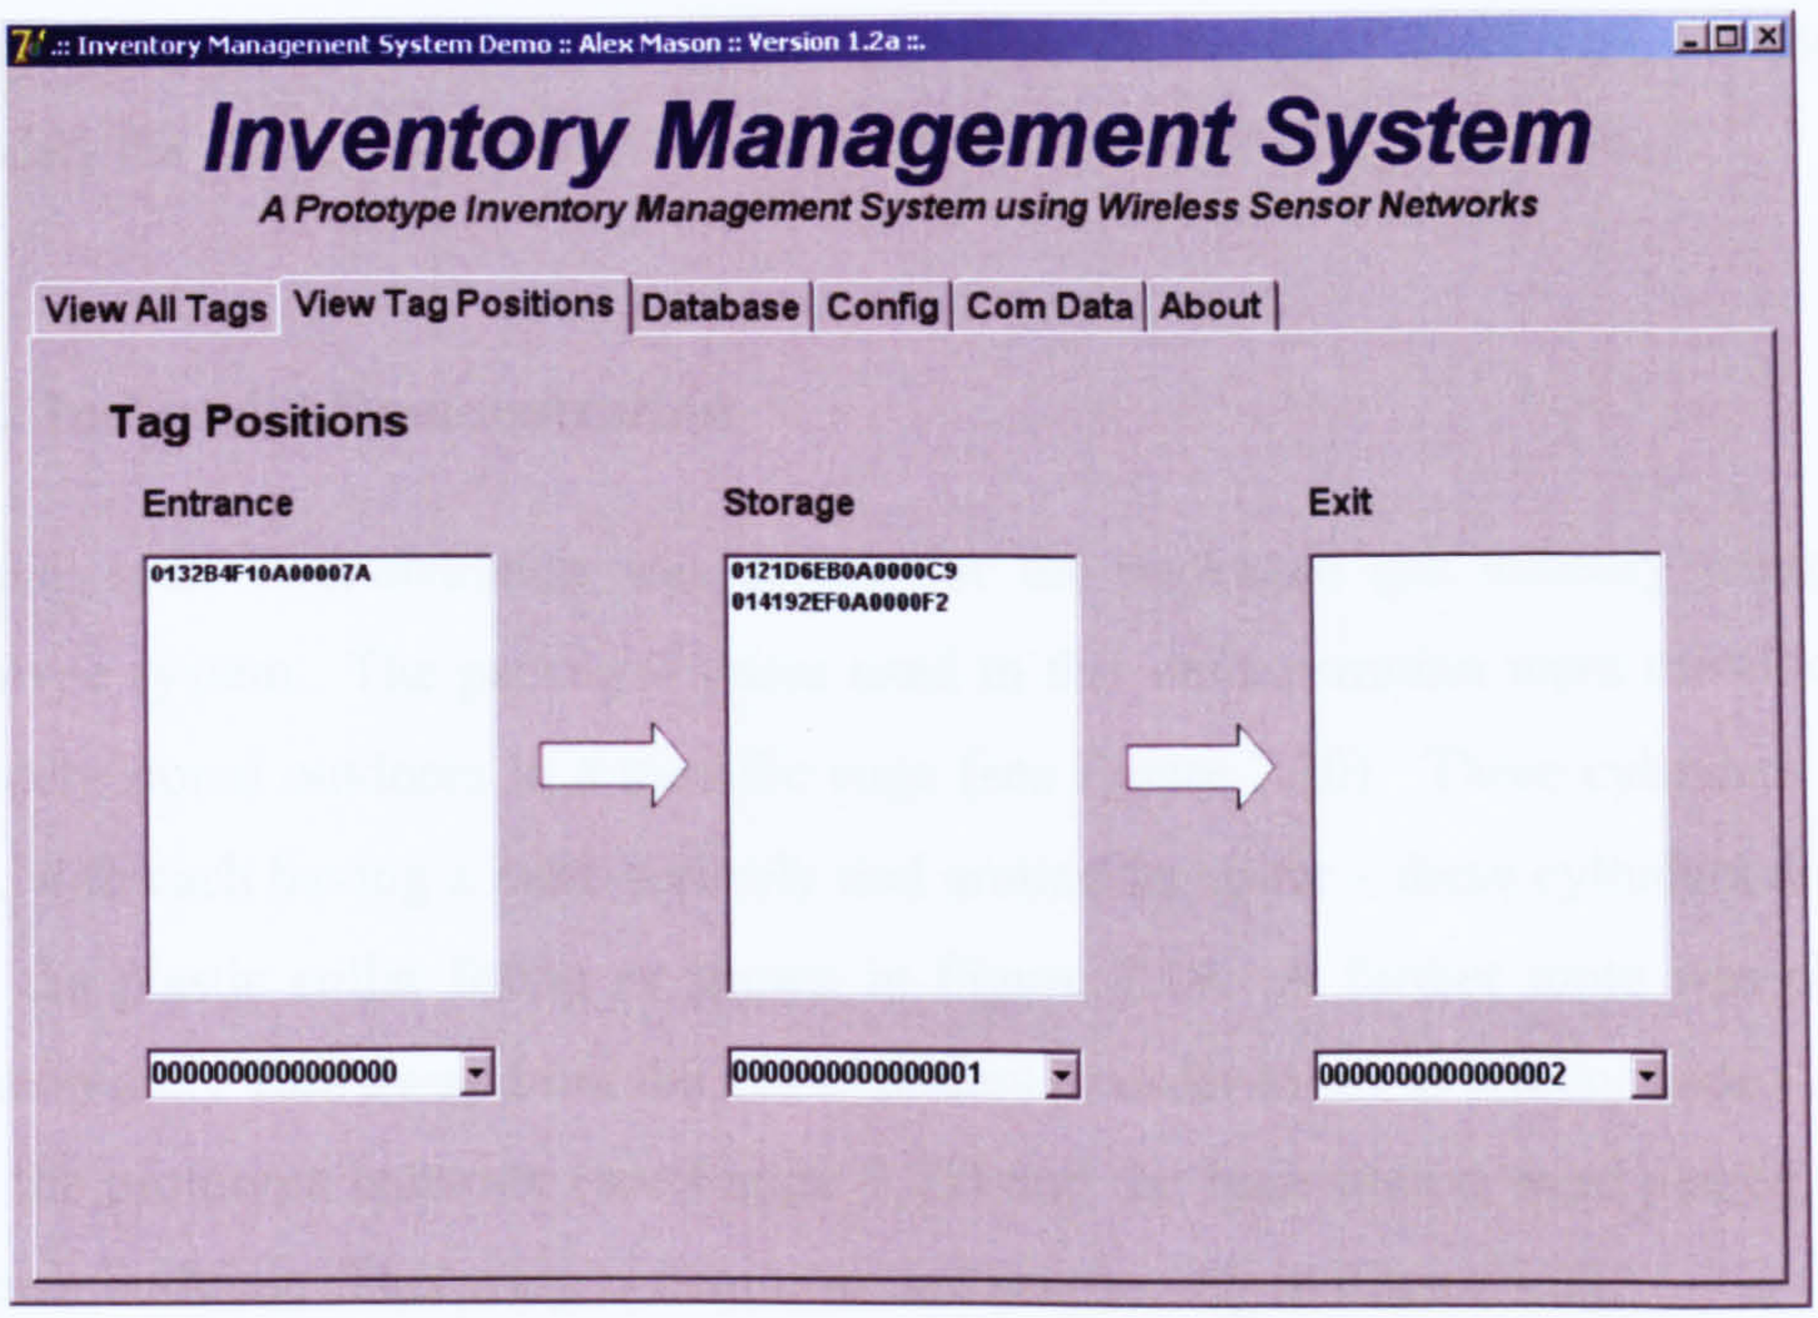

Figure 7.18 – Packaged gas automated compound position tracking

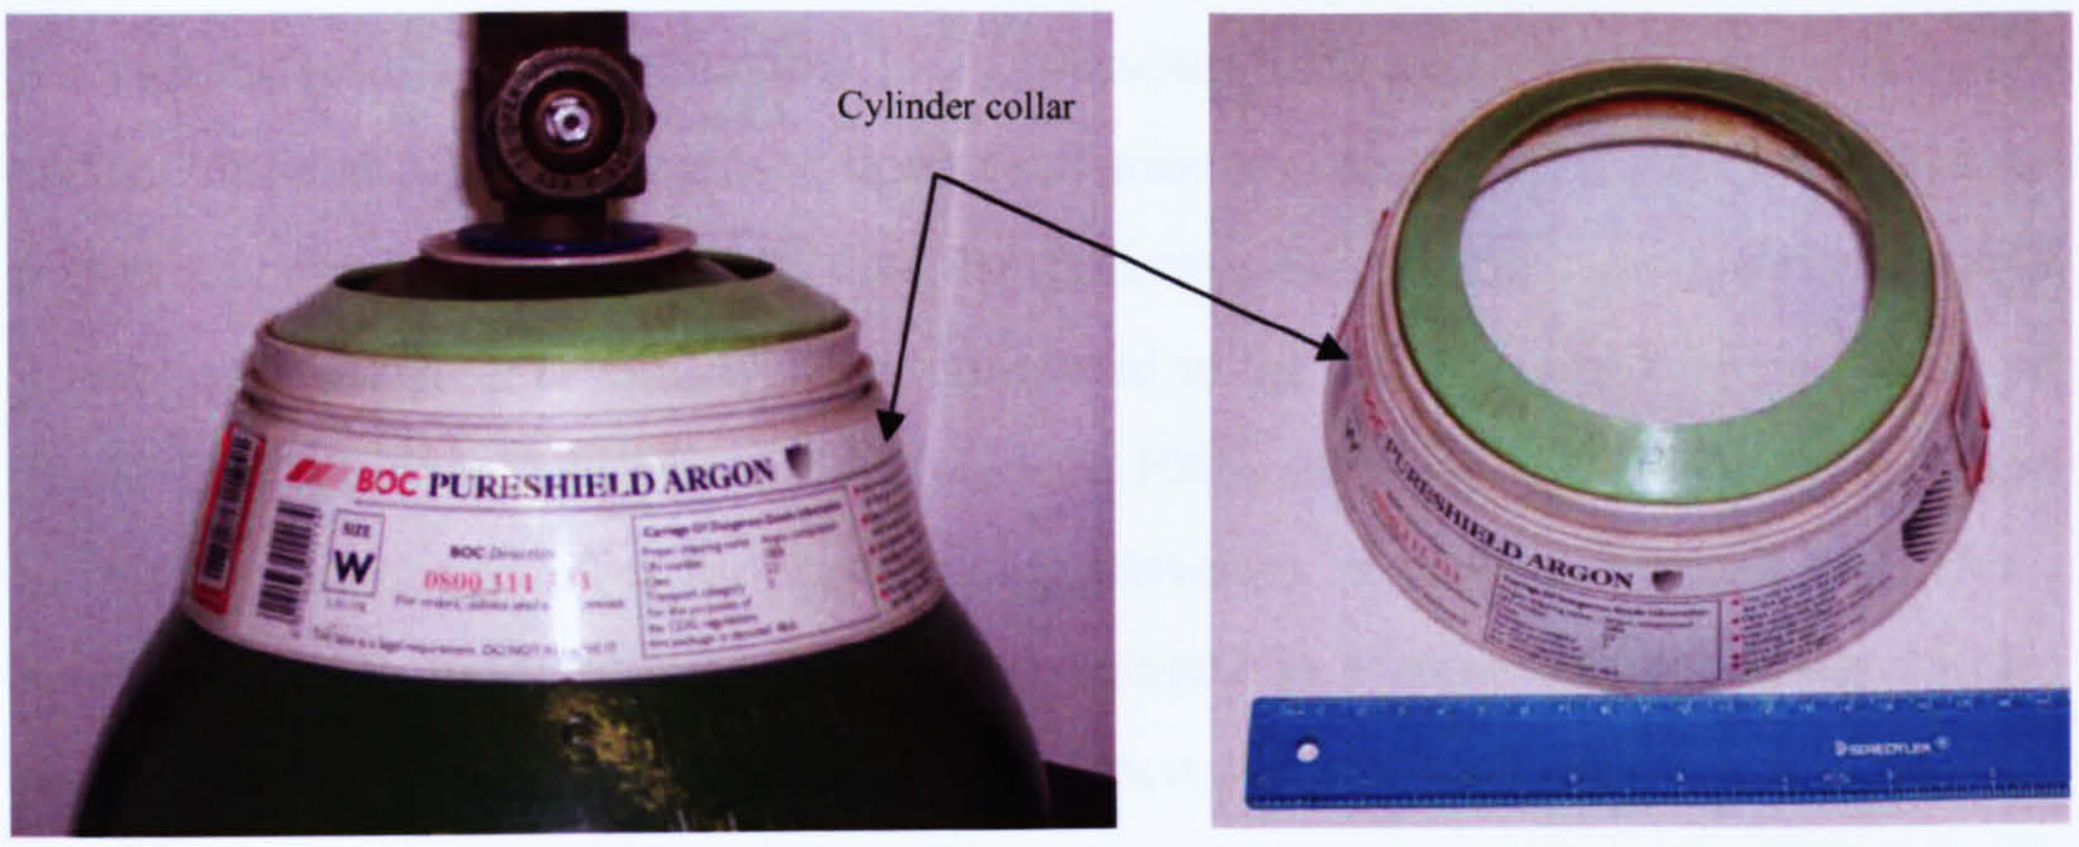

Figure 7.19 – Gas cylinder collar

With regards to actually attaching tags to gas cylinders, many modern cylinders have screw fitted plastic collars, as shown in Figure 7.19. The space between the inside of the collar and the the gas cylinder is large enough to fit a Mica2 mote PCB; the battery pack must be detached however. The fit is rather tight and Chapter 8 considers the future likelihood of much smaller motes becoming available.

7.4.2. Industrial Demonstration

A small scale demonstration was set up for the packaged gas industry using the prototype system. The packaged gases used in this demonstration were metallic gas cylinders stored outdoors in a metallic cage (see Figure 7.20). Three cylinders were used, with each having a mote securely tied around its collar – these cylinders did not have the plastic collar fitting as shown in Figure 7.19. A further mote was placed approximately 10m away from the gas cylinders in order to act as a relay node. A PC with the prototype software (see Figure 7.21) and the base station were placed in an adjacent building. This setup is demonstrated graphically in Figure 7.22.

In addition to the metallic storage cage, a mote on each cylinder was arranged so that it provided the worst possible operating conditions; non-LOS conditions were present, antennas were touching the cylinders they were not matched in terms of polarisation. As with the results taken in Section 7.2.4, 100 scans were performed with the software and the number of tags recognised upon each scan was recorded. Table 7.2 shows the results obtained, which indicate that the differing environment does have a small impact on the system reliability when compared to the ideal scenario results shown in Table 7.1. Whilst it is still possible to achieve 100% reads using the prototype system in this scenario, it requires further scanning iterations in order to achieve this. As a result a reliable delivery system may be required as a modification to the current routing mechanism – this is a topic which is discussed in Chapter 8.

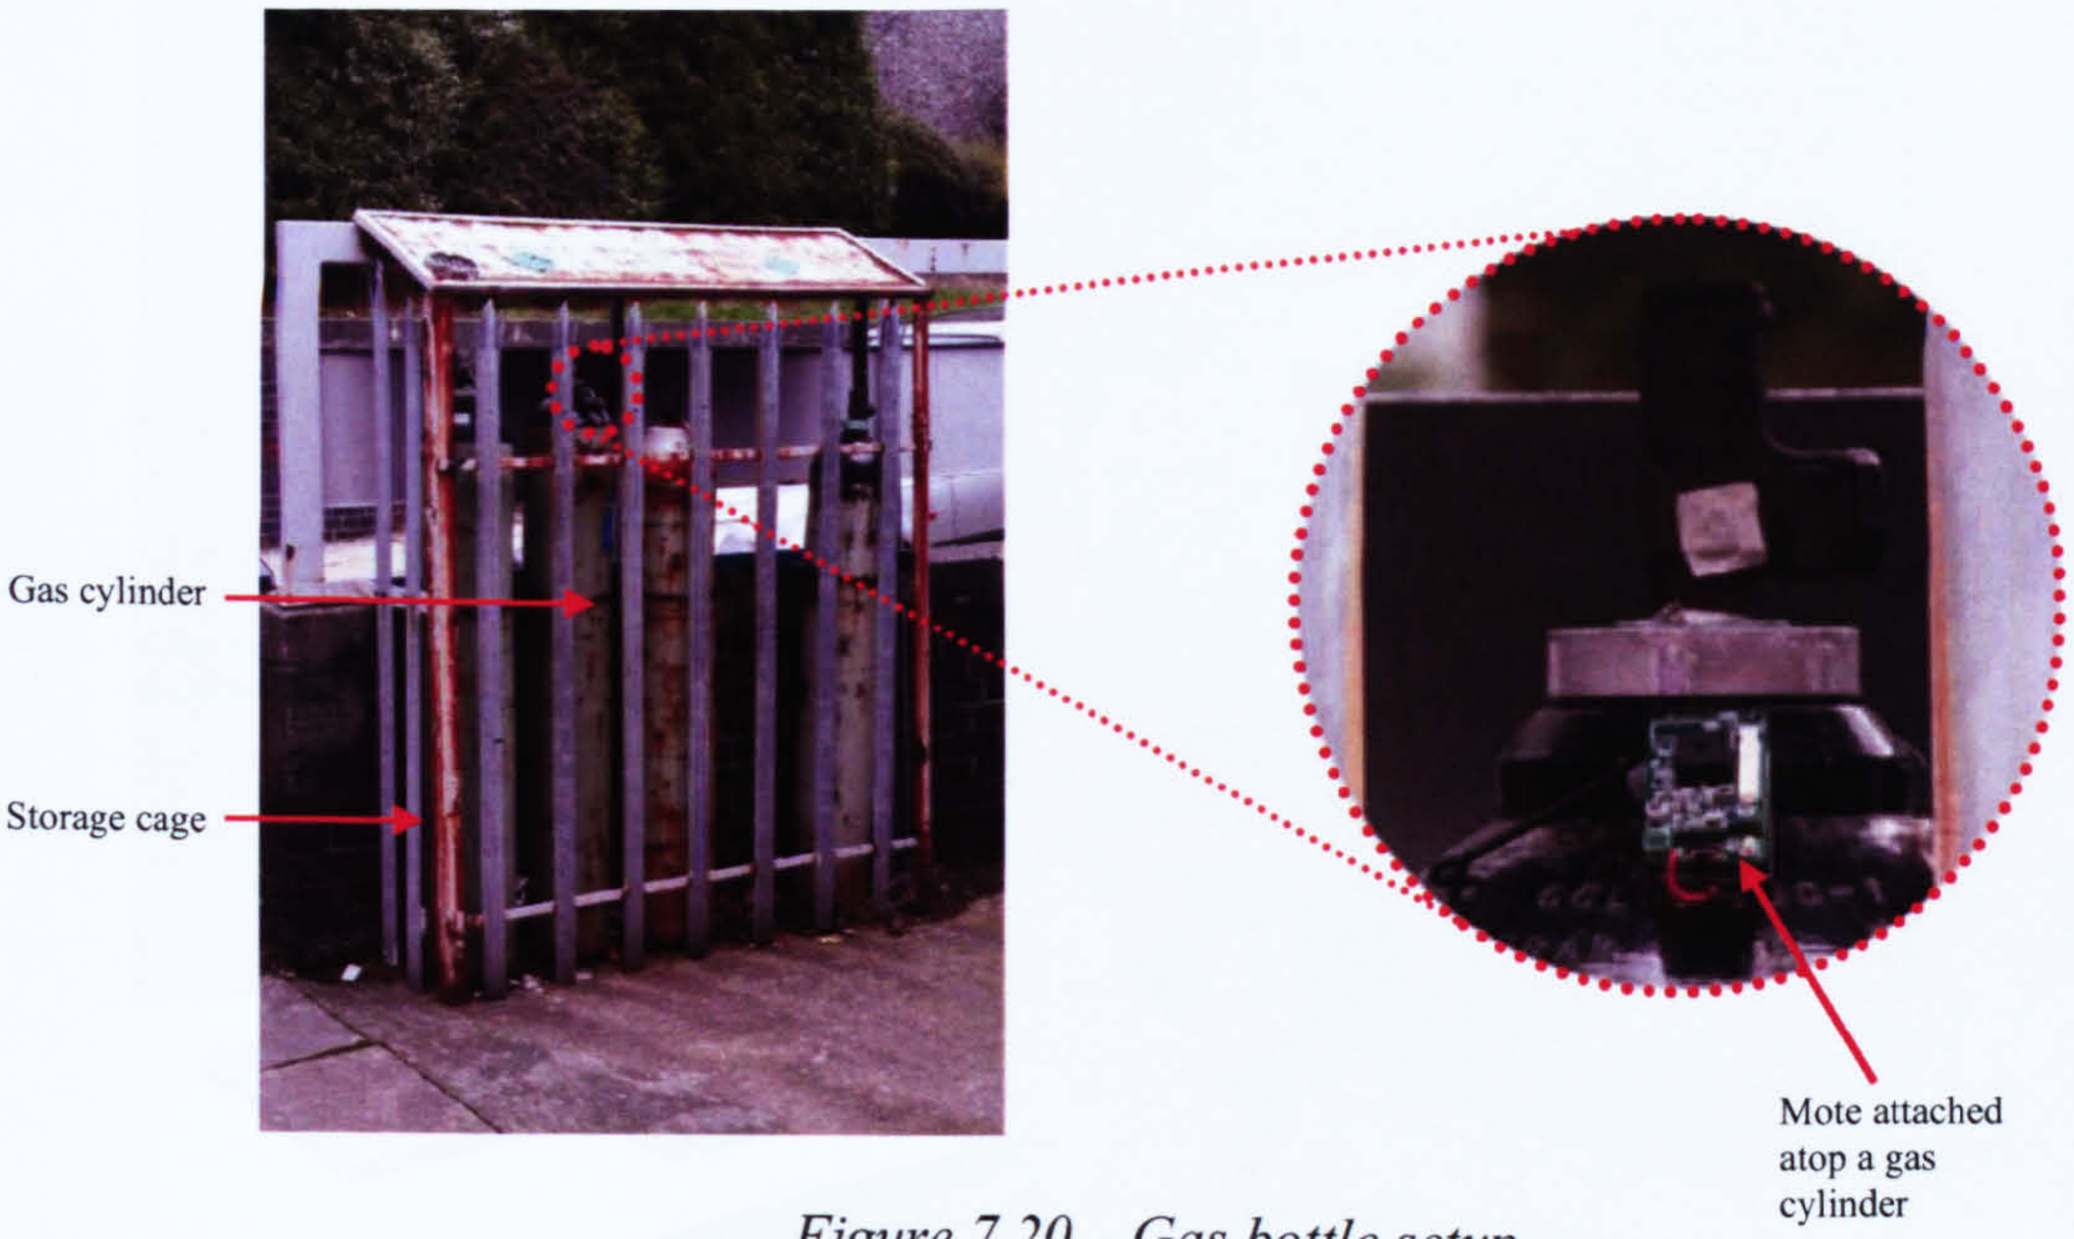

Figure 7.20 – Gas bottle setup

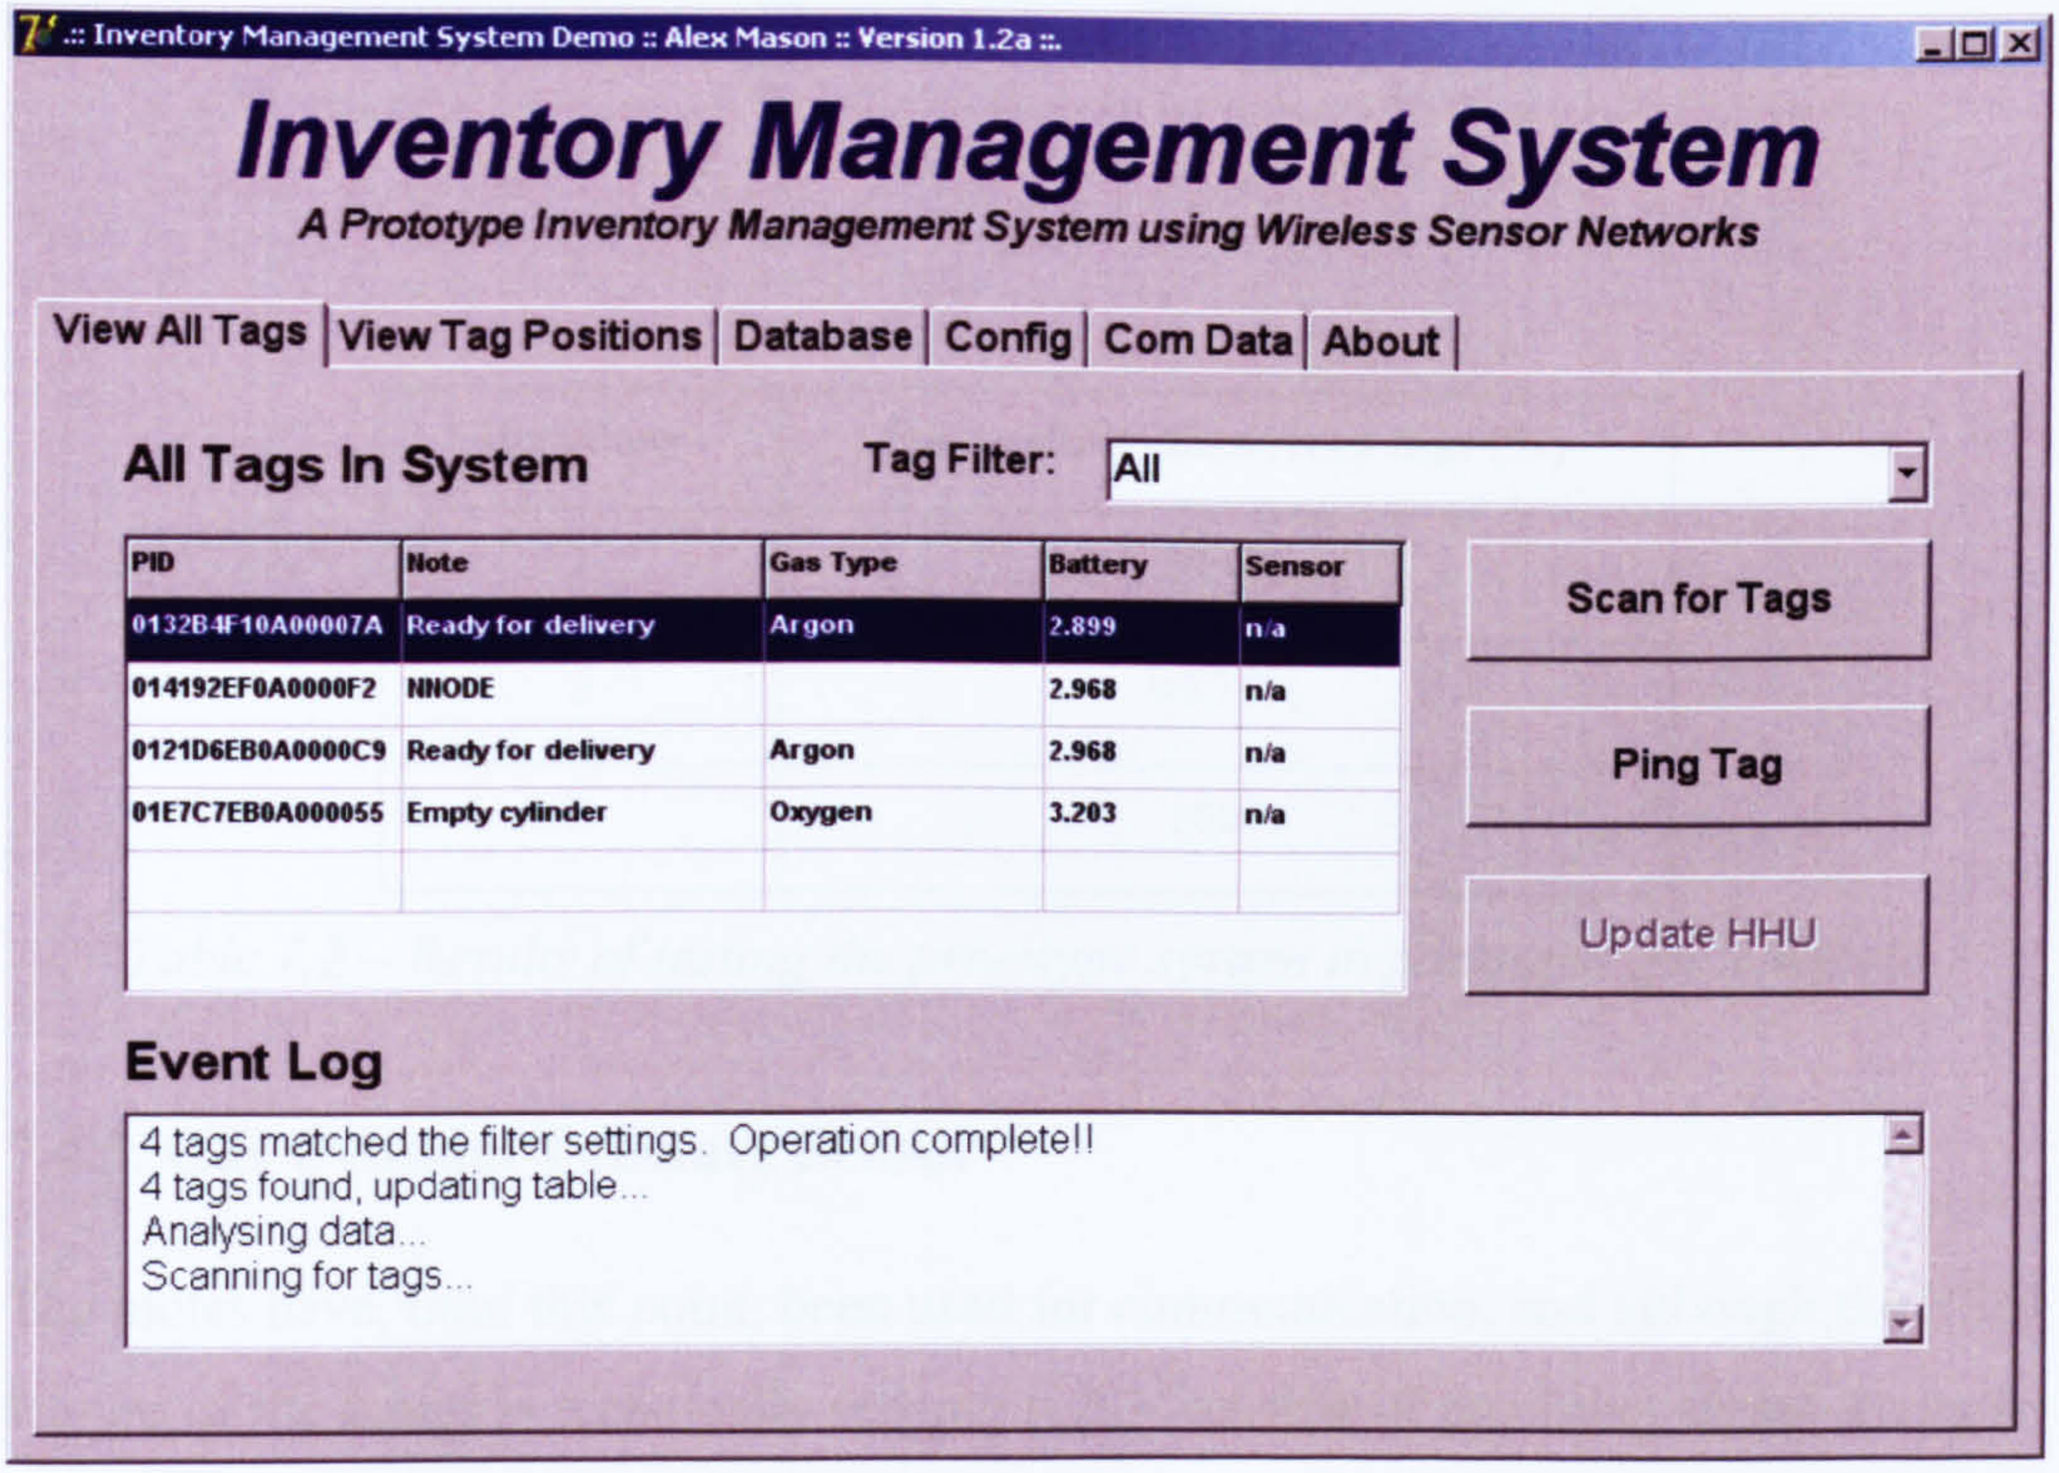

Figure 7.21 – Packaged gas inventory window

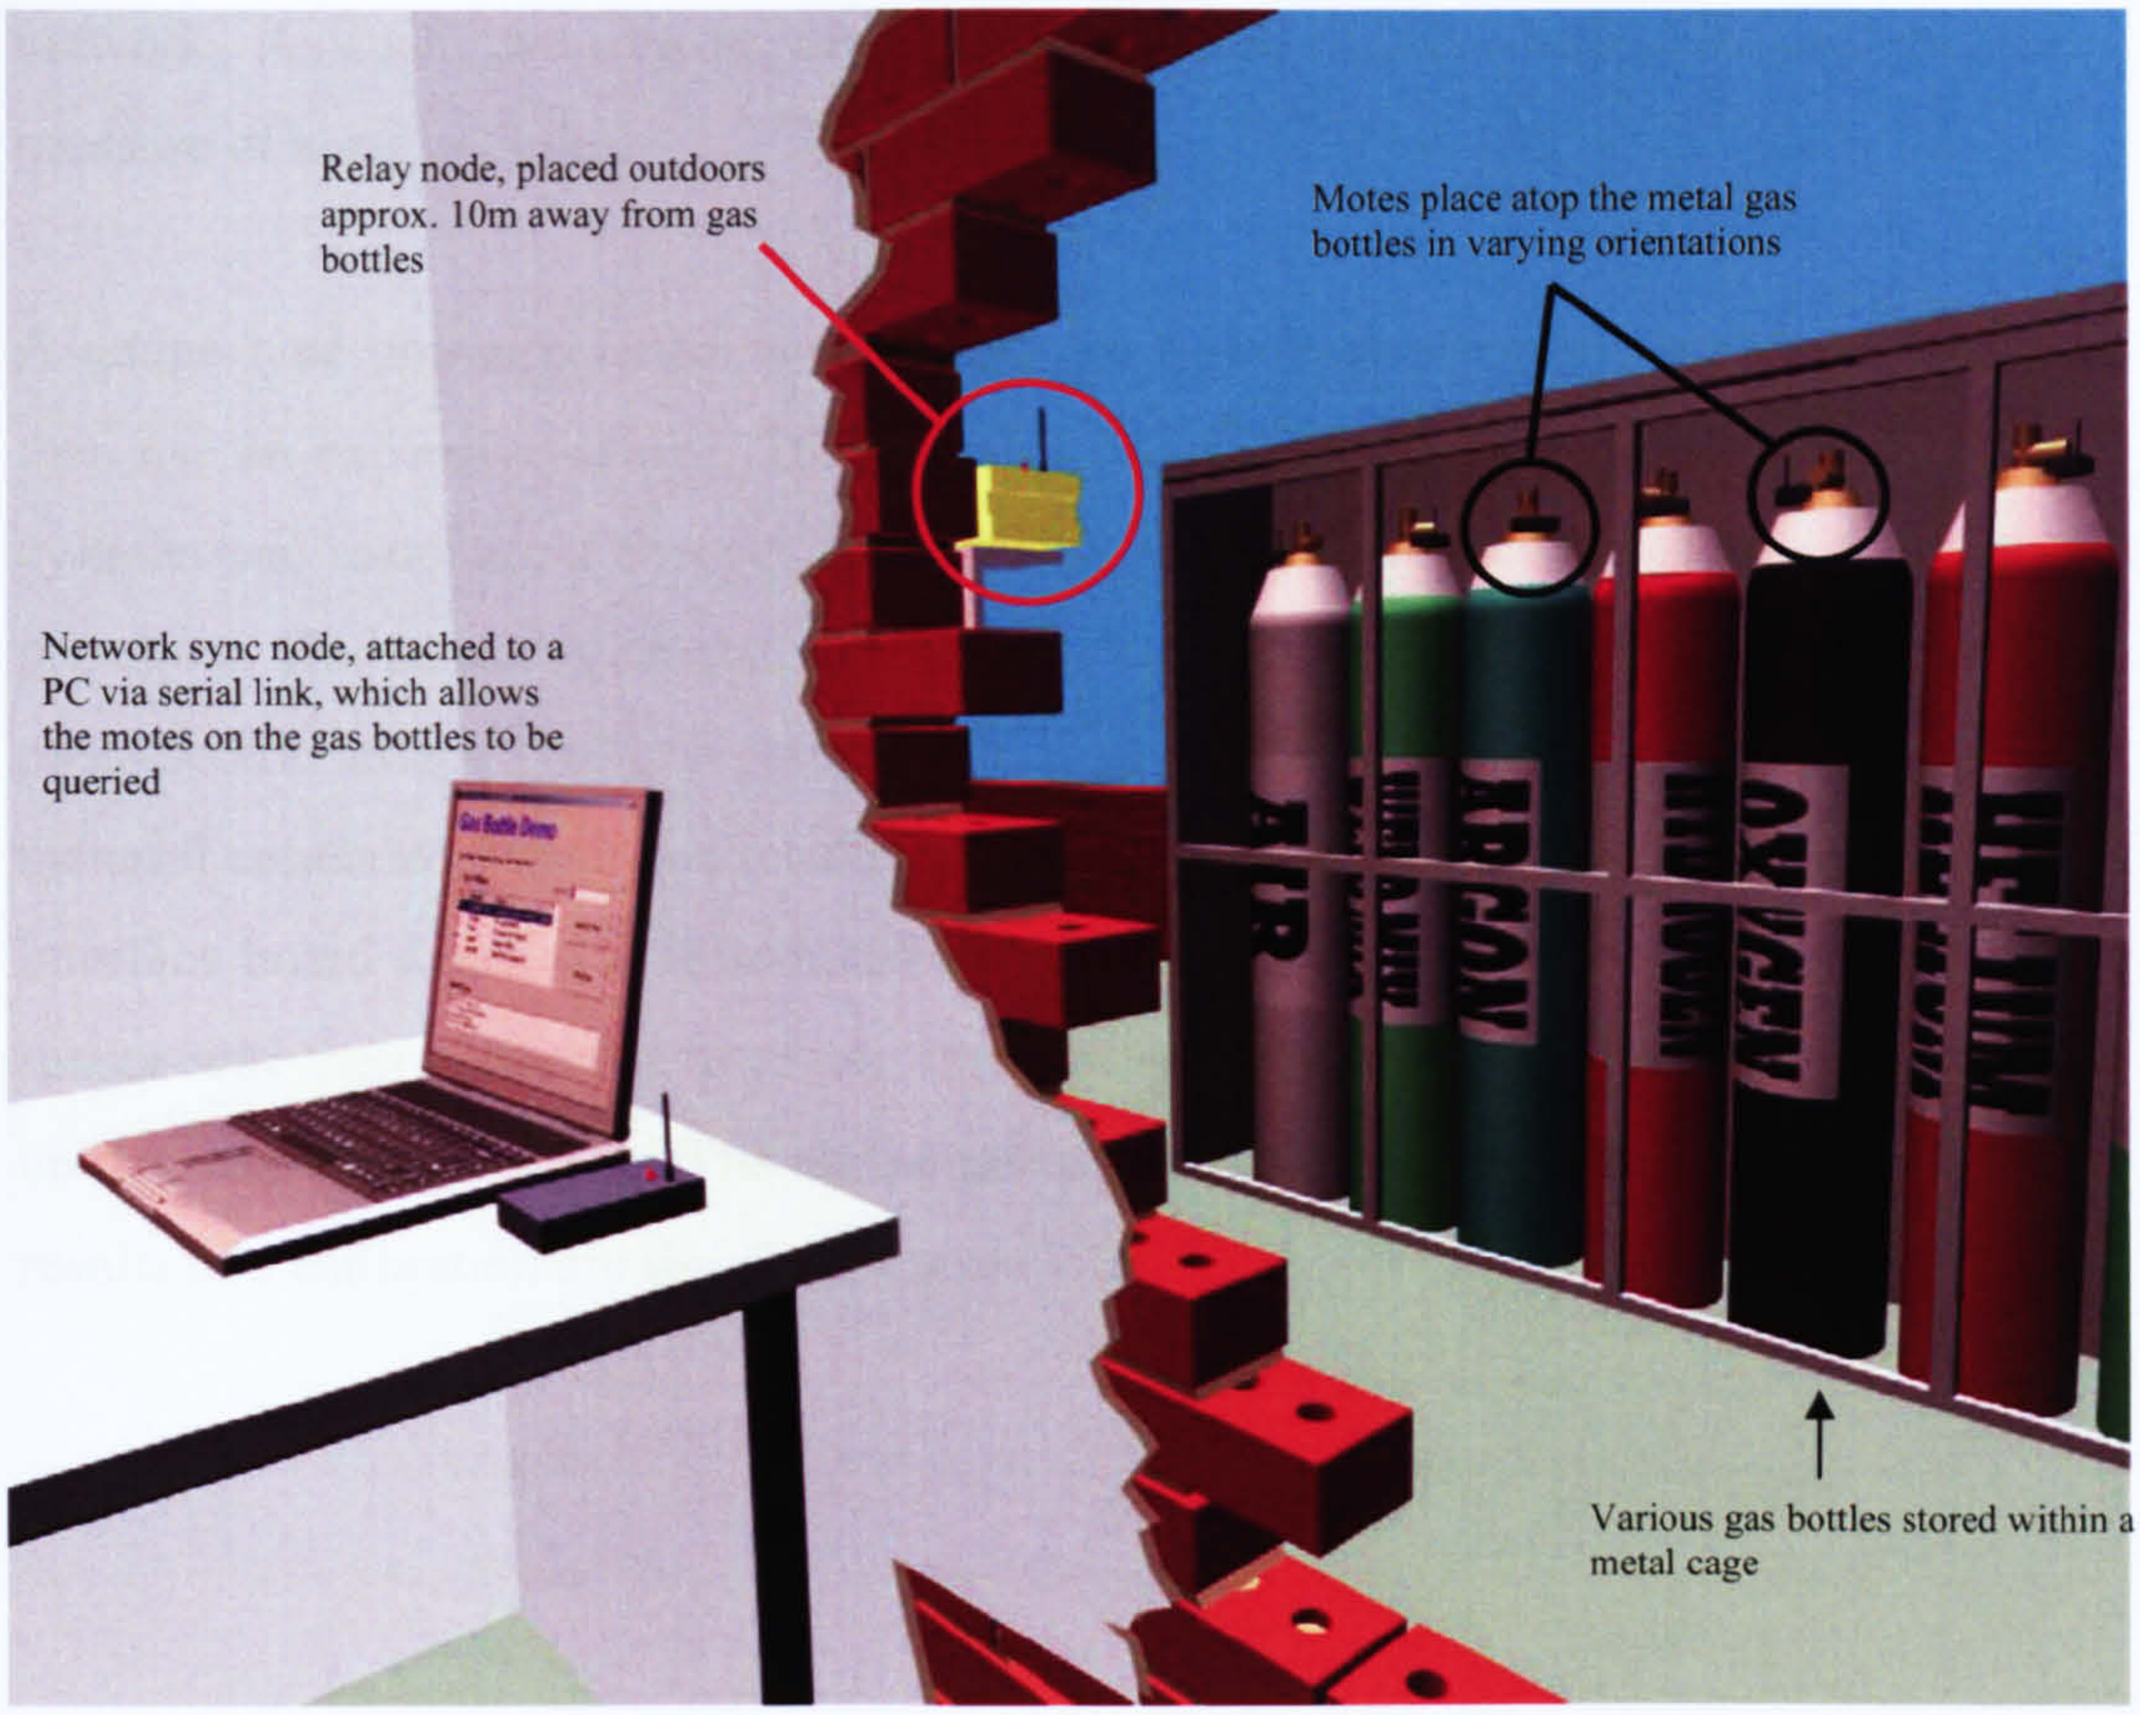

Figure 7.22 - Overview of prototype system setup

Iterations	Successfully discovered tags (%)
1	94.5%
2	99.75%
3	100%

Table 7.2 – Results of testing the prototype system in packaged gas scenario

7.4.3. Gas Cylinder Pressure Sensor

The motes have, until this point, been used for communication, and although the RSSI feature of the motes is technically sensing (i.e. – sensing of receiving signal strength) it is an integral feature. Therefore the aim was to look at external sensors which could be applied to demonstrate the capabilities of the motes as part of a true sensor

network. As such, two sensors were explorered with the first being for detecting the pressure of a gas cylinder.

A gauge type pressure sensor was chosen for this demonstration. However, rather than use an expensive sensor [165] capable of reading the high pressures that a gas cylinder can maintain, a cheaper option [166] (< £10) was found to achieve proof of principle. This sensor is capable of reading 0-100psi (0-6.9bar) and operates using the piezoelectric effect [167]; pressure at the sensor input port causes stress on the material contained within and results in an electric signal. It was attached to the mote interface board which was introduced in Chapter 4, as is shown in Figure 7.23. The sensor was then connected to an air compressor which allowed a pressure variation upto a maximum of 90psi. The air compressor setup is shown in Figure 7.24 and the results of a calibration are shown in Figure 7.25.

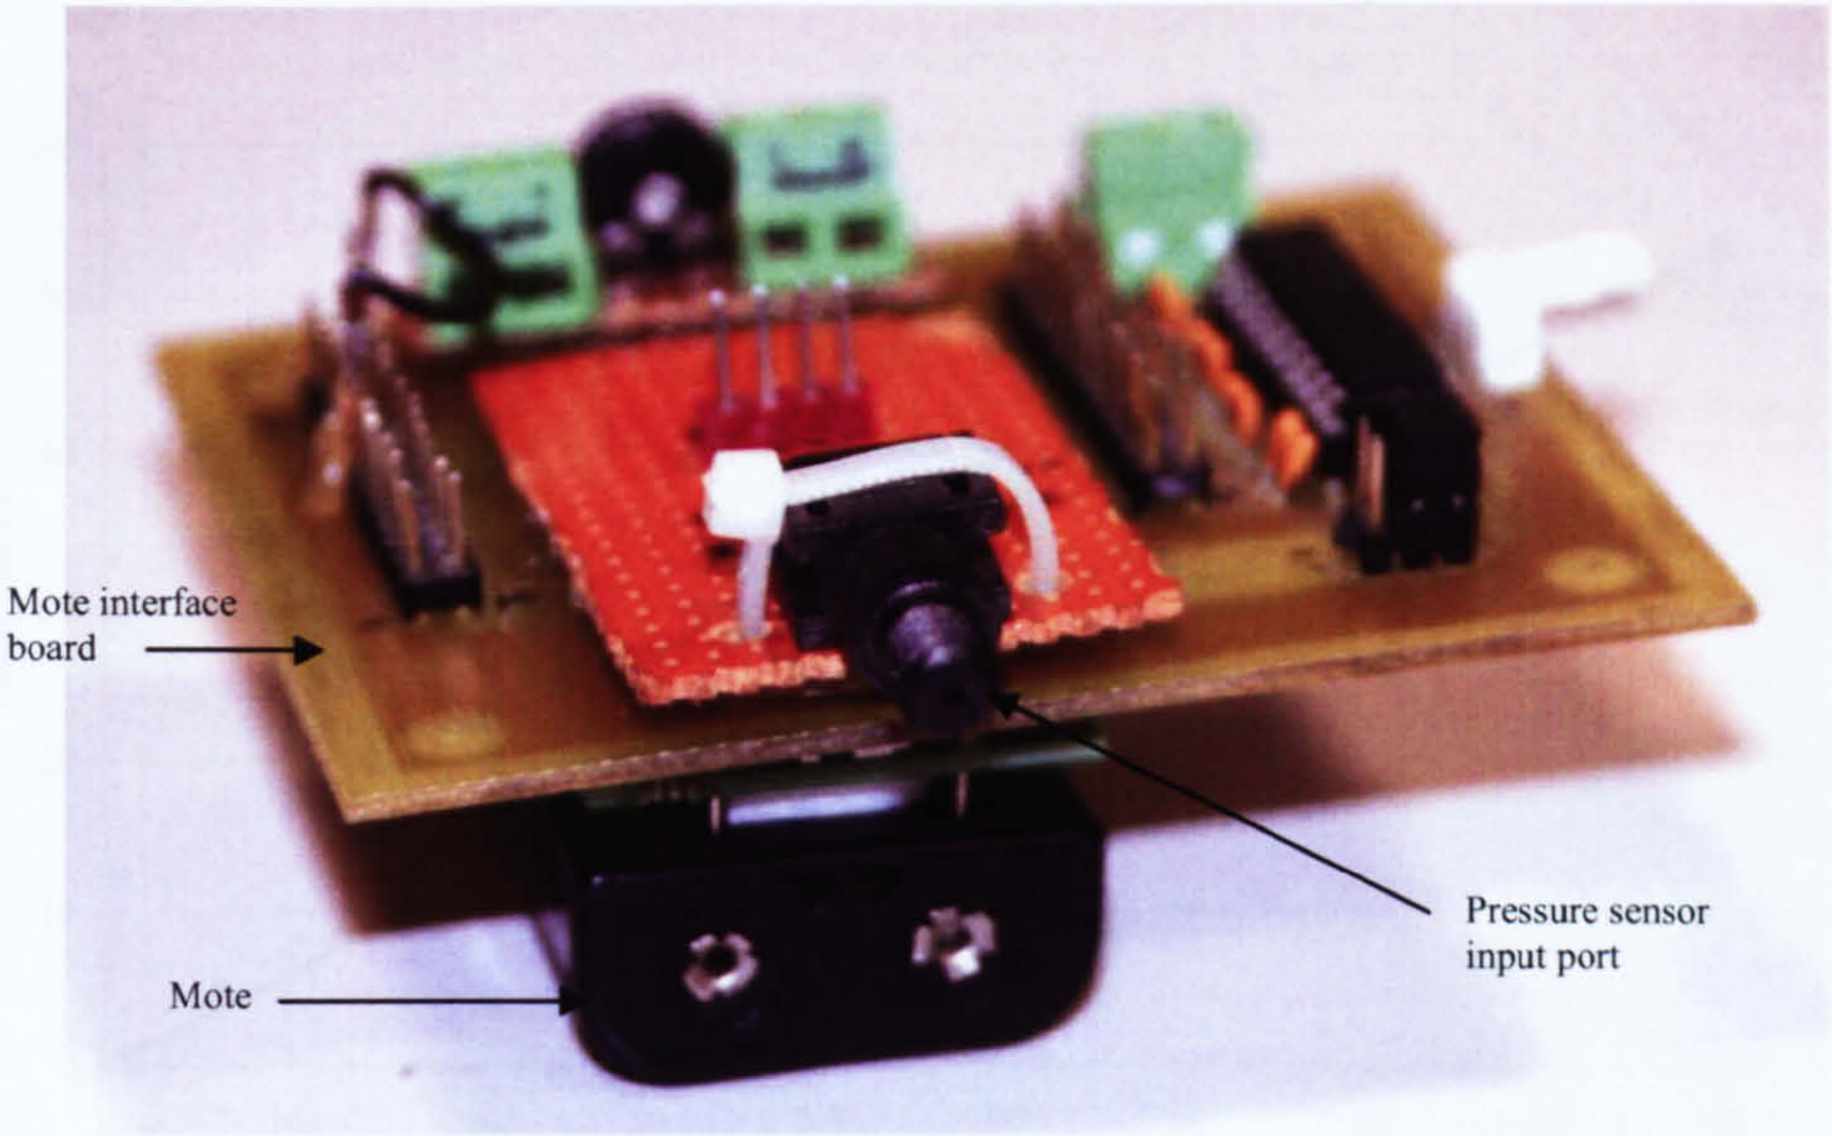

Figure 7.23 – Pressure sensor on mote daughterboard

Figure 7.24 – Pressure sensor calibration and test setup

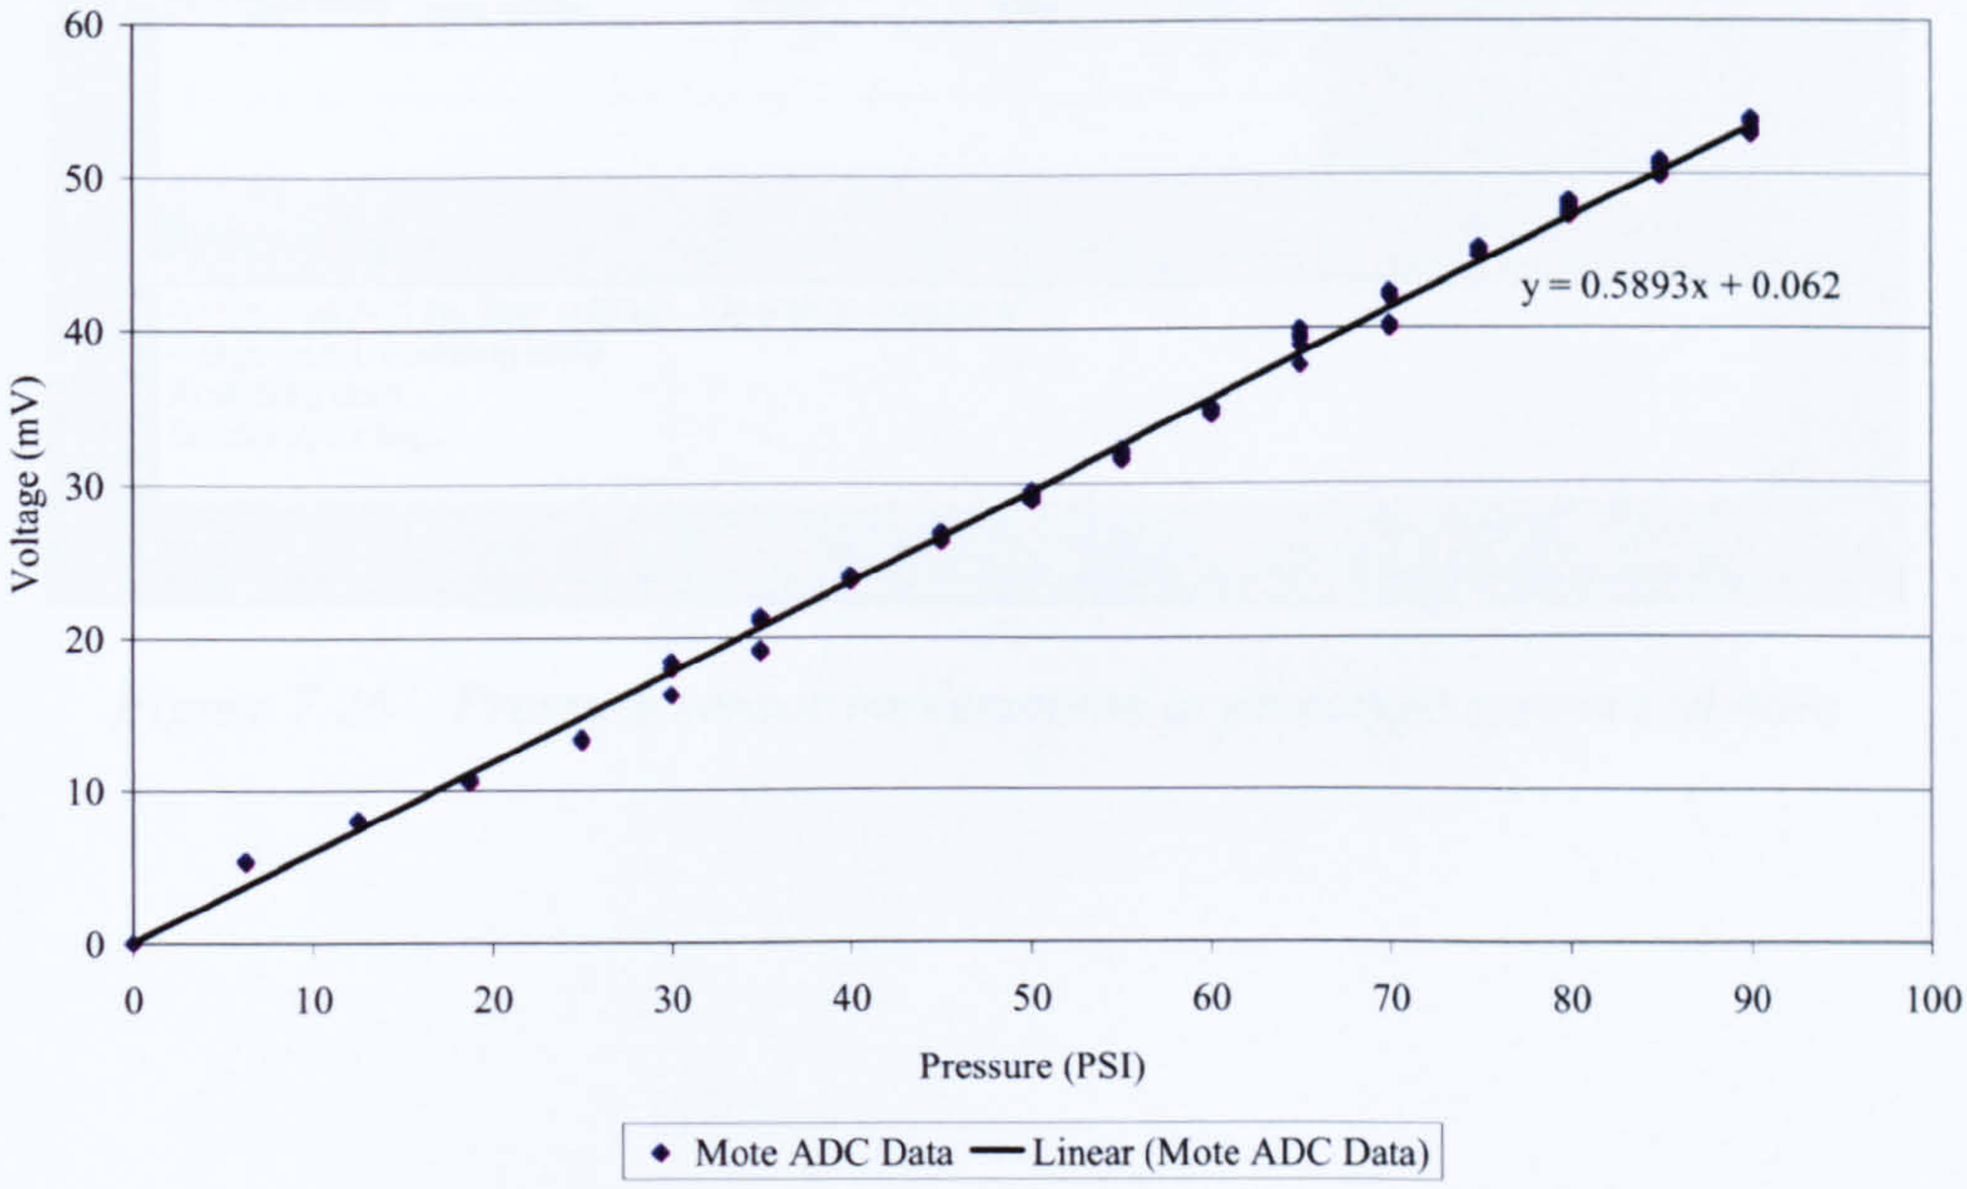

Figure 7.25 – Pressure sensor calibration results

Figure 7.25 shows that the output voltage of the pressure sensor increases linearly with pressure. Using this linear relationship it is possible to demonstrate that a mote is capable of measuring pressure. A cylinder pressure sensor could be useful for a number of purposes such as allowing automatic checking of leaks as well as automated alerts when a gas cylinder is running low; the later alert could even be triggered at the customer location in order to prompt a new delivery of gas. Figure 7.26 shows how the pressure sensor data can be integrated into the prototype system[21].

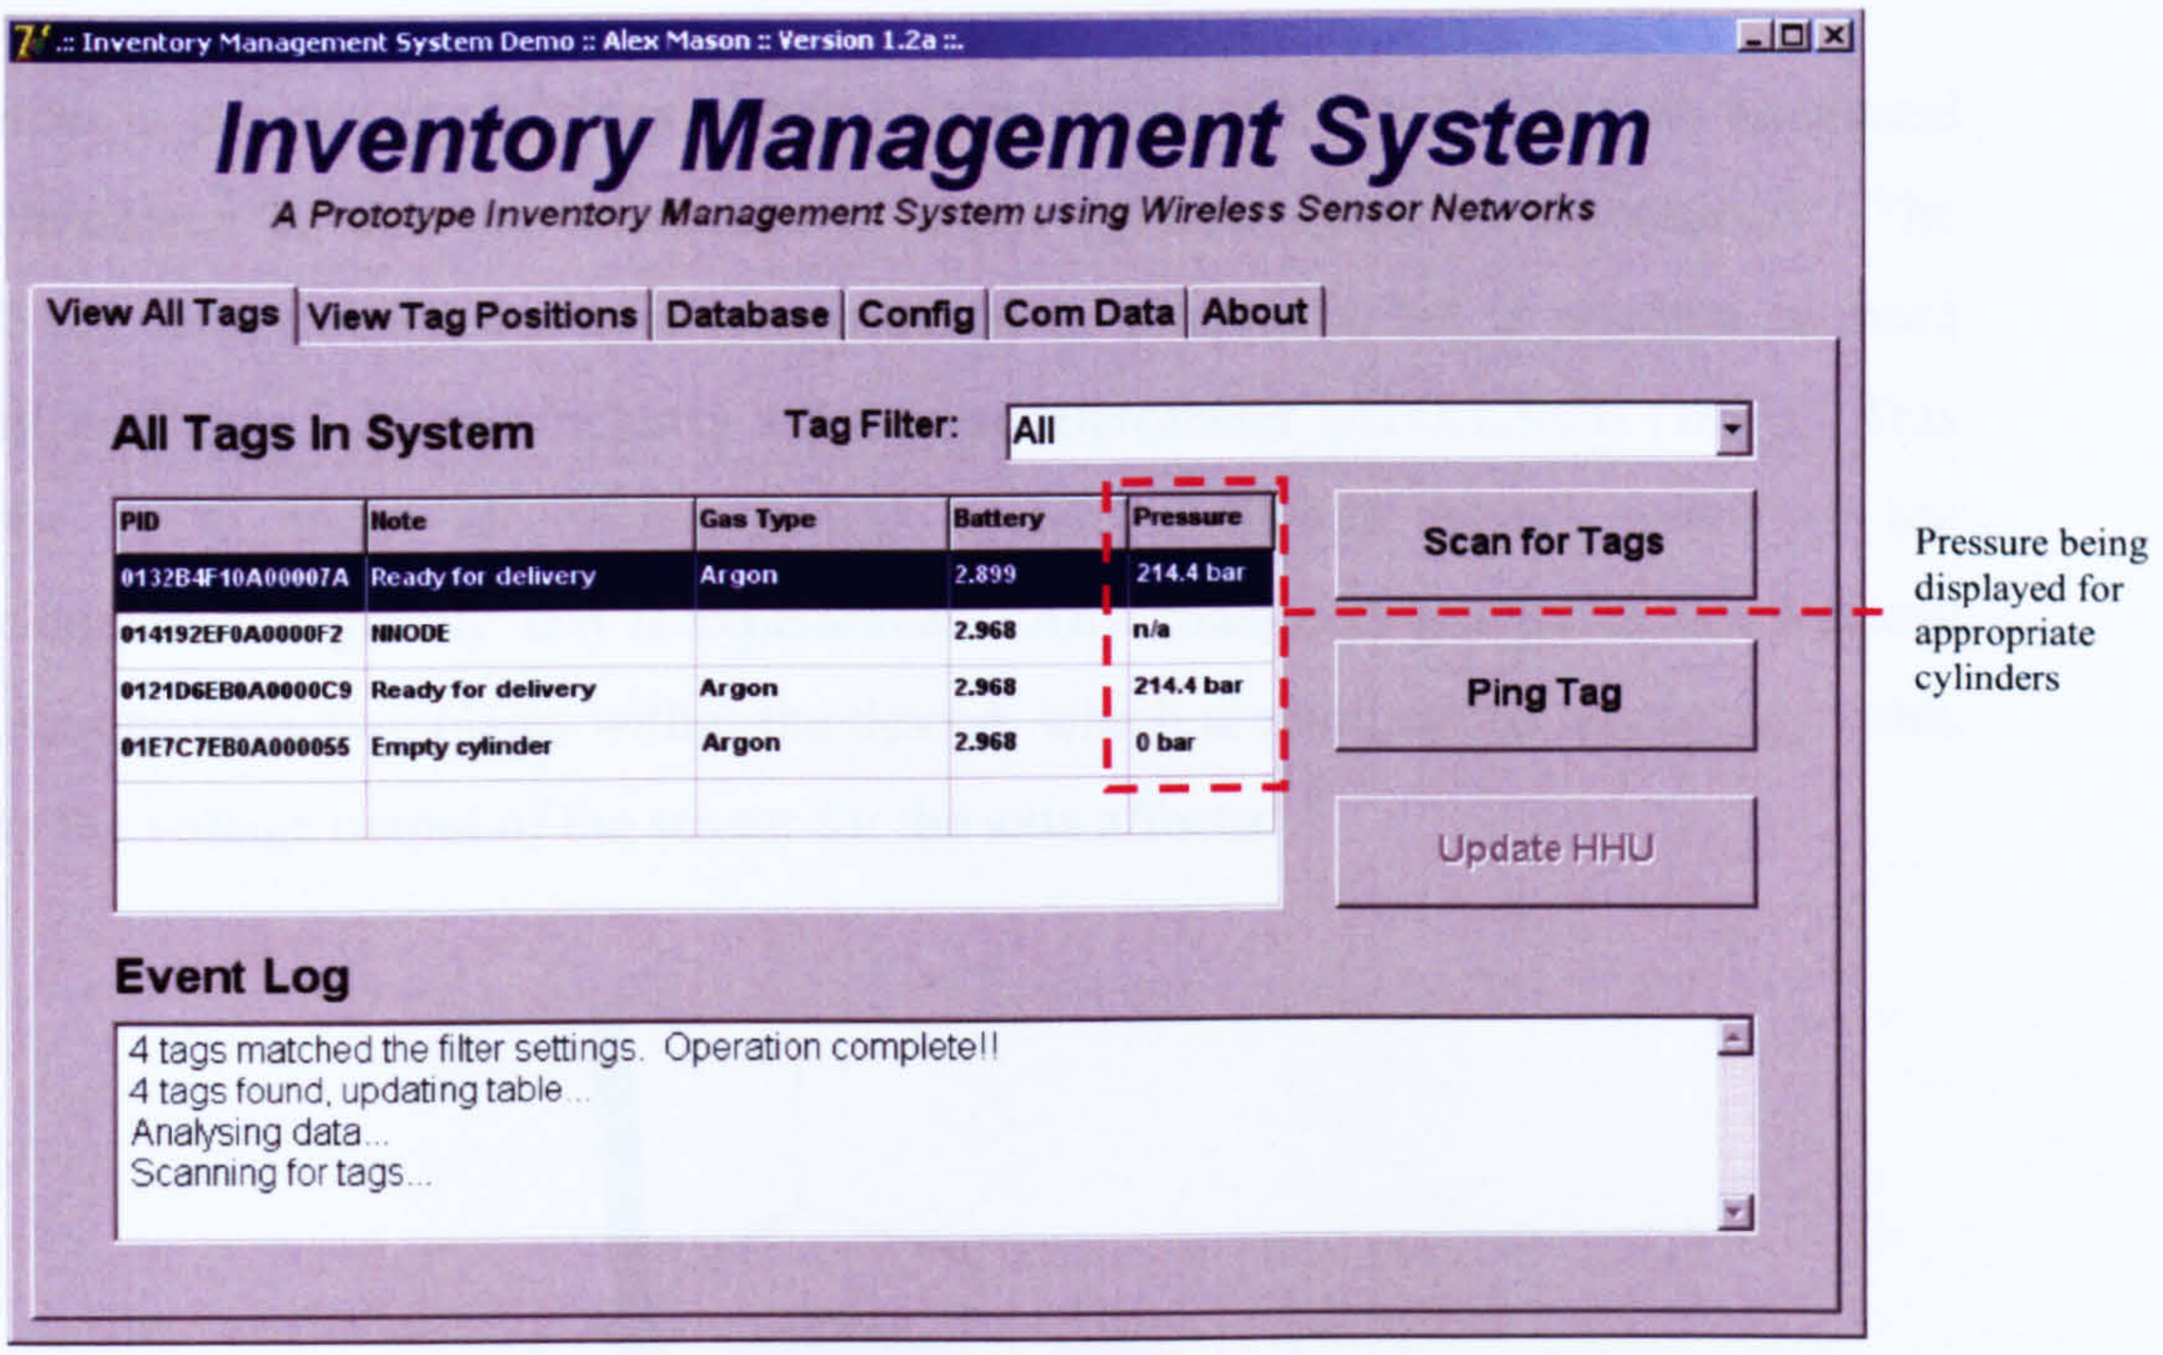

Figure 7.26 – Pressure sensor implemented in prototype system software

[21] Note that the values shown here represent real measured values, but they are scaled up by a factor of 32 in order to make them appear to be realistic gas cylinder pressures for demonstration purposes.

7.4.4. Detecting Cylinder Orientation

Whilst a pressure sensor could add value to a packaged gas inventory management system, it is not the only sensor which could be useful. Acetylene, a highly flammable gas which is used in oxyacetylene cutting and welding has strict storage guidelines [168]. The gas is unstable when pressurised, and so for storage in a pressurised cylinder it must be dissolved in acetone. The cylinder is also filled with a porous material which acts like a sponge to soak up the acetone. As a result of this arrangement, the cylinder must be kept upright in order to prevent acetone leaking from the cylinder – if such a leak does occur then a build up of acetylene gas may occur which could result in an explosion.

In order to provide an additional safety feature the mote tag could have an integrated accelerometer so that the orientation of a gas cylinder could be determined. The MTS420 sensor board was briefly introduced in Chapter 3, but is pictured in more detail in Figure 7.27 and includes a 2-axis accelerometer (ADXL202E [169]). This sensor is a micro electro-mechanical system (MEMS) which measures the acceleration and gravity that it experiences. Any change in gravity or acceleration repositions capactive plates within the device, which are suspended on springs – this alters the voltage output of the sensor for the axis affected.

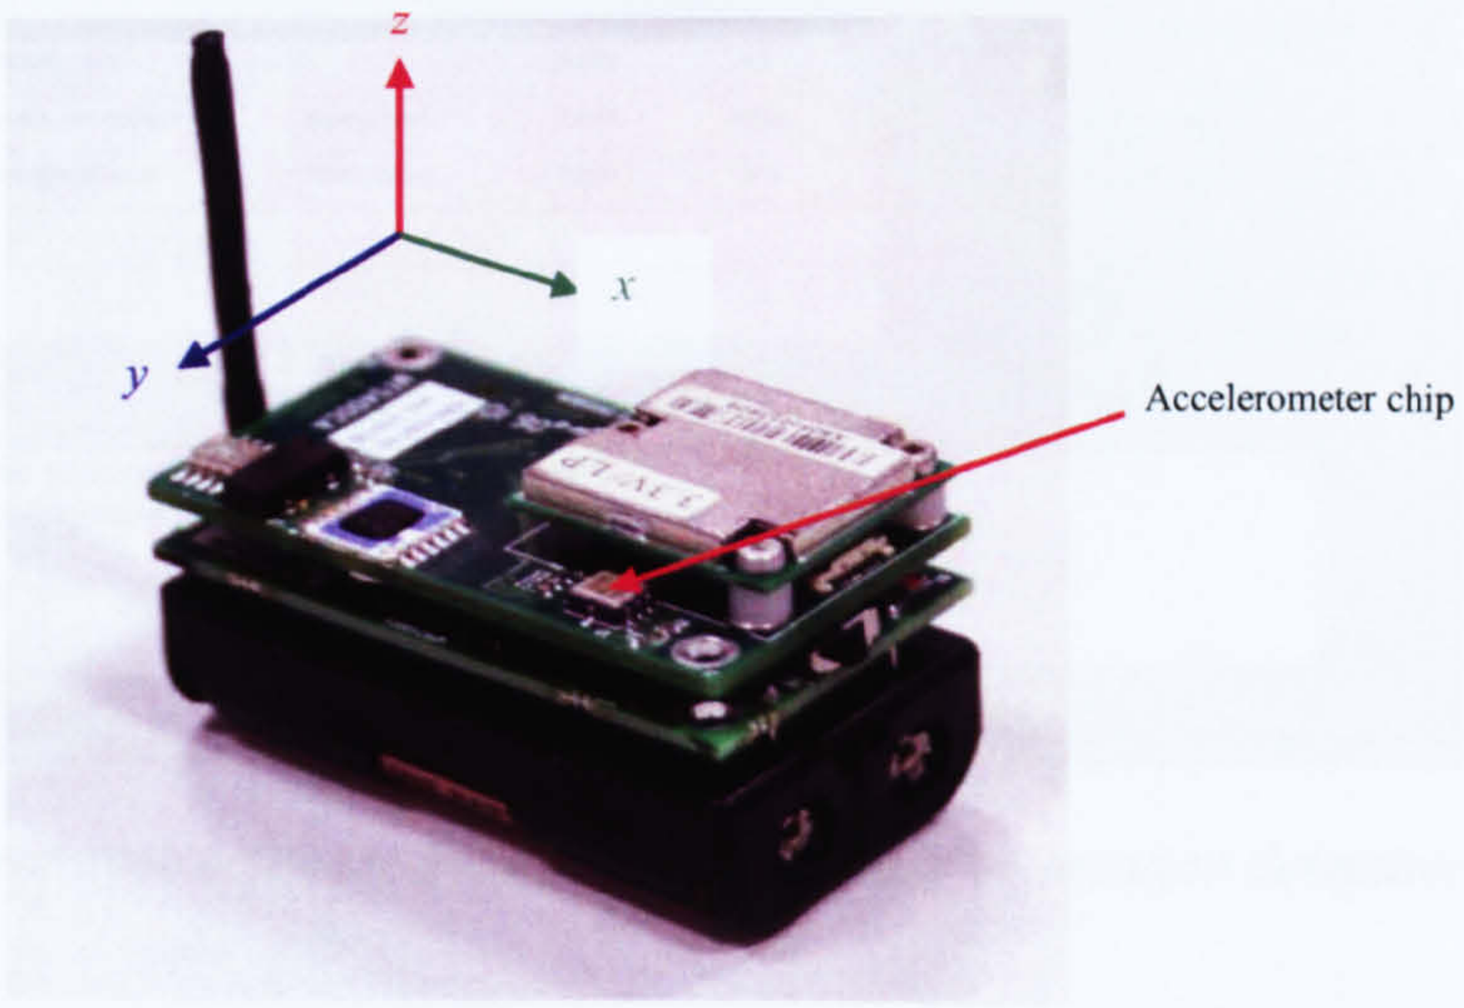

Figure 7.27 – Micax mote with MTS400 sensor board attached

Since the accelerometer covers only 2 axes, it can only detect rotations in the *xz* and *yz* planes. Through experimentation it was found that the accelerometer x-axis output varied over the range of 1.16-1.46V, whilst the *y*-axis output varied from 1.35-1.65V. In order to demonstrate this sensor in a packaged gas scenario, it was assumed that the mote would be orientated such that its antenna socket would be perpendicular to the earth and be pointing skyward. In this position the *x*-axis output is 1.31V and the y-axis output 1.65V. It was assumed that a deviation greater than ±0.05V on either axis would indicate that the gas cylinder was not upright. As such, this logic was incorporated into the prototype system software; since the tag is required to send two sets of sensor data it filled the `data1` message field with the *x* and *y* axis data on alternate scans. Once the GUI had both readings it sought to determine whether the gas cylinder was upright or not, as shown in Figure 7.28.

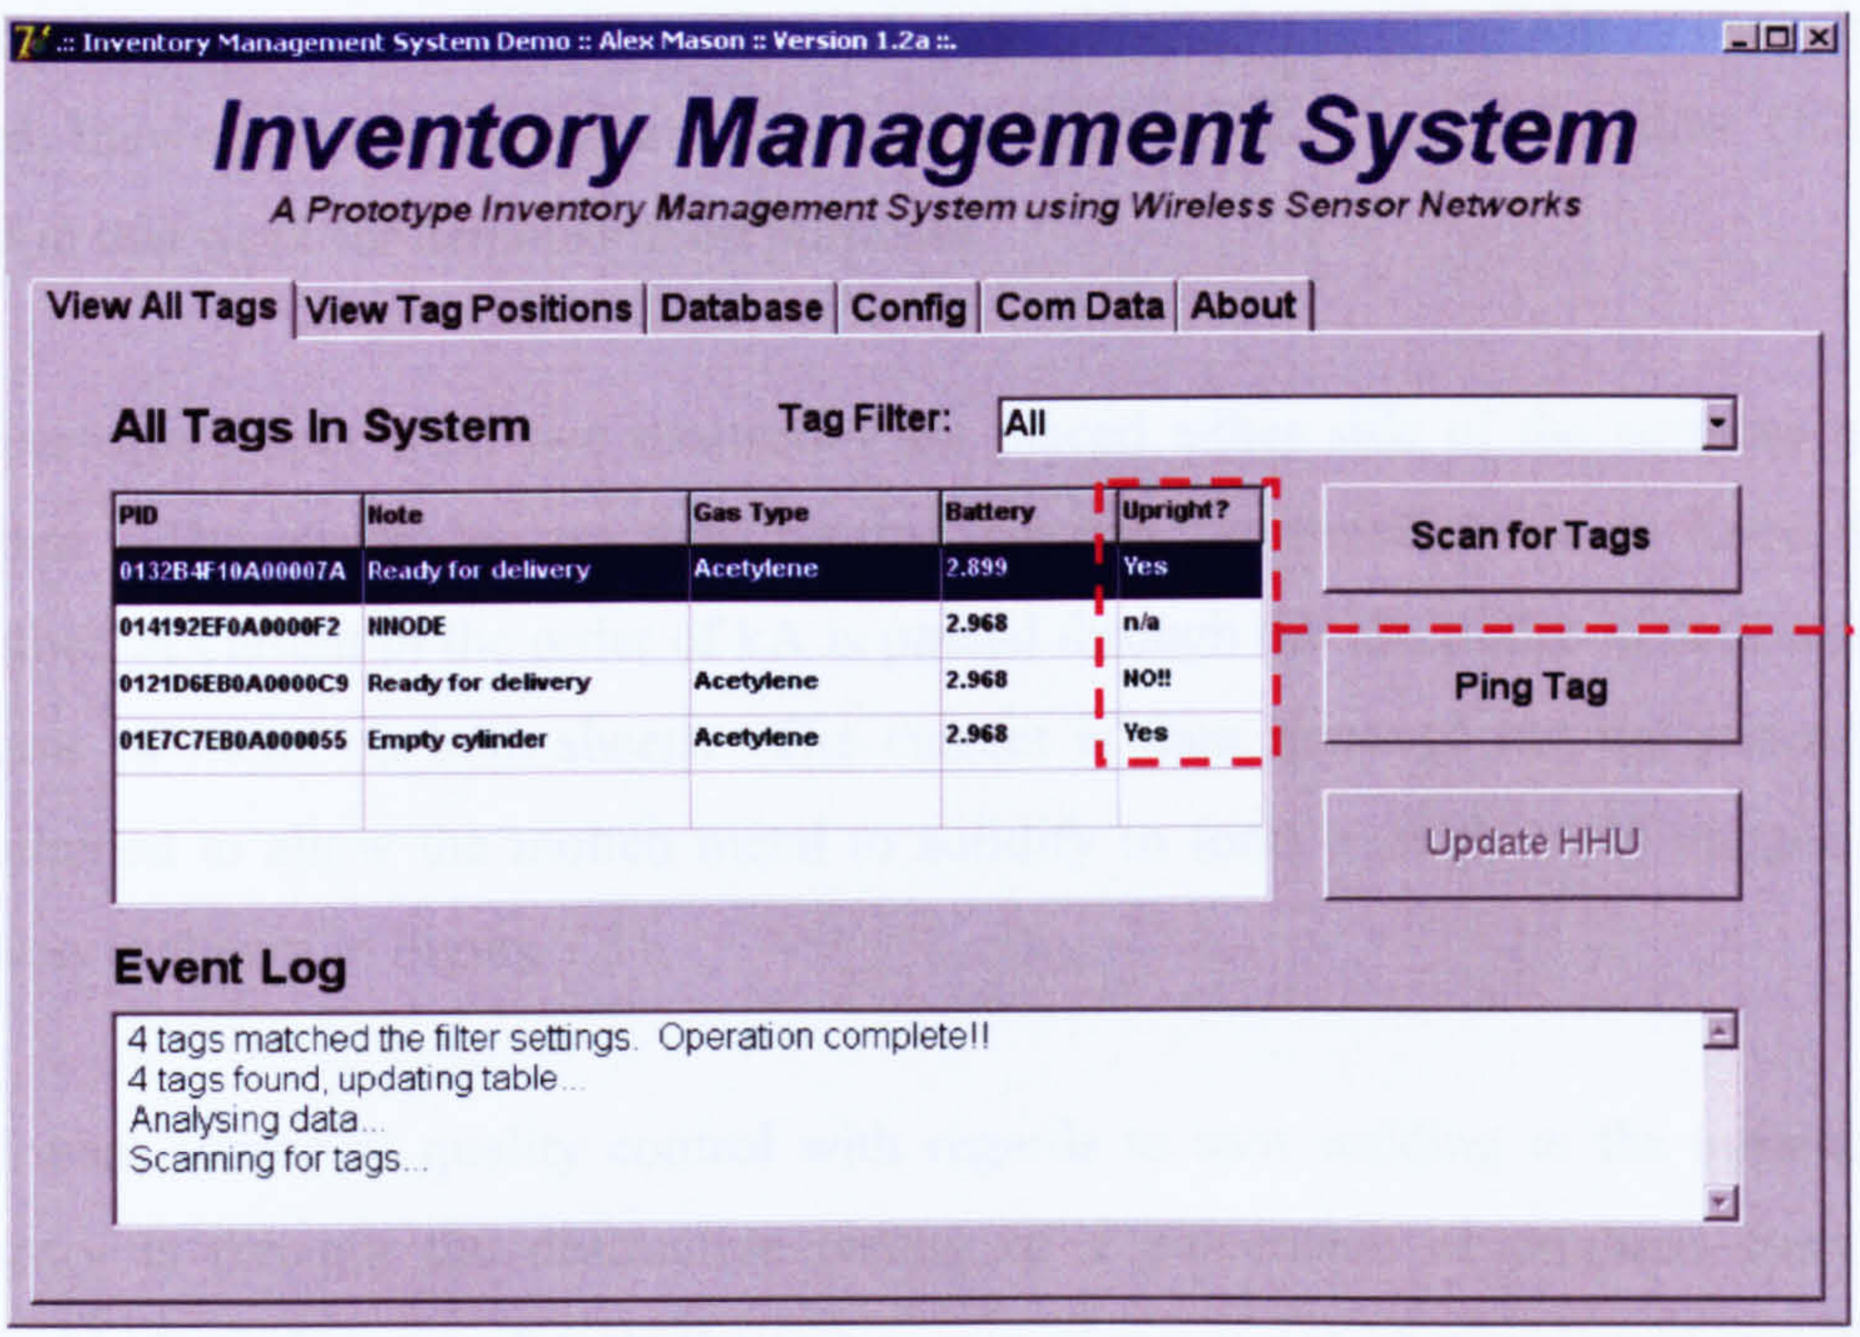

Figure 7.28 – Prototype software incorporating cylinder orientation detection

7.5. The Automotive Industry

7.5.1. A Brief Overview of Spot Welding

Thus far the discussion of WSN's for spot welding has been largely limited to Chapter 1 as it is not related to inventory management. Further industrial applications of WSN's were identified and an investigation into their use for spot weld monitoring was undertaken.

As previously mentioned spot welding is a quick and easy way of joining two, often metallic, materials [170]. The automotive industry relies heavily on it since it is useful in rapid, high volume welding applications. It can be used to join sheet metal of varying thicknesses and is adaptable to a variety of materials. In addition, it is extremely economical and the process can be largely automated. Figure 7.29 shows a robotic spot welder, whilst Figure 7.30 shows a pedestal welder. The later type of welder is still used in the automotive industry, although not as widely as the robotic types; they are useful for research too and a TECNA 4621 pedestal welder [171] is used in this work for demonstration purposes.

To perform a spot weld two electrodes are placed either side of the surfaces to be welded. The electrodes are then pressed against the metal sheets to force them together. A current in the order of kA is passed through the electrodes to melt a small volume of metal on both sheets. The current is then removed but the pressure is maintained to allow the molten metal to solidify to form a single weld nugget; this process is shown in Figure 7.31.

The main source of quality control with regards to spot welding in the automotive industry is through the destructive testing of a percentage of products from the assembly line. From such tests the failure rate of welds can be established and if it is significant then a whole batch of products may be rejected in order to maintain quality standards. Such testing is costly in terms of labour however and so in recent years the automotive industry has been looking for alternative methods of quality control which may be employed. To this end, research [5] has considered the application of multiple

sensors to monitor welding parameters during the spot welding process. The aim of this is the improve quality via a real-time monitoring and correction which would do away with the need for destructive testing.

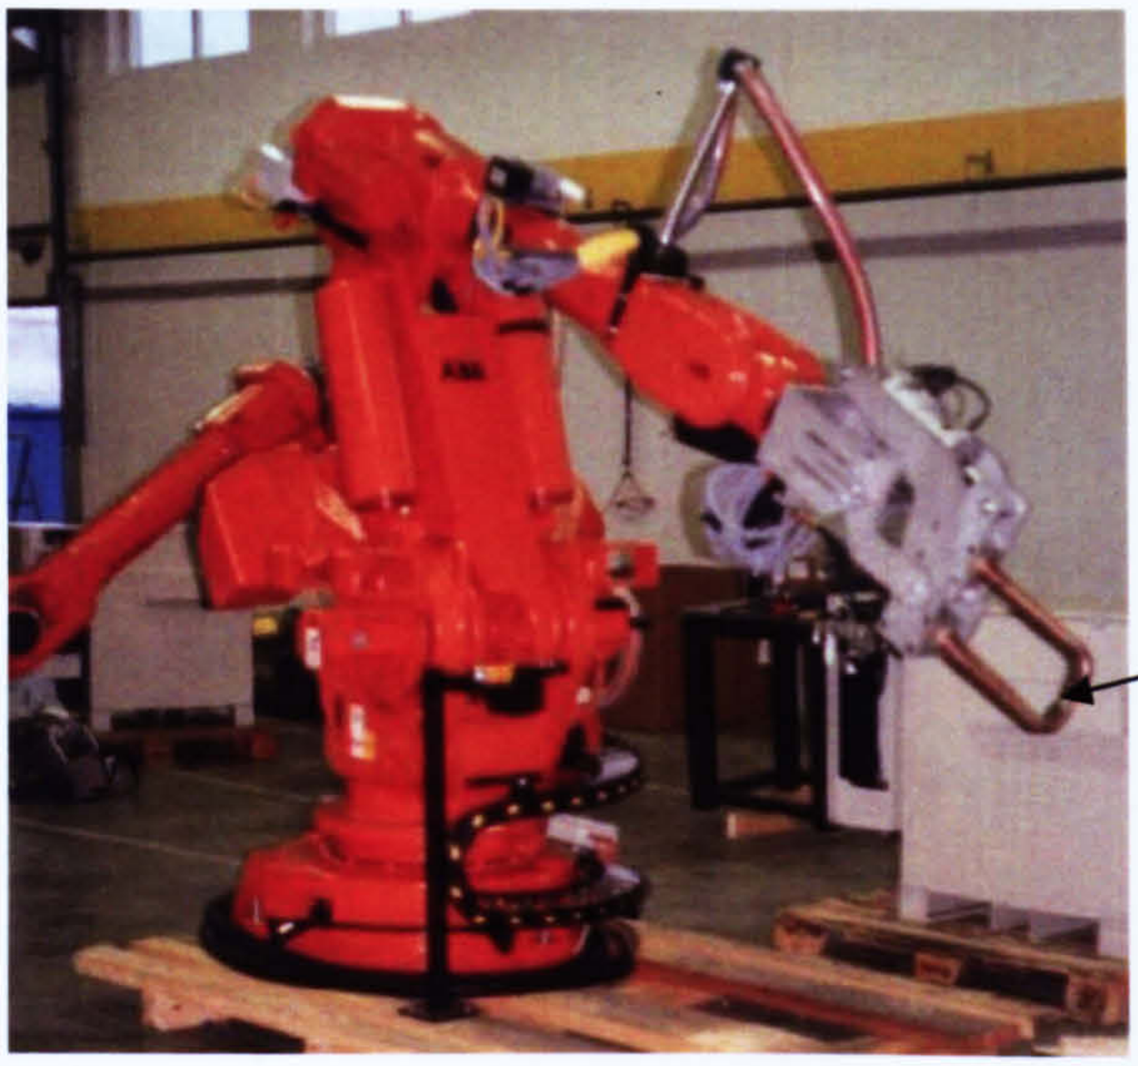

Figure 7.29 – Robotic spot welder

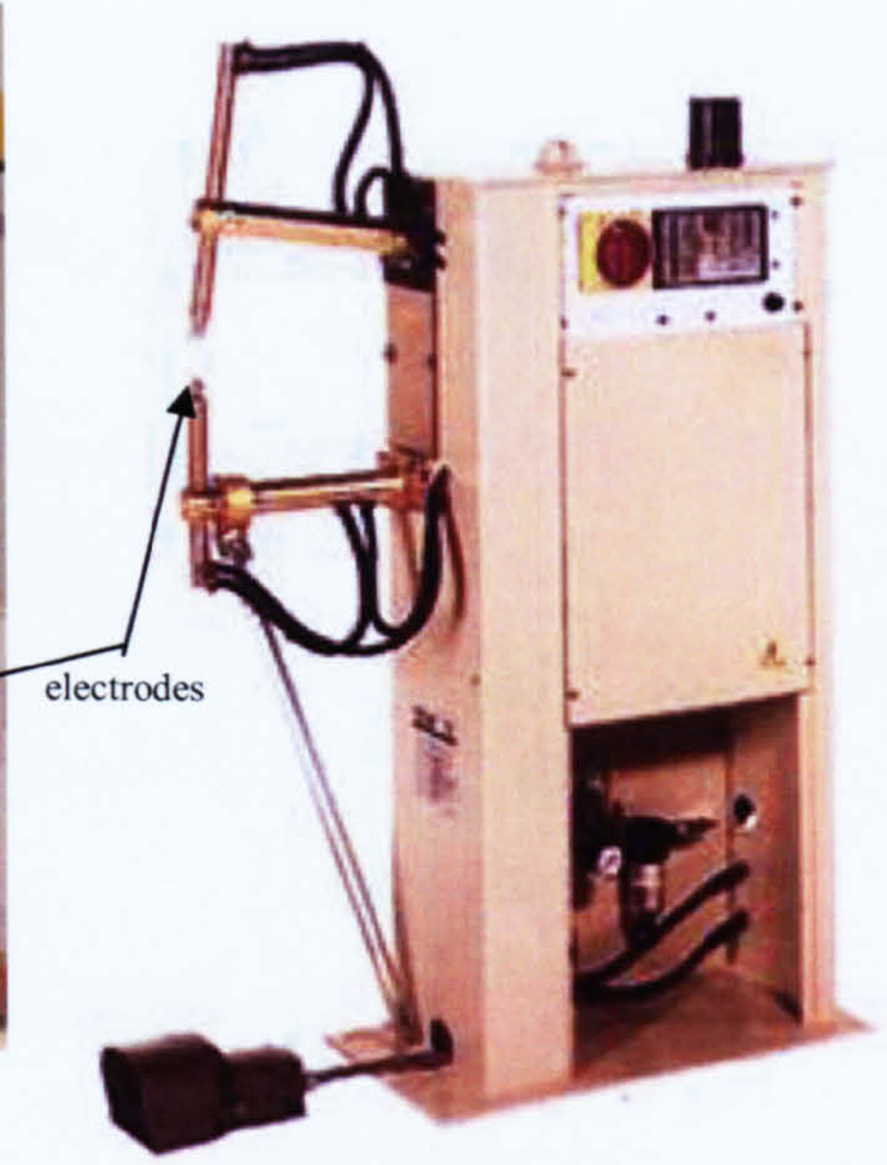

Figure 7.30 – Pedestal spot welder

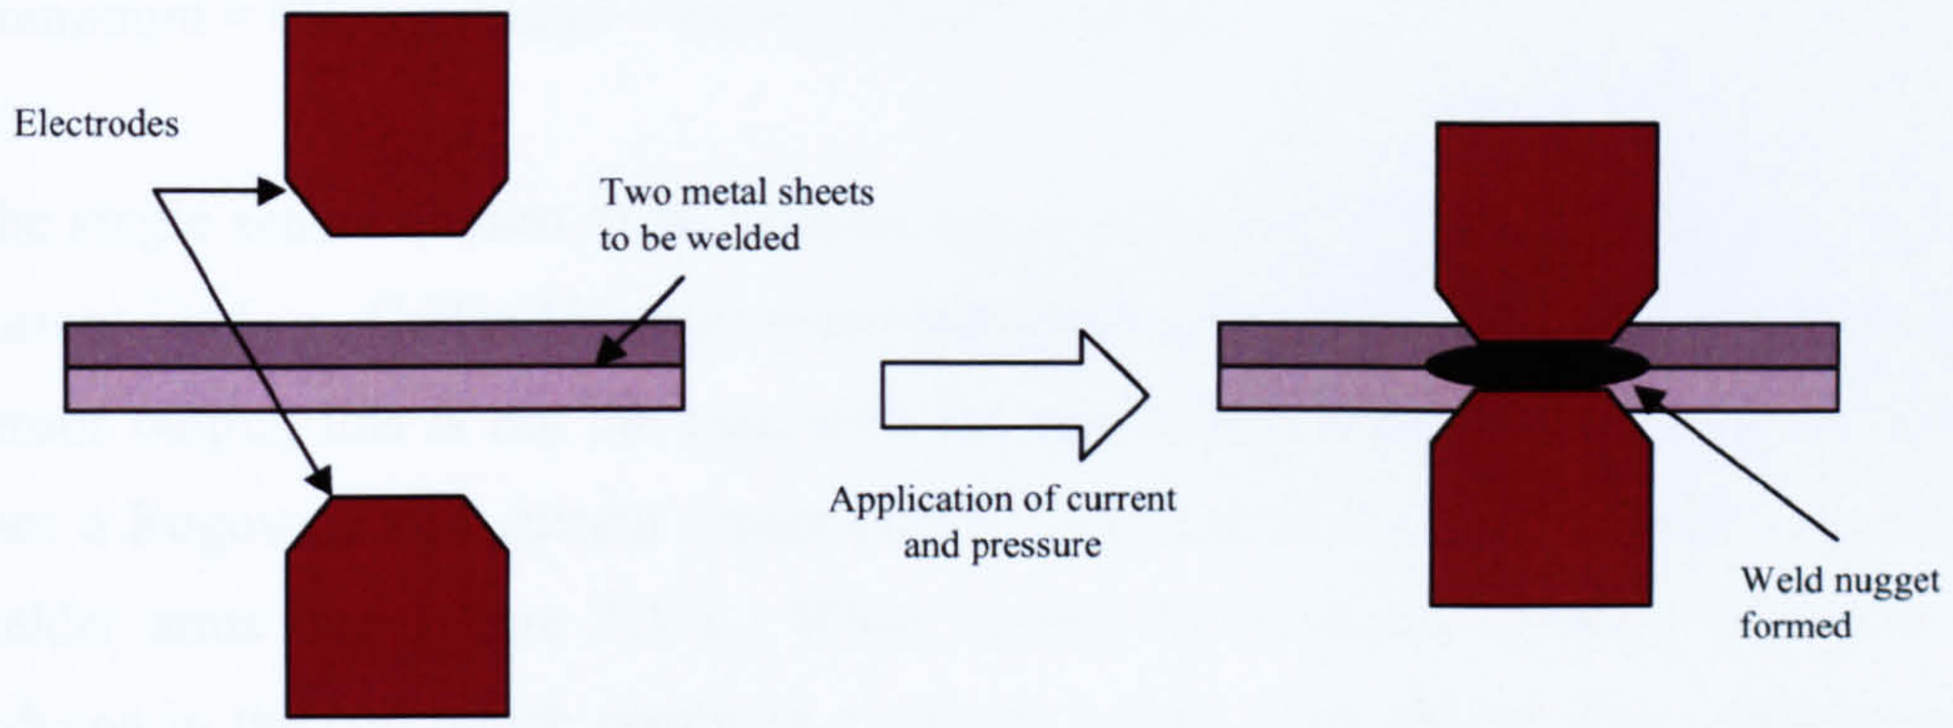

Figure 7.31 – Formation of a weld nugget using spot welding

Sensors included within the cluster collect current, voltage, infrared, acoustic and ultrasonic data. This is fed into a neural network in order to determine whether a weld is of acceptable quality. The function of each sensor is not important here, and so will not be discussed further.

As a further application, power could be saved by reducing the welding current to the bare minimum required [7]. Therefore the aim is not only to monitor the welding process, but also to attempt to minimise the welding current required.

In the existing system the sensors are wired to a computer. This is not feasible in an industrial situation since the highly mobile robotic spot welders (see Figure 7.29) could be hindered by such cabling in a retro-fit sensor system. Therefore the opportunity arose to apply a WSN to improve the applicability of the system.

7.5.2. Implementation of a Spot Welding Sensor

To demonstrate the principle of WSN operating in an automotive scenario, it was decided to utilise a single sensor rather than implement the entire sensor cluster previous mentioned, as each sensor would require some amount of conditioning logic for its output to be compatible with the Micax mote ADC interface. This is a result of many of the sensors having an output range greater than that of the mote ADC (minimum = 0V, maximum = current battery voltage).

The single sensor chosen to be adapted for connectivity to the Micax motes was for current sensing. Cullen [5] notes many industrial spot welders have an inbuilt current sensor output; this is not the case with the spot welder used in this research which uses a Rogowski coil current sensor [172]. This sensor is placed around one of the welder arms (see Figure 7.37). When current flows during welding a current is induced in the coil which results in a voltage output. It is driven from a 9V battery and its output varies between ±4.5V, which meant that a conditioning circuit was required.

A conditioning circuit was designed with the assistance of the National Instruments Multisim software package [173]. Multisim allows for rapid and inexpensive design of complex circuits, and even provides virtual tools for testing them. The circuit itself was designed to take the output of the current sensor and reduce the signal amplitude as well as shifting it such that it varied over the range of approximately 0.1-2.1V. Therefore, the circuit includes two pairs of resistors acting as voltage dividers; the first pair reduces the sensor output voltage to ±2.25V. The second pair of resistors further reduces the signal amplitude whilst the DC input (provided in reality by a 5V regulator) serves to shift the signal to its centre voltage of around 1.25V. Figure 7.32 shows the circuit as designed in Multisim, whilst Figure 7.33 shows the measured results from the circuit once it was built. Measurements were taken with a GDS-810S oscilloscope.

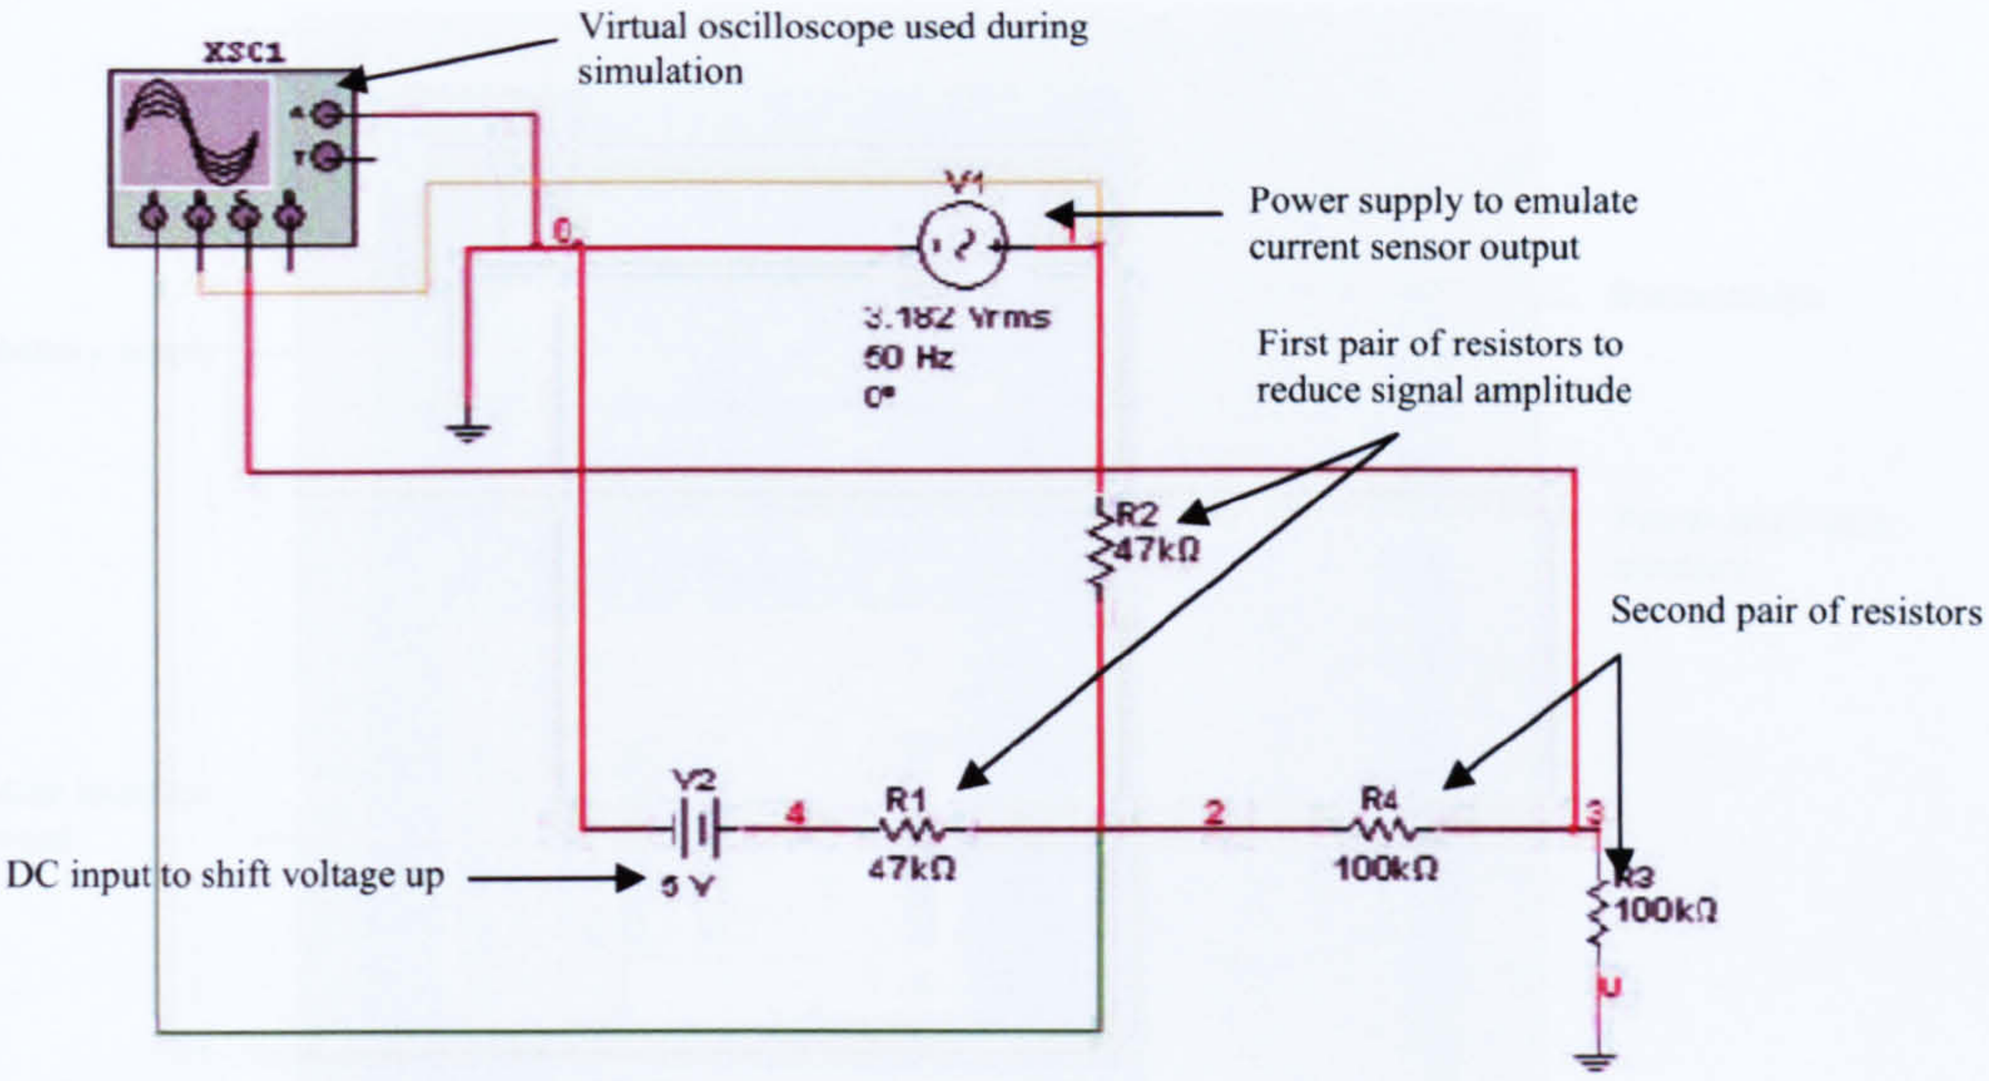

Figure 7.32 – Multisim conditioning circuit overview

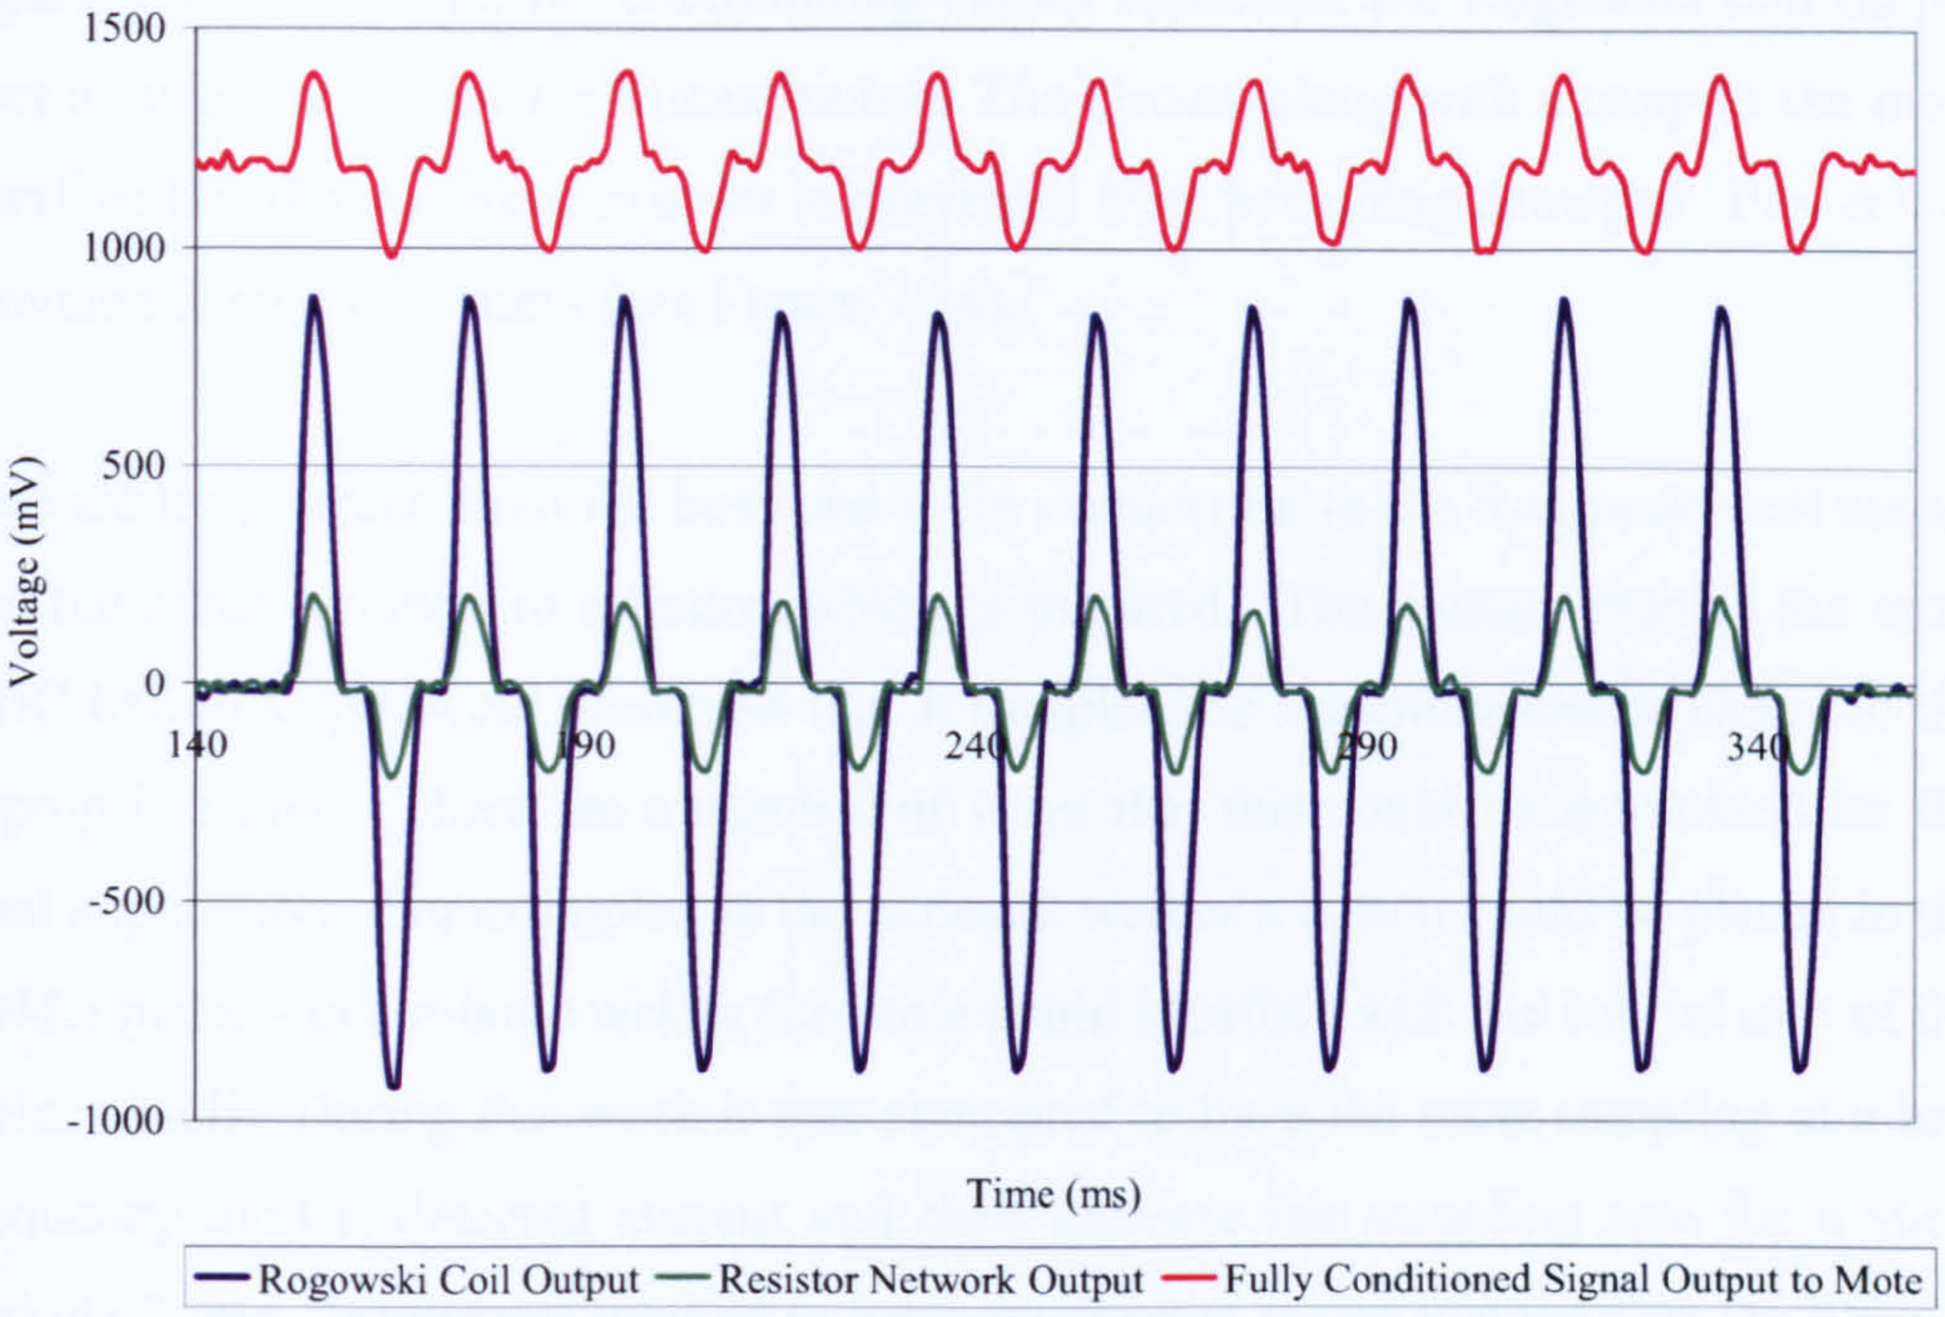

Figure 7.33 – Measured signal levels from sensor conditioning circuit

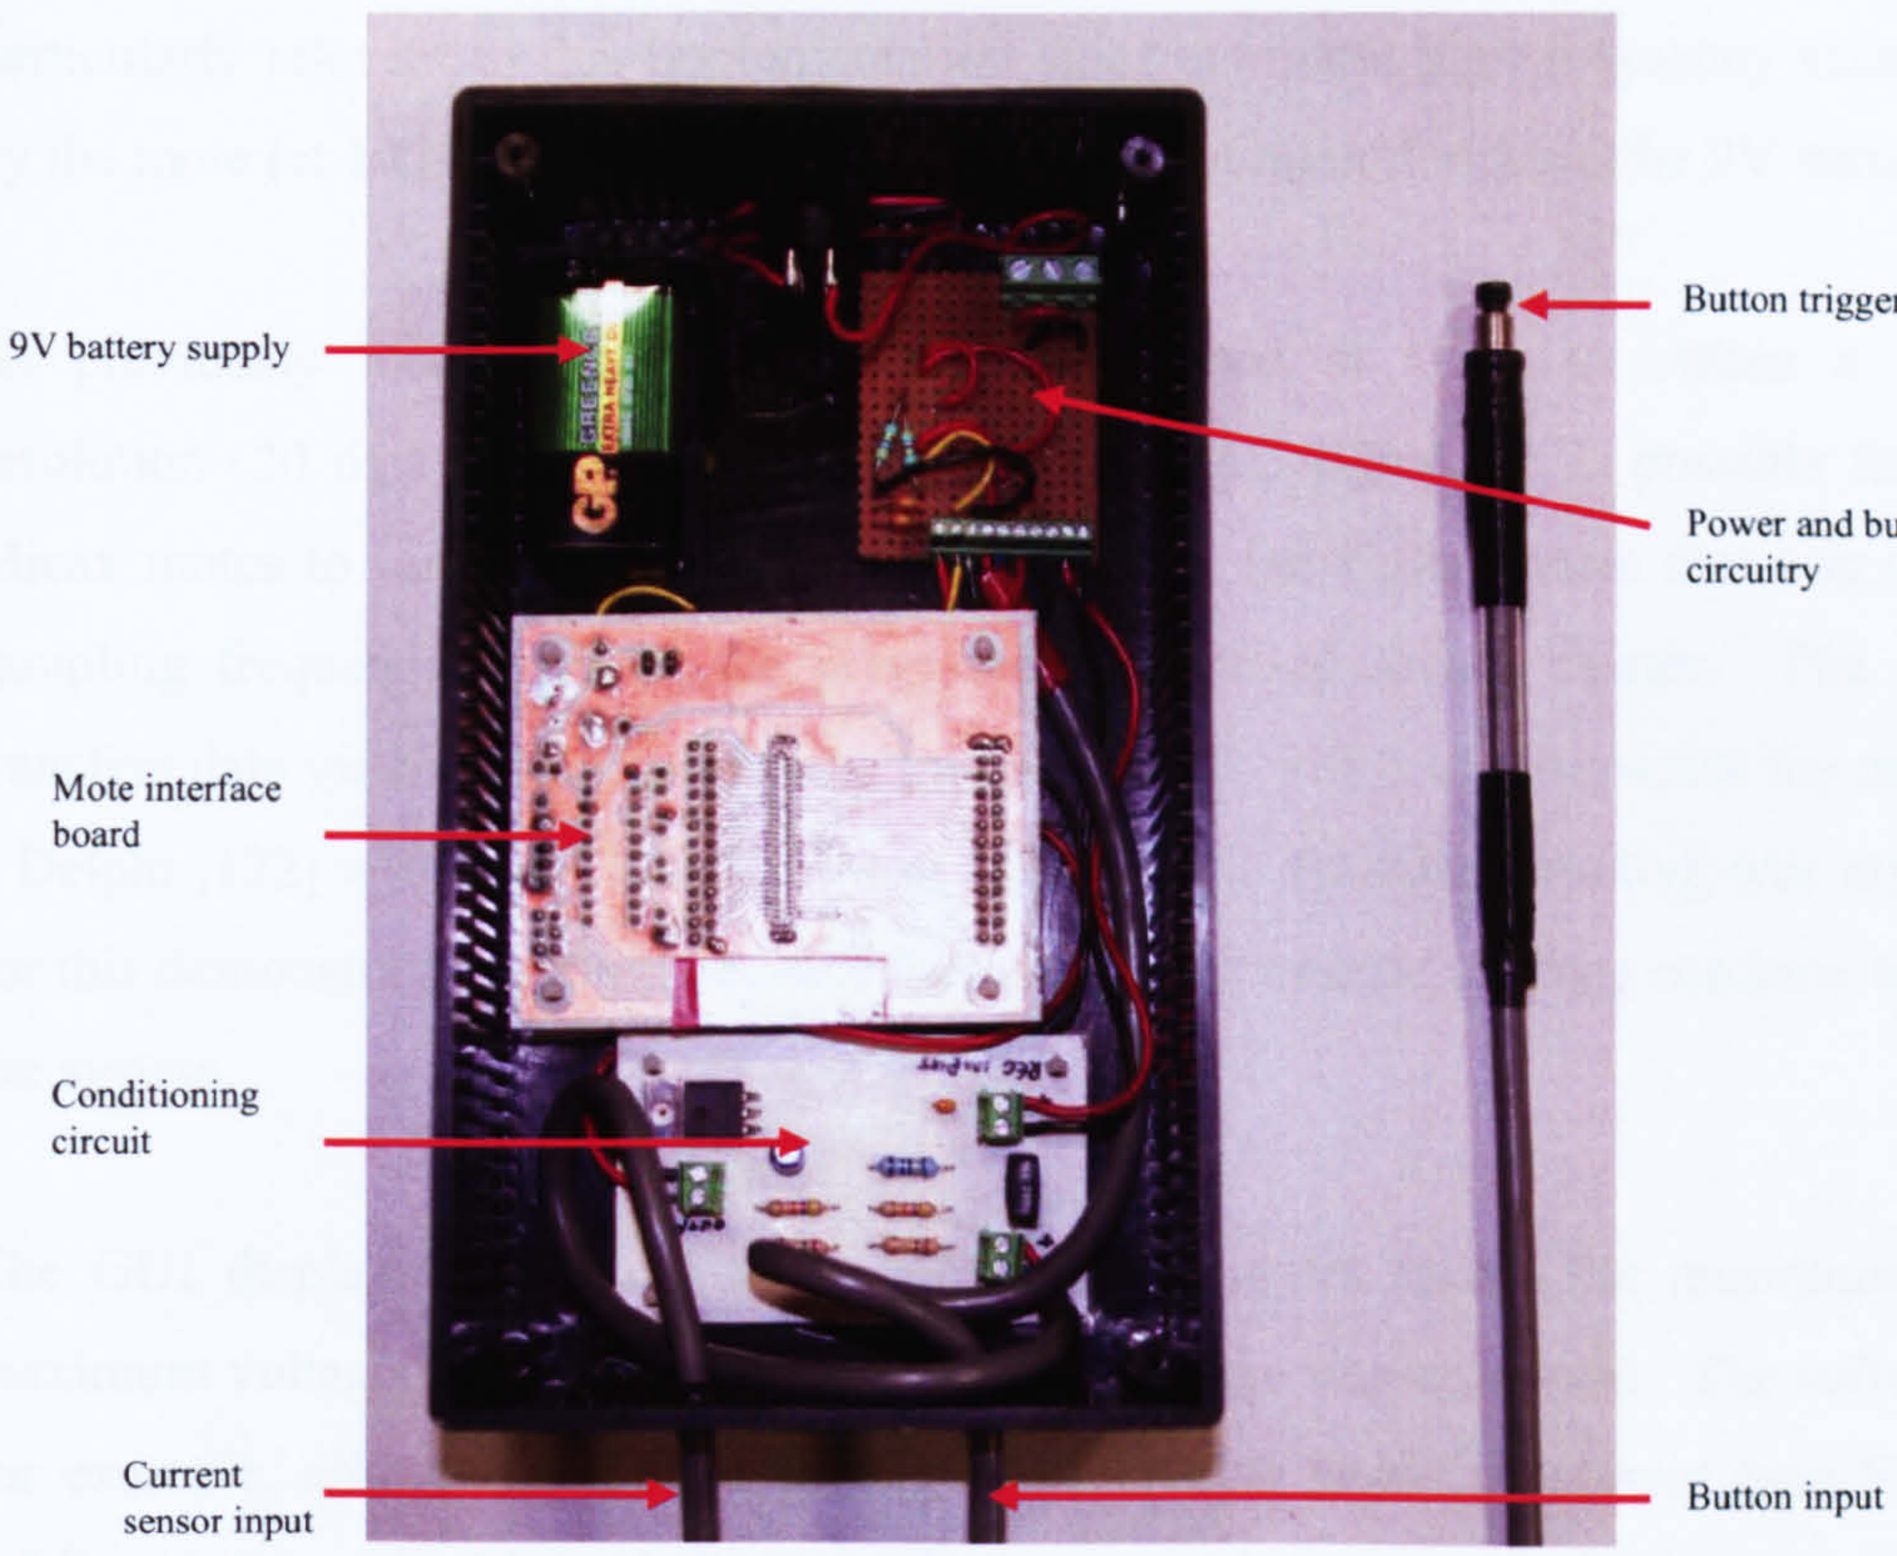

Figure 7.34 – Conditioning circuit and mote boxed to prevent damage

Figure 7.33 shows that the conditioning circuit replicates the Rogowski coil output over a range useable by the Micax motes. The circuit, along with a copy of the mote interface board, was boxed in order to prevent it from becoming damaged. Power was provided from a 9V battery (see Figure 7.34).

Two cables protrude from the box; one is for connection to the Rogowski coil sensor and the other connects to a button which is pictured. This button enables the mote ADC before a spot weld occurs so that it samples the incoming sensor signal at the appropriate time. There are a number of ways that this could be automated for the final application. For example, on the pedestal welder a button could be placed in the welder pedal – in a robotic welder the mote could interface with the control unit of the welder itself. During this work it was attempted to have the mote sampling at a low frequency until it detected current and then increase the sampling rate for a short period of time. Whilst this worked to a degree, it often meant that the first 10-20ms of the signal were not captured. Therefore it was thought that an external button trigger would be more appropriate so that the mote is not permanently sampling. This was particularly relevant to this implementation since continual high frequency sampling by the mote (at 1 kHz) resulted in only a few hours of operation from the 9V battery.

As previously mentioned, sampling was performed at 1kHz to obtain a good resolution (20 data points per cycle) of the 50Hz AC signal. It is possible for the Micax motes to sample at much higher frequencies but Cullen notes that this is the sampling frequency used for the previously mentioned sensor cluster. The mote transfers data via broadcast messaging to a base node, which in turn passes the data to a Delphi [122] written GUI, as shown in Figure 7.35. Broadcast routing was utilised for this demonstration simply because there was only a single welding machine to test the system.

The GUI displays data such as the number of samples taken, the minimum and maximum voltage readings as well as the mean voltage during a weld. The software, for example, allows the user to determine acceptable levels of current (see Figure 7.36) and then uses this to determine whether a weld was good or bad and shows an alert based upon this decision.

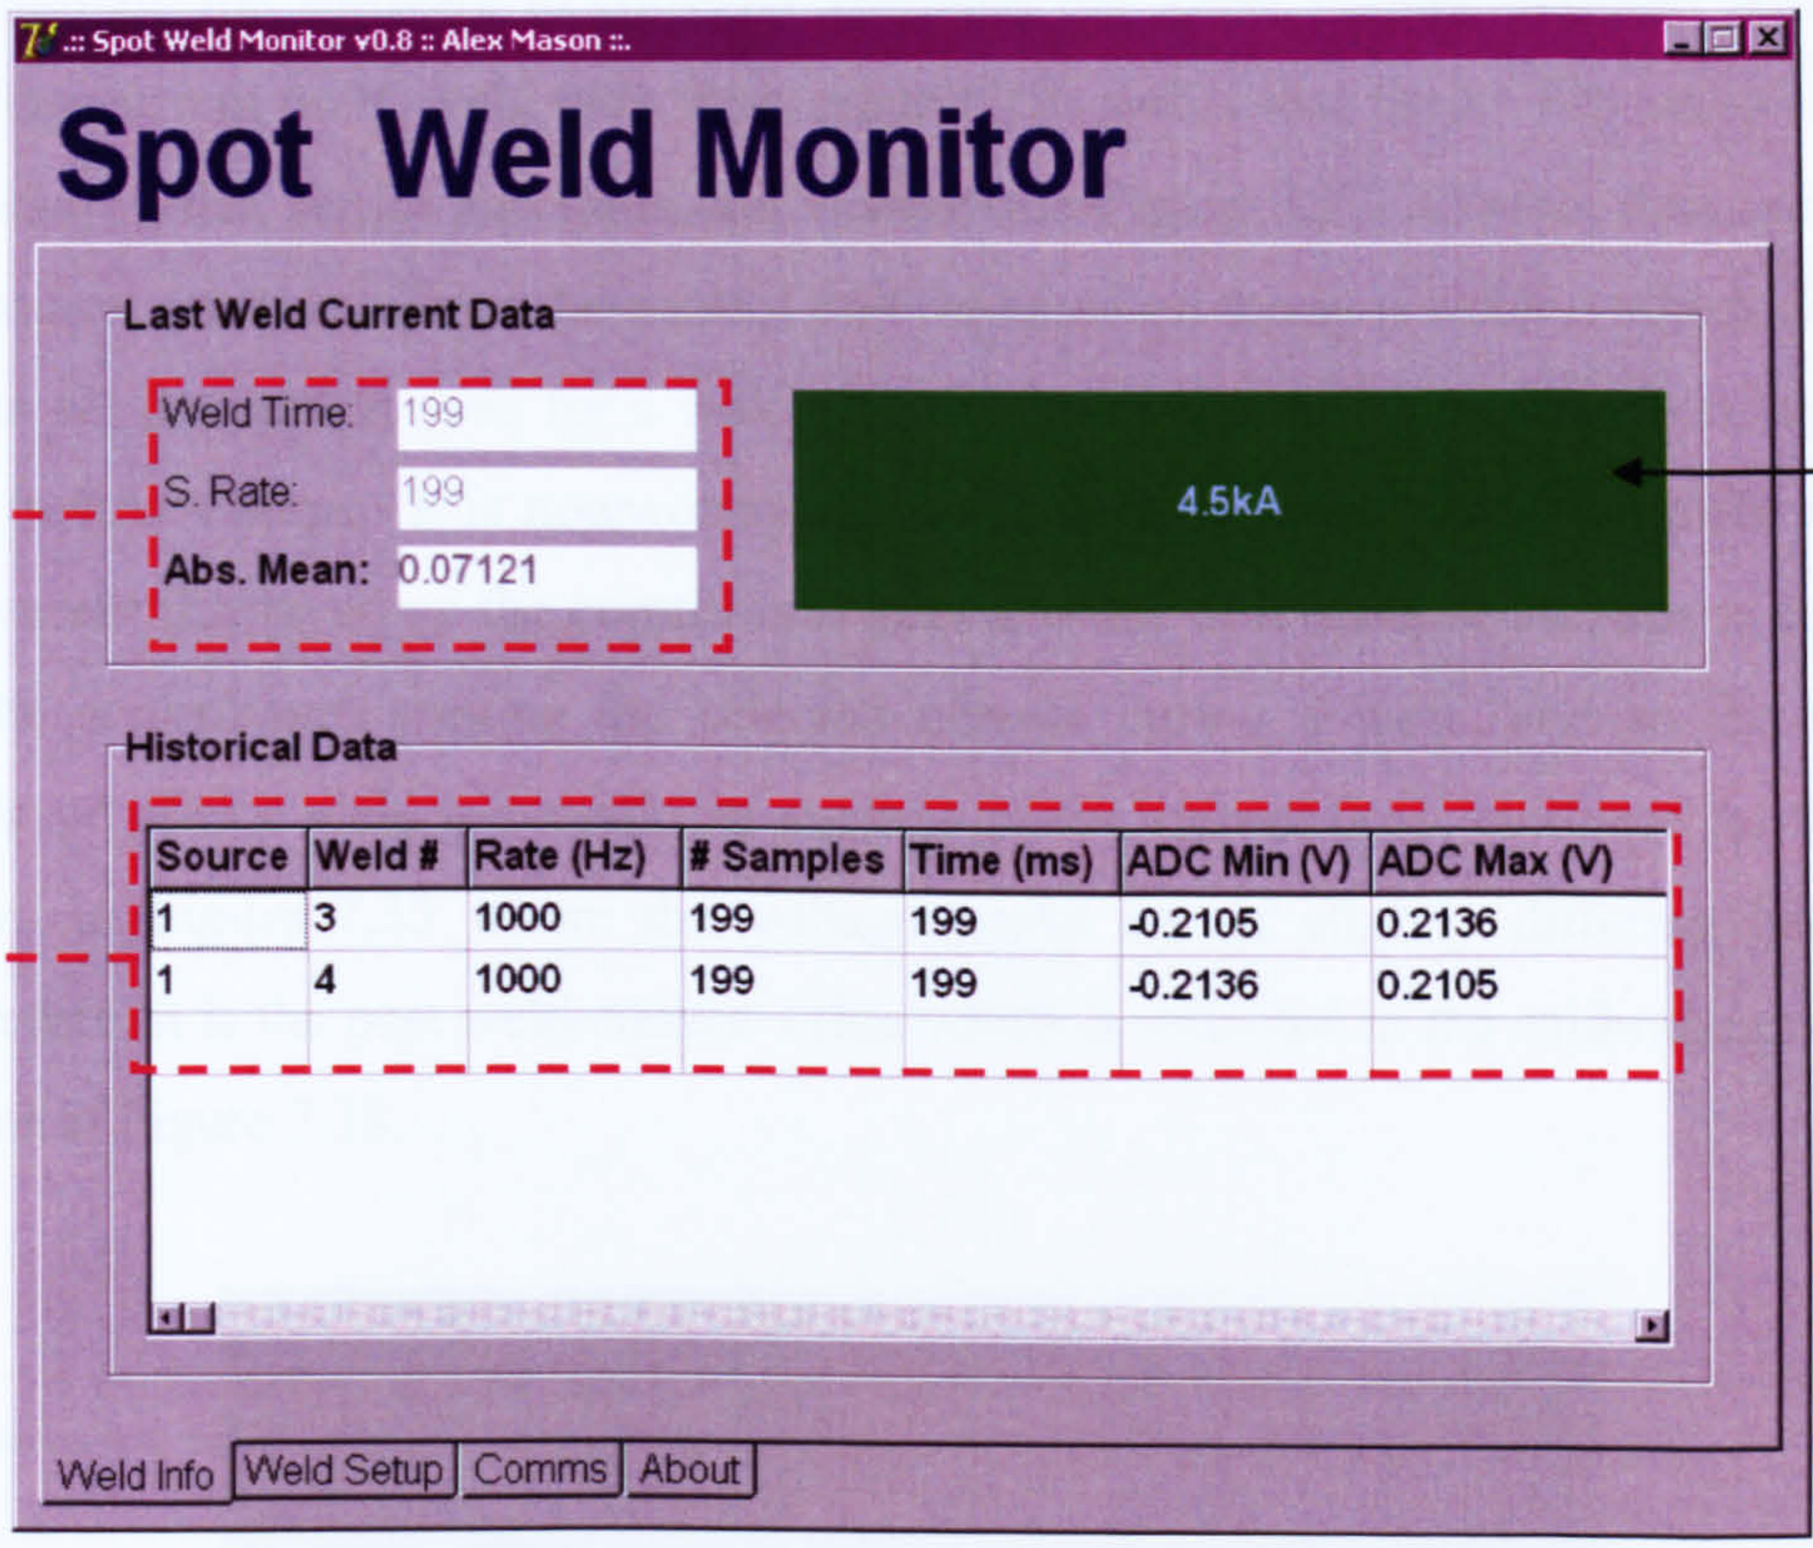

Figure 7.35 – Spot weld monitoring software

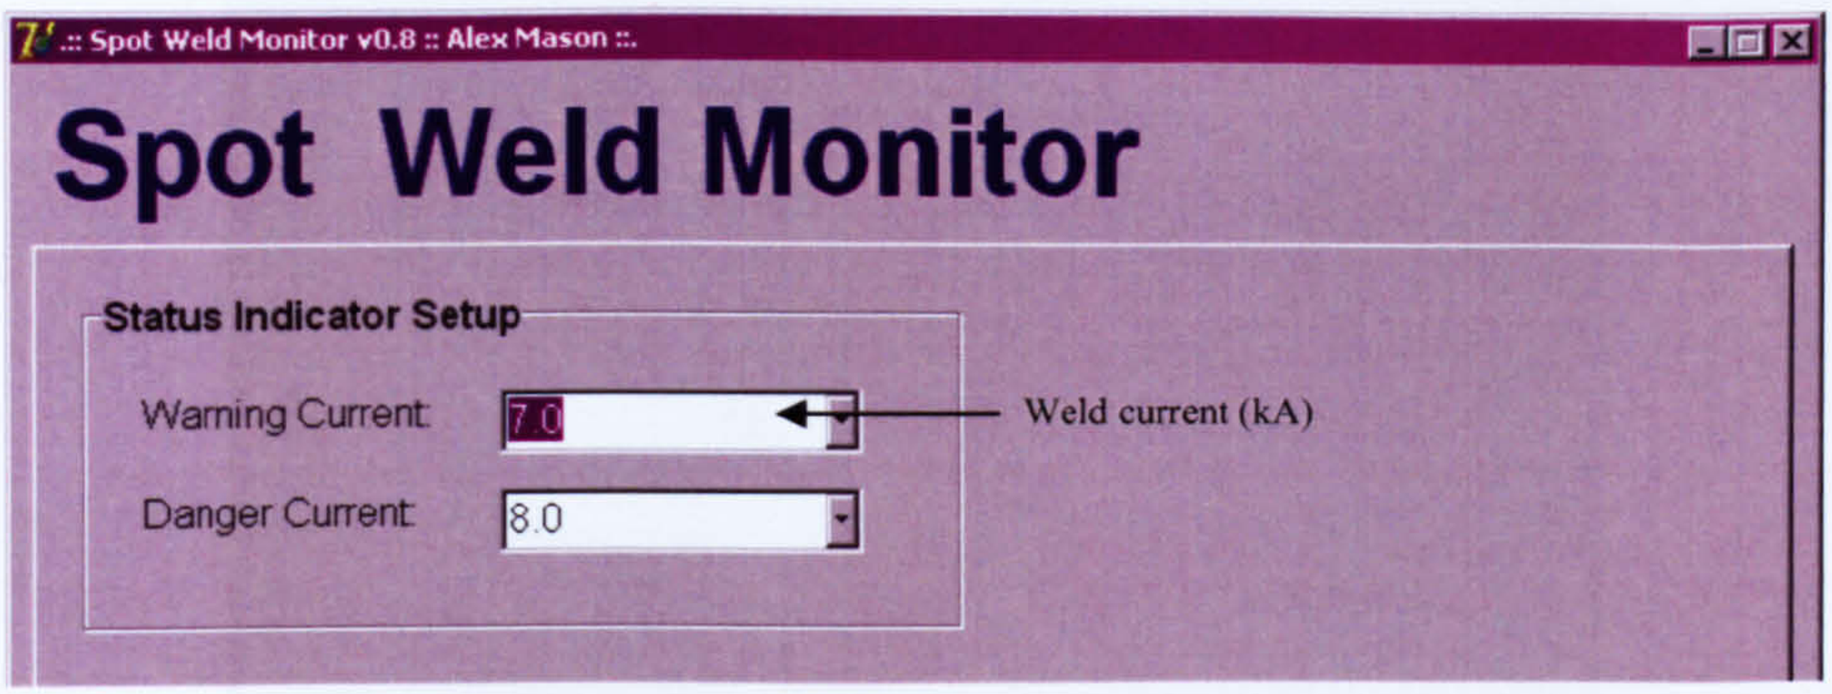

Figure 7.36 – Monitoring software current sensor setup

7.5.3. Calibration and Testing

To calibrate the software to correctly correlate the mean voltage detected with the welding current used, weld data from around 250 welds was used. The setup of the pedestal welder, sensor and mote unit is shown in Figure 7.37. Current data used for calibration was taken from the control unit mounted on the spot welder, which allows one to set the current used for a weld. The weld current was increased from 2kA to 8kA in 0.5kA steps. It is noteworthy that the selected current was often different to the current displayed on the control unit after a weld. It is thought that this is caused by the control unit seeking the selected current during a weld, and so the value displayed after a weld is actually an average of the current used; evidence for this is shown in Figure 7.33 where the voltage peaks are of slightly differing heights. Therefore, it is the post weld current value which is reflected in the calibration results shown in Figure 7.38.

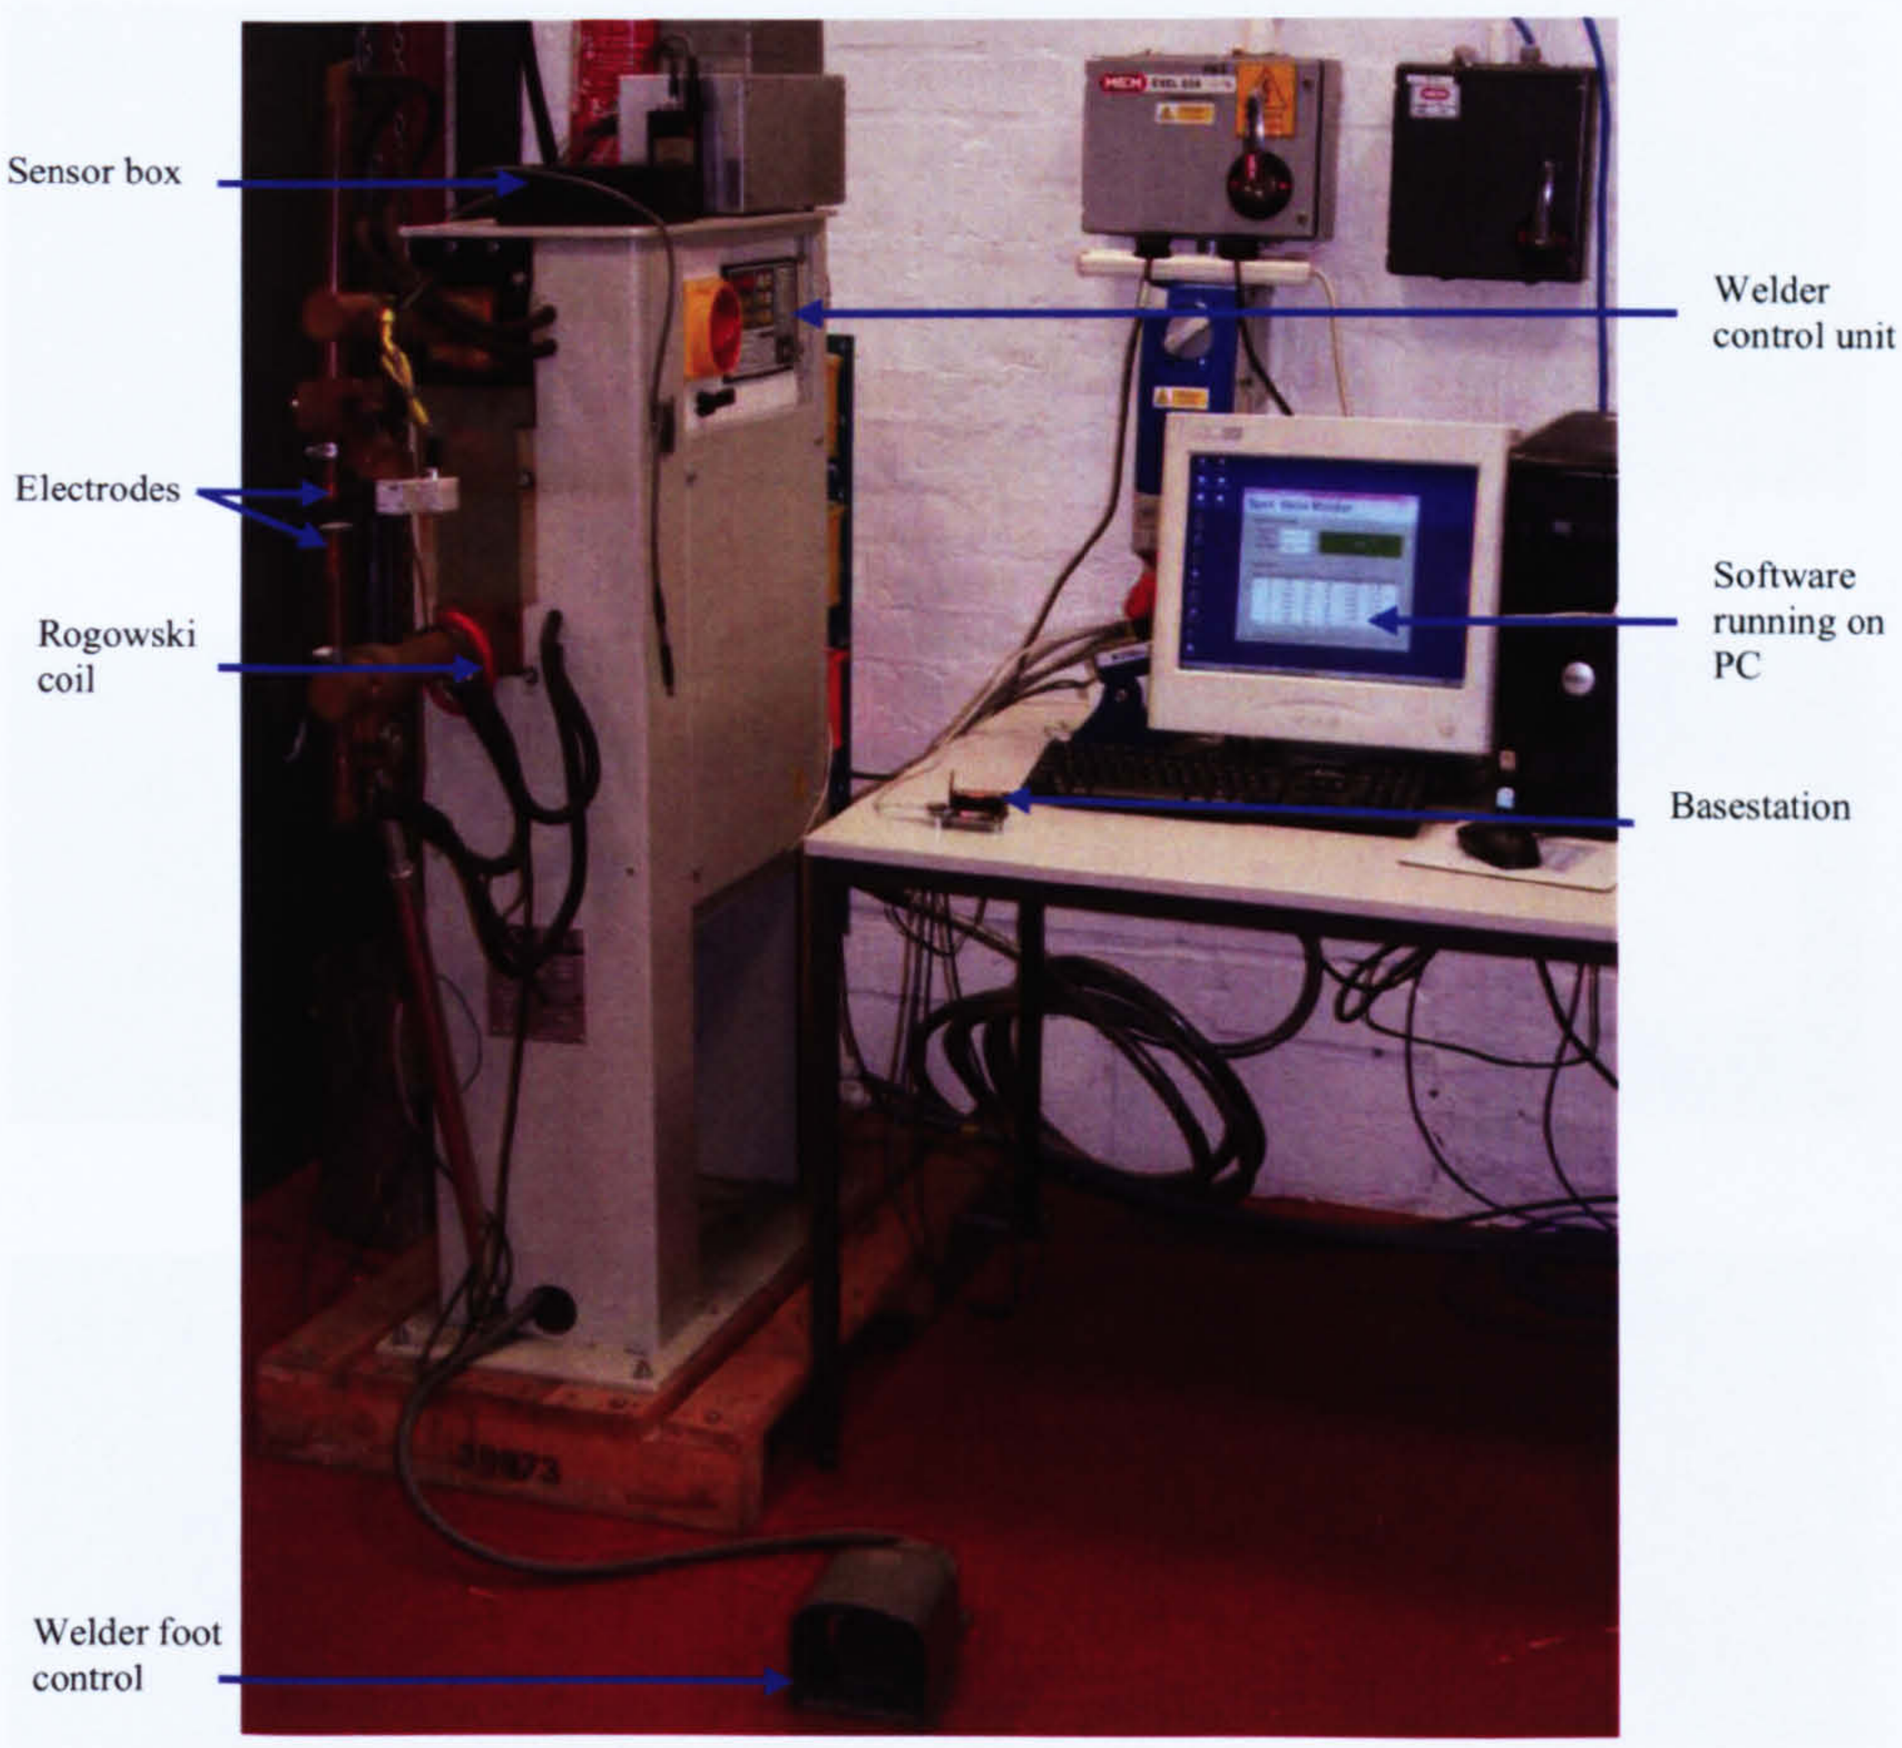

Figure 7.37 – Spot welding demonstration setup

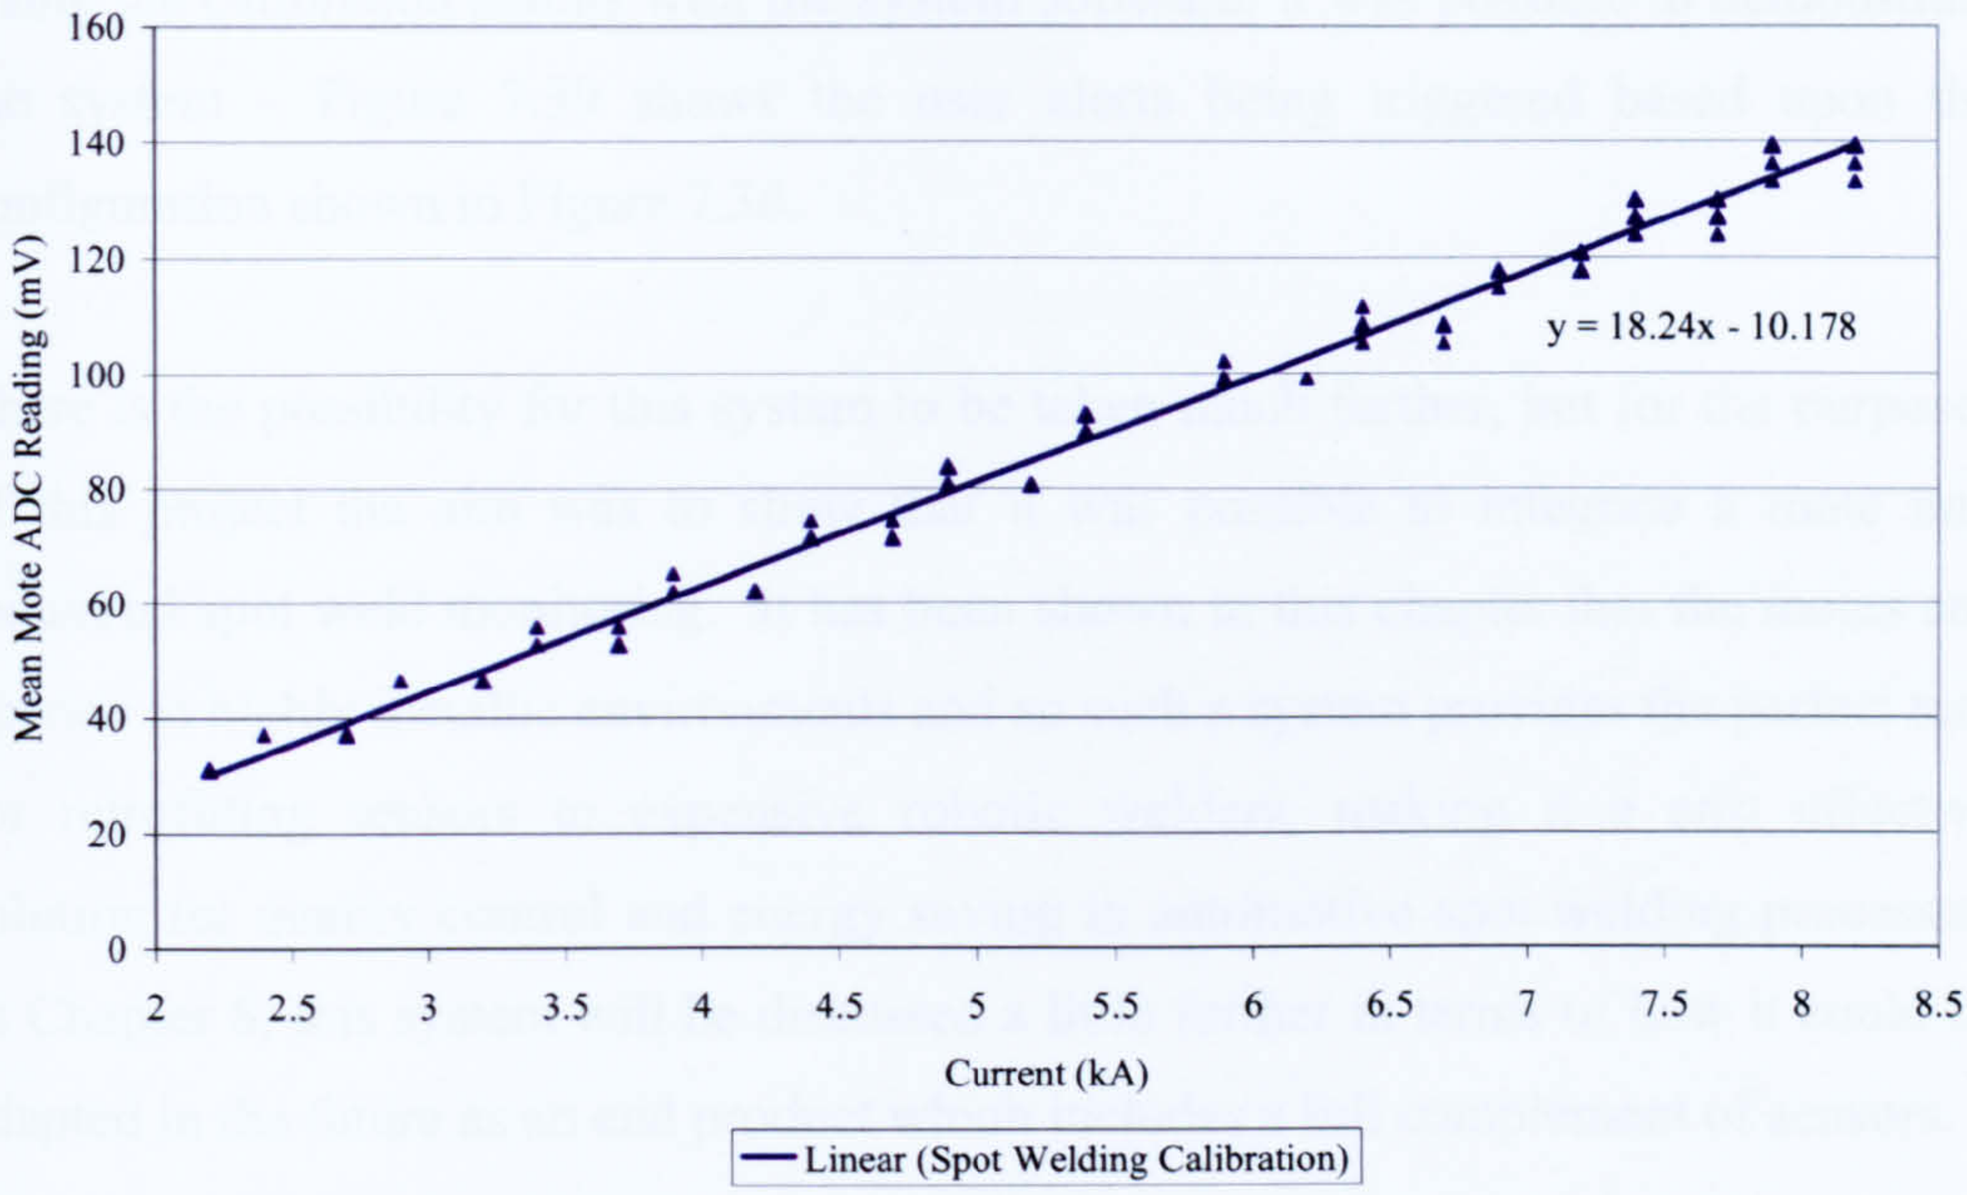

Figure 7.38 – Calibration of the spot welding sensor system

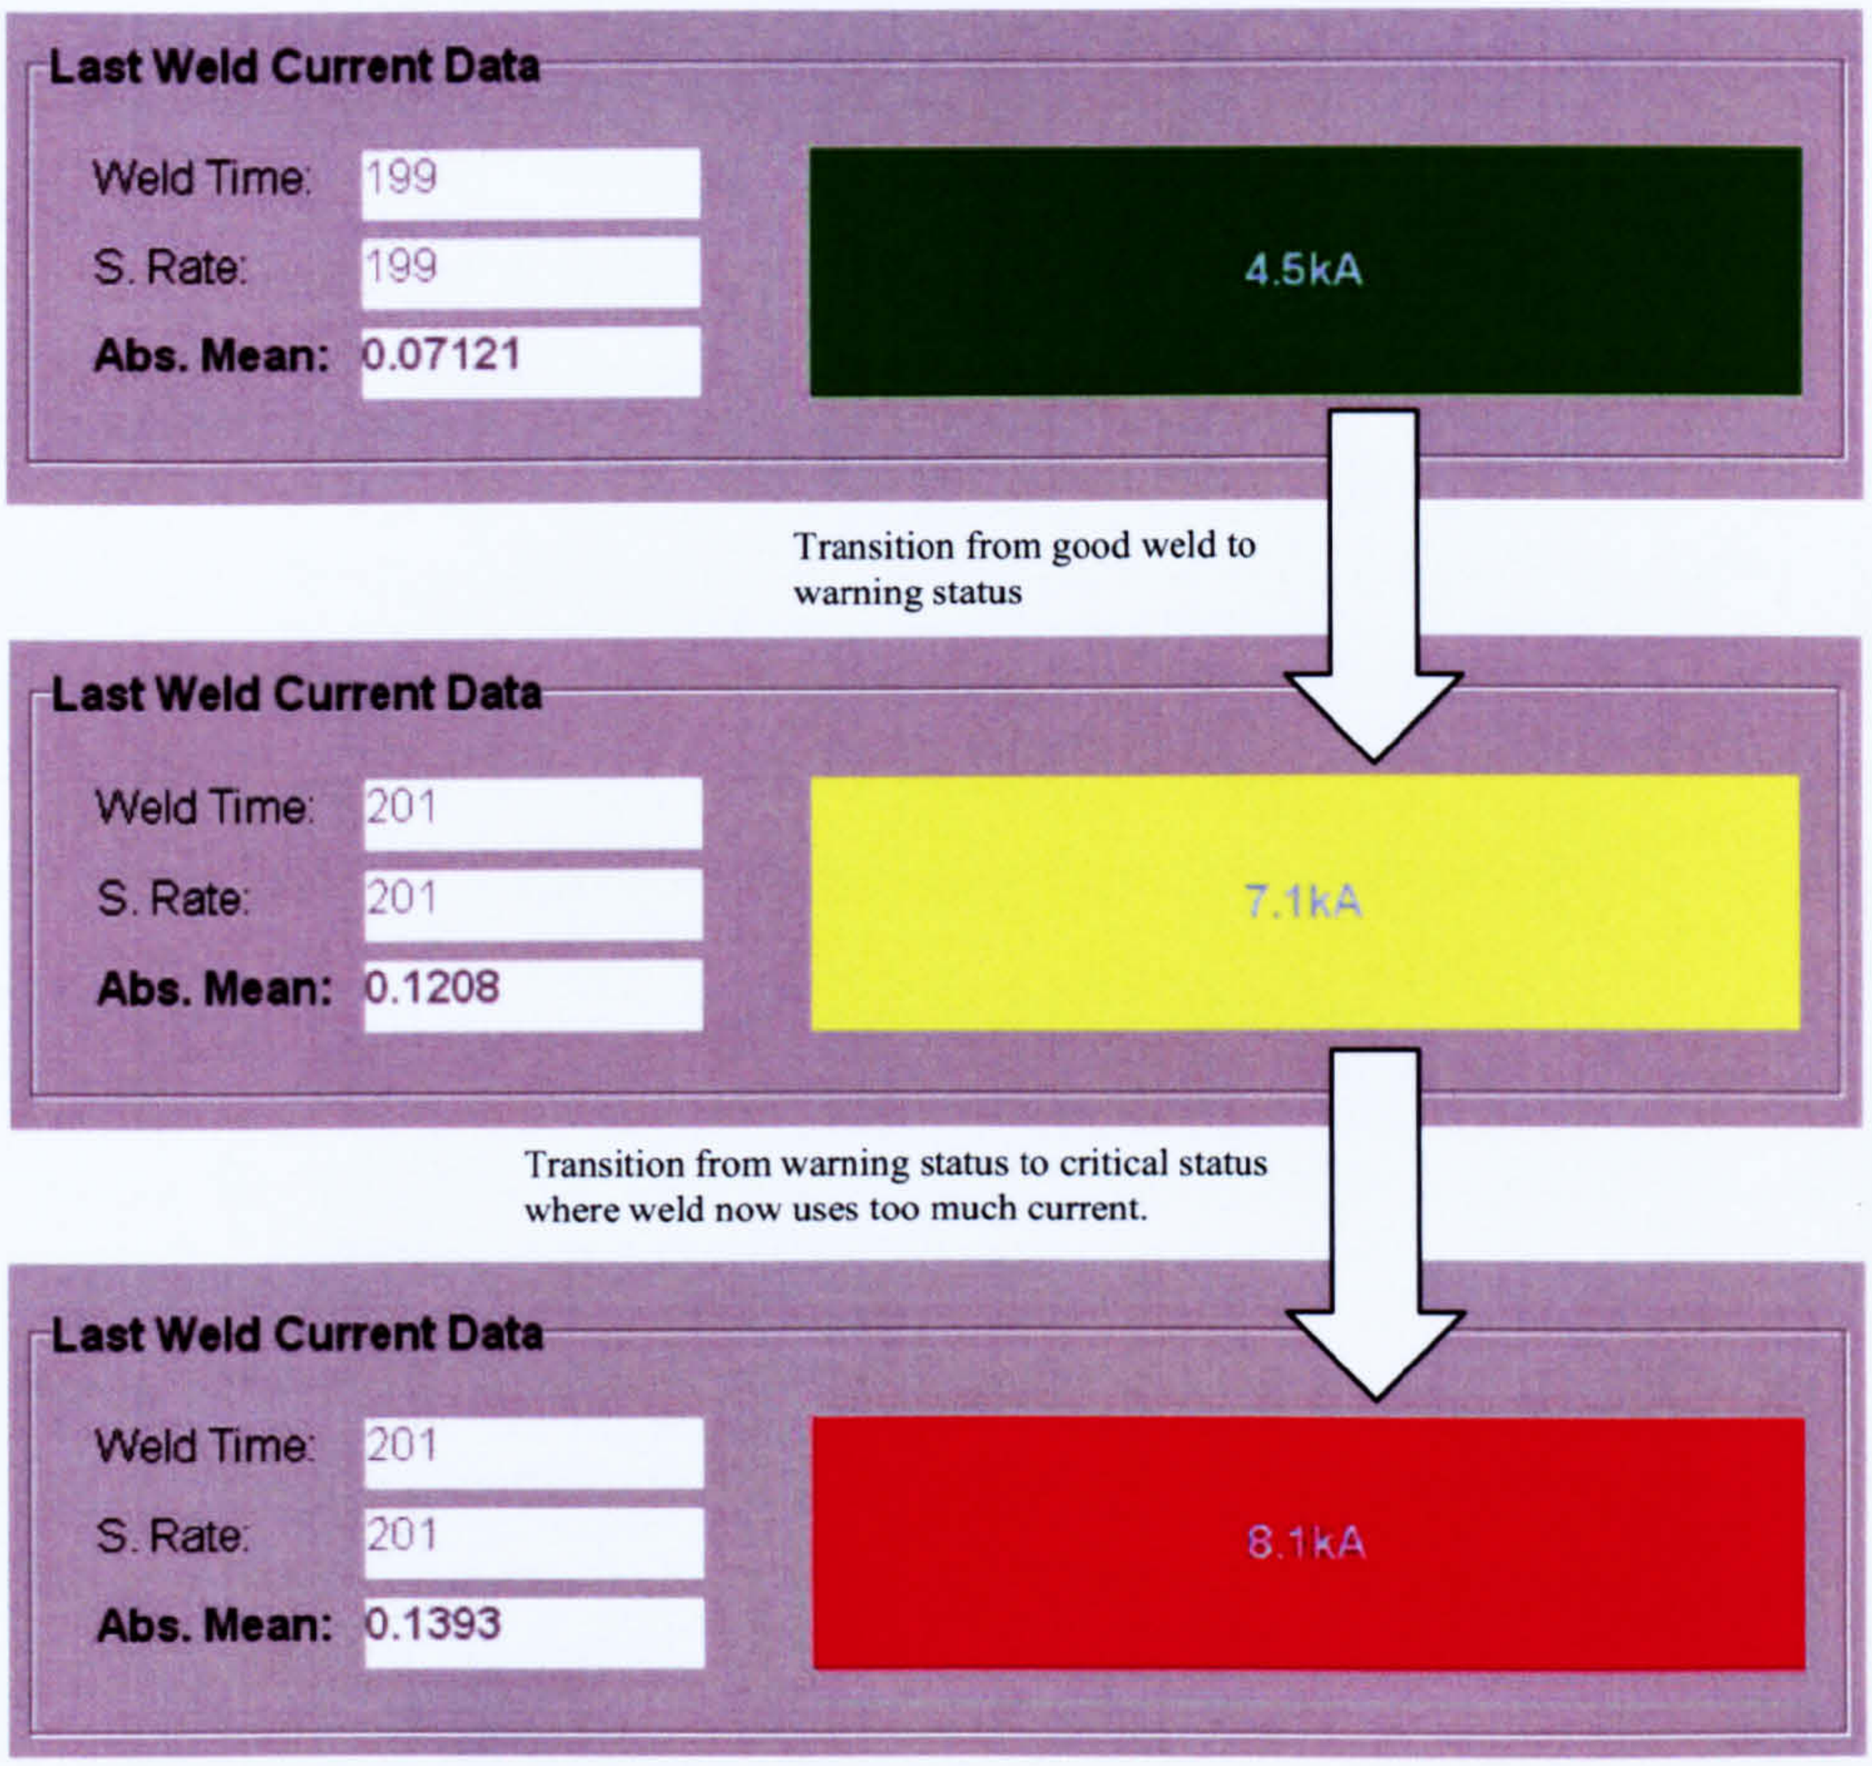

Figure 7.39 – Spot weld system user alert

Using the calibration results with the system software, it was possible to demonstrate the system – Figure 7.39 shows the user alerts being triggered based upon the configuration shown in Figure 7.36.

There is the possibility for this system to be taken much further, but for the purposes of this project the aim was to show that it was possible to integrate a mote into industrial spot weld monitoring. It has been shown in this chapter that the motes can operate in highly metallic environments and so such a system provides the perfect tool for retrofitting sensors to expensive robotic welders, making it a cost effective solution for quality control and energy saving in automotive spot welding processes. In Chapter 8, this system will be discussed a little further in terms of how it could be adapted in the future as an end product which includes a full complement of sensors.

8. Conclusions and Future Work

8.1. The Prototype Inventory Management System

A large part of this work has been concerned with inventory management in industry using WSN devices such as the Micax motes. Having a main focus for the work is not a negative point however since it has allowed many of the features of the motes, as discussed in Chapter 3, to be demonstrated in practise through the prototype inventory management system which was introduced in Chapter 4.

The prototype systems utilises a mesh network in order to attempt to overcome the issue of attenuation suffered by RF devices. The motes have been shown to be scalable through the reverse routing protocol which was shown to exhibit a linear increase in traffic (see Figure 4.6) when an increase in network nodes is experienced, which is far more suitable than the traffic increase exhibited by a purely broadcast based network. A positive side effect of this routing method is that, because it generates less traffic, it also saves energy. In addition, the size of packets used for communication are kept to a minimum, meaning that the mote transceiver is switched on for only a short period of time – this also ensures that usage of the shared bandwidth is kept to a minimum. This packet size could possibly be reduced in the future through the programming of proprietary software which does not incorporate the TinyOS structure. Whilst this would not offer the flexibility of rapid development, it may assist in further reducing energy consumption.

Further energy saving has been shown with the use of power management features implemented by the mote hardware and the operating system, such that high capacity batteries could possibly achieve a two year lifespan for a Micax mote. Additional lifespan gains could be achieved by reducing the duty cycle of the motes – this was discussed in particular with regards to the packaged gas industry where gas cylinders may lay dormant for many months at a time. It follows that the longer a tag battery lasts, the lower the maintainence costs of the system and the more attractive it is to

industry. As part of the prototype system, all tags report their battery voltage so that the system can alert operators to tags which may require maintenance.

The intelligent nature of WSN has been shown through the implementation of sensors – not only can the motes be used to identify assets uniquely they can also determine some characteristic of the asset. This has been shown with the gas cylinder application, where the motes have had both pressure and accelerometer sensors attached. This means that motes acting as tags can tell how much gas is left in a cylinder – this could be used for leak detection as well as allowing for automatic deliveries when a customer's gas supply is running out. The accelerometer sensor can be used to determine how an acetylene cylinder is orientated, and so could be a useful safety feature.

Expansion is demonstrated in multiple ways – the HHU was created as an application specific tool which could be developed further for other tasks. It was found that using the HHU as a location device was 80% reliable in limited testing, but was 100% reliable at identifying an item if it was held approximately 10-20cm away from it, leading to the conclusion that it may be most useful in its current state for close range asset discovery and authentication. Future work could consider other possible uses of the HHU; for example it could be used to discover assets within its vicinity rather than having a preloaded database of nodes available.

Further expansion of the motes has been demonstrated in the form of sensors for the packaged gas and spot welding industries, the later of which will be discussed later in this chapter.

In terms of testing the prototype system it was found that it was reliable provided some redundancy was added to the system through repeated scan requests. It is thought that the small amount of tags which are not read is due to data collisions which appear to occur despite MAC implementation and random messaging offsets. It is noteworthy that the Mica2 MAC protocol is implemented in software; these collisions may be further avoided if the MicaZ was utilised since it has MAC built into the transceiver as part of its Zigbee compliance. It is also possible that the Mica2

suffers from cellular network interference, since UK mobile phones operate around these frequencies and is the reason that, legally, 915MHz devices could not be deployed in the UK. This would appear to make the MicaZ motes favourable as the platform upon which to base future work, as they would be globally deployable.

It is proposed that future work should consider more extensive testing and development of the prototype system in terms of both simulation and industrial trials. Simulation has proven throughout this work to be a useful design tool – for network devices there is simulation software such as ns-2 [174] which can assist evaluation and aid the improvement of communication logic. An interesting comparison could be to see whether the service orientated structure that is in place now would be served better by a harvesting structure, which would remove the on-demand nature of the system presented in this work. Another aspect that could be evaluated is the effect asset movement has on the reliability of communication. It might also be useful to calculate the optimum arrangement of relay nodes, which provide additional network coverage.

Industrial trials will also be important however, since simulations may not take into account all situations which can occur in an industrial situation. Furthermore, industrial experience will allow for increased accuracy in the simulations. In order for such theoretical and practical testing to be feasible there needs to be industrial support and funding.

There needs to be an agreement on a communication standard for such devices from companies prepared to adopt the technology. In the work conducted so far, the devices do not conform to any particular standard. EPCGlobal Inc. define standards for a handful of RFID tags, but motes fall into class 4 (see Table 2.1) since they are intelligent and capable of ad hoc networking. It is extremely important that standards are agreed for this class of device so that future systems may be interoperable and do not create further expense for early adopters. It may be argued that there is a lack of demand for such standards to be drawn up, but WSN is a fast moving topic and it may only be a matter of years before there is a sufficient level of industrial expertise in this area.

Another argument may be that the technology is not ready for deployment. This however is untrue; many organisations now make money from the sale of WSN hardware which could be used for industrial purposes [175-176]. This proves that there is a market for such devices, so the aim now should be to reduce the cost so that they are feasible. One way in which this could be achieved is through removal of some unecessary components – the EEPROM memory on the Micax motes, for example, is not used in this work. In addition, there are SoC devices which combine both the microcontroller and the transceiver into a single package (see Figure 8.1) [177] – this may not only cut the cost of the device but also assist in size reduction. Whilst it would be preferrable to have a device which is smaller than the MicaZ mote, it is certainly not necessary that the technology get to $1mm^3$ before it is deployable. In most cases it really should be feasible for coin sized motes to be perfectly usable in most situations – the biggest challenge is not the hardware itself, but rather the power supply. Where possible it may be beneficial to consider a power scavenging method which may even mean that the motes do not require any maintenance.

One final point with regards to the improvements to the prototype system is that communication security should be considered to prevent rogue nodes from corrupting the inventory management system. The mote operating system includes some level of encryption for messages, but a bug in the version of TinyOS used appears to prevent it from operating. In this work, however, there is no private data being relayed (i.e. no passenger information in the case of airlines) so it is likely that an authentication based scheme would be sufficient rather than encrypting all data – the later would increase the processing time of the motes and therefore increase their power consumption.

Figure 8.1 – System on chip integrated circuit for future motes

8.2. Replacing the Standard Mote Antenna

Industrial fears over the mote antenna have been discussed and attended to through the design and implementation of a PCB antenna which exceeded initial specifications. It is operable at 2.45GHz (i.e. making it suitable for the MicaZ motes), has sufficient bandwidth to allow the MicaZ motes to utilise all of their transmission channels and is smaller than the MicaZ mote PCB, meaning that it could feasibly be integrated into future mote designs rather than being external like the existing monopole antenna. It also appears that a useful feature of the FCPW antenna is that it has two polarisations – this would seem to make it much more suitable than the existing monopole antenna for applications where the antenna orientation is not known or guaranteed. This is the case in spot welding, where the spot welding arm moves, as well as in inventory management where assets may be randomly arranged and stacked.

Future work for this antenna could involve attempting to establish a reliable value for its directivity, as this was something which was uncertain in this work due to a conflict between simulated and measured results. This may be possible through the use of an antenna testing facility which utilises a higher standard of measurement equipment than that used here, as well as an anechoic chamber. A prohibitive factor in this may be the cost of such an action however, so one has to consider whether the corroboration of directivity is actually of further benefit to the results already obtained and presented.

8.3. Industrial Testing

Testing of the Micax motes in situations emulating industrial environments appear to show that the MicaZ mote operated more effectively in metallic environments than the Mica2, which was particularly evident when the motes where almost completely surrounded by metal. Testing of other materials shows the MicaZ appears to be less effected by indoor surroundings, where the Mica2 mote displays rather more varied signal strength results probably due to EM phenomena relating to the operating wavelength. It was found that the Mica2 is affected less by brick walls than the

MicaZ. This is related to the fact that absorption due to skin depth increases with frequency, and so the higher frequency of the MicaZ serves as a disadvantage. That said however, communication was possible over at least 10m through a brick wall and so one could easily overcome any issues by increasing the density of relay notes near such objects.

A disappointing element of this work was the failure to manage real-life industrial testing - this has been previously mentioned as a topic for future work. Whilst efforts were made to set up tests which simulated industrial environments (i.e. the HHU test and industrial testing shown in Chapter 7), actually installing the technology in a real world situation for testing was not achieved. For the airline industry this has been a result of a lack of interest or understanding of what the technology has to offer – it is probable that recent tightening of airport security has also played a role meaning that airlines are wary of allowing the attachment of unknown devices to luggage. It is true to say that the airline industry as a whole appears happy to keep utilising bar code technology due to its cheapness; even relatively cheap passive RFID tags are failing to conjure a widespread positive response. There are reports of RFID being implemented in airports which are now over seven years old [178], yet still there appears to be little sign of bar codes being replaced. It is thought that one of the possible reasons for this is that luggage is not technically an asset to an airline. It has no actual value for the airlines who would prefer that passengers did not carry luggage at all. Therefore it is possible that the lack of investment in new technology for luggage handling systems is a result of there being no short term benefit for the airlines since any money they save as a result of losing less luggage will be taken up by the cost of implementing new systems. If cheap passive RFID tags, costing just a few pence each, cannot encourage serious interest for the airline industry then it is unlikely that a £50 mote will either! This may change in the future, but at the moment it is a serious hurdle to be overcome.

With the packaged gas industry, a different story has developed. Since many organisations still have either no automated identification or primitive bar code systems, there is the incentive for them to invest in new technology which will save them significant amounts of time. As such it appears that there has been recent

interest in WSN technology as a result of this work and it is hoped that in the near future a relationship will be established with industry. A unique attitude for the packaged gas industry is that they appear to recognise that motes and WSN are a growing technology which is determined to cut costs. After some discussion with Air Products it is noteworthy that they are interested in a solution that costs no more than £5 ($10) and as one may have noted from Chapter 3 mote manufacturers are aiming for a £0.50 price point ($1).

8.4. Spot Welding Application

The final part of this document discussed the idea of using WSN in spot welding. As previously mentioned, this shows the motes expandability. Furthermore it shows that there is potential for previous spot welding work to be implemented such that it is sufficiently small, wireless, and capable of overcoming difficulties presented by a metallic environment via multihop communication. It is thought that the Micax motes are actually small enough to be retrofitted to a robotic spot welding machine in their current form, and so a future work could be carried out to ensure that the sensors and intermediate conditioning circuits are also appropriately sized. In this way a complete sensor cluster could be implemented with a Micax mote at its heart. It may also be possible to have the mote monitor sensor data such that it can determine whether a weld is good or bad, as well as allowing it to alter welding parameters such as current to save power; a feature such as this would remove the need for external software.

Finally, it is important to note that the WSN concept is not just applicable to the industries mentioned here. There is also the intention of utilising some of the knowledge gained through this work to consider underwater WSN systems. Whilst this topic has not been considered in great depth at this point in time, water represents another substance in which RF propagation is difficult. It is therefore thought that WSNs could be applied such that they overcome this issue in a similar way to that used in this work to overcome difficulties in metallic surroundings.

9. References

[1] SITA: "Baggage claims and departure delays now costing $1 billion annually", Technical report, SITA Inc, 1 London Gate, 252-254 Blyth Road, Hayes, Middlesex, UB3 1BW, UK, http://www.sita.com/, 2005.

[2] Ginny McGrath: "BA and BMI top lost luggage poll", Times Online news article, http://travel.timesonline.co.uk/tol/life_and_style/travel/news/article3397831.ece, 19th February 2008.

[3] AEA: "Punctuality and Baggage Performance", Association of European Airlines punctuality information, Avenue Louise 350, B-1050, Brussels, Belguim, http://www.aea.be/research/performance/index.html, 2008.

[4] Global Aircraft: "Boeing 747 Specifications", http://www.globalaircraft.org/planes/b747.pl, 2008.

[5] Global Aircraft: "Boeing 737 Specifications", http://www.globalaircraft.org/planes/b737.pl, 2008.

[6] J. D. Cullen, A. M. A. El-Rasheed, A. I. Al-Shamma'a, W. Lucas: "Multi Sensor Fusion for On Line Monitoring of the Quality of Spot Welding in Automotive Industry", GARS-2005 Conference, Liverpool John Moores University, 22 June 2005.

[7] J. D. Cullen, N. Athi, M. A. Al-Jader, A. Shaw, A. I. Al-Shamma'a: "Energy reduction for the spot welding process in the automotive industry", Sensors and their Applications XIV, Journal of Physics Conference Series 76, 2007.

[8] The Oxford Dictionary and English Usage Guide, ISBN 0198613253, Oxford University Press, 1996.

[9] C. D. J. Waters: "Definition of Inventory Terms", Inventory Control and Management, Chapter 1, PP 4, ISBN 0471930814, John Wiley & Sons, 1992.

[10] C. D. J. Waters: "Just-In-Time Systems", Inventory Control and Management, Chapter 8, PP 295-322, ISBN 0471930814, John Wiley & Sons, 1992.

[11] Steven Shepard, "Bar codes... Up Close and Personal", RFID: Radio Frequency Identification, Chapter One, PP 26-41, ISBN: 0071442995, McGraw Hill, 2005.

[12] Steven Shepard, "A History of Bar codes", RFID: Radio Frequency Identification, Chapter One, PP 12-26, ISBN: 0071442995, McGraw Hill, 2005.

[13] Marian Visich, Jr: "Bar code Symbologies", Bar codes and their Applications, Monograph Series of the New Liberal Arts Program, Department of Technology and Society, State University of New York as Stony Brook, Available from (10/03/2008): http://www.math.dartmouth.edu/~mqed/NLA/Bar codes/Bar codes.phtml.

[14] Klaus Finkenzeller: "Bar code Systems", RFID Handbook, Chapter 1, PP 3, ISBN 0470844027, Wiley & Sons, 2003.

[15] GS1, Blue Tower, 326 Avenue Louise, BE 1050, Brussels, Belguim, http://www.gs1.org/

[16] Microscan: "Industrial Bar code Scanners, MS890", Microscan Corporate Headquarters, 1201 SW 7th Street, Renton, WA, 98057, United States, http://www.microscan.com/.

[17] Microscan: "Quadrus Imagers, MSQ-Basic", Microscan Corporate Headquarters, 1201 SW 7th Street, Renton, WA, 98057, United States, http://www.microscan.com/.

[18] Microscan: "Contract Manufacturers Use Data Matrix to Track Products", Microscan Corporate Headquarters, 1201 SW 7th Street, Renton, WA, 98057, United States, http://www.microscan.com/.

[19] Microscan: "Reading bar codes on tires ensures integrity of manufacturing process", Microscan Corporate Headquarters, 1201 SW 7th Street, Renton, WA, 98057, United States, http://www.microscan.com/.

[20] Microscan: "Bar code Readers Enable Automated Pallet Tracking", Microscan Corporate Headquarters, 1201 SW 7th Street, Renton, WA, 98057, United States, http://www.microscan.com/.

[21] RFID Journal: "The History of RFID Technology", Available from http://www.rfidjournal.com/article/articleview/1338/2/129/.

[22] Klaus Finkenzeller: "Animal Identification", RFID Handbook, Chapter 13, PP 372-5, ISBN 0470844027, Wiley & Sons, 2003.

[23] Klaus Finkenzeller: "The functionality of an immobilisation system", RFID Handbook, Chapter 13, PP 372-5, ISBN 0470844027, Wiley & Sons, 2003.

[24] Laurie Sullivan: "Wal-Mart Outlines RFID Expansion Plans", Information Week, http://www.informationweek.com/, 17 June 2004.

[25] David H. Williams: "The Strategic Implications of Wal-Marts RFID Mandate", Directions Magazine, http://www.directionsmag.com/, 29 July 2004.

[26] Klaus Finkenzeller: "Electromagnetic Backscatter Coupling", RFID Handbook, Chapter 3, PP 47, ISBN 0470844027, Wiley & Sons, 2003.

[27] Bill Glover and Himanshu Bhatt: "Inductive Coupling", RFID Essentials, Chapter 3, PP 66-7, ISBN 0596009445, O'Reilly, 2006.

[28] Klaus Finkenzeller: "Frequency, Range and Coupling", RFID Handbook, Chapter 2, PP 23, ISBN 0470844027, Wiley & Sons, 2003.

[29] N. C. Wu, M. A. Nystrom, T. R. Lin, H. C. Yu: "Challenges to Global RFID Adoption", PICMET 2006 Proceedings, PP 618-623, Istanbul, Turkey, 9-13 July 2006.

[30] Martina Gerst, Raluca Bunduchi, Ian Graham: "Current Issues in RFID Standarisation", RFID Standardisation Workshop, University of Edinburgh, High School Yards, Edinburgh, EH1 1LZ, United Kingdom, 22 February 2005.

[31] ISO Standard: "ISO 14223", Radio frequency identification of animals, Available from http://www.iso.org/.

[32] ISO Standard: "ISO 15693", Identification cards, contactless integrated circuit(s) cards, vicinity cards, Available from http://www.iso.org/.

[33] ISO Standard: "ISO 18000", Radio frequency identification for item management (Parts 1-7), Available from http://www.iso.org/.

[34] EPCGlobal Inc, GS1 UK, Staple Court, 11 Staple Inn Buildings, WC1V 7QH, London, United Kingdom, http://www.epcglobalinc.org/.

[35] EPCGlobal Inc, "EPC Radio-Frequency Identity Protocols Class-1 Generation-2 UHF RFID Protocol for Communications at 860MHz to 960 MHz Version 1.1.0", Specification for RFID Air Interface, GS1 UK, Staple Court, 11 Staple Inn Buildings, WC1V 7QH, London, United Kingdom, http://www.epcglobalinc.org/

[36] Jim Harper: "RFID Tags and Privacy: How Bar-Codes-On-Steriods Are Really a 98-Lb Weakling", Executive Summary, Competitive Enterprise Institute, 1001 Connecticut Avenue NW, Suite 1250, Washington, DC 20036, 21 June 2004.

[37] Dave Piasecki: "RFID Update: The Basics, The Wal-Mart Mandate, EPC, Privacy Concerns and More.", RFID News Website, http://www.inventoryops.com/RFIDupdate.htm

[38] ZDNet.co.uk: "Marks & Spencer tags shirts with RFID", News Website, http://www.zdnet.co.uk/, 17 October 2003.

[39] Scott Bradner: "Security Issues Swamp RFID", Tech World, http://www.techworld.com, 7 February 2005.

[40] Intermec: "The Guide to RFID Tag Selection", Product Brochure, Intermec Corporate Headquarters, 6001 36th Avenue West, Everett, Washington 98203, http://www.intermec.com/.

[41] RFIDShop: "RF2ID On-Metal Multi-Purpose Tag", RFID Equipment Store, http://www.therfidshop.com/product_info.php?products_id=251.

[42] SmartCode: "SmartCode Corp Announces the Worlds First 5 cent RFID Tag", SmartCode Corporation, 245 Park Avenue, New York, NY 10167, United States, http://www.smartcodecorp.com/newsroom/01-05-06.asp, January 2006.

[43] Ronan Clinton: "RFID.Bomb", Heavey RF Blog, http://www.heaveyrf.com/news_and_events/rfidbomb.453.364.html, 20 July 2007.

[44] Ann Bednarz: "Leeway found in Wal-Mart's RFID mandate", Network World, http://www.networkworld.com, 29 November 2004.

[45] Christian Annesley: "UPS rules out RFID tagging for small packages", Computer Weekly Magazine, PP 10, 20 June 2006.

[46] P. Bahl and V. Padmanabhan: "RADAR: An In-Building RF-based User Location and Tracking System", Proceedings of IEEE INFOCOM, PP 775-784, Tel-Aviv, Israel, March 2000.

[47] AeroScout: "AeroScout Visibility Overview", Datasheet, 1450 Fashion Boulevard, Suit 510, San Mateo, CA 94404, http://www.aeroscout.com/.

[48] Ekahau: "Ekahau Positioning Engine 3.1", Data Sheet, 12930 Saratoga Avenue, Suite B-8, Saratoga, CA 95070, http://www.ekahau.com/.

[49] PanGo Networks Inc: "PanGo Locator", Data Sheet, 959 Concord Street, Suite 100, Framingham, MA 01701, http://www.pangonetworks.com/.

[50] Aeroscout: "Vehicle Tracking at American Port Services", Customer Case Study, 1450 Fashion Boulevard, Suit 510, San Mateo, CA 94404, http://www.aeroscout.com/.

[51] Ubisense: "Ubisense Hardware Datasheet", St Andrews House, 90 St Andrews Road, Chesterton, Cambridge, CB4 1DL, http://www.ubisense.net/.

[52] Robert J.Fontana and Steven J. Gunderson: "Ultra-Wideband Asset Location System", Proceedings of IEEE Conferences on Ultra Wideband Systems and Technologies, PP147-150, May 2002.

[53] Paul Hansell, Selcuk Kirtay: "Ultra Wide Band (UWB) Compatibility", Final Report to the Radiocommunications Agency, Aegis Spectrum Engineering, 30 Anyards Road, Cobham, Surrey, KT11 2LA, UK, http://www.aegis-systems.co.uk/.

[54] OFCOM: "The Wireless Telegraphy (Ultra-Wideband Equipment) (Exemption) Regulations 2007", Electronic Communications, No. 2084, http://www.ofcom.org.uk/, 13 August 2007.

[55] Visible Assets Inc: "t-Tag", Technology, 195 Bunker Hill Avenue, Stratham, New Hampshire, 03885, United States, http://www.visible-assets.com/.

[56] Visible Assets Inc: "Construction Visibility Network I-Beam Discovery Protocol", Analysis, 195 Bunker Hill Avenue, Stratham, New Hampshire, 03885, United States, http://www.visible-assets.com/.

[57] Visible Assets Inc: "RuBee Visibility Pilot TopRack Network State of Colorado", Presentation, 195 Bunker Hill Avenue, Stratham, New Hampshire, 03885, United States, http://www.visible-assets.com/.

[58] Connexions Direct: "Airport Baggage Handler", Job Article, jobs4u Database, http://www.connexions-direct.com/jobs4u/index.cfm?pid=64&catalogue ContentID=668, 14 March 2008.

[59] Eric Rydar: "Baggage tracking system enhances air travel safety", iStart, New Zealand's e-Business Portal, http://www.istart.co.nz/, November 2005.

[60] Sarah Tapley and David Riley: "Baggage handling in narrow-bodied aircraft: Identification and assessment of musculoskeletal injury risk factors", Health and Safety Executive Report, January 2005, Available from: http://www.hse.gov.uk/airtransport/msdreport.pdf

[61] Bob Brewin: "Delta Says Radio Frequency ID Devices Pass First Bag-Tag Test", Computer World, Mobile & Wireless News, http://www.computerworld.com/, 22 December 2003.

[62] Harry Hutchinson: "every bag in place", Mechanical Engineering Magazine, Feature Article, http://www.memagazine.org/, April 2004.

[63] HKIA: "HKIA Boosts Baggage Handling Efficiency with RFID Technology", Hong Kong International Airport (HKIA) Press Centre, http://www.hongkongairport.com/eng/pr/pr_914.html, 15 January 2008.

[64] Bill Ray: "Heathrow to trial RFID tags", The Register, Enterprise security news, http://www.theregister.co.uk/, 6 July 2007.

[65] RFID Journal: "First Multifrequency Chip Unveiled", Report on Toppan Forms RFID transponders, http://www.rfidjournal.com/, 15 March 2005.

[66] Air Products PLC: "ACTS – Cylinder Tracking System", Proprietary cylinder tracking system, http://www.airproducts.co.uk/services/acts.htm.

[67] Pollution Control Industries: "Gas Cylinder Management", Cylinder tracking and management system, http://www.pollutioncontrol.com/services/cylinders/ GasCylinders.html.

[68] Dataweld Inc: "AcuTrax – Cylinder Tracking", Bar code tracking system, 1909 Citizens Bank Drive, Bossier City, Louisiana 71111, United States, http://www.dataweld.com/.

[69] Jim Broughton: "Side by side testing of Trovan and Motorola RFID Gas Cylinder Tracking Systems", Report by Dataweld Inc – Comparison of read distances, http://www.rfidnews.com/side.html.

[70] Feng Zhao, Leonidas Guibas: "Roles of Sensor Nodes and Utilities", Wireless Sensor Networks: An Information Processing Approach, Chapter 5, PP 139, ISBN 1558609148, Elsevier Inc, 2004.

[71] Feng Zhao, Leonidas Guibas: "Energy-Minimizing Broadcast", Wireless Sensor Networks: An Information Processing Approach, Chapter 3, PP 83, ISBN 1558609148, Elsevier Inc, 2004.

[72] Alec Woo, Terence Tong, David Culler: "Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks", Sensys '03, November 5-7, 2003.

[73] Andrew S. Tanenbaum: "Connection Oriented and Connectionless Services", Computer Networks Fourth Edition, Chapter 1, PP 32-33, ISBN 0130384887, Prentice Hall, 2003.

[74] Ekram Hossain, Kin K. Leung: "Introduction", Wireless Mesh Networks: Architectures and Protocols, Chapter 1, PP 2, ISBN 9780387688381, Springer Science, 2008.

[75] Lamont V. Blake: "Characteristics of Electromagnetic Waves", Chapter 1, PP10-35, ISBN: 0471079286, Wiley & Sons, 1966.

[76] Chipcon (now Texas Instruments): "CC1000 Single Chip Very Low Power RF Transceiver", Chipcon as Smart RF Datasheet (rev 2.3), 2005.

[77] Chipcon (now Texas Instruments): "CC2420 2.4 GHz IEEE 802.15.4 / Zigbee Ready RF Transceiver", Chipcon as Smart RF Preliminary Datasheet (rev 1.2), 2004.

[78] National Frequency Planning Group: "United Kingdom Frequency Allocation Table", Issue Number 14, Annex B, PP 194, Available from: http://www.ofcom.org.uk/radiocomms/isu/ukfat/, 2007.

[79] Neil Reid, Ron Seide: "Medium Access Method", 802.11 (Wi-Fi) Networking Handbook, Chapter 5, PP 100, ISBN 007226234, McGraw Hill, 2003.

[80] Alan Mainwaring, Joseph Polastre, Robert Szewczyk, David Culler, John Anderson: "Wireless Sensor Networks for Habitat Monitoring", WSNA '02, Atlanta, Georgia, September 28, 2002.

[81] L.E.Cordova-Lopez, A. Mason, S. Ross, J. D. Cullen, A. Shaw, A. I. Al-Shamma'a: "A GIS Based Wireless Sensor Network fro the Remote Monitoring and Mapping of Vehicle Exhaust Pollution", Geographical Information Systems Conference 2006, Liverpool John Moores University, Liverpool, L3 3AF.

[82] Zhongmin Wang, Zhen Song, Peng-Yu Chen, Anisha Arora, Kevin Moore, YangQuan Chen: "MASmote – A Mobility Node for MAS-net (Mobile Actuator Sensor Networks)", Center for Self-Organizing and Intelligent Systems, Utah State University, Logan, USA.

[83] VIA Technologies: "VIA Embedded Mainboad Catalogue", http://www.viaembedded.com, 2007.

[84] William Merril, Katayoun Sohrabi, Lewis Girod, Jeremy Elson, Fredric Newberg, William Kaiser: "Open Standard Development Platforms for Distributed Sensor Networks", Sensoria Corporation, 200 Corporate Point Suite 100, Culver City, CA, 90230, USA.

[85] Ralph Kling, Robert Adler, Jonathan Huang, Vincent Hummel, Lama Nachman: "Intel Mote: Sensor Network Technology for Industrial Applications", Intel Corporation, 3600 Juliet Lane, Santa Clara, CA, 95051, USA.

[86] SiliconFarEast.com: "System on a Chip (SoC)", Semiconducter Online reference website, http://www.siliconfareast.com/soc.htm, Cited 01/03/2008.

[87] uPart: "uPart 014xilmt", Preliminary Datasheet, http://particle.teco.edu/upart/, Cited 01/03/2008.

[88] BTnode: "BTnode rev3 – Product Brief", Swiss Federal Institute of Technology, Zurich, http://www.btnode.ethz.ch/, Cited 01/03/2008.

[89] Anthony Rowe, Rahul Mangharam, Raj Rajkumar: "Firefly: A Time Synchronised Real-Time Sensor Networking Platform", Department of Electric and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA.

[90] Joseph Polastre, Robert Szewczyk, Cory Sharp, David Culler: "Commercial Off the Shelf Components (COTS)", The Mote Revolution: Low Power Wireless Sensor Network Devices, Oral Presentation, Hotchips 2004.

[91] Doug Steel: "Smart Dust", Technology Briefing, University of Houston, Information Systems Research Centre, March 2005.

[92] J. L. Hill, "Spec takes the next step toward the vision of smart dust", Spec project website, http://www.jlhlabs.com/jhill_cs/spec/, 2005.

[93] Crossbow: "MPR/MIB Users Manual", Crossbow Technology Inc, 4145N First Street, San Jose, California, http://www.xbow.com/.

[94] Willow Technologies Ltd, Shawlands Court, Newchapel Road, Lingfield, Surrey, RH7 6BL, http://www.willow.co.uk/.

[95] ATMEL: "8-bit AVR Microcontroller with 128K Bytes In-System Programmable Flash", ATmega128L datasheet, http://www.atmel.com/.

[96] Zigbee Alliance, Official Website, http://www.zigbee.org/en/about/, Cited 01/03/2008.

[97] ATMEL: "AT45DB011B DataFlash", Data sheet, http://www.atmel.com/.

[98] Harwin: "M40-600 Series Connectors", Harwin Customer Information Sheet, , http://www.harwin.com/.

[99] Crossbow: "MTS/MDA Sensor Board Users Manual", Crossbow Technology Inc, 4145N First Street, San Jose, California, http://www.xbow.com/.

[100] Philip Levis: "TinyOS: Getting Started", A TinyOS Tutorial, March 2003.

[101] TinyOS Homepage, http://www.tinyos.net/.

[102] Sourceforge: "TinyOS", TinyOS Current Versioning System (CVS) Repository, http://sourceforge.net/cvs/?group_id=28656.

[103] Philip Levis: "TinyOS Programming", TinyOS 2.0 programming guide, Available from http://csl.stanford.edu/~pal/pubs/tinyos-programming.pdf, 2006.

[104] David Gay, David Culler, Philip Levis: "nesC Language Reference Manual", www.tinyos.net/api/nesc/doc/ref.pdf, September 2002.

[105] Jeff Thorn: "Deciphering TinyOS Serial Packets", Octave Tech Brief #5-01, Octave Technology, http://www.octavetech.com/pubs/TB5-01%20Deciphering%20TinyOS%20Serial%20Packets.pdf, 10/03/2005.

[106] Stefano Basagni, Marco Conti, Silvia Giordano, Ivan Stojmenovic: "Destination Sequenced Distance Vector Routing", Mobile Ad Hoc Networking, Chapter 10, PP277-8, ISBN: 0471373133, 2004.

[107] Stefano Basagni, Marco Conti, Silvia Giordano, Ivan Stojmenovic: "Ad Hoc On-Demand Distance Vector Routing", Mobile Ad Hoc Networking, Chapter 10, PP281-3, ISBN: 0471373133, 2004.

[108] Sze-Yao Ni, Yu-Chee Tseng, Yuh-Shyan Chen and Jang-Ping Sheu: "The Broadcast Storm Problem in a Mobile Ad Hoc Network", Proceedings of Mobicom' 99, PP151-162, Seattle, United States, 17/08/99.

[109] Dallas Semiconductor: "DS2401 Silicon Serial Number", Data sheet, http://www.maxim-ic.com/getds.cfm?qv_pk=2903.

[110] Instek: "GDS-810 Oscilloscope", Test and Measurement Product Suppliers, 3661 Walnut Avenue, Chino, United States, http://www.instek.com/.

[111] Mason. A, Shaw. A, Al-Shamma'a. A. I and Welsby. T, "RFID And Wireless Sensor Network Integration For Intelligent Asset Tracking Systems", 2nd GERI Annual Research Symposium GARS-2006, Liverpool, UK, 15 June 2006.

[112] Bader. E, Shaw. A and Al-Shamma'a. A. I, "Using A Powerful, Accurate Speech Compression Algorithm In Low-Power Consumption Wireless Devices", 3rd GERI Annual Research Symposium GARS-2007, Liverpool, UK, 27 June 2007.

[113] Maxim: "MAX3223 RS-232 Transceiver with AutoShutdown", Data Sheet, http://www.maxim-ic.com/.

[114] National Semiconductor: "LM117 800mA Low Dropout Linear Regulator", Data Sheet, http://www.national.com/.

[115] Cadsoft: "Eagle PCB CAD Software", Cadsoft Online, http://www.cadsoft.de/, 2008.

[116] Everbouquet International: "MC16024X-Series LCD Display", Data Sheet, http://www.everbouquet.com.tw/.

[117] Microchip: "PICF87xx 28/40 pin 8-bit CMOS Flash", Data Sheet, http://www.microchip.com/.

[118] Fairchild Semiconductor: "LM78xx 3 Terminal 1A Positive Voltage Regulator", Data Sheet, http://www.fairchildsemi.com/.

[119] Watson: "Watson WSMA889 Handheld Antenna", Waters and Stanton Catalogue Webpage, http://www.wsplc.com/acatalog/Watson_Handheld_Antennas.html.

[120] Microchip: "MPLAB Users Guide", PIC Development Library, http://ww1.microchip.com/downloads/en/DeviceDoc/51519a.pdf, 2005.

[121] Hitachi: "LM016L" Liquid Crystal Character Display Modules Catalogue, PP42-3, Hitachi Europe Ltd, Whitebrook Park, Lower Cookham Rd, Maidenhead, Berkshire, SL6 8YA, Printed March 1994.

[122] Borland Corporation: "Borland Delphi 7", http://www.borland.com/, 2001.

[123] Sun Microsystems: "MySQL", http://www.mysql.com/, 1995.

[124] Joseph J. Carr: "Radio waves", Practical Antenna Handbook (Fourth Edition), Chapter 2, PP 8, ISBN: 0071374353, 2001.

[125] Fawwaz T. Ulaby: "The Electromagnetic Spectrum", Fundamentals of Applied Electromagnetics, Chapter 1, PP 22, ISBN: 0130115541, 1999.

[126] Lamont V. Blake: "Characteristics of Electromagnetic Waves", Antennas, Chapter 1, PP 10, ISBN: 0471079286, 1966.

[127] Fawwaz T. Ulaby: "Plane Wave Propagation in Lossy Media", Fundamentals of Applied Electromagnetics, Chapter 7, PP 277-9, ISBN: 0130115541, 1999.

[128] Health & Safety Executive: "Specification for high-strength seamless steel gas cylinders in a modified chromium molybdenum steel", Specification Number HSSS-2, http://www.hse.gov.uk/cdg/pdf/standard/hsss2.pdf, June 1992.

[129] Klaus Finkenzeller: "Electromagnetic Backscatter Coupling", RFID Handbook, Chapter 3, PP 47, ISBN 0470844027, Wiley & Sons, 2003.

[130] Joseph J. Carr: "Fading Mechanisms", Practical Antenna Handbook (Fourth Edition), Chapter 2, PP 35, ISBN: 0071374353, 2001.

[131] Stefano Basagni, Marco Conti, Silvia Giordano, Ivan Stojmenovic: "Multipath Fading and Shadowing", Mobile Ad Hoc Networking, Chapter 8, PP235-6, ISBN: 0471373133, 2004.

[132] Lamont V. Blake: "Interference", Antennas, Chapter 1, PP 26-7, ISBN: 0471079286, 1966.

[133] Fawwaz T. Ulaby: "Faradays Law", Fundamentals of Applied Electromagnetics, Chapter 6, PP 231-3, ISBN: 0130115541, 1999.

[134] Joseph J. Carr: "Transmission Lines Characteristic Impedance", Practical Antenna Handbook (Fourth Edition), Chapter 2, PP 67, ISBN: 0071374353, 2001.

[135] Fawwaz T. Ulaby: "Quarter wave monopole antenna", Fundamentals of Applied Electromagnetics, Chapter 9, PP 358, ISBN: 0130115541, 1999.

[136] Constantine A. Balanis: "Loop Antennas", Antennas Theory Analysis and Design (Second Edition), Chapter 5, PP 203, ISBN: 0471592684, 1997.

[137] Constantine A. Balanis: "Loop Antennas", Antennas Theory Analysis and Design (Second Edition), Chapter 5, PP 210, ISBN: 0471592684, 1997.

[138] Lamont V. Blake: "Antenna Parameters", Antennas, Chapter 3, PP 122, ISBN: 0471079286, 1966.

[139] Joseph J. Carr: "Near field and far field", Practical Antenna Handbook (Fourth Edition), Chapter 19, PP 395, ISBN: 0071374353, 2001.

[140] Constantine A. Balanis: "Field Regions", Antennas Theory Analysis and Design (Second Edition), Chapter 2, PP 33, ISBN: 0471592684, 1997.

[141] Fawwaz T. Ulaby: "Reciprocity", Fundamentals of Applied Electromagnetics, Chapter 9, PP 341, ISBN: 0130115541, 1999.

[142] Lamont V. Blake: "Directive vs Power Gain", Antennas, Chapter 3, PP 126, ISBN: 0471079286, 1966.

[143] Constantine A. Balanis: "Bandwidth", Antennas Theory Analysis and Design (Second Edition), Chapter 2, PP 63-4, ISBN: 0471592684, 1997.

[144] Ramesh Garg, Prakash Bharitia, Inder Bahl, Apisak Ittipiboon: "Broadbanding of Microstrip Antennas", Microstrip Antenna Design Handbook, Chapter 9, PP 534-5, ISBN: 0890065136, 2001.

[145] Arie Voors: "NEC based antenna modeller and optimizer", 4nec2 antenna modelling software, http://home.ict.nl/%7Earivoors/, March 2008.

[146] W7EL: "EZNEC Antenna Software", EZNEC modelling software, http://www.eznec.com/.

[147] Ansoft: "HFSS", 3D Full-wave Electromagnetic Field Simulation, 225 West Station Square Drive, Suite 200, Pittsburgh, PA 15219, USA, http://www.ansoft.com/products/hf/hfss/, 2008.

[148] EM Software & Systems: "Feko", Langley Research Park, 144 Research Drive, Hampton, VA23666, USA, http://www.feko.info/, 2008.

[149] Lamont V. Blake: "Radiation Resistance and Efficiency", Antennas, Chapter 3, PP 138-9, ISBN: 0471079286, 1966.

[150] Constantine A. Balanis: "Microstrip Antennas", Antennas Theory Analysis and Design (Second Edition), Chapter 14, PP 722, ISBN: 0471592684, 1997.

[151] Neil Reide and Ron Seide: "Modulation Schemes", 802.11 (Wi-Fi) Networking Handbook, Chapter 3, PP 46-7, ISBN: 0072226234, 2003.

[152] Constantine A. Balanis: "Microstrip Antennas", Antennas Theory Analysis and Design (Second Edition), Chapter 14, PP 727-730, ISBN: 0471592684, 1997.

[153] Ramesh Garg, Prakash Bharitia, Inder Bahl, Apisak Ittipiboon: "Effects of Finite Size Ground Plane", Microstrip Antenna Design Handbook, Chapter 4, PP 293-296, ISBN: 0890065136, 2001.

[154] Ansoft: "Probe Feed Patch Antenna Example", HFSS Users Guide, Section 5.3, 225 West Station Square Drive, Suite 200, Pittsburgh, PA 15219, USA, http://www.ansoft.com/products/hf/hfss/, 2008.

[155] Anritsu: "VNA Master MS2024A", Americas Sales Region Headquarters, 1115 East Collins Boulevard, Suite 100, Richardson, TX 75081, http://www.us.anritsu.com/products/ARO/North/Eng/showProd.aspx?&ID=742, 2006.

[156] K. Nithisopa, J. Nakasuwan, N. Songthanapitak, N. Anantrasirichai and T. Wakabayashi: "Design CPW Fed Slot Antenna for Wideband Applications", Piers Online, Vol. 3, No. 7, P1124-7, http://piers.mit.edu/piersonline/, 2007.

[157] Ansoft: "Port Tutorial Series: Coplanar Waveguide (CPW)", HFSS v8 Training, http://web.doe.carleton.ca/~mmariani/Thesis/port_tutorial_CPW.ppt.

[158] John D. Kraus, Ronald J. Marhefka: "Folded Dipole Antennas", Antennas for All Applications (Third Edition), Chapter 16, PP 593-7, ISBN: 0072321032, 2002.

[159] John D. Kraus, Ronald J. Marhefka: "Directivity and Gain", Antennas for All Applications (Third Edition), Chapter 2, PP 23-6, ISBN: 0072321032, 2002.

[160] John D. Kraus, Ronald J. Marhefka: "Anechoic Chambers and Absorbing Materials", Antennas for All Applications (Third Edition), Chapter 24, PP 841-4, ISBN: 0072321032, 2002.

[161] Lamont V. Blake: "Waveguides", Antennas, Chapter 2, PP 94, ISBN: 0471079286, 1966.

[162] Constantine A. Balanis: "Long Wire Antenna Polarisation", Antennas Theory Analysis and Design (Second Edition), Chapter 10, PP 496, ISBN: 0471592684, 1997.

[163] Energizer: "Temperature Effects on Performance", Alkaline Manganese Dioxide Handbook and Application Manual, http://www.energizer.com/, 2008.

[164] Mason. A, Al-Shamma'a. A. I, Shaw. A, Irven. J and Wiktorowicz. R, "Intelligent Wireless Asset Tracking Of Packaged Gases", Sensors And Their Applications XIV Conference (SENSORS07), Liverpool, UK, 11–13 September 2007, Journal of Physics: Conference Series, ISSN: 1742-6588, doi:10.1088/1742-6596/76/1/012037, Vol. 76, IOP Publishing, Article No. 012037, 2007.

[165] Farnell: "Pressure Transmitter 0-1000 bar", http://uk.farnell.com/7056473/industrial-controls-automation/product.us0?sku=vdo-imt-3396088001.

[166] Farnell: "Pressure Sensor 0-100 PSI", http://uk.farnell.com/731729/industrial-controls-automation/product.us0?sku=honeywell-s-c-24pcffm6g.

[167] Jim Breithaupt: "Piezoelectricity", Understanding Physics for Advanced Level (Third Edition), Chapter 13, PP 178, ISBN: 0748715797, 1995.

[168] Air Products: "Acetylene", Safetygram-13, http://www.airproducts.com/Responsibility/EHS/ProductSafety/ProductSafetyInformation/Safetygrams/safetygram13.htm, 1994.

[169] Analog Devices: "ADXL202E Low-cost 2G Dual-Axis Accelerometer", Product Datasheet, http://www.analog.com/en/prod/0,,764_800_ADXL202,00.html, 2000.

[170] A. L. Phillips: "Fundamentals of Welding", Welding Handbook 6th Edition (Volume 1), Chapter 3, PP3.1-3.10.

[171] Tecna: "Rocker Arm Pneumatic Spot Welders", Corporate Website, http://www.tecna.net/pr46xx.html.

[172] Rocoil: "Precision Rogowski Coils",5 Almsford Avenue, Harrogate, North Yorkshire, HG2 8HD, http://www.rocoil.cwc.net/.

[173] National Instruments: "NI Multisim", http://www.ni.com/academic/multisim.htm , 2008.

[174] The Network Simulator: "ns-2", http://nsnam.isi.edu/nsnam/index.php/User_Information, 2008.

[175] Crossbow: "Wireless Sensor Networks", 4145 N. First Street, San Jose, CA 95134, USA, http://www.xbow.com/Home/wHomePage.aspx, 2008.

[176] Microstrain: "Wireless Sensors", 310 Hurricane Lane, Suite 4, Williston, VT 05495, USA, http://www.microstrain.com/, 2008.

[177] Texas Instruments: "CC2431 System on Chip for 2.4GHz Zigbee / IEEE 802.15.4 with Location Engine", CC2341 PRELIMINARY Datasheet (Rev 1.01), 2006.

[178] Nicki Hayes: "Airport Security – Can Wireless Help?", Wireless Developer Network, http://www.wirelessdevnet.com/channels/wireless/features/newsbyte36.html, 20th November 2001.

10. Workplan

Task	1/9/05 to 30/11/05	1/12/05 to 28/2/06	1/3/06 to 31/5/06	1/6/06 to 31/8/06	1/9/06 to 30/11/06	1/12/06 to 28/2/07	1/3/07 to 31/5/07	1/6/07 to 31/8/07	1/9/07 to 30/11/07	1/12/07 to 28/2/08	1/3/08 to 31/5/08	1/6/08 to 31/8/08
Literature review of RFID and WSN												
Familiarisation with motes and TinyOS structure												
Develop mote software to utilise onboard signal strength monitor and conduct extensive testing of motes in different environments												
Create standalone WSN based tracking prototype												
Theoretical studies and understanding of antenna fundamentals												
Design, simulated and build an alternative mote antenna												
Research the method for luggage location and setup a laboratory experiment												
Industrial trials; both laboratory setup and industrial testing (if possible).												
Spot welding trials for the automotive industry.												
Thesis writing												

11. Related Work

A. Mason, A. Shaw, A. I. Al-Shamma'a, *'Intelligent Radio Frequency Identification'*, IET Speak Out for Engineering 2006, The University of Liverpool, May 2006.
- (Awarded a prize for being the 'Best Young Researcher').

Mason. A, Shaw. A, Al-Shamma'a. A.I, "Asset tracking: Beyond RFID", PGNET, 7th Annual PG Symposium on the Convergance of Telecommunications, Networking and Broadcasting, 26-27 June 2006, ISBN 1902560139, pp:267-272, 2006.

Mason. A, Shaw. A, Al-Shamma'a. A. I and Welsby. T, "RFID And Wireless Sensor Network Integration For Intelligent Asset Tracking Systems", 2nd GERI Annual Research Symposium GARS-2006, Liverpool, UK, 15 June 2006.

Cordova-Lopez. L. E, Mason. A, Cullen. J. D, Shaw. A and Al-Shamma'a. A. I, "Online Vehicle And Atmospheric Pollution Monitoring Using GIS And Wireless Sensor Networks", Sensors And Their Applications XIV Conference (SENSORS07), Liverpool, UK, 11–13 September 2007, Journal of Physics: Conference Series, ISSN: 1742-6588, doi:10.1088/1742-6596/76/1/012019, Vol. 76, IOP Publishing, Article No. 012019, 2007.

Mason. A, Al-Shamma'a. A. I, Shaw. A, Irven. J and Wiktorowicz. R, "Intelligent Wireless Asset Tracking Of Packaged Gases", Sensors And Their Applications XIV Conference (SENSORS07), Liverpool, UK, 11–13 September 2007, Journal of Physics: Conference Series, ISSN: 1742-6588, doi:10.1088/1742-6596/76/1/012037, Vol. 76, IOP Publishing, Article No. 012037, 2007.

Mason. A, Shaw. A, and Al-Shamma'a. A.I, "Intelligent Radio Frequency Identification Positioning Using Wireless Sensor Networks", IEEE Antenna and Propagation Conference, Loughborough, UK, 2-3 April, 2007, ISBN:1424407753, doi:10.1109/LAPC.2007.367452, pp 145-148, 2007.

Mason. A, Shaw. A and Al-Shamma'a. A. I, "Intelligent Radio Frequency Identification Using Wireless Sensor Networks", 3rd GERI Annual Research Symposium GARS-2007, Liverpool, UK, 27 June 2007.

Mason. A, Shaw. A and Al-Shamma'a. A. I, "An Industrially Suitable Mote Antenna", 4th GERI Annual Research Symposium GARS-2008, Liverpool, UK, 27 June 2008.

Appendix A – Inventory Management System Code

Due to the volume of this appendix it is included on the below CD in PDF format.

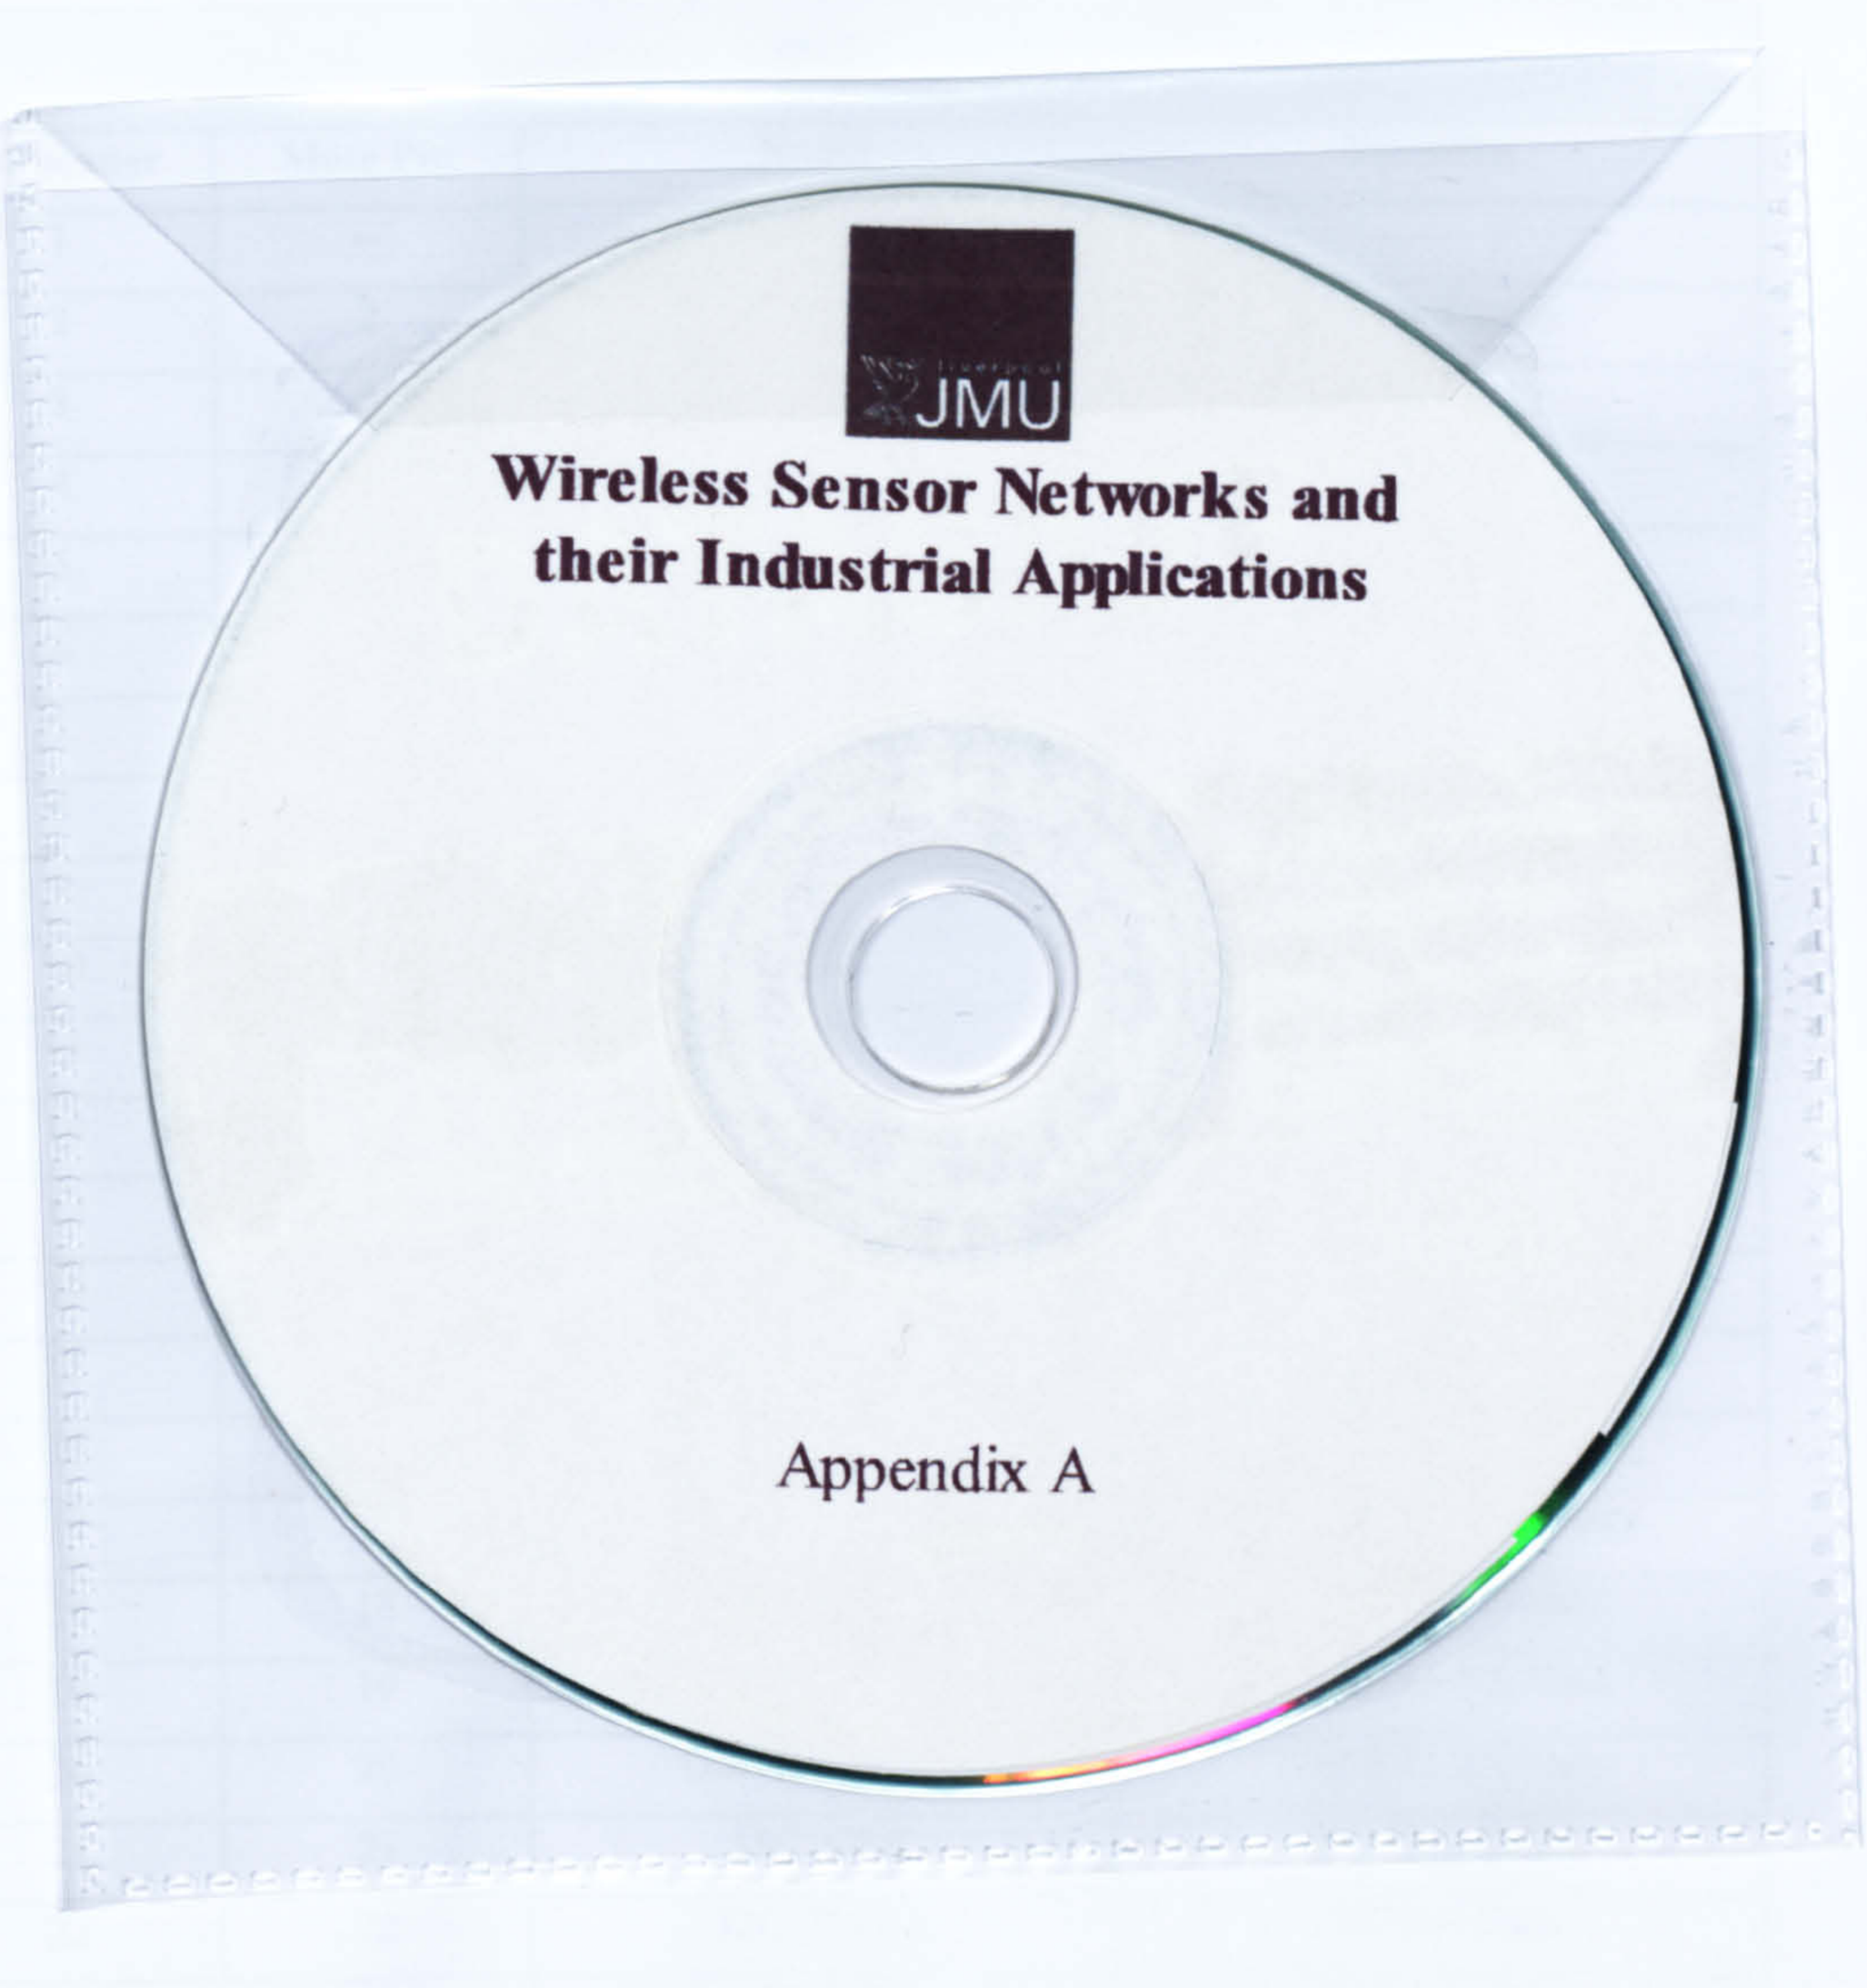

Appendix B – Circuit Schematics and Related Material

B.1. Mote Interface Board Pin Mapping

SKT1			
Pin header	**Mote Pin**	**Name**	**Function**
1	--	GND	
2	2	VSNSR	Sensor Supply
3	3	INT3	GPIO
4	4	INT2	GPIO
5	5	INT1	GPIO
6	6	INT0	GPIO
7	7	BAT_MON	Battery Monitor Enabled
8	8	LED3	LED3 YELLOW
9	9	LED2	LED2 GREEN
10	10	LED1	LED1 RED
11	11	RD	GPIO
12	12	WR	GPIO
13	13	ALE	GPIO
14	14	PW7	Power Control 7
15	15	USART1_CLK	USART1 Clock
16	16	PROG_MOSI	Serial Program MOSI
17	17	PROG_MISO	Serial Program MISO
18	18	SPI_SCK	SPI Serial Clock
19	19	USART1_RXD	USART1 RX Data
20	20	USART1_TXD	USART1 TX Data
21	21	I2C_CLK	I2C Bus Clock
22	22	I2C_DATA	*I2C Bus Data*
23	23	PWM0	GPIO/PWM0
24	24	PWM1A	GPIO/PWM1A
25	25	AC+	GPIO/AC+
26	26	AC-	GPIO/AC-

SKT2			
1	--	See Jumper	--
2	--	See Jumper	--
3	29	PW0	Power Control 0
4	30	PW1	Power Control 1
5	31	PW2	Power Control 2
6	32	PW3	Power Control 3
7	33	PW4	Power Control 4
8	34	PW5	Power Control 5
9	35	PW6	Power Control 6
10	36	ADC7	ADC Input 7 / Batt monitor
11	37	ADC6	ADC Input 6
12	38	ADC5	ADC Input 5
13	39	ADC4	ADC Input 4
14	40	ADC3	ADC Input 3
15	41	ADC2	ADC Input 2
16	42	ADC1	ADC Input 1
17	43	ADC0	ADC Input 0
18	44	THERM_PWR	Enable Temp. Sensor
19	45	THRU1	Thru Connect 1
20	46	THRU2	Thru Connect 2
21	47	THRU3	Thru Connect 3
22	48	RSTN	Reset Neg
23	49	PW1B	GPIO/PWM1B
24	50	VCC	Digital Supply
25	--	GND	Ground
26	--		

B.2. Mote Interface Board Schematic

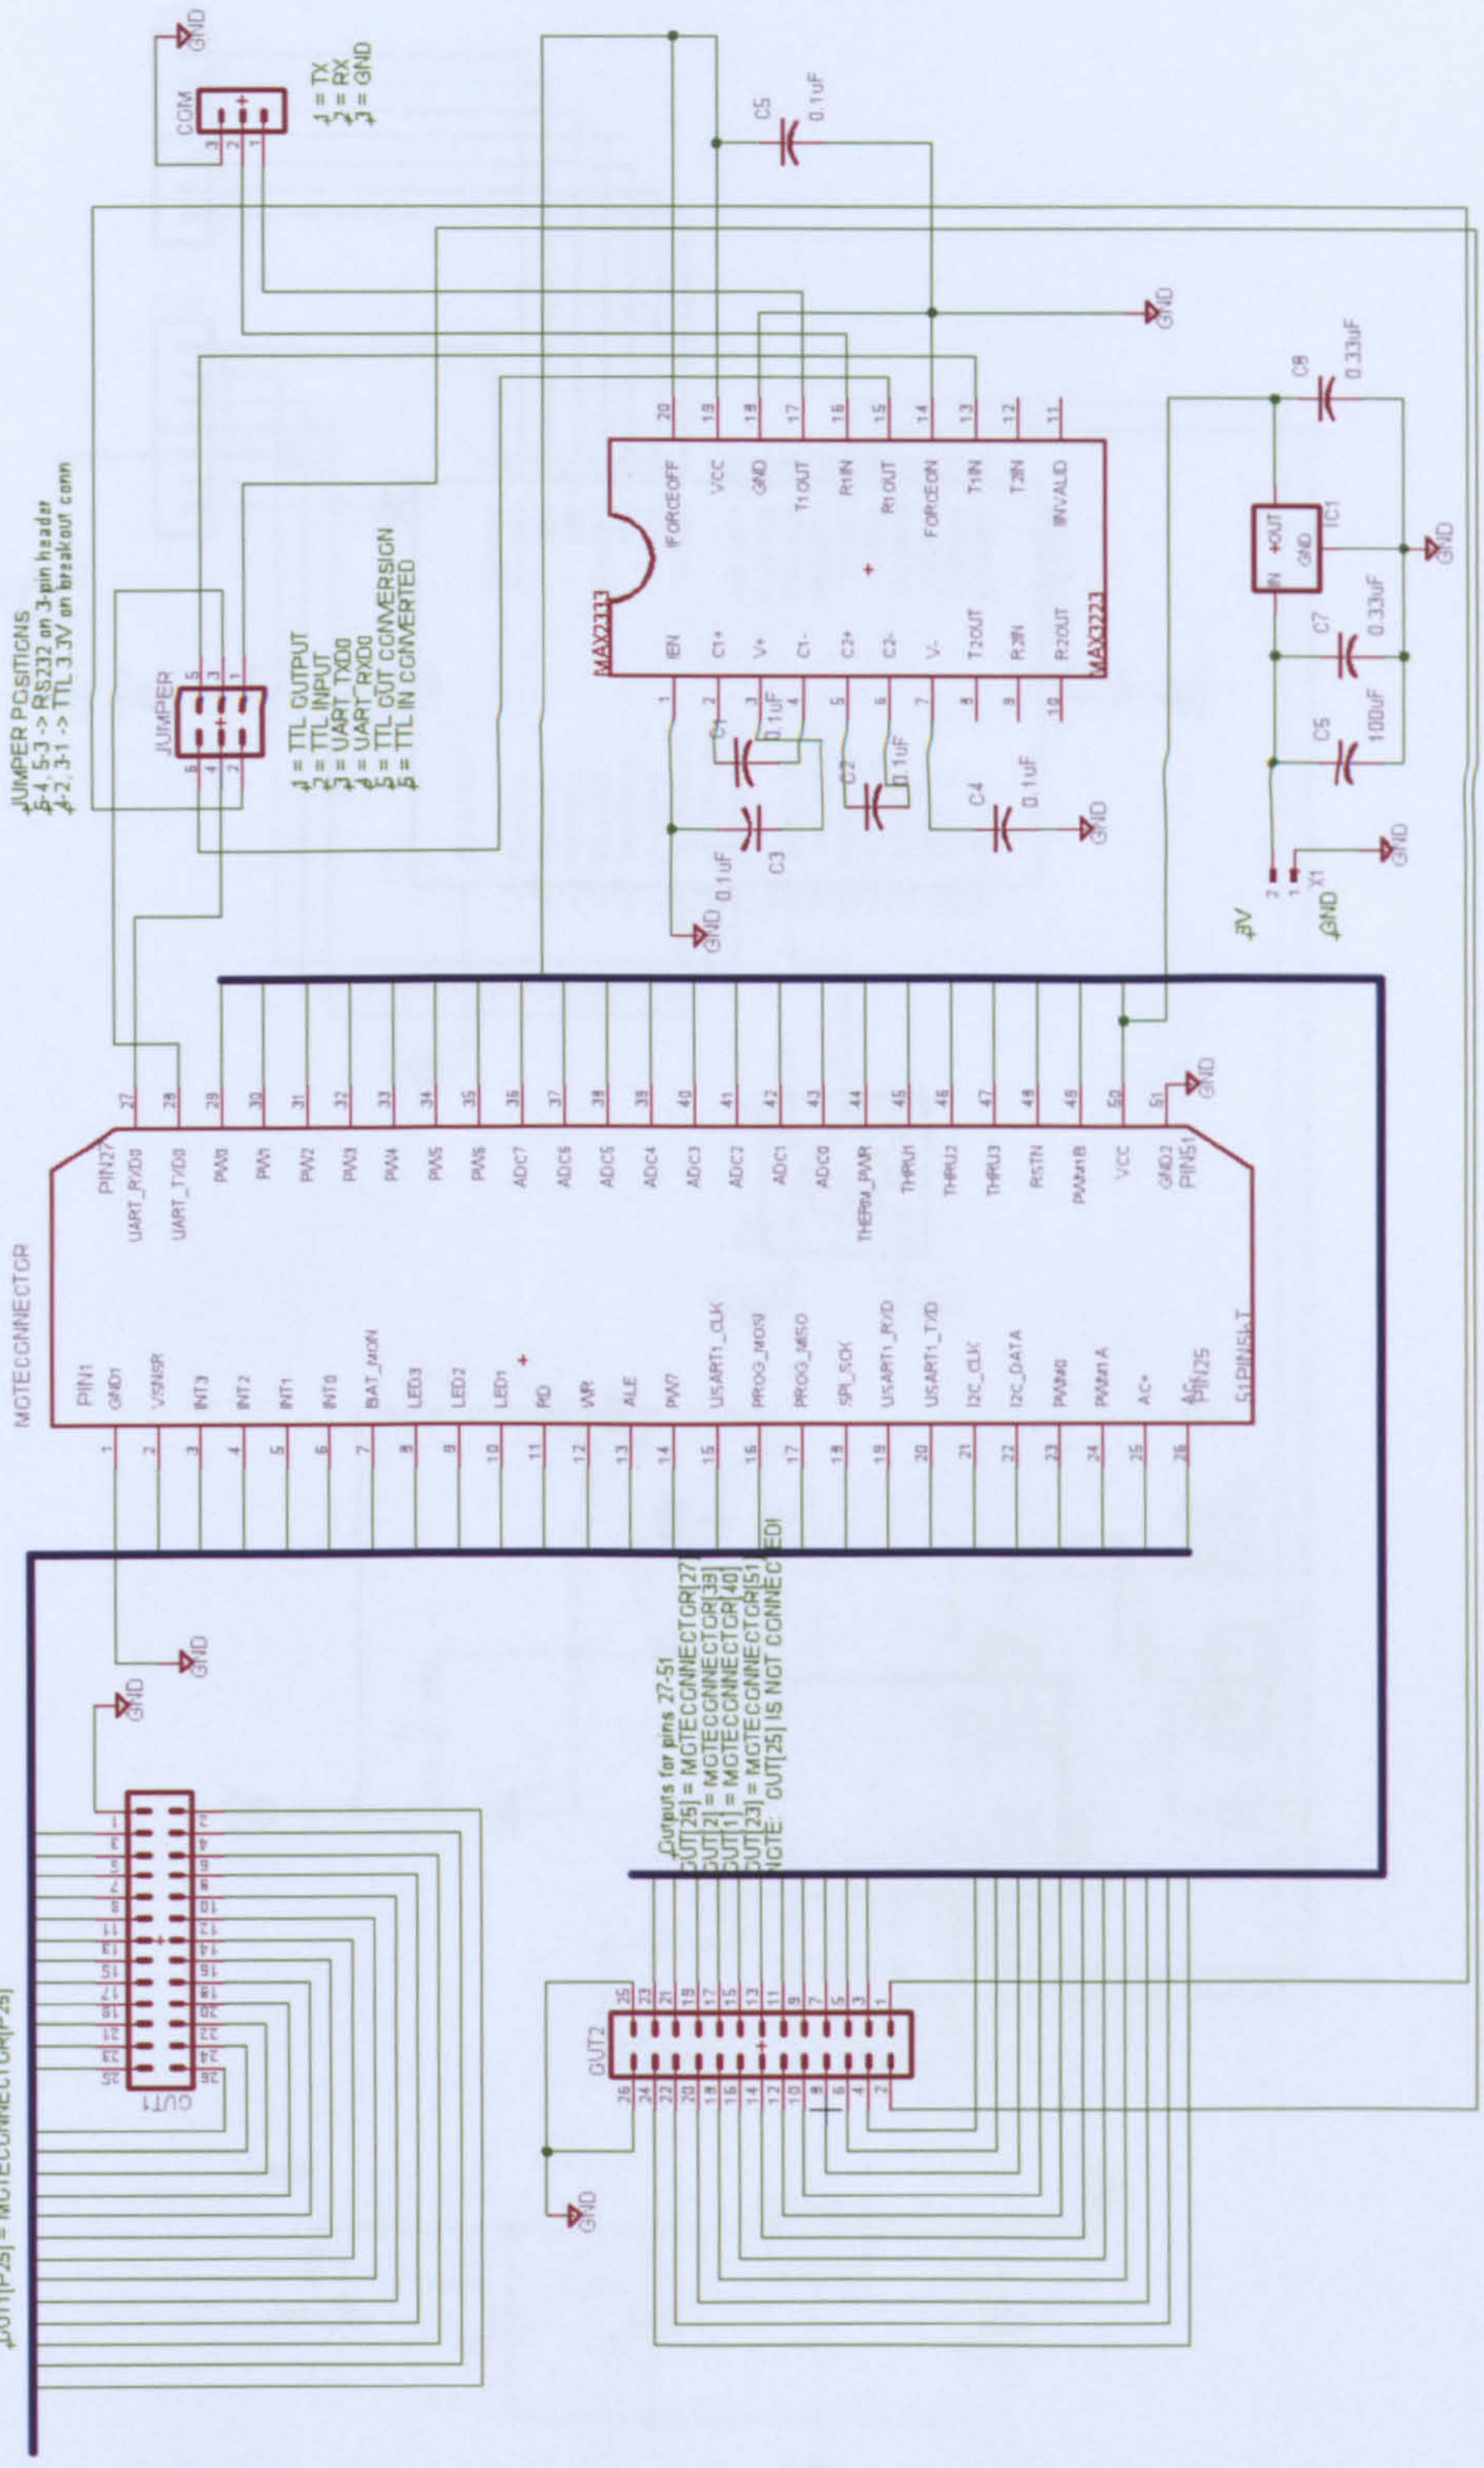

B.3. LCD Interface Board Schematic

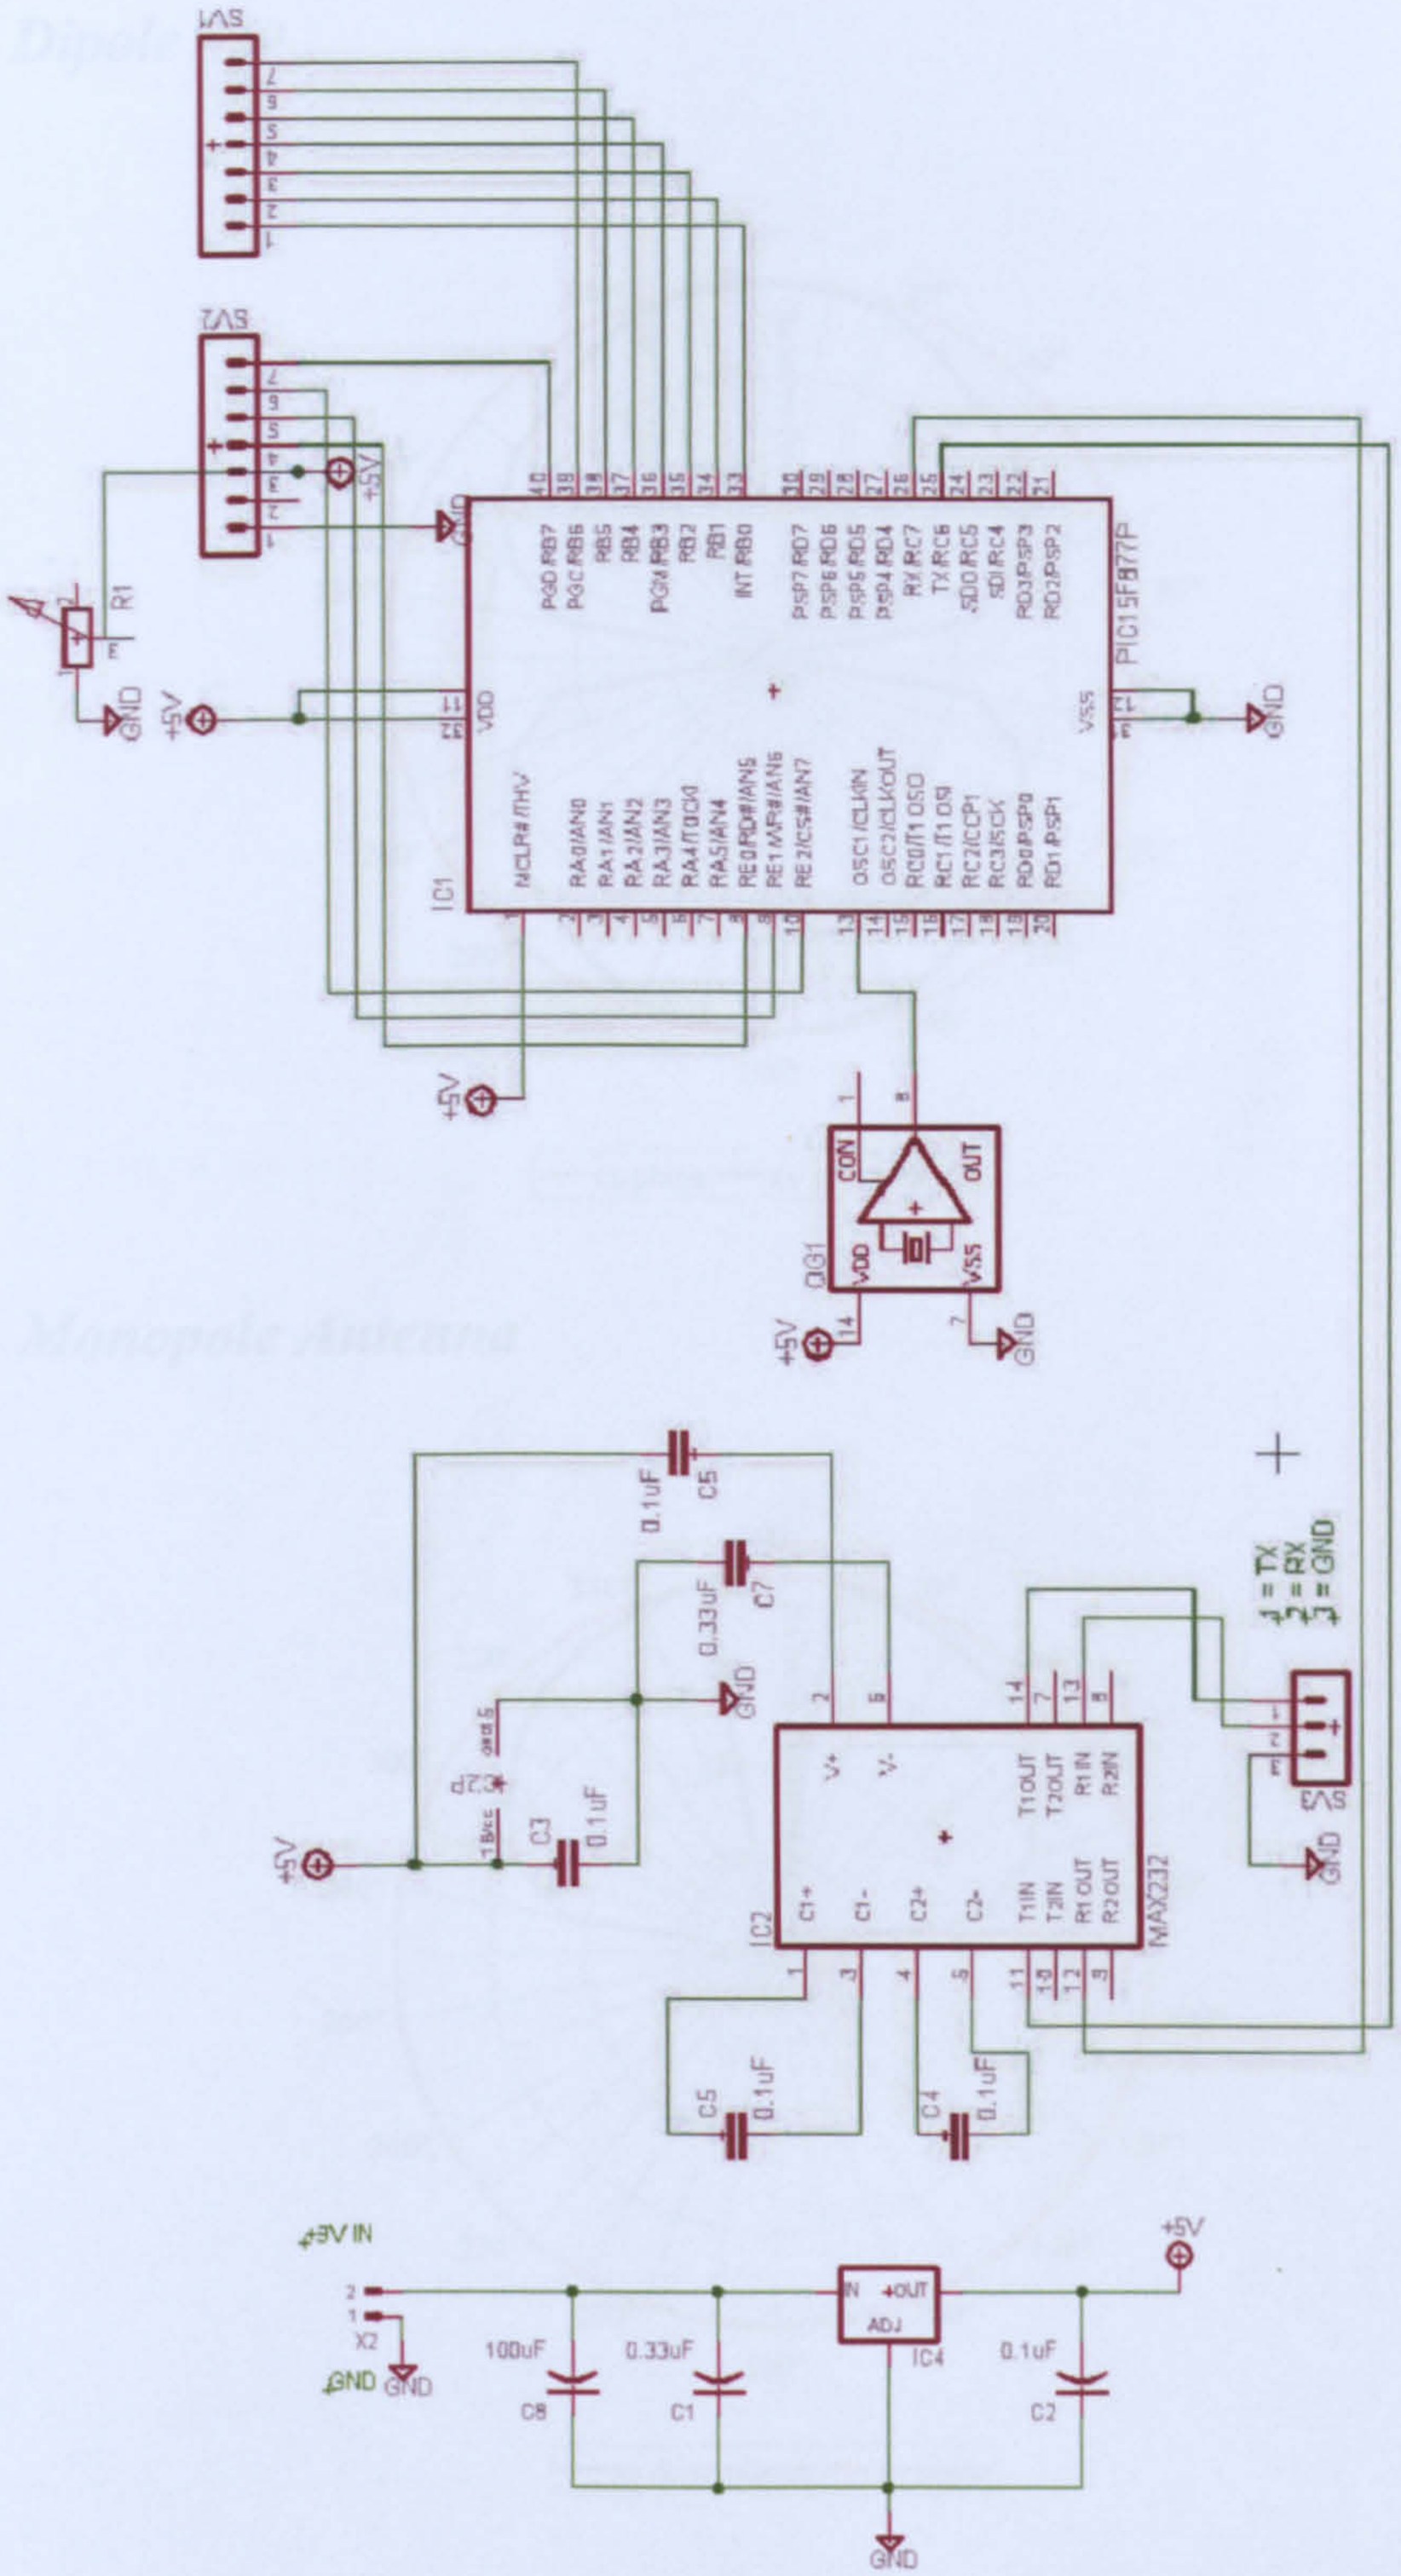

Appendix C – Typical Radiation Patterns

C.1. Dipole

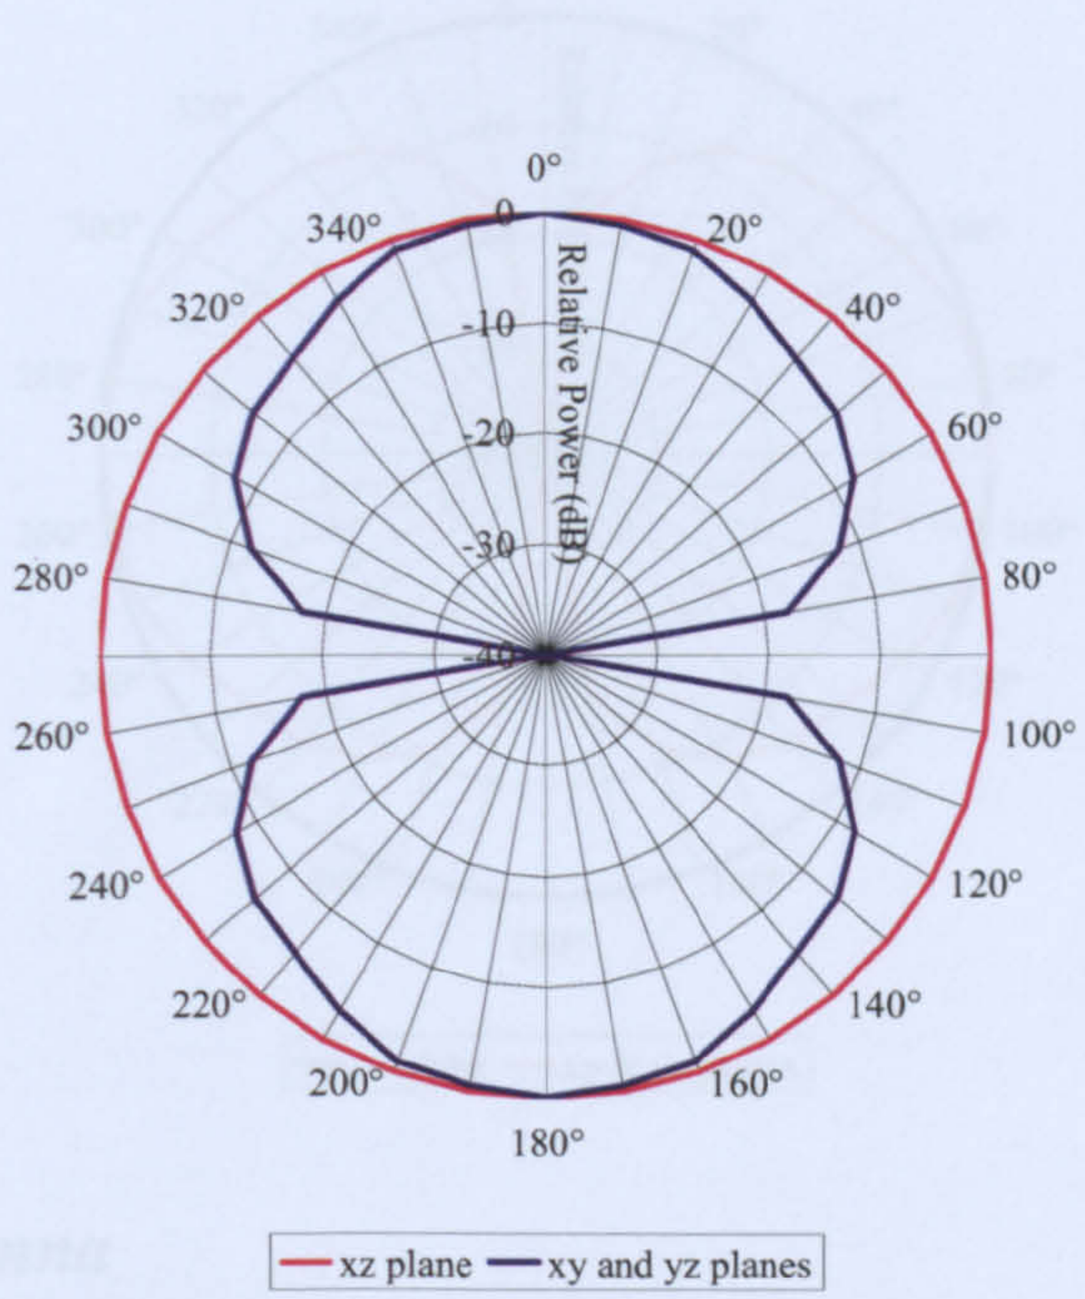

C.2. Monopole Antenna

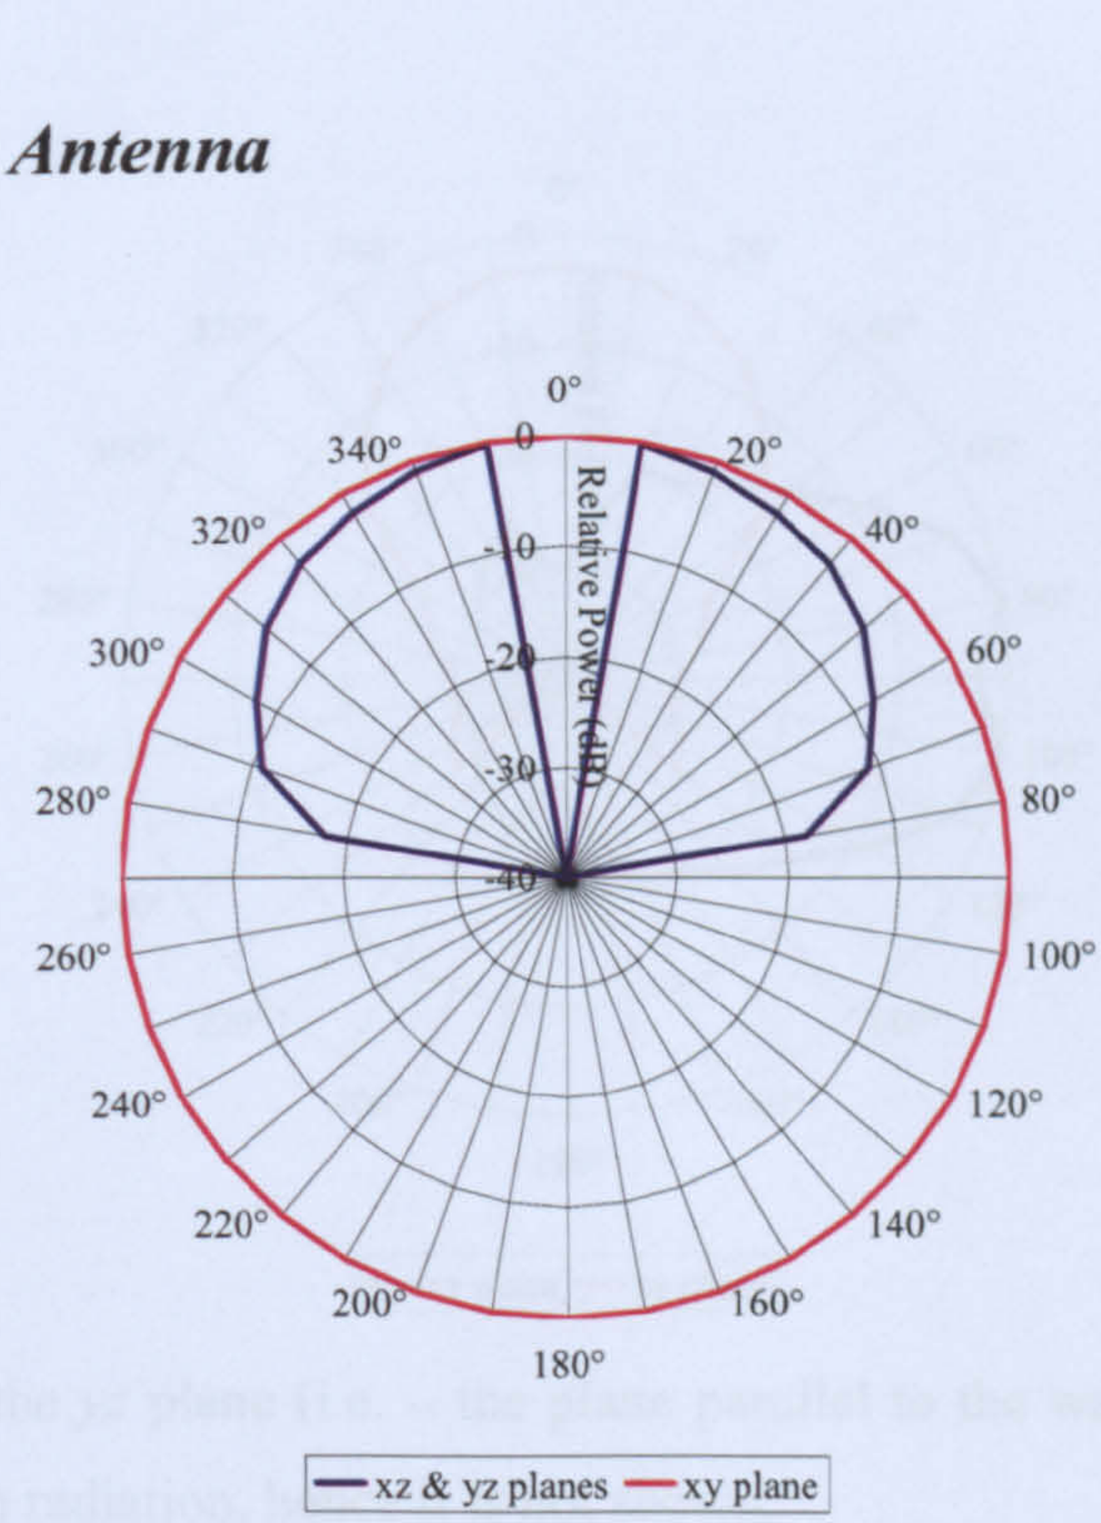

C.3. Loop Antenna

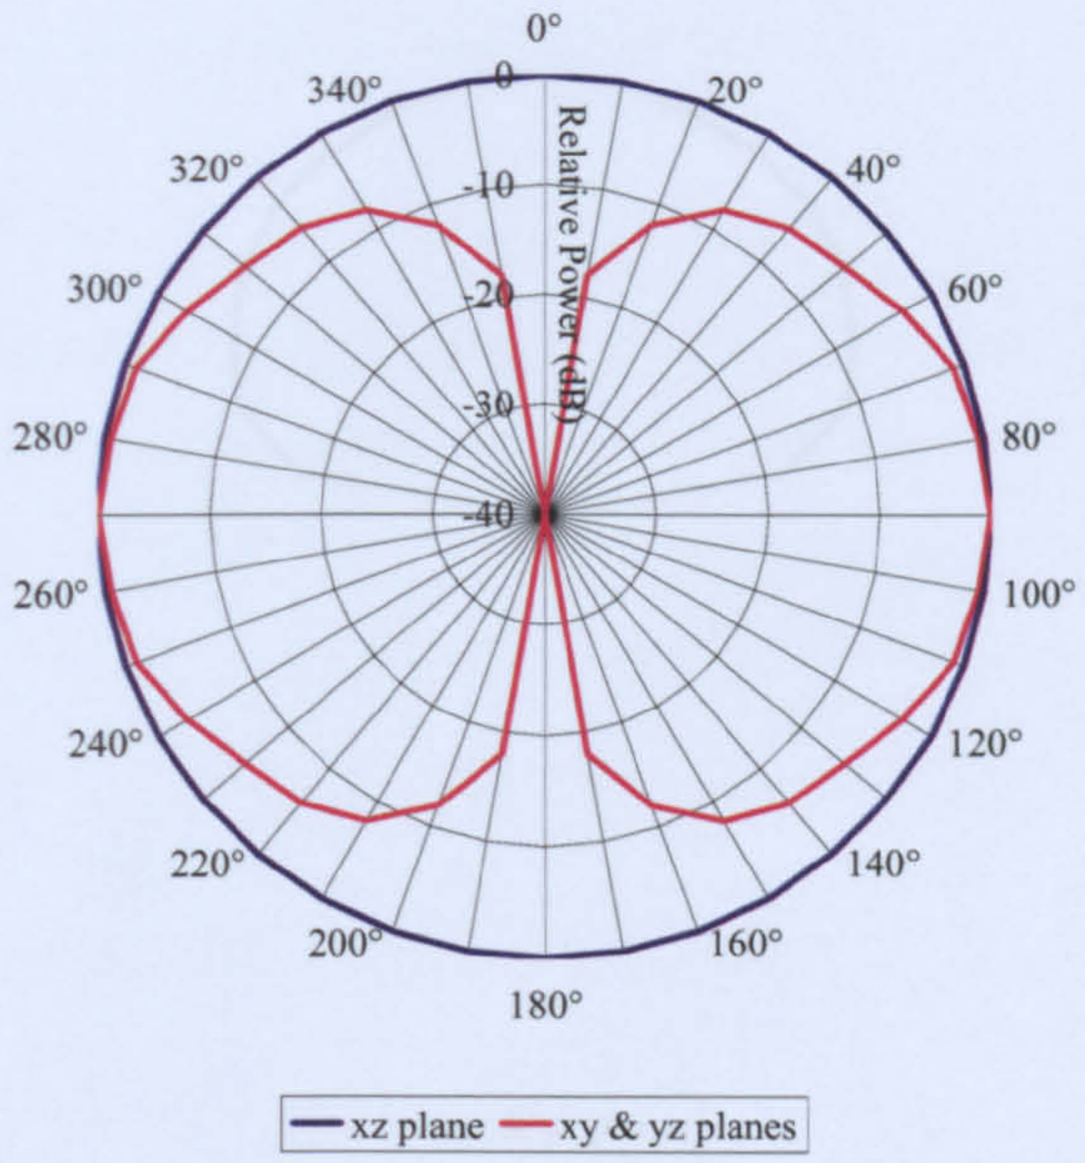

C.4. Horn Antenna

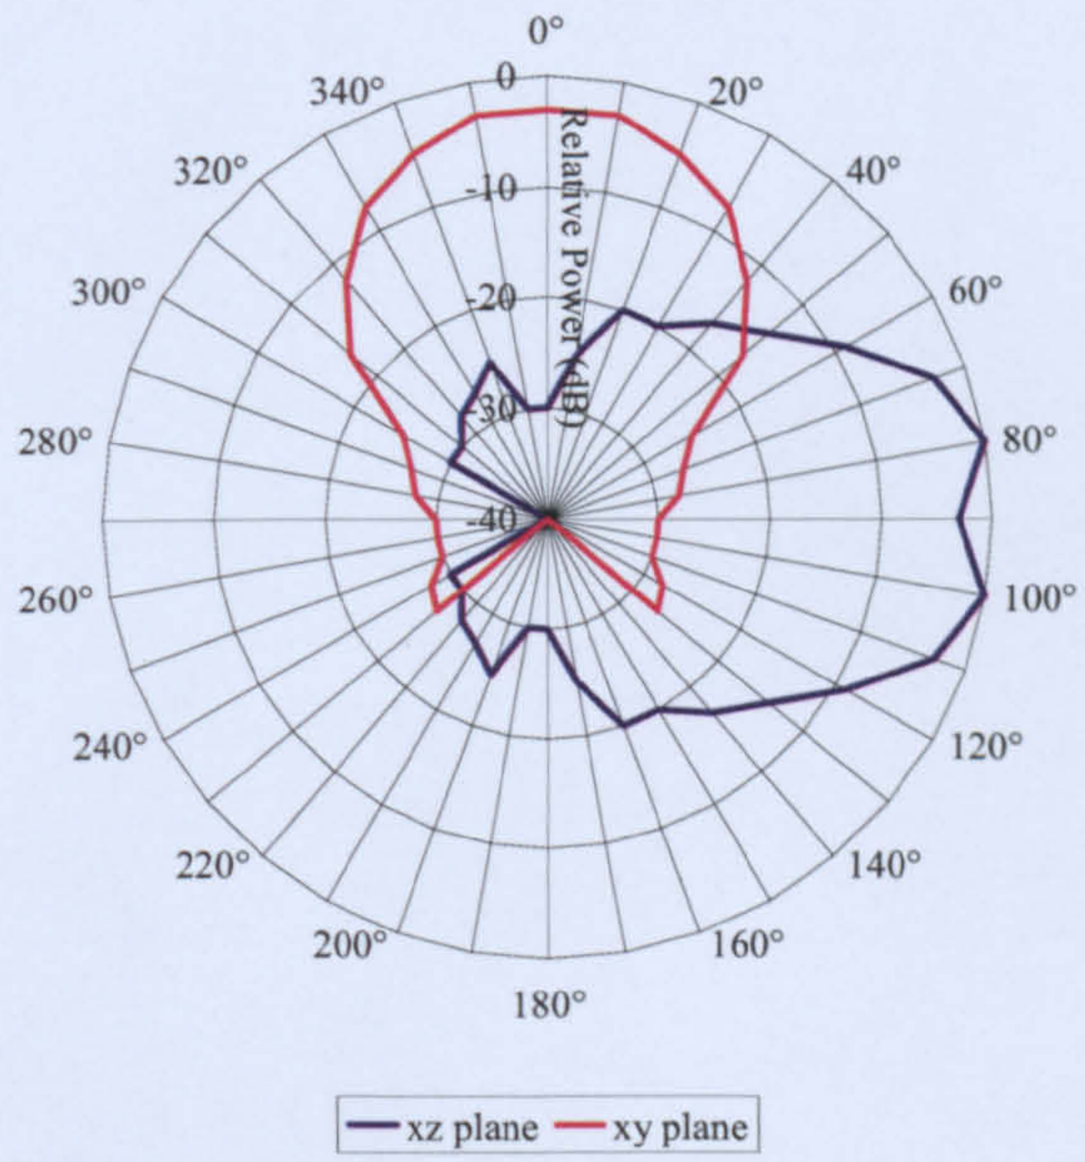

Note: In this plot the *yz* plane (i.e. – the plane parallel to the waveguide opening) is assumed to have no radiation, hence it is not shown.

C.5. Patch Antenna

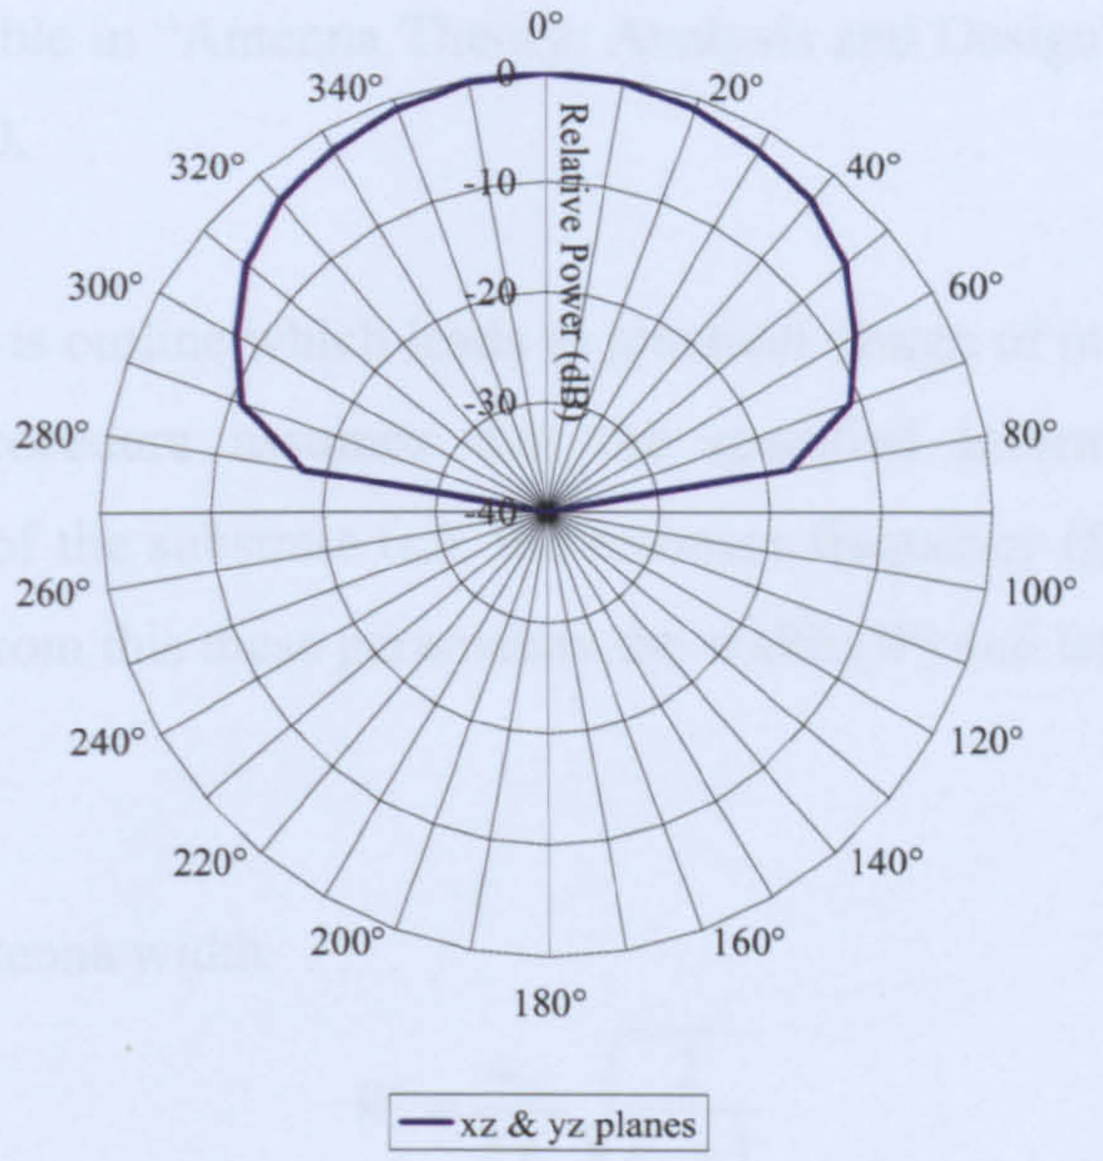

Note: In this plot the *xy* plane (i.e. – the plane parallel to the PCB) is assumed to have no radiation, hence it is not shown.

Appendix D – Simple Patch Antenna Dimensions Guide

This guide is available in "Antenna Theory: Analysis and Design" by Constantine A. Balanis on Page 730.

A design procedure is outline which leads to practical design of rectangular microstrip antennas. The procedure assumes that the specified information includes the dielectric constant of the substrate (ε_r), the resonant frequency (f_r), and the height of the substrate (h). From this these parameters the width (W) and length (L) of the patch can be calculated.

1. Calculate the antenna width

$$W = \frac{v_o}{2f_r}\sqrt{\frac{2}{\varepsilon_r + 1}}$$

2. Calculate the effective dielectric ε_{reff} constant; since the EM energy propagates in air and in the dielectric substrate ε_{reff} takes both mediums into account.

$$\varepsilon_{reff} = \frac{\varepsilon_r + 1}{2} + \frac{\varepsilon_r - 1}{2}\left[1 + 12\frac{h}{W}\right]^{-\frac{1}{2}}$$

3. Now calculate the reduction of L, ΔL, which will take into account fringing effects.

$$\Delta L = 0.412h\frac{\left(\varepsilon_{reff} + 0.3\right)\left(\frac{W}{h} + 0.264\right)}{\left(\varepsilon_{reff} - 0.258\right)\left(\frac{W}{h} + 0.8\right)}$$

4. Finally calculate L.

$$L = \frac{1}{2fr\sqrt{\varepsilon_{reff}}} - 2\Delta L$$

Appendix E – Reference Monopole Simulation

E.1. Simulation Setup

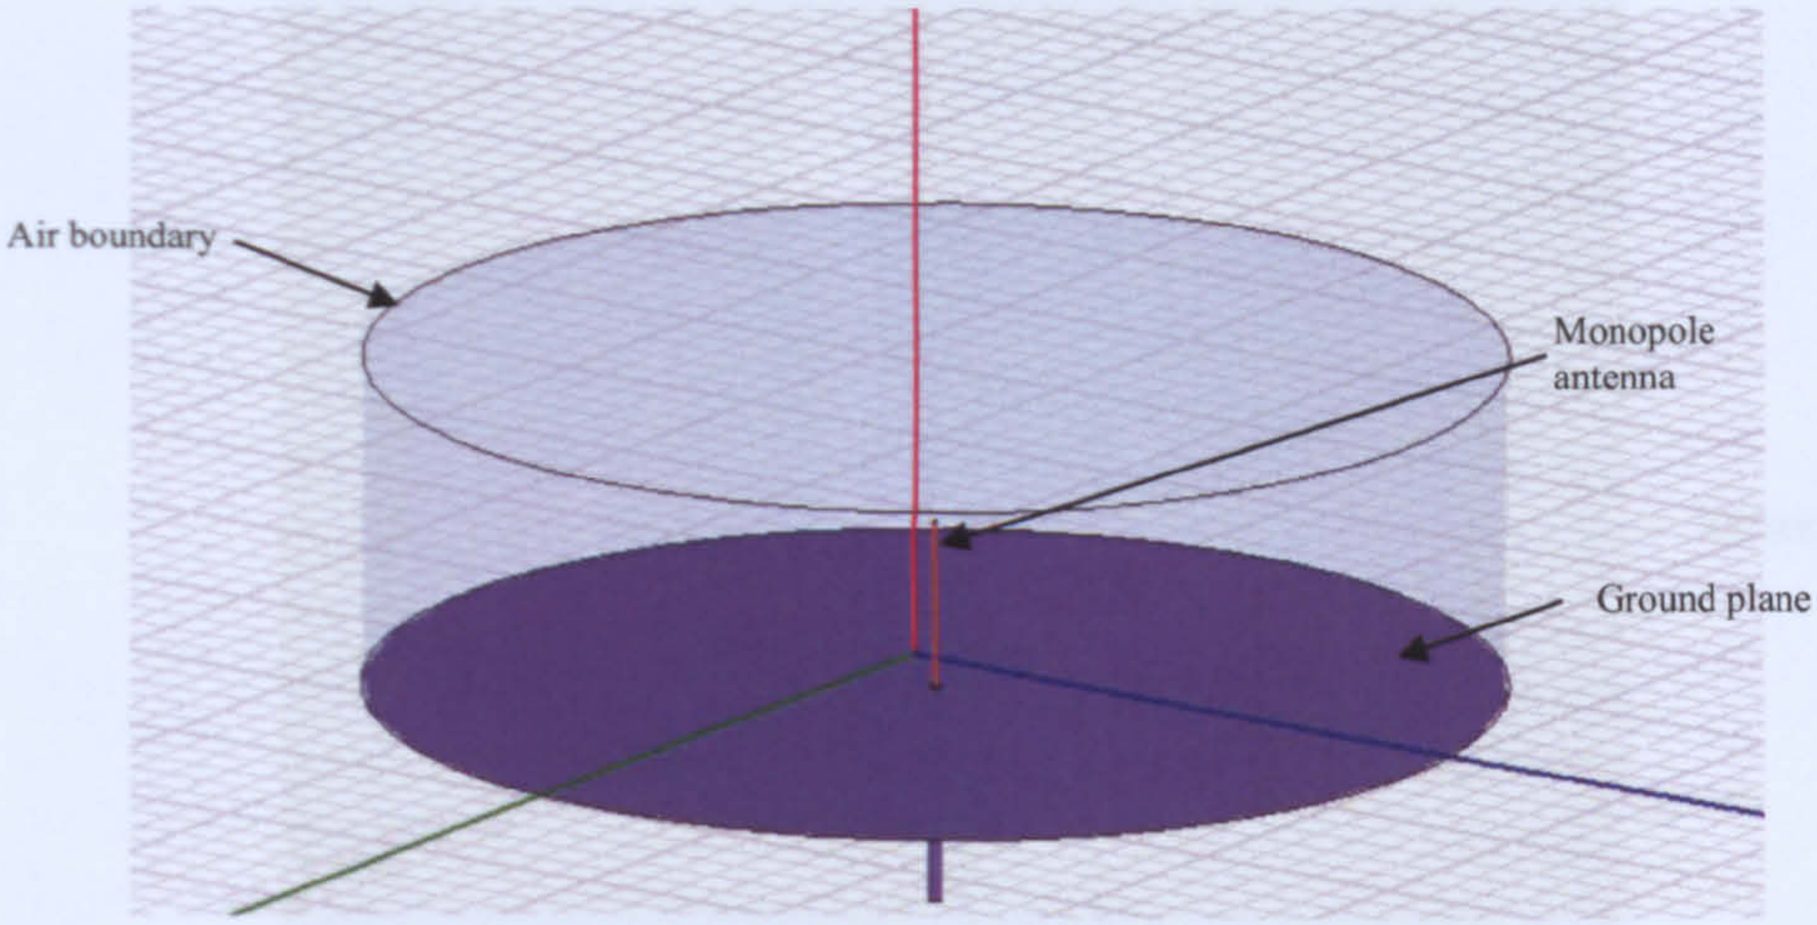

E.2. Directivity Results

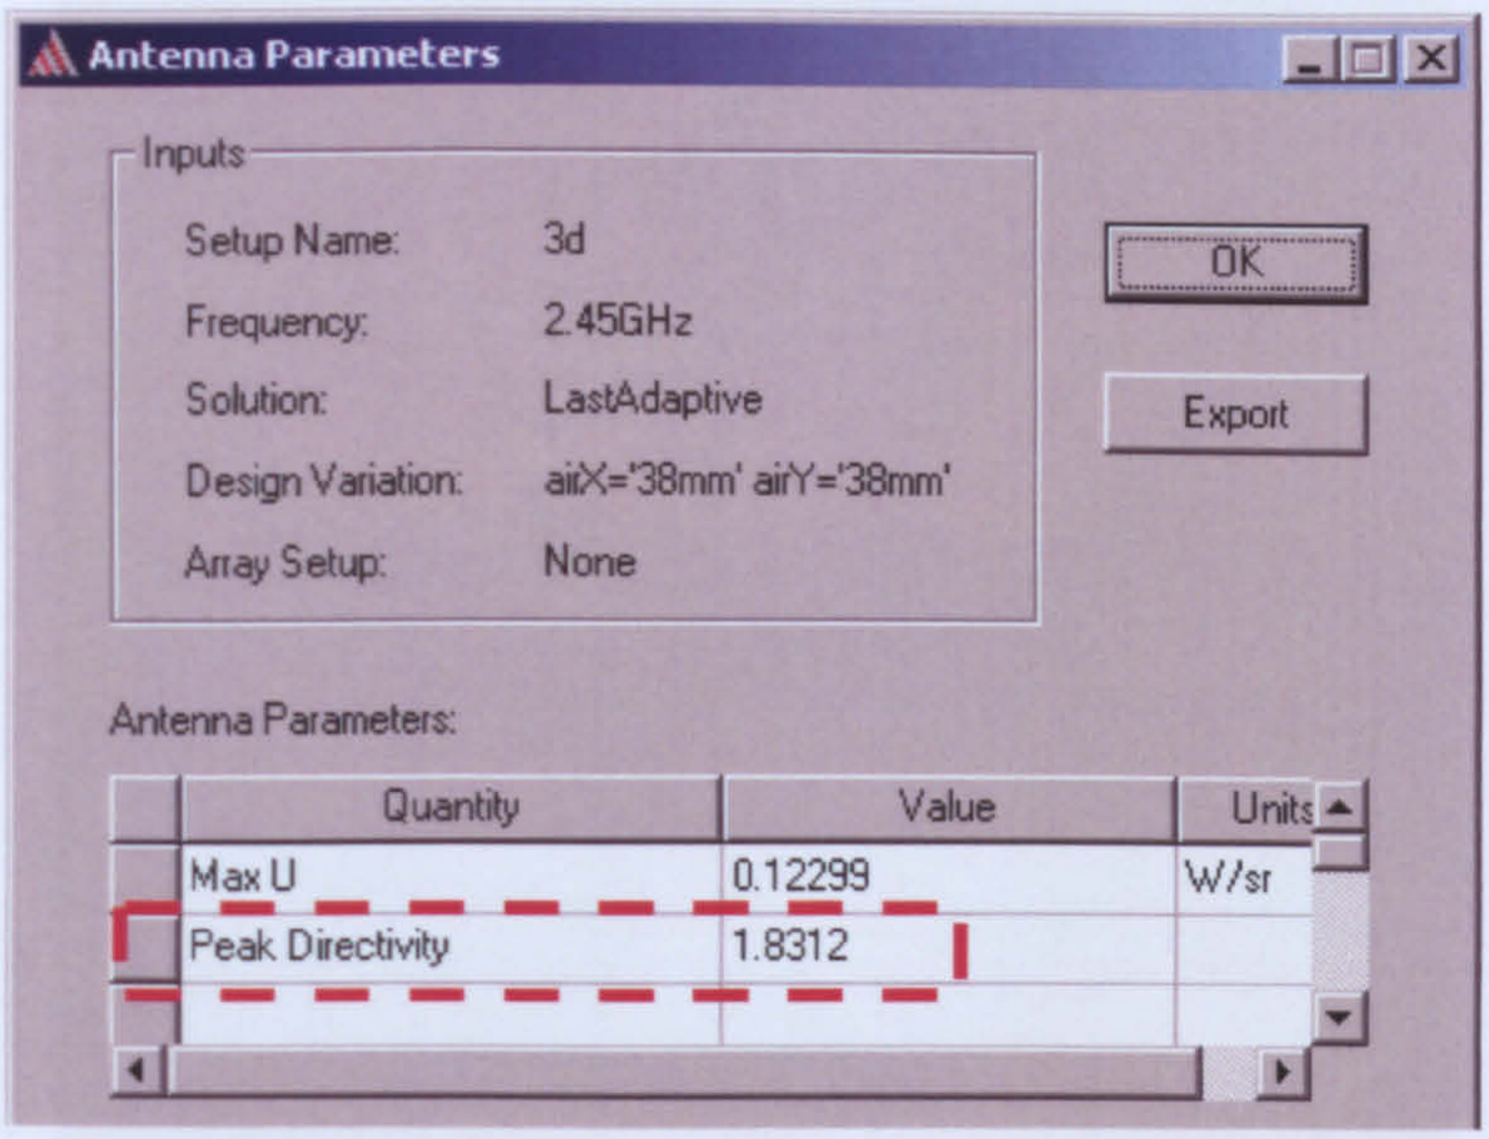

Where:

$$D(dB) = 10\log_{10} 1.8312 = 2.627dB$$

Appendix F – Material Effects on Transmission

F.1. Metal Tin (Lid Off and On)

Approximate dimensions: 220mm wide, 150mm tall, 0.25mm thick

F.2. Glass

Approximate dimensions: 220mm tall, 9mm diameter, 1.5mm thick glass

F.3. Wood

Approximate dimensions: 470mmx290mmx255mm (15mm thick)

F.4. Cardboard

Approximate dimensions: 310x240x130mm (6mm thick corrugated cardboard)

F.5. Plastic

Approximate dimensions: 400x300x200mm (1mm thick plastic)

F.6. Cloth

Approximate dimensions (of bag): 480x240x150mm (material varies in thickness due to som parts being polyester and plastic; 1-2mm thick).