Development of Condition Monitoring Robots for High Voltage Equipment

A thesis submitted to The University of Manchester for the degree of

Doctor of Philosophy

in the Faculty of

Engineering and Physical Sciences

2011

Chithambaram A. Veerappan

School of Electrical and Electronic Engineering

Contents

This thesis contains 59,379 words

This page is intentionally blank.

List of Figures

This page is intentionally blank.

List of Tables

This page is intentionally blank.

Acronyms

6DOF Six Degrees of Freedom.

ADC Analogue to Digital Converter.

AM Amplitude Modulated.

API Application Programming Interface.

ARP Address Resolution Protocol.

ASCII American Standard Code for Information Interchange.

ASP Active Server Page.

BCTC British Colombia Transmission Corporation.

BGA Ball Grid Array.

BIOS Basic Input Output System.

CARPI Conference on Applied Robotics in the Power Industry.

CCA Composite Cross Arm.

CDC Communication Class Device.

CGI Common Gateway Interface.

CMOS Complementary Metal Oxide Semiconductor.

CPU Central Processing Unit.

CRC Cyclic Redundancy Check.

cRIO Compact Real-time Input Output.

CSL Chip Support Library.

CSV Comma Separated Value.

DCE Data Communication Equipment.

DMA Dynamic Memory Access.

DSP Digital Signal Processor.

DSP/BIOS Digital Signal Processor/Basic Input Output System.

DTE Data Terminal Equipment.

DVSDK Digital Video Software Development Kit.

EFDM Electrical Field Distribution Method.

EIRP Effective Isotropic Radiated Power.

FEC Forward Error Correction.

FFC Flat-Flexible Connect.

FPGA Field Programmable Gate Array.

FSR Feedback Shift Register.

FTP File Transfer Protocol.

GFSK Gaussian Frequency Shift Keying.

GIS Gas Insulated Switchgear.

GPIO General Purpose Input/Output.

GPP General Purpose Processor.

GPRS General Packet Radio Service.

GPS Global Positioning System.

GRP Glass Reinforced Plastic.

HCI Human Computer Interface.

HV High Voltage.

I²C Inter-Integrated Circuit.

ICSP In-Circuit Serial Programmer.

IDE Integrated Development Environment.

IEEE Institute of Electrical and Electronic Engineers.

IEPE Integrated Electronic PiezoElectric.

IETF Internet Engineering Task Force.

IMU Inertial Measurement Unit.

IREQ Hydro-Quebec Research Institute.

ISDN Integrated Services Digital Network.

ISM Industrial, Scientific, Medical.

JPEG Joint Photographic Experts Group.

LLM Live-Line Maintenance.

LQI Link Quality Indication.

LSB Least Signification Bit.

MEMS Micro Electro-Mechanical System.

MIPS Mega Instructions Per Second.

MJPEG Motion JPEG.

MPEG Motion Pictures Experts Group.

MSB Most Significant Bit.

MSK Minimum Shift Keying.

OAIS Open-Air Insulated Substation.

OGIS Open-Air GIS.

OMAP Open Multimedia Application Processor.

PCB Printed Circuit Board.

PRN PseudoRandom Number.

PRNG PseudoRandom Number Generator.

RAID Redundant Array of Inexpensive Disks.

RF Radio Frequency.

RFC Request for Comments.

RFS Root File System.

RGB Red, Green, Blue.

RIV Radio Interference Voltage.

RSSI Received Signal Strength Indication.

SD Secure Digital.

SDHC Secure Digital High Capacity.

SPI Serial Peripheral Interface.

SRD Short Range Devices.

SSD Solid State Drive.

SSE Scottish and Southern Energy.

SSH Secure SHell.

SSSI Site of Special Scientific Interest.

TCP Transmission Control Protocol.

TCP/IP Transmission Control Protocol/Internet Protocol.

TDMA Time Division Multiple Access.

TFTP Trivial File Transfer Protocol.

TI Texas Instruments.

UART Universal Asynchronous Receive Transmit.

UDP User Datagram Protocol.

UGIS Underground GIS.

USART Universal Synchronous Asynchronous Receive Transmit.

USB Universal Serial Bus.

UV Ultra-Violet.

VGA Video Graphics Array.

VI Virtual Instrument.

WAP Wireless Access Point.

WLAN Wireless Local Area Network.

WPAN Wireless Personal Area Network.

XOR eXclusive OR.

Abstract

Society has an every increasing thirst for electrical energy; this is only set to increase as the 21st Century progresses. In order to sustain this increasing demand, the power industry needs to consider a number of factors; adding generation capacity and maintaining the transmission and distribution networks that connect the producers to the consumers. This work focuses on the development of systems to aid maintenance operations.

Parts of the transmission network in the UK date back to the 1950's and 60's, consisting of over 22,500 circuit km of overhead lines. The monitoring of this network is a significant ongoing task and needs to locate potential problems prior to failure. Numerous assessment techniques are presented in literature which discuss the examination of line components from the air or ground using the visual, infrared or ultra-violet spectrums.

Of particular interest in this work is the live-line inspection of composite insulators; thereby aligning with other ongoing work at The University of Manchester. While existing techniques have proved adequate to date, not all insulator surfaces can be appropriately seen. The ideal solution would be a device capable of photographing all insulator surfaces from a camera mounted on the insulator itself. While a number of live-line robotic systems are both in development and use around the world, operation and performance information is lacking; possibly due to commercial sensitivity issues.

This work aims to clarify this situation, in particular focusing on the nature of broadband communication from, and survivability of complex electronics in areas of intense electric field strength and partial discharges. These areas are explored through the development of a technology demonstrator, a robot capable of imaging composite insulator surfaces in real-time and transmitting them to a ground

station. Knowledge gained can then be adapted to create systems for other high-voltage monitoring situations.

A systems level approach is taken whereby the technology demonstrator is divided into its constituent functional components. The requirements of each are assessed and research and development needs are detailed. Literature is reviewed to collate existing knowledge and enable comparison with the envisaged requirements.

Prototype systems are developed to test the selected communication mechanism under high voltage conditions, while designs are created and fabricated for imaging and mechanical needs. The separate systems are then combined into the technology demonstrator and examined as a single unit under energised conditions.

The author presents extensive results on the capability and nature of broadband radio frequency communication from areas of high electric field strength and partial discharges. They show that high data rates from such environments is possible up to a certain point at which high enhanced shield and antenna protection needs to be considered. They additionally demonstrate the transmission of live video from an energised composite insulator. This knowledge can be used to both improve the current system and as a basis to create additional monitoring solutions for high voltage situations. As such a new method of electric field distribution assessment is proposed.

Lay Abstract

Society has an every increasing thirst for electrical energy; this is only set to increase as the 21st Century progresses. In order to sustain this increasing demand, the power industry needs to consider a number of factors; adding generation capacity and maintaining the transmission and distribution networks that connect the producers to the consumers. This work focuses on the development of systems to aid maintenance operations.

Parts of the transmission network in the UK date back to the 1950's and 60's, consisting of over 22,500 circuit km of overhead lines. The monitoring of this network is a significant ongoing task and needs to locate potential problems prior to failure. A number of methods are in use around the world using visible light as well as infra-red and ultra-violet light imaging techniques. Assessment is usually conducted from the air using helicopters or from the ground with the aid of zoom lenses.

Increase energy demand combined with limited transmission capacity means that deactivating a line for maintenance purposes is becoming impractical. This therefore means that a safe method of detailed assessment of transmission line infrastructure is required. While a number of commercial systems which examine the transmission line itself exist, the same is not true for the assessment of the insulators that separate the line from the steel tower.

This work documents the development of a robot capable of examining insulator surfaces while the line is energised, through the wireless transfer of images to operators at a ground station. This presents a number of challenges, such as the transmission of large volumes of digital data and the protection of complex electronics in high voltage conditions.

It is concluded that high bandwidth communication is possible in such situations

up to a certain limit at which point enhanced shielding needs to be considered. This knowledge can therefore be used to both improve the existing system and create additional monitor/assessment solutions for other areas of the high voltage network.

Declaration

I hereby declare that no portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

This page is intentionally blank.

Copyright Statement

The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the "Copyright") and s/he has given The University of Manchester certain rights to use such Copyright, including for administrative purposes.

Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made only in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of any such copies made.

The ownership of certain Copyright, patents, designs, trade marks and other intellectual property (the "Intellectual Property") and any reproductions of copyright works in the thesis, for example graphs and tables ("Reproductions"), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.

Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy (see http://documents.manchester.ac.uk/DocuInfo.aspx?DocID=487), in any relevant Thesis restriction declarations deposited in the University Library, The University Library's regulations (see http://www.manchester.ac.uk/library/aboutus/regulations) and in The University's policy on Presentation of Theses.

This page is intentionally blank.

Acknowledgements

The author would like to acknowledge the help and support of a number of people without whom this Ph.D would not have been possible.

The author would like to thank Mr Peter R. Green and Professor Simon M. Rowland, the author's supervisors for their support, encouragement and advise during the research and writing up process.

The author would like to thank their parents who have also supported and encouraged them throughout this research programme.

Thank you to Chris Kirham and John Bramwell from the school's mechanical workshop for refining the robot's design and its construction.

Thank you to Frank Hogan, Simon Hayes, Graham Bruce and Dr Antonious Tzimas, Christos Zachariades for their help during experimentation in the high voltage laboratory. The author would like to thank Dr Vidyadhar Peesapati for his assistance and expertise in modeling the author's equipment in field modelling software.

An additional thank you to Keith Williams for his assistance showing the author how to appropriately use the school's spectrum analyser.

Finally a thank you to everyone who the author has not mentioned by name but have assisted the author during this Ph.D.

This page is intentionally blank.

The Author

Chithambaram A. Veerappan has studied at The University of Manchester since September 2002 when it was The University of Manchester Institue of Science and Technology (UMIST). He undertook an undergraduate degree in Computing and Communication Systems Engineering with Industrial Experience. He graduated in 2007 with a First Class Honors degree.

His team's forth year project was selected as the university's entry in a competition organised by The Royal Society of Engineering, the project subsequently was named the winner in the category of Best Integrated System. He was awarded The Siemens Excellence Award 2007 for his contribution to the university.

Chithambaram has undertaken industrial expeience at 3Com Europe Limited as a software engieneer working on thier remote network configuration software and at BT Group plc as a software engineer working with 'Web 2.0' techologies. He has undertaken periods of work at The University of Manchester as an electronic engineer exploring and developing teaching tools for the undergraduate Microcontroller Engineering module. He has additionally worked as a teaching assistant working with the Protel PCB Design environment and Microchip Assembly and C Laguages.

This page is intentionally blank.

Chapter 1

Introduction

The population of the world has an ever increasing thirst for electric power. Combined with the ageing infrastructure of the electricity transmission and distribution networks in many western countries, complex planning regulations and high installation costs associated with new routes, it is becoming more important to fully utilise existing infrastructure. However as utilisation increases the network as a whole becomes less tolerant to failures as less capacity exists for re-routing.

One aspect to maintaining energy security is to ensure the network's infrastructure components are properly maintained and fixings[1] are replaced prior to failure. Figure 1.1 shows the basic components of a typical electricity transmission line. In order to appropriately schedule maintenance operations the utility operator needs to be aware of potential problems; this is possible through regular condition monitoring of the all parts of the network. This is a vast undertaking for any utility company and any device which could increase assessment speed while reducing costs would be beneficial.

At the same time, little is known about the survivability of complex electronics and the possibility of broadband wireless communication from areas of intense electric field strength. As such this research aims to answer the following questions:

- Is it possible to establish and maintain a reliable high bandwidth communication link from a device located in an area of high electric field strength and partial discharges to a device located some distance away with relatively low

---

[1]Fixings include components such as spacers, clamps, dampeners.
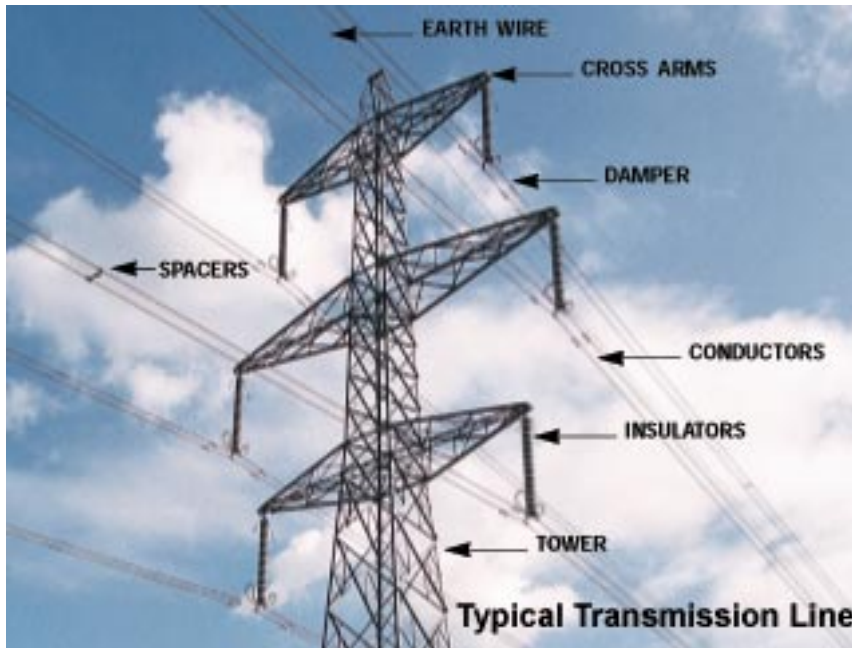
Figure 1.1: Basic components of an electricity transmission line[1]

electric field strength?

- How complex electronic circuits that include processors, transceivers and digital camera modules cope in intensive field strength environments with minimal shielding?

Furthermore, while in research environments components under test can be removed and examined in detail 'off-line', the same does not hold true for the real networks were service interruptions must be kept to a minimum. This means that in-service components have to be regularly examined for signs of fatigue or imminent failure. As discussed in the literature review in Chapter 3, there are a number of options for this in the event that unenergised inspection is not possible[2].

Therefore the development of a robot incorporating the above knowledge and assessment needs will form a technology demonstrator by which the research outcomes can be discussed and knowledge disseminated to the wider community. It will also form a platform which, through refinement, can be used in industry for

---

[2]From what can be ascertained, live-line working by company operatives is not undertaken by the transmission operator in the United Kingdom.

maintenance operations, as an asset management tool and in research for on-line investigation of new high voltage concepts.

There are a number of other robotic systems mentioned in literature and in use in the field. Systems mentioned are either remotely operated or controlled via a optical umbilical cord; a number also provide real-time video to operators and experts. However, there are a number of issues with the systems noted in literature as will be discussed in Chapter 3: the majority appear to be in the prototype/theoretical stage, with little or no empirical evidence to as their real-world operability. Of those systems which have been experimentally tested, little or no information is provided as to wireless communication methodology or environmental protection.

An opportunity arose to attend and present at the inaugural Conference on Applied Robotics in the Power Industry (CARPI), in Montreal in October 2010. Here a number of robotic systems both in development and in use were discussed. The most interesting of these were LineROVer and LineScout[2, 3] developed by the Hydro-Quebec Research Institute (IREQ) and the 'AApe-C1' insulator monitoring robot developed by Shenyang Institute of Automation, Chinese Academy of Sciences (SIACAS)[4]. However the issue with the former systems is that they operate on either the ground wire or the conductors themselves, an environment different to that around an insulator. The latter operates specifically on porcelain tension insulators and does not appear to transmit live data. Due the commercial sensitivity of all systems mentioned, very little information is available on the mechanism by which they operate, with particular reference to the transmission of real-time video and protection of the electronics.

1.1 The Development Scenario

A single high voltage condition monitoring scenario was selected, with the standpoint being that any technology developed could then be re-worked for other monitoring scenarios and requirements. The scenario selected was based on other work being undertaken at The National Grid High Voltage Research Centre at The University of Manchester and the interests of researchers at the facility. One area of interest is research into the ageing and performance characteristics of composite insulators for use on transmission lines.

Therefore one of the primary goals was to develop a system to assist with research on the ageing of composite insulators through real-time close-range optical analysis. A simultaneous goal was to provide the National Grid with a mechanism for monitoring insulators on their network, thereby allowing them to discover potential problems with a higher degree of accuracy and with sufficient time to rectify them. Therefore this work aims to initially develop a robot capable of being inserted onto an energised composite insulator possibly via hot-stick, which then can provide an expert visual information on the insulator's surface condition.

Specifically the robot will need be inserted at the tower side of a composite insulator and operate on a 400 kV transmission line, providing sufficient communication bandwidth for the transfer of real-time video. The device will need to be able to capture a combination of video and photographs of each surface of a composite insulator and be able to move along an insulator string when commanded to do so. This data will then need to be transmitted to a operator at ground levelor stored on-board for later analysis. The robot will need to support itself and cannot damage the insulator's mechanical or electrical integrity while doing so. Based on a review of live-line working techniques by Looms[5], it is apparent that the weather can have a major bearing on whether it is considered safe to operate live-line. As such, the system developed should only expect to be operated under ideal conditions, i.e. dry surfaces, with a weather outlook for calm conditions.

1.2 Thesis Structure

The following is a quick summary of the contents of each chapter in this thesis.

Chapter 2: Requirements Analysis and Specification, converts the project's vague goals into a set of defined requirements that could be used to design and build a robot for one situation to enable the technology to be showcased.

Chapter 3: Literature Review details exiting work into radio frequency interference from high voltage lines, high voltage inspection systems and live-line inspection techniques. It discusses the various failings of existing material and the reasoning behind the activities undertaken within this research.

Chapter 4: Robot System Level Overview takes the requirements and charts the

envisaged robotic system. It discusses the planned spliting of the robot into an number of subsystems and how they were to be interconnected.

Chapter 5: Preliminary Communication System Development details the system built to assess the quality of the 2.4 GHz license free frequency band and the testing undertaken to ensure the systems stability and quality prior to energised testing.

Chapter 6: Preliminary High Voltage Testing discusses the test scenarios used and associated results for detailed assessment under high voltage conditions

Chapter 7: Vision System Platforms details the different approaches that could have been taken to construct a vision system. It discusses the system that was selected for development and platforms that had potential but whose development complexity was too high for the time available.

Chapter 8: Communication System Development details the development and testing of the high data rate communication system and evaluates and considers the effects of external and ambient radio energy in the test environment.

Chapter 9: Vision System Development details the implementation of a video streaming and capture system and the steps taken to improve image quality.

Chapter 10: Mechanical System Development discusses the various mechanisms required by the robot and a number of possibles devised to achieve them. It then discusses the development of the selected option.

Chapter 11: System wide Communication and Control charts the system wide communication requirements and how they were achieved. It subsequently discusses the various control strategies developed for the control of the robot.

Chapter 12: This chapter presents a case study into future needs that could be fulfilled by robotic investigators. The study focuses on the Composite Cross Arm technology that is being developed at The University of Manchester and how the technology developed within this research could be applied.

Chapter 13: Conclusions and Future Work summarises the results of the systems created and what work needs to be carried out before they can be carried forward into the commercial arena. It additionally proposes a new system that could be created for the monitoring of insulators based on discoveries made during the research.

1.3 Future Possibilities

The knowledge obtained through the design and construction of a robot capable of transmitting real-time video while operating in an electrically harsh environment has multiple future prospects. If shown to be viable, the technology can be applied to other areas of the power industry; examples include the monitoring of other transmission line hardware, in particular the cables themselves.

Furthermore, the core research questions are not directly tied to the development of this robot, rather it is to discover if high bandwidth communication is possible, therefore video and photographic data are just two potential high data rate sources. Likewise, the video camera module and its associated processing components are also just two examples of complex electronic components which are susceptible to electrostatic discharge and damage from high electric field strength environments.

Therefore, if it can be proved that these systems work in these harsh conditions, other monitoring sensors could also be used in similar situations. For example, the visible light camera could be replaced with modules sensitive to either infra-red or ultra violet radiation; the former allowing heating to be identified while the latter could identify partial discharge and organic growth[6].

Another possible extension to the system would be the collection of surface samples for later analysis. This is something that cannot be achieved by other means of energised inspection, bar possibly live-line working. This facility would be advantageous for researchers who would be able to ascertain what biological growth and pollution is building up on composite insulators regularly over a period of time. Transmission utilities on the other hand would be able to assess the rate of insulation performance degradation (if any) in relation to contaminant build-up.

While initially movement will be semi-automatic with the operator instructing the robot to perform each action, this could then be extended to allow each shed or entire insulator to be visually inspected automatically with the issue of a single instruction. It is important to note that in the context of this research, the term 'automatic' refers to the control mechanism of the robot itself rather than any sort of image processing based insulator surface categorisation algorithm. However, should it be proved that the robot can survive in its intended operating environment, an image processing algorithm could be produced to undertake this

task; this could be produced to function either on the robot itself or on a separate system.

1.4 Contribution to the Subject

This research has investigated the use of the 2.4 GHz license free band for the transmission of data from areas of high field strength and partial discharges. Evidence shows that in the majority of situations these parameters do not effect reliable transmission.

Comparison between the simulated electric field strength distribution across a representative suspension insulator and the gathered data suggests a link between position on the insulator and the achievable data rate when using Wireless Local Area Network (WLAN)[7] technology. This has then lead to a proposal for a potential new tool for the assessment of electric field distribution based on achievable data rates.

Omnidirectional, cylindrical waveguide transition and Yagi antennas have been used during this research. It has been discovered that omnidirectional antennas have limited use in areas of high field strengths and partial discharges. Cylindrical waveguide transition were shown to be more resilient at the detriment of beam angle.

The use of an omnidirectional-to-omnidirectional antenna link under energised conditions was shown to be unreliable when high data rates were required. The use of a Yagi antenna at the base-station substantially improved the stability of the link. Further assessment revealed that the Yagi antenna's directivity was of more benefit than its gain.

The knowledge gathered during this research has wider applicability than just the technology demonstrator that has been constructed. The findings presented can be used in the creation of other monitoring devices, not limited robots, which need to operate wire free in areas in close proximity to or in direct contact with high voltage equipment.

This page is intentionally blank.

Chapter 2

Requirements Analysis and Specification

The key requirement was the creation of a robot capable of being attached onto a live-line composite insulator via hot-stick. The robot would then need to capture high quality images and real-time video of the surfaces of each insulator shed prior to moving to the next. These images were to be transmitted to a technical expert on the ground a safe distance away. After this process, the robot was required to return to its starting point for retrieval via hot-stick.

The envisaged inspection system is to initially operate as master-slave system, as illustrated by Figure 2.1. The operator is to issue commands via a Human Computer Interface (HCI) consisting of either a laptop computer or custom hardware solution. These commands would be sent to the robotic investigator which would then carry out the requested operation, returning a status code indicating the success or failure the request. Image and video data would be update and streamed respectively on request.

A fail-safe mechanism was also required, which would allow the robot to detach itself from the insulator string and drop to the ground. It was not categorically stated if the robot was required to survive this scenario; therefore it was decided that a 'fall' would not have to be endured. One reason for this fail-safe requirement was to avoid the need to de-energise a transmission line in order to retrieve a failed robot.
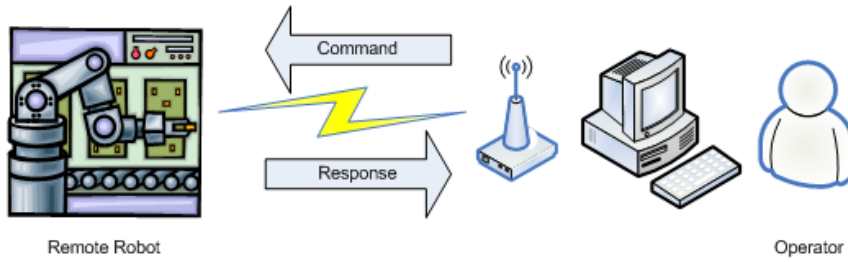
Figure 2.1: Intended interaction between operator and robotic investigator

In addition there were a number of implementations that had to be avoided.

- The creepage length[1] of the insulator must not be degraded by the robot to an extent that the insulator no longer provides sufficient electrical clearance and therefore cause flash-over.

- A wireless transmission system must be used. Fibre optics are not sufficient even if the required creepage length is exceeded due to potential pollutants on the outer sheath reducing this creepage length.[2]

- The robot needs to be predominantly constructed of non-conductive materials to reduced the likelihood of an electrical equipotential forming along the length of an insulator string and thereby reducing the creepage length and flash over safety margin. However, it is envisaged that a Faraday Cage of some description would be required to protect any on-board electronic systems.

Barring the few safety requirements and implementation 'no-goes' listed above, further specifications were not provided; as a result this research could explore a wide range of solutions. Based on the key requirement and with the safety concerns in mind, a number of questions needed to be answered in order to define a useful technology demonstrator. These are detailed next:

---

[1]This is the distance that leakage current flows along from one end of an insulator to the other. It is the shortest surface path between the two end of an insulator.

[2]Even though this stance has been taken within this research, it should be noted that other researchers have developed boom truck applications which utilise optical fibres for the transmission of video streams captured in close proximity to high voltage equipment and the manipulation of robotic manipulators[8, 9].

- What wireless communication method should be chosen: unguided optical or radio frequency?

- Both surfaces of the insulator need to be examined, should this be done with a single camera or dual cameras? What overhead will this decision have on processing and transmission requirements? What effect does this decision have on the time it takes to examine each insulator shed?

- What quality video is required? Is the same quality required all the time or just when interesting phenomena appear on the insulator surface? Is video actually necessary, would moderate to high resolution photographs be sufficient for day-to-day insulator monitoring? Is only visible light photography and video needed? What about infra-red or ultra-violet photography and video?

- What is the maximum weight the robot can be? What is the best shape to avoid over stressing the local electric field strength?

- What is the best way to record images and/or video of the surface of a shed?

These questions are answered in the following sections by examining possible solutions and the likelihood of their successful & robust implementation. However before doing so access limitations should be considered as this will potentially drive certain aspects of the design. Early work by Looms in 1981 [5] notes that there are two ways in which live-line work can be undertaken, 'hot-stick' and 'bare-hand'. The former is where the operative remains at earth potential and works on live elements through the used of insulated equipment, while in the latter situation the linesman is brought to the line potential and highly insulated from earth.

Looms further indicates that there are two types of safety distance: the 'minimum working length of surface' and the 'minimum approach distance'. The former is associated with 'hot-stick' access approach, sometimes refered to as a 'distance' technique and indicates the minimum length of the insulating tool used. The latter is applicable in 'bare-hand' operation and provides the minimum distance "between a linesman at line potential and (any) object at earth potential" [5]. Figures provided suggest that 'minimum working length of surface' for system (phase-to-phase) voltages of up to 400 kV is greater than 3.2 m while the 'minimum approach distance' is approximately 2.75 m.

More recent information is included within Safety Rules issued by the National Grid[10], which suggest the safety distance from exposed high voltage conductors at 400 kV is 3.1 m, with an additional clearance of 300 mm required between a worker an any part of the supporting insulator which is outside the safety distance. This document additionally suggests that this safety distance applies from the end of any tool or object that is being used by the linesman, however it is not specified if this only applies to non-insulates tools and objects.

This then suggests that even if linesmen are able to access the steel cross-arm, they or any implement they use may not be able to directly contact the supporting insulator. This then has a bearing on the initially proposed concept of using a hot-stick for mounting the device, particularly for UK operations.

2.1 Communication Mechanisms

From the perspective of the communication system, the video stream and photographs are forms of digital data. Therefore, the task can be seen as the need to transfer large volumes of data with low latency and high accuracy; a common requirement for modern communication systems. No special consideration was taken for the fact that data being transferred is visual in nature. Initial thoughts revealed two possible mechanisms for a communication system which could provide enough bandwidth for live video streaming; unguided (freespace) optical and radio frequency.

The basis of a unguided optical link is similar to that of a fibre-optic system in that light is used as the transfer mechanism, but without the fibre to guide the light. Guided systems are used throughout the world to deliver cable television and Internet services and therefore an optical link, in principle, would have sufficient bandwidth. However, the challenge would be to make the jump from a guided to unguided solution without detrimentally degrading performance.

The key to creating a reliable wire-free optical link is maintaining a clear line-of-sight between the transmitter and receiver. This is an almost trivial task when both devices are stationary. However, if one or both parties can move, the issue of maintaining alignment arises; requiring the development of a tracking system

for either the primary receiver or transmitter[3]. Such a tracking mechanism would probably best be suited to the ground-based device to avoid adding weight and complexity to the robot being developed. Although achievable, the tracker itself could be considered a complex undertaking, both mechanically and electronically and might therefore detract from the main intention of developing a robot to examine composite insulators.

The use of radio frequency communication has seen dramatic growth over the last decade, with similarly expansive growth in technologies and devices which exploit its capabilities. This is most prevalent in the expanding mobile broadband and Internet capable mobile phones markets, with speeds similar to the maximum achievable by fixed line broadband services only a few years ago. However, the most relevant technology for this situation are the ubiquitous and license free wireless network standards ratified by the Institute of Electrical and Electronic Engineers (IEEE).

The most recent standard, IEEE802.11n, is able to achieve theoretical maximum data speeds of 300 Mbits/s, while the more common IEEE802.11g[7] can theoretically achieve 54 Mbits/s. Considering Internet video streaming services, such as BBC iPlayer, can deliver reasonable quality video over a 2 Mbits/s broadband connection, the video bandwidth required here would be significantly less than the theoretical maximums on offer.

Implementation of this technology on the robot is not without its own concerns though. Radio frequency communication transceivers have not regularly been deployed in such close proximity to energised transmission lines, in either industrial or research contexts, as confirmed by Chapter 3. This situation then requires the frequencies of interest, the 2.4 GHz license free Industrial, Scientific, Medical (ISM) band[4], to be thoroughly examined for potential interference from line related phenomena.

It is clear that both systems possess sufficient data bandwidth and are theoretically realisable; however neither is without their own issues. The unguided optical

---

[3]This link needs to be bi-directional in order to allow commands to be sent to the robot, as such both it and the ground-based device are in fact transceivers. However, as the robot is the main data source it will be termed the primary transmitter and the ground-based control and viewing unit will be termed the primary receiver.

[4]The 2.4 GHz ISM band begins at 2.4000 GHz and extends to 2.4835 GHz.

solution is fraught with complexities; primarily maintaining a line-of-sight 'lock' between the transmitter and receiver when one party is attached to a suspended insulator, some tens of meters in the air.

Meanwhile, the use of radio frequency communication shows more promise; for example only minimal modifications are required, in particular the protection of the antenna which cannot be shielded. There is no need to adapt existing technology for the change in transmission media as required for an optical link. Furthermore, research presented Chapter 3 suggests that interference in the frequencies of interest may not actually exist to detrimental levels.

Based on the complexities of maintaining constant line-of-sight and the associated possibility of detracting away from the main research objective by trying to solve the issue, it was decided that radio frequency communications be researched first. If it was shown to be neither realisable nor resilient the optical solution would be investigated.

2.2 Sensing and Optical Systems

In order to decide on an appropriate condition monitoring mechanism, it is important to understand the ways in which composite insulators can degrade. Degradation can be in the form of a loss of hydrophobicity, electrical erosion and intentional or accidental damage. This degradation can lead to an increased risk of flash-over or mechanical failure[11].

Specific degradation can occur in the form of surface crazing[5], light and heavy erosion, tracking[6], chalking[7], punctures, spliting or cutting of the composite material and corrosion to the metal end fittings. Physical vandalism and bird bites can additionally affect the performance of a composite insulator. All these degradation mechanisms change the physical appearance of an composite insulator string and hence can be capture/assessed through visible light imaging techniques.[13, 12]

Corona activity can lead to a loss in hydrophobicity in composite insulators. This can subsequently lead to dry band arcing which will cause localised heating of the

---

[5]Micro-fractures on the surface[12]

[6]The development of conductive paths on the surface of the insulating material[12]

[7]Rough or powdery surface due to exposure of the filler particles from the housing materials[12]

insulator surface. The arcing itself can lead to an increased loss in hydrophobicity and erosion of the surface; this can eventually lead to either electrical or mechanical failure. This surface heating can be detected through infra-red imaging and can provide a qualitative view on the current level of dry band arcing[14].

As corona activity is known to lead to the loss of surface hydrophobicity[14], it is useful to ascertain the current level of corona activity on insulator surfaces. Corona discharges while sometimes visible at the blue end of the visible-light spectrum, it additionally produces light components in Ultra-Violet (UV) spectrum. Therefore, use of UV sensitive imaging equipment will allow the detection and assessment of the level of surface corona activity.

It is important to note that while visual techniques have been highlighted in detail above, they are not the only possibilities. Other techniques include acoustical[14, 15] and an electric field distribution method[14, 16, 17]. However as one of the requirements was a visual conditional monitoring system, such a system is the major focus of this work.

As such, the vision system had the largest number of questions that need to be resolved(in comparison to either the communication or mechanical systems). No indication was provided as to what was meant by 'high quality' and this was difficult to quantify. For example, 320x240 pixel video from a distance of a few centimetres may (neglecting use of optical zoom lenses) actually provide a higher quality than 800x600 pixel video from a few hundred meters, especially on the basis of pixels per square centimetre of insulator surface. As pixel information can be removed more easily than it can accurately estimated (digital interpolation), it is more appropriate to use a larger imaging sensor than required and either scale or crop the image.

The question of whether full quality is required all the time is linked to minimum quality expected and aimed to reduce the load the compressed video data places on communication bandwidth available. If not required permanently, then there is less emphasis on the need to create a communication link which exceeds peak load[8]. This is somewhat mitigated by the choice of imaging sensor as it dictates the highest resolution available. Given the use of the highest resolution imaging components within the system, a firmware/software solution can alter the

---

[8]The bandwidth needs to exceed to peak load to account for communication protocols.

rate/resolution at which videos/photographs are captured.

Another software/firmware issue is the case of whether video itself is actually required. A video is essentially a sequence of photographs captured and displayed at a rate which allows movement to be conveyed. A composite insulator string is a solid object and hence the question arises of what added benefit video would bring to any assessments that were to be made as there are no 'moving parts'; any damage should also be visible on photographs. However, as the switch between video and photographs can be made in software/firmware this can be changed as the need arises.

As discussed in Chapter 3, intra-red and ultra-violet photography can play a part in assessing insulators for damage and earlier signs of failure. However, as mentioned in Section 2.1, video streams and photographs are just forms of data. Therefore, if it is shown that it is possible to transfer visible light images from the robot to the ground based operative, the camera element can be easily replaced by one which is sensitive to either infra-red or ultra-violet radiation.

2.3 Mechanical Considerations

In order to appreciate the complexity of the mechanical design requirements, knowledge of the form of a composite insulator is needed. It is important to note that there are multiple shapes of composite insulator, designed for different situations and to alter the localised electric field in different ways. As such it is not feasible to attempt to design a single solution to operate on all cases. Therefore a typical insulator was chosen as the design target.

Each insulator shed is approximately 15 cm in diameter and is fixed to a central core of approximately 2.5 cm diameter. The lower surface is flat and is almost perpendicular to the central shaft. The upper surface consists of a increasing gradient from the circumference towards the central shaft. The shed spacing is approximately 4 cm at the circumference, reducing to about 1 cm at the central shaft.

Insulator strings with sheds spaced in this manner can also vary in length depending on the system voltage. One end is attached to the tower while the other

is attached to the energised line; the latter can sometimes be referred to as the 'hot-end'. Strings can be either mounted vertically (on suspension towers) or horizontally (in tension towers). They can be found individually or in a set of two or four; again based on system voltage[9].

This multitude of mounting orientations and styles again suggests that it is not feasible to design a 'one-solution-fits-all' device. As such, a single insulator situation will be targeted. This also has two possible configurations as noted above; vertical and horizontal. The former can, in part, rely on gravity for maintaining alignment between the insulator and robot. The latter, however, may need additional supports to ensure that gravity does not distort this alignment. Horizontal insulator configurations are used by deviation towers which are used to change the direction of a line, shown in Figure 2.2. Vertical insulator configurations are used on suspension towers which support the line, shown in Figure 2.3. Based on this, a vertically orientated insulator will initially be targeted.
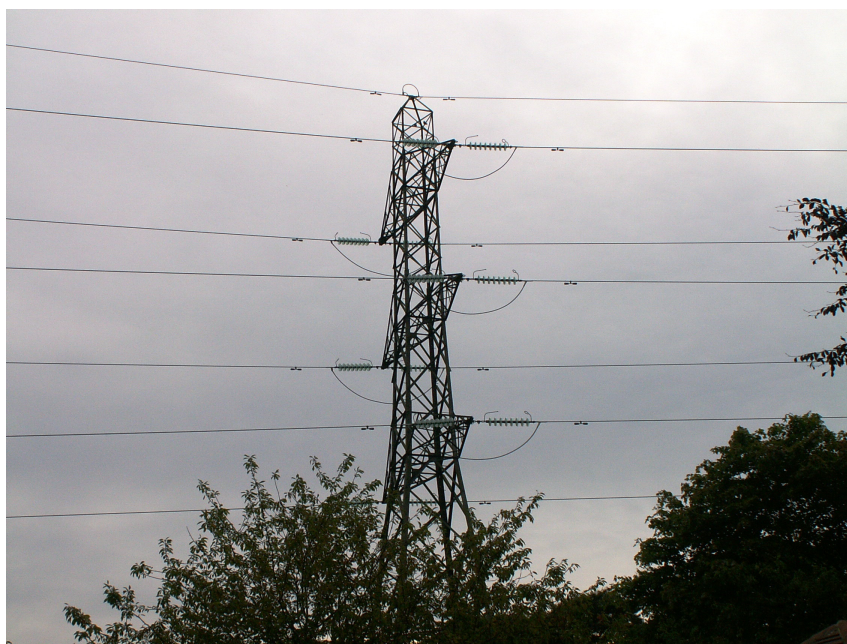
Figure 2.2: Deviation Tower - Horizontal Insulators

In a similar manner to the existence of many insulator configurations, the National Grid also has many types of transmission tower structure within it's network. Common tower types used at 275 kV and 400 kV include the L2, L6 and L12

---

[9]This is not an exhaustive set of circumstances.

Figure 2.3: Suspension Tower - Vertical Insulators

designs[18, 19]. Dimensional information of these tower types can be found in Table 2.1 and can be referenced to Figure 2.4.

Code	Description	L2	L6	L12
a	Tower Height	33.26	41.51	39.22
b	Bottom Phase Height (typical)	12	12	12
c	Middle Phase Height	19.85	21.00	20.70
d	Top Phase Height	27.62	31.54	30.00
e	Bottom Phase Width	6.10	8.44	7.12
f	Middle Phase Width	5.72	10.45	9.12
g	Top Phase Width	5.49	6.98	6.30
h	Ground Clearance (typical)	12	12	12

Table 2.1: Dimension of common 400 kV (suspension) transmission towers used in by the National Grid[19]

Other mechanical questions are more general and would apply regardless of the targeted situation. Possibly the most important question regards the maximum weight of the robot. There is actually not a definitive answer to this barring 'as light as possible'. The mounting point between the tower and the insulator would already have to support the insulator itself, at least half the weight of any
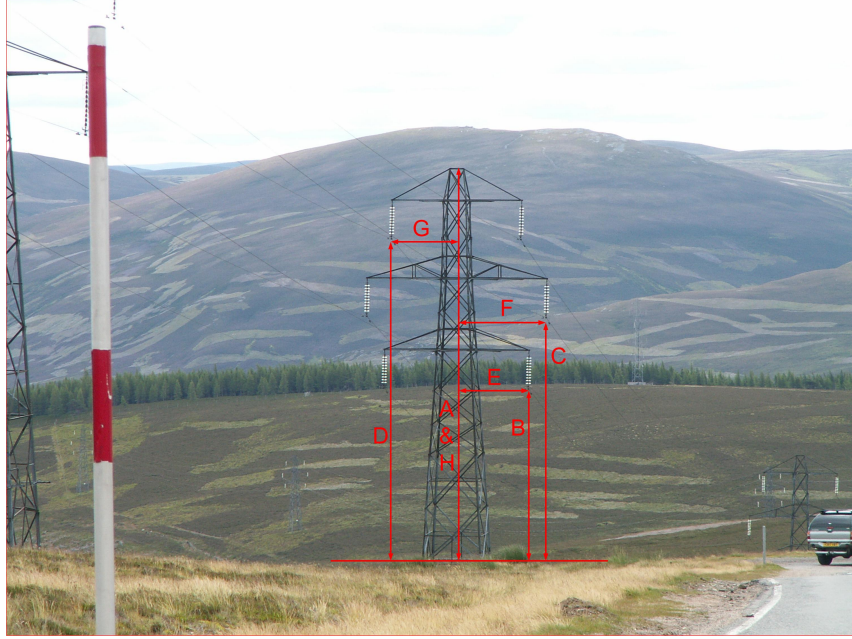
Figure 2.4: Definition of dimension in Table 2.1

attached line, plus a margin for factors such as ice loading and an additional safety margin. However, a bigger concern would be to avoid damaging the surface; this being dependent on the 'clamping' mechanism in addition to the robot's weight. Initially a maximum weight of 5 kg was suggested.

As noted at the start of this section, insulator strings are also designed to provide a certain creepage length between the energised line and the earthed tower so that the breakdown strength of the medium surrounding the insulator, in this case air, is not exceeded. Therefore it important to firstly avoid the use of metals wherever possible in the construction of the robot, coupled with minimising its physical height.

Furthermore, it is important to avoid any sharp edges or points as these features can enhance the local field strength and lead to partial discharge, which is know to cause radio frequency interference[20]. Based on this, the best shape for the robot would be cylindrical, with rounded edged at the ends of the cylinder.

Finally, the mechanism of capturing surface images needs considering. Due to the spacing of the insulator sheds, a single angled view of the surface of each shed is not sufficient and not actually possible because of the central column. Furthermore, by

imaging from an angle, features towards the central column will be captured in a lower relative resolution compared to foreground features. Additionally, those features towards the central column will also contain some level of optical distortion. A more detailed description of the problem is presented in Section 10.1.1.

Therefore, a better solution would be to capture images from directly above the surface. Dependent on the implementation, this should provide images which are undistorted and of similar relative resolutions. However this solution requires the camera to scan and track over the entire surface of each shed and subsequently individual images need to be digitally stitched together if complete photographs are required.

2.4 Cross Subsystem Considerations

A number of points have to be considered across the range of subsystems and are strongly interrelated. Both surfaces of each shed needs to be examined. This can either be done with a single camera which rotates through 180° to enable both surfaces to be examined or two cameras which capture simultaneously.

Based on the description of a typical insulator string in the previous section, it is clear that the portion of the camera assembly which 'scans' the surface needs to be less that 1 cm in height. While a dual camera arrangement can be built to meet this requirement, a single camera system is simpler to prototype; especially when lens focal length requirements are additionally considered.

Furthermore, a dual camera system would require twice the processing power for image capture & compression and twice the communication bandwidth for reliable video streaming compared to a single camera solution. Additionally, the receiver hardware would also need to cope with this increased data stream; again needing twice the processing power to decode and display the video streams.

However a dual camera system does have one benefit: a reduction in the time needed to examine each insulator when compared to a single camera solution. Although, as this system's main purpose is to test the concept, it is debatable whether examination time is actually of major concern.

It therefore seems apparent that a single camera solution appears to have more

benefits over a dual camera system, with the major benefit of the latter not really being relevant in this situation. Therefore it was decided that a single camera system should be initially developed; with the belief that a single camera system should still provide enough scope to prove the concept of real-time live-line monitoring.

2.5 Power System

So far the power requirements and methodology for the proposed system has not been considered. While a dedicated system for power provisioning is outside the scope of this research, it is important to consider one during development, if only to provide a route by which such a system could be integrated in the future.

Ideally, the implemented power system, either via power scavenging or integrated battery cells, should have the ability to power the device for one working day or at least for the assessment of one insulator string. The former is prefered solution, which would allow utilities to under take one full day's assessment. In this situation it would be possible to have integrated battery cells which could be re-charged overnight.

Alternatively, the minimum solution would require the device to assess one complete insulator before battery capacity diminishes[10]. This minimum requirement is essential to prevent the device failing due to power problems part way thorough an assessment sequence. In the event of failure in this manner recovery becomes problematic and may involve de-energising the line: something that the proposed systems actually aims to avoid. In such as system, the batteries would need to be quickly replaceable so that they may be changed easily during the course of a working day.

2.6 Summary

This chapter has considered the basic requirements that the robot needs to fulfil, through which a development strategy has been outlined. Two communication

---

[10]There should additionally be a safety margin to the required battery capacity.

mechanisms have been detailed of which radio frequency communication will initially be explored. The mechanical and imaging systems have been shown to be interrelated to some degree and are additionally related to system complexity and the monitoring time required. The ambiguity of the term 'high quality video' has been debated with the conclusion that, if designed correctly, video can be added through software if required.

Due to the exploratory nature of this work a single camera system has been chosen for implementation, thereby simplifying fabrication and reducing development costs. A single composite insulator design has also been selected and detailed for this phase. Knowledge gathered can then be applied to integrate other sensors into the existing system or aid in the development of tools for other high voltage monitoring situations.

Chapter 3

Literature Review

As detailed in previous chapters the aim of this research was to design and build a robot capable of examining the surfaces of composite high voltage insulators under live line conditions by means of real-time video transmission. This task presents many interesting and distinct challenges including high bandwidth communication, video capture and robust, minimalist & protrusion-free mechanical design. This work intended to use industry standard video and photographic compression systems, such as Motion Pictures Experts Group (MPEG)[21] and Joint Photographic Experts Group (JPEG)[22], to reduce the required bandwidth prior to transmission or storage; as such alternative compression techniques have not explored.

Over the years numerous researchers have investigated various phenomena resulting from high voltage assets; in particular substations and high voltage transmission lines. During assessment of existing material it became apparent that research falls into two fields of study; 1) engineering phenomena and 2) health effects of high voltage equipment, e.g. pylons. The latter field of research is irrelevant to this work and was not considered further.

This chapter starts by examining current research into wireless communications in high voltage environments and acknowledges the lack of quantitative data in the area of the transmission frequencies of interest. Past and present insulation monitoring techniques & methods are subsequently investigated and the benefits of the proposed system is reinforce.

A number of live line robots do exist around the world and a number are currently

in development, some of which were presented at CARPI 2010. A discussion of similarities, differences and research omissions will be presented, therefore re-iterating the need for this research. Finally international standards of interest are noted.

This literature review splits discovered works into seven areas of interest;

1. Radio Interference as a Concern
2. Radio Interference as a Benefit
3. Non Specific Wireless & Power System Co-location Studies
4. HVDC Line Radiometric Phenomena
5. Insulator Assessment Techniques
6. Existing Robotic & Comparable Systems
7. International Standards

3.1 Radio Interference as a Concern

Research in to radio interference from high-voltage lines at times was found to consider the effects on the broadcast frequencies, such as those used for television and radio. Other researchers have considered the topic without a particular focus on a particular communication or transmission system type; for example Van Atta & White[23], Adams[24] Mather & Bailey[25] and Pakala & Chartier[26].

Van Atta and White summarise some of the problems faced with radio interference from insulators in their 1930 paper, detailing the cause of interference to be corona and brush discharge from/between points on insulator disks. Suggesting that discharges are caused by a combination of poor insulator design, damage and pollution, they discuss and evaluate various solutions offered by manufacturers and utilities[23]. It is also noted during discussion of the paper that the insulator alone may not be causing all the interference, with the conductor size and system voltage also playing a part.

In 1934 Langton and Bradshaw conducted laboratory and field experiments looking at pin-type, cap-and-pin-type and interlink insulators under various voltages and environmental conditions. While their conclusions align well with later researchers observations, they too only look at a small range of radio frequencies: 230 to

300 kHz and 600 to 1000 kHz, however they do note that other frequencies could be affected. In essence they conclude that small air gaps that exist either between the line and mounting point on the insulator or between the insulating material and its external fixing are the cause of radio interference and by 'filling' these gaps in most cases reduces observed interference[27].

Adams in 1958 looked at the effects that the design of a line itself has on radio interference. He suggests that through differences in design, two systems at the same voltage can have appreciable differences in radio interference. Results were obtained through calculation and showed the effect that different line configurations could have. It is shown that lateral distance, phase spacing and phase layout all affect the 'relative interference level'; as the former two increase, the relative interference experienced decreases. It is also suggested that the spacing between conductors in a bundle can also affect interference, with neither a very large nor very small spacing being preferential. Adams also notes that in some circumstances conductor spacing can lead to a decrease in interference at long distances but simultaneously increase it close to the line. It is further suggested that the majority of interference is due to the centre conductor[24].

Reichman and Leslie[28] looked at the the radio interference from two test lines at Coldwater, Ontario, Canada. What sets these researchers out from others is their use of long term tests, whereby measurements were recorded from autumn 1959 to summer 1960. Their conclusions agree with other researcher in that positive corona has a larger bearing on radio interference production, additionally noting the effect the weather and the proximity of the line to the tower structure has. However efforts were mainly concentrated at lateral distances of 100 feet and at frequencies of <2 MHz.

In their 1963 paper entitled 'Radio Interference from High Voltage Lines', Mather & Bailey[25] investigated the radiometric effects of an energised transmission line at multiple voltage levels and with various conductor diameters, thereby examining the theories suggested by Van Atta & White in 1930[23]. Experiments were conducted on both six-span isolated test sections and on normal lines. Readings were taken at varying distances from the operational line but results were converted to a common slope distance for comparison purposes.

Like Reichman and Leslie before them, Mather & Bailey made use of long du-

ration experiments. When data collected over a three year period was analysed, the results indicated that various factors affect the level of radio noise seen. It was revealed that over the assessment period the average voltage on the line had increased by approximately 5%[1] and manifested itself in the results as an increase in the noise level.

Atmospheric conditions also affected the level of radio noise experienced; the paper indicates that noise levels increase in later summer and autumn; rain additionally affects values. It is noted that the type of conductor used is another factor; for example a double 'Drake' conductor bundle was shown to create less noise than a single conductor of 'Chukar' type[25]. However, these findings could be refuted by research in 1994 by Chatier et al[29] which suggested that the addition of a conductor to the former scenario could have lead to the decrease in noise.

However, while Mather & Bailey showed that interference exists at various distances from the line and the factors that influence the level of interference, they do not state outright the frequencies at which this interference is strongest. This situation is partially resolved by Pakala & Chartier[26] who investigated radio noise at frequencies up to 10 GHz and at line voltages of up to ~800 kV.

In 1969 Warburton, Laio and Hoglund[30] investigated the effects of noise from power lines on television reception, indicating that interference above 15 MHz is due to sparks or micro-sparks. This paper does not deal with the actual noise-frequency profile at average TV transmitter/receiver separations, instead making reference to certain frequencies and the perceived interference at those frequencies. It should be noted that the authors have primarily focused on the causes of and mitigation techniques for radio frequency noise, suggesting that historically poor line design and currently ageing, are the causes of sparks and micro-sparks; the former in part concurring with Adams[24].

In 1971 Pakala & Chartier[26] used an 'artificial gap-type radio noise generator' to produce broadband noise that was subsequently measured. This could therefore be considered similar to the air-gaps which Langton & Bradshaw concluded were the cause of interference[27], while the device can be considered to produce micro-sparks similar to those investigated earlier by Warburton, Laio and Hoglund[30]. The pair state that this "artificial gap was used to increase the radio noise level

---

[1]More specifically between 1959 and 1960/1.

from a line so that the field strength could be measured at relatively large lateral distances and at the higher frequencies with the instrumentation available"[26]. Therefore it follows that the presented results should relate to the worst case scenario; as this increase noise level would not be present without the artificial gap device.

Of particular relevance are the results of experimentation in the 1-10 GHz range. Results indicated that the noise level in dB peak above 1 $V/m/MHzBW$ falls as the frequency increases; either as a single linear relationship, or multiple relationships depending on the frequency band[26]. The authors additionally indicate that the noise-frequency profile of conducted and radiated interference are different, a fact that is confirmed by international standards that will be discussed later.

Pakala & Chartier's results indicate at a frequency of approximately 2.4 GHz, the radio noise is between 40 and 60 dB peak above 1 $V/m/MHzBW$ for various configurations of high voltage line, line voltage and distance from the line. It is additionally shown that as the lateral distance from the line increases, the radio noise decreases. However, it shown that this attenuation is less pronounced as the frequency increases; changing from approximately 0.013 dB/m in the 25-50 MHz range to approximately 0.0085 dB/m in the 2.5-10 GHz range[26]. While these results provide a useful basis there were two missing elements; the noise level at very close proximity to the line was not provided and the possible effects to reliable transmission are not discussed[2].

Moreau and Gary[31] in 1972 investigated if it was possible to predetermine the radio interference level produced by high voltage transmission lines. They relate the interference level to a number of factors, some of which can be easily determined while others cannot, such as climate and conductor surface state. They note the influence of the latter factors are diminished under certain weather conditions, namely heavy rain, quantified as >1 mm/hour. They go on to state that the worst interference is normally experienced under such conditions and therefore it is this level that should be used to characterise the line. The effect of weather has previously been empirically assessed by Mather & Bailey in 1963[25].

They make use of a parameter termed the 'excitation function' which they claim is independent of the conductors' capacitance and is therefore a better "measure

---

[2]Though this was not the focus of the paper.

of the cause of interference". This paper mainly focuses of the variations of noise with respect to voltage gradient, conductor radius and number of conductors in the bundle under heavy rain. It is then shown how these 'reference' charts can be modified to give an indication of noise levels under other weather conditions[31].

In 1975, Maruvada & Trinh, while at IREQ looked at radio interference with the intention of finding a basis by which limits could be set for future transmission lines. They succinctly summarise the cause of radio interference to be corona discharges on the line and associated fixings, further detailing that the pulsing nature of the discharge leads to wideband noise. Although they focus their efforts on interference to Amplitude Modulated (AM) radio transmissions, their observations on the effects of weather, distance and transmission frequency appear to be similar to observations made by other researchers discussed in this chapter. Importantly they note that the radio interference cannot be treated in isolation and has to be considered with received signal strength. Finally, as some data from which these observations were drawn were sourced from various utility companies and a local AM radio broadcaster, details on measurement equipment and techniques were not available[32].

In 1994 Chartier et. al. investigated the effect of bundle orientation on produced audible and radio noise levels, additionally examining the effects of conductor arrangement. Primarily looking at the future use of four conductor arrangements, they compared trapezoidal 'Hood' conductors to threaded 'Bunting' conductors. It was revealed that neither the orientation, vertical or horizontal for double, 'vee' or inverted 'vee' for triple and square or diamond for quad conductor, nor the conductor type had an effect on radio or audible noise. The one exception to this was the single conductor configuration under 'aged' testing, where the 'Hood' and 'Bunting' conductors produced different amounts of audible noise when stressed to over 250 kV[29]. This then refines Mather & Bailey's earlier suggestions on the relationship between conductor type and interference[25].

A further study by Chartier et. al.[33] in 1996 on the radio effects of T2 conductors on single, twin, triple and quad formations revealed that conductor type in most cases does not effect the level of noise produced at 990 kHz. In this investigation Partridge T2 conductors were tested against Dove conductors with an equivalent current capacity. The 'T' indicates that the conductor's formation is of twisted nature and the '2' indicates that two elements are used in the twisting process.

Chartier et. al. additionally examined the stand-alone Partridge conductor (i.e. without twist-paring); results suggest that under dry conditions at line to ground voltages of over 300 kV the noise level of a quad conductor Partridge line is up to 14 dB higher than a quad conductor Dove line. Noise increases are also seen in other scenarios for Partridge conductors but are less pronounced. Results for Partridge T2 and Dove conductors show very close correlation throughout. However, as this 'stand-alone' Partridge conductor appears to be non-standard, associated noise results can be ignored.

Greater interest was placed on publications that related to the use of Global Positioning System (GPS) receivers for fault monitoring on power transmission systems. This was for the reason that the GPS receiver and associated antenna would need to be located on power transmission equipment. Silva & Olsen[34] initially discuss the operating frequencies of the GPS system: Link 1 at 1575.42 MHz and Link 2 at 1227.60 MHz, each with a 20.46 MHz bandwidth. It was stated that at frequencies above 300 MHz, the majority of noise is generated internally in the receiving system; the noise generated externally is termed the 'antenna noise temperature', the portion of which is related to electromagnetic interference is of importance.

Silva & Olsen suggest that power lines can cause two type of interference to GPS signals; 1) when noise from the line is greater than the 'equivalent thermal noise', and 2) when the electromagnetic scattering is significant, in which case the received signal may become reduced. They continue to state, in common with other researchers, that there are two mechanisms of interference from power lines; conductor corona and gap or spark discharges. They additionally note the effects of weather and altitude on the level of observable noise[34].

Silva & Olsen experimented under a number of double circuit twin conductor 345 kV transmission lines; the purpose being to ascertain the effect of the conductors on the signal received. They noted no practical change in received carrier to noise ratio as the path of the line was traced. As foul weather can lead to higher levels of corona the same experiment was conducted in poor weather; again showing no appreciable degradation to receiver operation. As the noise could not be directly measured with their equipment, assessment was made by monitoring the number of satellites that the GPS unit 'locked' onto. It was indicated that the unit maintained a lock on a large number of satellites even in poor weather

thereby showing that near conductor corona or scattering does not affect GPS reception.[34].

However there are a number of issues relating to the work of Silva & Olsen, some of which are noted by the researchers themselves[34]. Firstly, measurements were not taken within the lattice steel towers. Additionally measurements were undertaken at near ground level, which is some distance from the location of both conductor corona and gap discharges.

While many authors have undertaken experimental approaches, Nayak and Thomas discuss simulated results and suggest that interference levels reduce as conductor diameter is increased, distance from the line is increased and the height of the line is raised; all within certain limits[35]. The suggestions on conductor diameter concur with Van Atta & White[26] and those researchers who have investigated multiple conductor types[25, 29].

Huang & Ruan[36] also look at audio and radio performance in relation to a proposed 1000 kV transmission line. Using formula from The Westinghouse Electric Corporation and the Bonneville Power Administration they assess three line configurations at varying heights. They show that both the phase configuration and height above sea-level effects the noise level experienced. Radio noise was assessed at a reference frequency of 0.5 MHz.

Huang, Ruan & Huo extended their work in 2009 by investigating the effect of bundle spacing[37]. This work focused solely on double circuit lines and compared two Chinese 1000 kV tower configuration with one from Japan. Calculated results mirror the authors' previous findings, with the additional conclusion that bundle spacing also effects radio interference. A key point being, neither a very small or very large spacing of sub-conductors is inherently better, with the authors' calculations suggesting that a spacing of approximately 0.3 m yields the best results. A similar sub-conductor spacing observation was made by Adams in 1958[24].

3.2 Radio Interference as a Benefit

Partial discharges are another cause of electromagnetic interference; there are many active researchers in this field and work stretching back over a decade exists. This

domain of work is based on the fact that partial discharges can be used as an indicator for the degradation of electrical insulation. Thus, by being able to accurately detect and locate such events, it is be possible to prevent electrical equipment from failing catastrophically.

Early work by Hikita et al[38] investigated the electromagnetic spectrum produced in air and SF_6 gas by partial discharges. This team's goal was to separate the radio noise of partial discharges in air (and background noise) from those occurring within Gas Insulated Switchgear (GIS). This research used a sphere-plate electrode to create partial discharge and a bi-conical antenna to detect radiation at a distance of 1 m. Their paper presented results in both SF_6 and air; it summated that greater radiation was experienced at lower frequencies, specifically in the range from 30 to 150 MHz when assessed in SF_6. When comparing average positive charge to the gain of the radio emission results suggested that the frequency band from 35 to 70 MHz was the most sensitive to an increase in charge for both insulation media.

This work was supplemented a year later and showed the distance characteristics of partial discharges on the electromagnetic spectrum. In this investigation, the previously used sphere-plane electrode geometry was replaced with a needle-plane geometry. A similar bi-conical antenna was used to detect emissions, while the distance was varied between two and eight meters. When the background noise spectrum was compared to the partial discharge spectrum, it was seen that they were nearly identical for frequencies above 60 MHz. For frequencies in the band 30 to 50 MHz the partial discharge spectrum was up to 10 dB greater than background levels. It was additionally seen that the distance from the partial discharge source affects the produced frequency components in different manners; higher frequencies showed 'pockets' of greater interference after initial decreases[39].

Hakita et al subsequently applied similar principles to three types of substation: an Open-Air Insulated Substation (OAIS), an Open-Air GIS (OGIS) and an Underground GIS (UGIS), with the results for OAIS being most relevant to operating environment of interest. Similar equipment to previous investigations were used, while the distance to the partial discharge source was not explicitly specified. Measurements were only made at frequencies up to 300 MHz. Results indicate that (in OAIS) the entire range suffers from interference; levels of interference were shown to decreases at higher frequencies[40].

These experimental results were in part confirmed by Babnik et al[41]. This team proposed detection and identification methods based on both time and frequency domain measurements. Results related to air discharges indicated that such events cause energy in frequencies below 150 MHz with a spectrum trace indicating large amounts of energy at approximately 50 and 110 MHz. Results additionally indicate a brief energy spike at approximately 220 MHz.

Moore, Portugues & Glover[20] are a group of particularly active in the field of partial discharge detection. They have shown that by using four antennas in various arrangements and calculating the time difference between the waves arriving at each antenna, the location of the partial discharge can be discovered. Measurements made by the team suggest that the majority of radio energy caused by partial discharge activity is focused below 1 GHz; they furthermore showed that stronger dielectrics showed greater interference than weaker ones.

Simulations by Xiao et al[42] for one scenario suggested that radiation was most noticeable between 1.7 GHz and 3 GHz; being potentially disastrous for communication in the 2.4 GHz IMS band. Testing under experimental conditions revealed this to be incorrect, additionally confirming previous knowledge that partial discharge activity creates most interference at frequencies below 1 GHz; in particular noting that most energy was focused below 200 MHz.

Both the simulation and laboratory results were captured by custom diskcone antennas at three locations within an sealed 'tank'; each at the same distance from a spark source at the centre. Both results indicated that the majority of the radio energy occurred from the direction of the spark gap, the x-axis in presented experiments[42]. While the paper's aim was the isolate the best location at which faults could be located via radiometric means, it also provides information on zones that should be avoided for reliable wireless communication; as in the case of the robot being developed.

3.3 Non Specific Wireless & Power System Co-location Studies

While previous researchers either took the standpoint that interference was a concern or a benefit, others looked at whether wireless communication was actually possible. This section looks at work in this domain.

In 2002 Kezunovic et al [43] produced a report for the Power Systems Engineering Research Centre which discussed wireless and system-wide spread-spectrum communication in substation environments. The authors noted that signal level depended more on the location of the substation than on the power delivery of that substation. They noted that while research has been undertaken into the effect of switching transients on various frequencies, effects due to substation layout have never been considered.

Kezunovic et al note that various factors effect interference in wireless channels and the fact that the channel utilisation is continually changing requires regular assessments to be made. They add that the type of channel, interference or bandwidth limit has a direct baring on the modulation scheme and device settings chosen[43].

Kezunovic et al made field measurements for both 900 MHz and 2.4 GHz systems; the former being of minimal interest to the current work. They noted the importance of assessing the modulation quality rather than spectrum analysis or single frequency measurement. The latter, while accounting for areas of diminished signal, cannot account for the spreading in frequency applied by spread spectrum modulation schemes. Using commercially available 2.4 GHz Wireless Local Area Network (WLAN)[7] devices for the transmitter and receiver, the team was able to gather information on packet error rates for the case when the receiver was placed in the substation environment over a period of 14 days[43].

It is believed that given the information gathered by Kezunovic et al for the 2.4 GHz frequency band, an adequate analysis of the spectrum (in terms of detrimental effect to signal and modulation quality) has not be provided. During analysis, there were a few references to the 2.4 GHz test data of notable concern; the correlation between the transformer load and the assessed radio quality. The results show moderate inverse correlation in the early part of the day. Kezunovic

et al, even for the 900 MHz data, appeared to be more concerned with the correlation to environmental factors and channel utilisation by other frequency band users than the effect of the substation itself.

Liang et al[44] focus mainly on a video transmission system, specifically the time duration of the capture, compression, transmission, decompression and display of images. Their interest is in the consistency of video quality and they reasoned that parallel processing of the video transmission stages combined with fixed duration stages aid the issue. While they make use of a WLAN system, they have not addressed radio interference or antenna protection issues nor does the system appear to have been tested under live-line conditions.

Between 2007 and 2008 Shan et al[45, 46] assessed the use of wireless technologies operating in the 2.4 GHz license free band in electricity substations. Technologies investigated were WLAN and Wireless Personal Area Network (WPAN)[47], with emphasis on the effects of partial discharges and impulse noise respectively. ZigBee (based on IEEE802.15.4[48]) was the WPAN technology under investigation. Due to similarities with the proposed technologies to be used within this research, Shan et al's work could potentially be used as a basis for further investigations. However a substation environment is different from the environment surrounding insulator string, therefore results will not be directly applicable. Although it could be stated that a substation environment is far more challenging than on an insulator string, thus representing a worst case scenario.

Shan et al initially discuss the sources of partial discharges within transmission substations before discussing the basic transmission methodology of ZigBee[45] and WLAN[46]. The authors used field trials for the former while both laboratory tests and field trials were used to investigate WLAN interference.

Shan et al's ZigBee transmission system could transmit at up to 250 kbit/s with a maximum transmit power of 3 dBm. Data transfer experiments were conducted with raw and error control coded data. At the point the article was written, field trials were under way. 'Proof-of-system' results for raw data with a length of 114 bytes revealed that byte positions below 50 and above 70 suffered a greater error rate than the centre 20 bytes. Using stated error figures, the average byte error rate can be calculated as $3.59 \times 10^{-3}\%$ [45].

There are however some severe shortcomings with this research work. The authors

did not specify the ZigBee channel used for the laboratory 'proof-of-system' experiment. Further the authors state that the outcome of field tests would be presented at a conference, however, the name of the conference was omitted, thus the crucial results from the field trials could not be located.

The groups' next paper on the effect of impulsive noise on WLAN performance appears to have more substance. The laboratory tests included both noise free and controlled impulse noise experiments. In each case the transmission medium was replaced by a microwave coaxial cable and microwave attenuators. A 'power combiner' was used to combine the signal and the interference in the latter experiment. The actual payload length was 2304 bytes, the maximum allowed by the transmitter used. Both 11 Mbps and 54 Mbps transmission rates were assessed[46].

Using a 10 ns 0.316 V impulse, frame error rates were recorded as the impulse repetition frequency was increased. When the transmission rate was 11 Mbps, the frame error rate reduced as the repetition frequency increased to 10,000 impulses a second and climbed thereafter. With 54 Mbps transmission, frame error rates increased as the repetition frequency approached 10,000 impulses a second and fell beyond this. The average frame error rate was 4.43×10^{-5} and 2.52×10^{-3} for air-data-rates of 11 Mbps and 54 Mbps respectively[46]. As with the previous paper on ZigBee[45], the channel used for experimentation was not specified. Field tests were mentioned as 'being planned' but again subsequent work incorporating the results could not be located.

The data that this group's research aimed to produce would have provided an ideal basis for performing similar analysis for pylon mounted insulators. Given the dates that the papers were published, reasoning leads to the assumption that both works should be complete; however further evidence of this could not be located.

3.4 HVDC Line Radiometric Phenomena

Many researchers have looked at the effects of ultra, extra and normal HVDC transmission lines in relation to radio noise and electromagnetic interference[49, 50, 51].

Bailey in 1970 performed a number of tests on DC test lines in association with

the Bonneville Power Association[52]. This study tested combinations of single and double conductored sections at various pole spacings, mainly at +/- 400 kV under various weather conditions. A number of interesting points were concluded. It was revealed that radio noise was approximately 10 dB lower for double conductor configurations compared to single conductor configurations. In clear weather, reducing the pole spacing also reduced noise levels; reductions were also seen under rainy conditions. A 10% increase in line-to-ground voltage translated to a 4 dB increase in noise. Bailey additionally mentioned that for AC situations, the generally accepted figure for a 10% increase in line-to-line voltage is 6 dB.

3.5   Insulator Assessment Techniques

While the use of radio communication in the high voltage environment around an insulator string of transmission pylons has not been considered in great detail, assessment techniques of such insulators have been analysed in the previous decade or so. This research was two-fold; firstly, in determining a faster mechanism of defect detection in classical insulators and secondly into damage & defect identification of composite insulators. Despite the varying nature of each researcher's objective, the techniques used to achieve those objectives are similar.

In the case of the classical porcelain cap and pin type insulators, Lang, Allen and Zhou[15] considered various remote detection solutions for the costly and time-consuming traditional manual test methods used at the time. They detail a number of possible remote diagnostic solutions, of which thermal imaging and ultra-sonic discharge location had been in use for a number of decades. However it was noted that while these techniques were used in air-borne surveys of line hardware, their application for detecting defects on individual insulator elements represented a significant challenge.

Later research placed emphasis on the detection of pollution and defects on composite insulators. Wallström et al[6] investigated the effects of micro-organisms on the hydrophobic surface of the insulator and the subsequent change in the electrical properties of the insulator as a whole. While Giorgio et al[17] investigated the effect of insulation defects with a view to stimulate a wider debate and to invoke the introduction of international Live-Line Maintenance (LLM) standards

(e.g. procedures and approach distances) for transmission lines using composite insulators.

Lang, Allen & Zhou[15] determined that remote assessment techniques for installed insulation on energised lines fall under three categories: electrical, acoustic and visual. It was noted that electrical methods are most commonly relied upon; however they need a hands on approach. It is this hands-on approach that concerned Giorgio et al who note that an Electrical Field Distribution Method (EFDM) provides a quantitative view into dielectric strength while visual methods provide a qualitative view of damage[17].

Giorgio et al suggest that without a quantitative view it cannot be determined whether it is safe to carry out LLM in particular, carry out the suggested EFDM. However, it should be noted that this stance appears to be a flaw in their discussion; in order to do LLM a form of EFDM needs to be preformed, however in order to perform EFDM, LLM needs to be possible[17].

Despite this stance, work carried out by Giorgio et al[17] supports earlier work by Lang, Allen & Zhou[15], Jaensch, Hoffmann & Markees[16] and Bologna, Mahatho & Hoch[53].

According to Lang, Allen & Zhou infra-red imaging trials can be dated back to 1971 when thousands of kilometres of circuit were examined with the discovery of 515 defects. The authors additionally note that while the technique of hot-spot detection is well established, its use in the monitoring of insulator strings is more complex[15]. This is supported in part by Jaensch, Hoffmann & Markees[16], whose work focuses on the monitoring of the transmission line without particular reference to the insulators. They suggest that thermography should only be used for special cases and not regular inspection; although an everyday alternative is not explicitly stated. In reality this has proven not to be the case, as many utility companies use aerial infra-red techniques to detect the early signs of defects, as noted by Bologna, Mahatho & Hoch[53].

Lang, Allen & Zhou additionally indicate that not all conditions are suitable for infra-red imaging, with humidity, ambient temperature and the emissivity of the surface greatly affecting results; stating that a humid overcast day provides the best possible results. Furthermore, they note that heat 'signatures' need to be analysed in the context of adjacent insulator elements; they show that defects can

result in both lower and higher temperatures[15].

Although greatly discussed and widely used, infrared imaging is not the only technique available for the detection of defects; additional methods include eye-sight and ultra-violet measurements. Visual assessment has been noted by Jaensch, Hoffmann & Markees[16] and Wallström et al[6] although for different primary purposes. The former team suggests visual methods for defect detection, both from the ground and in the air, with the possibility of aerial checks being supplemented by photography and videography[3]. Conversely the objective of the study by Wallström et al intended to use visual methods to estimate the extent of biological growth.

A final visual based method focuses on the detection of UV radiation from corona, which has been assessed by both Bologna, Mahatho & Hoch[53] and Giorgio et al[17]. It is noted by Bologna, Mahatho & Hoch that night-vision technology is used to assess this aspect. However, due to the presence of UV radiation at frequencies below those seen in natural daylight, cameras which filter out these background (daylight) UV components can be used for daytime observation.

Whilst numerous visual techniques have so far been discussed, alternative solutions do exist; Lang, Allen & Zhou[15] briefly mention an acoustical system. In a similar stance to UV detection, this system relies upon the detection of corona sources. They indicate that a sound insulator will not experience many discharges while a defective one will. In order to overcome other noise sources within the system, they suggest the use of highly directional 'corona guns'.

However there are certain shortcomings with the development of these monitoring and analysis techniques. All the techniques discussed above revolve around the imaging/detection system being a number of meters away from the target, i.e. the insulator string/shed. It is intended for the system being developed to operate on the insulator itself and not a distance from it. Nevertheless, the techniques noted in these research papers could still be used in developing an analysis system.

One defect identification system that has not been discussed in detail relates to the use of electric field measurements. The principle behind this is that faulty insulators and equipment give rise to fluctuating and sporadically occurring electric fields[16] or an abnormal field distribution[17]. It may be possible to use such a

---

[3]The supplemental documenting allowing detailed examination at a later time.

technique as an extension to the required robot. However the presence of the robot itself will alter the distribution of the electric field, consequently invalidating results that may be obtained. This problem exists with any type of field measurement system, in that the field probe itself will alter the field it is aiming to measure.

3.6 Existing Robotic & Comparable Systems

The creation of robotic systems for use on high voltage equipment is not a new idea; over the last few decades institutions and utilities around the world have been designing, building and experimenting with remote controlled and automated robots for the monitoring and maintenance of power system equipment. Projects have the primary intention of increasing operator safety and system reliability as well as reducing costs[9].

While there are many examples of such assistive robots, either due to commercial sensitivity or research practices and procedures, literature (both commercial and academic) seems to be sparse. Research material that was located mostly seemed to lack essential details. This section looks at robotic systems that have been discovered in literature and at The 1st Conference on Applied Robotics in the Power Industry (CARPI) in October 2010.

One invention, patented in modified forms in US Patent 4,808,917[54], European Patents 125,050[55] and 314,850[56] relates to a device for monitoring various parameters of a power line conductor while the line is energised; for example conductor and ambient temperature, current and voltage. The device is mounted to a hot-stick and is either clamped into place for extended duration monitoring or used pressed against the conductor for 'one-off' measurements. The measurements are transmitted to a local concentration unit, e.g. on a transmission tower, before being sent to a central monitoring station. Suggestions are made for both radio transmission and fiber-optic connection solutions between the devices and the concentration unit.

However no data is provided on operational performance. It is noted that a single frequency in the 902-928 MHz US ISM band will be used for downlink and a second frequency for an uplink. This uplink serves to aid a Time Division Multiple Access (TDMA) scheme detailed in the later patents. It is also briefly noted that

a transmission rate of 20 kBuad would be used[54, 55, 56].

There are a number of issues with this invention; firstly, this bandwidth is not adequate to transmit live video and secondly the selected transmission band is not defined for use as ISM in the United Kingdom. Therefore, even if concrete evidence existed into the successful use of the system, it would merely provide a sound basis for the investigation of a higher transmission frequency, rather than a springboard for further development of the presented system.

Nonetheless, there are certain features of the invention that can be possibly adapted. The line mounted unit (module) is toroidal in shape and is comprised of two haves and hinged towards one side, imitating a jaw. Prior to mounting on the conductor, the module is attached to a hot-stick via a slot on the module; it is subsequently pressed against the conductor. If the hot-stick is then rotated (while still attached to the module) the 'jaw' of the module begin to close, thereby clamping onto the conductor[54, 55, 56].

For completeness, it is sensible to discuss the proposed method of power provision. The above invention suggests electromagnetic induction as a method of power scavenging with a battery backup for when the line voltage is below a certain threshold. An alternative US patent filed by Lau et al[57] presents a similar sensor device but utilises photovoltaic cells for power generation and additionally removes the complexities of the clamping mechanism for a freely supported device.

Hirose and Aoki[58] have proposed and demonstrated a design for an insulator washing robot. This is a two part device consisting of a unit which can grip an insulator shed and rotate around while cleaning it, along with a system to move the former from shed shed. The latter is achieved through the use of a supporting pole situated adjacent to the insulator being washed. This system is moved into the correct location via a custom railing system fitted to each tower. As such the system was clearly intended for use on un-energised systems.

Murakami et al[59] suggest that the Kyushu Electric Power Company have been using robots of some description for live line distribution work since 1989. One of their robots is part autonomous and part remote controlled by a ground based operator. The remote control element comprises of directing the robot to the correct position, after which autonomous routines identify the target location, calculates a movement path and subsequently performs the maintenance task.

The maintenance task documented involves the replacement of the insulators on distribution networks. Monitoring and remote control is aided by the use of two cameras. However, this paper fails to document the control and transmission mechanisms. The only experiments that were conducted were in laboratory off-line on network 'mock-ups'.

While Murakami et al developed a robot to (semi)automatically replace insulators on distribution networks, other institutes have developed robots which move along transmission lines and even automatically navigate line obstacles. Liang et al[44] have proposed such a system which is primarily automated but can be remotely with the aid of a video feed transmitted over the Internet. Specifics about the inspection systems are not provided, while the control system consists of multiple computers and Digital Signal Processor (DSP)s. The primary focus of this work appears to be obtaining consistent images over the Internet rather than the design and control of the robot itself.

Jiang et al[60] propose a similar obstacle navigating transmission line monitoring system. This system has also been designed to self navigate where possible and transmit condition and picture data to the ground up to a distance of 2 km. However, like Liang et al, they too do not mention details about interference issues surrounding high voltage transmission lines.

The systems developed by Liang et al[44] and Jiang et al[60] have only been tested in laboratory environments in unenergised conditions. While documented as functional, both neglect the potential radio frequency interference surrounding high voltage transmission lines. Only Jiang et al noted the need for investigation into the electromagnetic shielding of the robot itself.

One device which appears to be functional is 'LineROVer', developed by IREQ[61]. Originally intended as a de-icing tool, its abilities were quickly adapted for wider application. It can withstand electromagnetic interference and has been repeatedly tested on a '315 kV circuit with a 800 A load'. Its worth was also proven in practice when Hydro-Quebec needed an assessment of a damaged ground wire prior repair. LineROVer was then developed into a basic mobile platform onto which numerous sensors could be attached.

Further innovation and practical application was seen at CARPI in Montreal in October 2010. This included LineROVer's successor LineScout[2, 3, 62] and an in-

sulator crawling device (the 'AApe-C1') by The Shenyang Institute of Automation, Chinese Academy of Sciences[4].

IREQ has a long tradition of developing and operating robotic systems for high voltage monitoring and maintenance applications. These include devices for operation in underground distribution systems, dam inspection (Maski and Neptune), operation in nuclear facilities and the aforementioned transmission network devices.

While most robotic equipment has seen little or no (documented) field use, LineScout has not only been used extensively on Hydro-Quebec's own lines, but that of other utilities. Toch, Pouliot & Montambault[3] discuss LineScout's use on the British Colombia Transmission Corporation (BCTC) network and the collaborative partnership formed between the two companies.

Through this partnership BCTC gained valuable data and photographs which were not possible by other means, while Hydro-Quebec gained an insight into what utilities may need and want. This lead to the implementation of a splash-proof coating, enhanced geo-location, the addition of infra-red cameras and improved image stabilisation & iris control systems. The importance of line-of-sight for communication was also highlighted, as in one case operators had to relocate their base-station. The authors note that the success of this collaboration paves the way for LineScout to be offered to utilities around the world, pinpointing the device's proven use within standard live line practices, its transportability and robustness[3].

AApe-C1[4] was another robot identified at CARPI and noted in Chapter 1. Operating on double porcelain tension insulator strings, the device uses twin probes to measure the resistance of each shed in turn. No wireless system is documented, which suggests that this robot is fully automated and stores results for later offline processing. The system uses a DSP for control and appears to be inserted and removed from the tower side connection of the insulator. No information was provided as to whether live-line tests were conducted with this robot.

Although not discussing a particular system, an interesting paper was presented by Montambault & Pouliot in which they discuss the future of power line robotics[62]. While ground-wire operation is useful they note that operation on energised conductors provides the broadest scope. They state that new devices need to fit within

existing practices and not require new or specialist tools in order to reduce the learning curve and avoid the need for certification respectively. The authors detail the importance of designing equipment for operation on current assets without modification rather than a specifically designed line. They proceed to note the conditions that robots may be expected to operate in, the importance of developing an easily transportable device, issues surrounding maintaining line-of-sight for communication purposes and planning retrieval methods in the event of device failure.

The authors then discuss future applications of power line robotics and the benefits they may provide, but note the importance of considering the trade-off between development cost, surveying time and quality of information. They discuss future problems that researchers need to resolve, for example power budgets, image processing, sensor technologies and data volumes generated before suggesting that complementary robotic platforms will eventually work cooperatively to achieve a common goal. They conclude that at present high-value resides in systems which have the dexterity to undertake complex maintenance operations in harsh and hard-to-reach environments in the presence of unknown obstacles and damaged components[62].

3.7 International Standards

While constructive, destructive and non-specific issues of interference have been assessed, the thoughts of the standards agencies have not been considered. A number of standards do exist pertaining to the production and assessment of radio interference from power lines and high voltage equipment. Most notable is British Standard 5049: Radio Interference Characteristics of Overhead Power Lines and High Voltage Equipment: Part 1: Description of Phenomena[63]. This standard agrees and confirms numerous findings detailed in the previous studies; a selection of these conformities are detailed below.

This standard covers the frequency range 0.15 MHz to 300 MHz, suggesting that noise from power lines above this frequency are not sufficiently large to affect television signals[4]. The standard continues to discuss the causes of such interference,

---

[4]Note how primary concern focuses on broadcast bands.

stating them as conductor corona, stressed parts of insulators and imperfect or loose contacts; as discovered by previous researchers.

It is noted that at low frequencies, noise is caused by propagation along the line, while at higher frequencies (>30 MHz) noise is due to radiation. The former dominates at close proximity from the line while that latter dominates at further distances. In addition the effect of weather on the production of corona noise is documented; in particular conductor corona is worse in rain and fog, while bad contacts appear worse in fine conditions.

Other factors such as the conductor diameter, time after energising and overall line ageing effect the level of radio energy produced in varying ways. As the former increases, radiation levels increase, however as time progresses, the levels drop. With particular reference to insulators, it is stated that a clean and dry insulator will produce radiation at frequencies up to 30 MHz although at low levels. It should be noted that results revealed in this document pertain to a measurement height of two meters at distances of up to 200 m; hence is somewhat removed from the target sphere of operation.

3.8   Literature Review Summary

Literature shows that extensive work has been conducted into the radiometric effects of high voltage power lines and equipment. Research is split into negative and beneficial aspects as well as studies on specific radio technologies in specific high-voltage environments. Further literature examined current insulator examination techniques and existing robotic systems.

Research indicates that a combination of line and insulator design and quality lead to varying levels of radio interference. In particular aspects such as conductor diameter, number of conductors in a bundle and inter-bundle spacing all potentially affect levels of interference. It has been suggested that a large-diameter multi-conductor bundle at a 'medium'[5] inter-bundle separation produces the least interference.

It was first suggested that the conductor type effects interference[25], however

---

[5]It was suggested that neither a small or large spacing is best but a position in between.

the comparison on which this conclusion was made is slightly flawed and later researchers made the suggestion that this is only valid in single conductor lines[29]. There is also a common consensus that spark gaps (from defects in insulators and gaps in line fittings) and corona around edges are the primary cause of interference. However, their effect on actual communication at frequencies above those in the broadcast bands have not been thoroughly analysed.

When all radiometric research is combined, there is a common suggestion that interference from high voltage equipment seems to be predominantly limited to sub 1 GHz frequencies. However, this is at a distance from the radiation source and there is only limited work focusing at the 2.4 GHz public band. Researchers using this band within their own work simply give the suggestion that their transmission systems work with no empirical evidence to prove the fact neither do they discuss proposed/implemented data protection or correction techniques used in their systems.

Of those studies which have either considered the frequency of interest or similar area of operation, no statement nor analysis has been made of actual realisable data rates or the data correction techniques that have been used to achieve those rates[54, 55, 56].

Many insulator assessment techniques have also been identified, however the most predominant focus on visible light photography with supplementary infra-red or ultra-violet analysis for the detection of specific damage. Visible light photography allows quick assessment of surface condition while infra-red allows insulator heating to be detected and corona & organic growth being identifiable through ultra-violet analysis. Acoustic detection has also been noted, but it is possible that a robot's mechanical systems may interfere with such a detection technique.

Numerous robotic systems were seen in literature and at CARPI. Some robotic proposals reviewed have only been tested under off-line conditions, thus their feasibility in real world applications is unknown. Others such as IREQ's LiverROVer[61] and LineScout[2, 3, 62] are operational devices but minimal literature concerning detailed low level operation could be located, most likely due to commercial sensitivity. Furthermore, these two devices operate on the line rather than an insulator; therefore presenting a different set of challenges. AApe-C1[4] while operating directly on ceramic insulators, appears to have a 'collect then analyse' methodology

rather than live transmission.

In conclusion, research suggests that the transmission frequency under consideration should be relatively interference free. This is confirmed in part by the fact that GPS receivers work in the presence of corona discharges. These findings provide a good basis for further exploration of the 2.4 GHz ISM band to assess suitability for broadband communication at close proximities to high voltage equipment. The 2.4 GHz band is also known to provide sufficient bandwidth to support the high data rates required by proposed sensors, which is unlikely to be matched by frequencies previously published in literature.

While this research has multiple monitoring techniques to choose from, visible light will initially be used to prove if the complete monitoring concept is viable. It has been shown that this is a common technique for monitoring insulator and tower structures, however no system offers a view as detailed as what is expected from the robot being proposed. As such the robot will be a useful asset management tool for utilities and a research aid for The University.

While numerous robots have been discussed and interesting movement actions seen, none appear to be compatible with the requirements of the proposed robot. Suggestions can be drawn from earlier work[54, 55, 56] however most movement actions either need to draw inspiration from other technology areas or be extrapolated from basic singular actions.

Chapter 4

Robot System Level Overview

This chapter looks at the robot from a systems level perspective. Initially each subsystem will be explored in turn, re-iterating comments made in Chapter 3 (if required) and detailing research work that needs to be carried out prior to being able to create a working prototype. The layout of the initially proposed system will be presented before detailing the final system design and the reasons for the modifications.

4.1   System Components

The robotic demonstrator being build was ideal in that, once its requirements had been specified, its development could be split into three distinct parts: communication, vision and mechanical, shown in Figure 4.1. As each part had a high degree of independence from the others, development on each subsystem could, in theory, progress in parallel. However as detailed in Section 4.3, with the discovery of advanced, powerful and small computer platforms, more integration was possible than originally envisaged.

4.1.1   The Communication System

As discussed in Chapter 2, there were two main options for wireless communications: unguided optical and radio frequency. The latter was chosen due to not
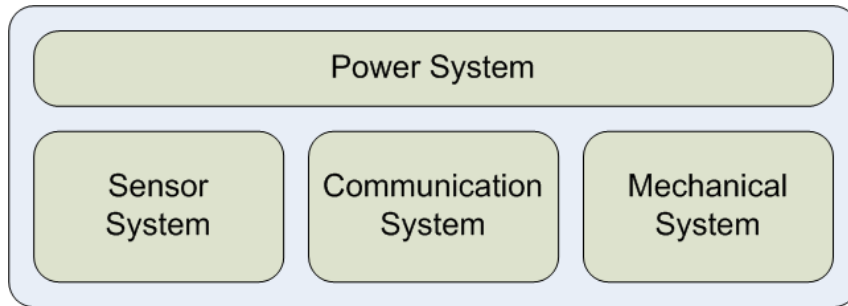
Figure 4.1: High Level Functional System Diagram

needing high-precision alignment between transmitter and receiver. In particular the license free 2.4 GHz ISM band was selected. However the question remained as to whether communication in this band would actually be possible in an environment containing partial discharges and high electric field strengths.

Chapter 3 went part way into revealing an answer to this question. This chapter discovered that a number of researchers were using radio frequency emissions to locate defects at substation sites[20]. This research revealed that radio frequency interference from partial discharge activity tends to be most prevalent at sub 750 MHz frequencies. While between 750 MHz and 1 GHz radio energy seems to drop dramatically. There is one potential 'flaw' though; this research did not examine frequencies above 1 GHz, therefore it is not known if partial discharges would cause any radio frequency emissions in the frequency band under consideration.

The literature review also discussed various researchers' developments in creating live-line robots. The major 'faults' with these developments, from the point of view of realisation and reporting, is the fact they were only tested in off-line laboratory conditions. Furthermore, for the few that were tested under live-line conditions, no mention was given as to data transmission statistics, communication protocols implemented and data correction methodologies.

Therefore, the first objective was to thoroughly assess the 2.4 GHz ISM band under high voltage, high field strength conditions. This was to be achieved by firstly using low cost low data rate radio frequency modules. If these low data rate experiments functioned appropriately, then the final system could substitute the single low data rate module with either a WLAN compliment module or a

bank of low data rate modules functioning in parallel. Details of the test system's implementation, the data collection strategy & experimentation, subsequent data processing and results can be seen in Chapter 5.

4.1.2 The Sensor/Vision System

It is important to note that this robot aimed to provide real-time live-line monitoring to an expert rather than it possessing automatic diagnostic and reporting functionality. As such, there was no need to investigate image analysis techniques in literature. However, should this requirement be needed in future, based on the final implementation, its addition should be possible. Therefore, the main task was to assess the various routes by which this subsystem could effectively be realised.

In Chapter 2, it was decided that a single camera should be used for image capture in order to alleviate complexities with limited shed clearance and to reduce the burden on the processing and wireless transmission systems. The following questions then arose:

- To use a pre-built module complete with lens, a customised camera system or a USB 'web cam'?
- How to process the imaging data and prepare it for transmission over a wireless communication link?
- How to decode the video stream at the operator's end?

Details of this analysis, early Printed Circuit Board (PCB) design efforts and final choice for the image capture device and subsequent processing can be found in Chapter 7.

4.1.3 The Mechanical System

Based on the requirements specified previously, there are essentially three mechanical tasks that need to be established. A number of motions are then needed to achieve these tasks. The required tasks are detailed below:

1. Clamping onto the insulator string

2. Movement along insulator string

3. Scanning surface of an insulator shed

As seen in Chapter 3, there are numerous examples of monitoring robots for high voltage systems; most examples however relate to movement along power lines and the navigation of obstacles[2, 3, 44, 61, 62]. Although examples were seen that could wash[58] and replace[59] insulators on distribution lines. However none of these previous attempts has great relevance to the movements that this robot needs to make. As such an innovative movement strategy needs to be conceived for both the movement along the insulator and for scanning the surface.

The clamping onto and removal from the insulator can however take guidance from previous work. Specifically the patents filed by Fernandes et al relating to live line monitoring hardware[54, 55, 56]. The documented device is circular in shape with a hinge at a point on the circumference, with a break 180° further along. This hinge is then mechanically operated by rotating a hot-stick in a port also located on the circumference. Details of the proposed mechanics to achieve these tasks will be reported in Chapter 10. This chapter also includes details of the mechanical systems which were not further developed due to time constraints.

4.1.4   Power System

The remit for this work did not include the research and implementation of a power scavenging system. As such the primary focus was to develop the live-line monitoring robot using high capacity batteries as the power source, with the intention to consider a power scavenging system after all other development tasks were completed. However as the primary task itself, in mechanical terms, was scaled back due to time concerns, it was not possible to consider a power scavenging system.

4.2   Initial System Design

The initial system level design can be seen in Figure 4.2. This design creates a highly modular system by requiring each subsystem to have its own processing

and control intelligence; noted by the presence of a Central Processing Unit (CPU) device in each functional area in the diagram.
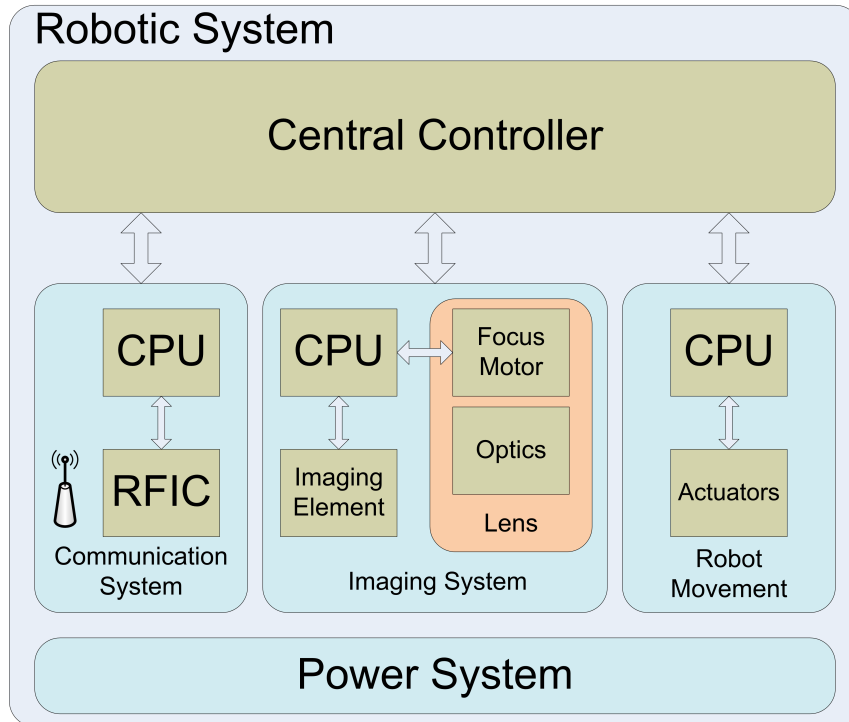
Figure 4.2: Initial Architecture

These individual modules are then interfaced to a 'master' controller. It would be this controller's responsibility to co-ordinate the operations of all attached modules. Furthermore, by defining standard interface requirements, new modules could be designed and should function without modification of the existing hardware. This 'plug and play' mechanism would prove useful for quickly changing the communication mechanism or sensor technology, for example.

While modularisation provides flexibility and in some cases the ability to use cheaper & simpler technology[1], it does have a number of negative aspects. In particular, the number of individual systems to maintain and the space needed to mount the multiple PCBs increases.

Despite these negative aspects, this was still a valid system architecture and hence initially considered. However research into different video capture and processing systems, discussed in Chapter 7, revealed a range of sophisticated, powerful and

---

[1]This is possible because there is less demand on individual processors.

small processors. This discovery lead to the system architecture transforming into that seen in the next section.

4.3   Final System Design

Without repeating too many details found in Chapter 7, one of the processor families[2] discovered contain a General Purpose Processor (GPP), a graphics accelerator and in some cases a DSP. These processors are available as part of pre-built hardware solutions, which are better termed miniature-computers.
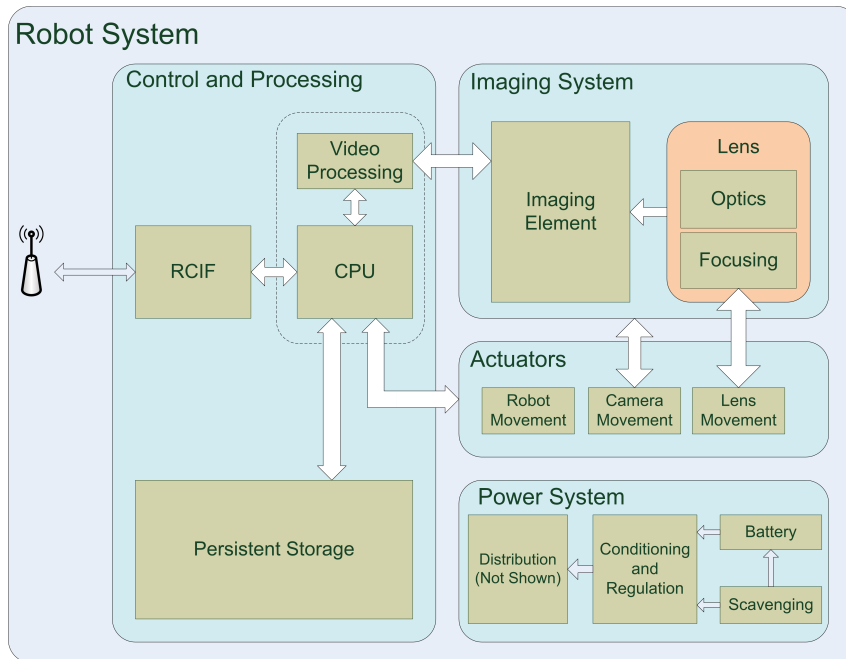
Figure 4.3: Proposed System Architecture

These processors also contain an image capture interface; therefore image capture, compression (via DSP), motor control and communication (via GPP) duties can all be handled via a single hardware platform as shown in Figure 4.3. This miniature computer has the ability to run a Linux based operating system; therefore simplifying some of the programming requirements; for example things like concurrency would be handled by the operating system. Furthermore, if a WLAN module is

---

[2]The family in question is the Open Multimedia Application Processor (OMAP) - OMAP35xx.

selected for use as apart the communication link, then a Linux operating system should readily be able to accept and use it.

To reiterate comments made in Chapter 2, the envisaged device is to be designed to initially operate as master-slave system, as illustrated by Figure 4.4, but with the flexibility to allow mostly automated operation in the future. The operator is to issue commands via a HCI consisting of either a laptop computer or custom hardware solution. These commands would be sent to the robotic investigator which would then carry out the requested operation, returning a status code indicating the success or failure the request. Image and video date would be update and streamed respectively on request.
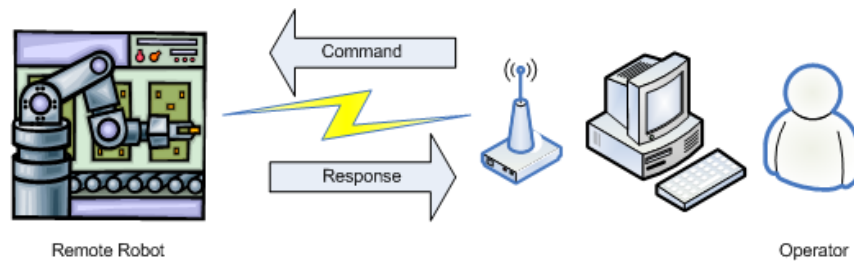
Figure 4.4: Intended interaction between operator and robotic investigator

4.4 Summary

Splitting the robotic device into three subsystems and using knowledge from the literature review (Chapter 3), this chapter detailed the research that still needed to be conducted and questions that needed to be answered for the development of each subsystem. The frequencies of interest for communication have not been studied in detail and no mechanical system of the type envisaged has been previously documented. The mechanical subsystem was partitioned into three separate motions for onward parallel investigation. Based on the Chapter 2 decision to use a single visible-light camera, a number of questions were posed with regards to its implementation and associated processing methods. An initial proposal for distributed intelligence and control was put forward before being revised to include highly versatile integrated processors. The system will be designed to initially operate using master-slave methodology, but have the ability to be used for highly automated inspection in the future.

This page is intentionally blank.

Chapter 5

Preliminary Communication System Development

It has been previously mentioned that there are two possible communication systems that could be implemented: free space optical link or radio frequency. The latter is the system that is favoured because a precise direct line of sight is not required in order to maintain communication. However there is one possible problem that needs further investigation; the effects of high field strength and partial discharges on stable and reliable radio frequency communication at 2.4 GHz.

It was seen in the literature review that there was some conflicting information. On one hand it has been suggested that radio frequency noise can be used to detect fault locations in transformers within substations [20, 64], while another study seemed to suggest that radio frequency communication can be used for monitoring on high-voltage lines[54, 55, 56]. This needs to be resolved prior to confirming whether or not radio frequency communication is feasible for the intended application.

5.1 Wireless Communication

5.1.1 Radio Frequency Use in the United Kingdom

The United Kingdom has a number of frequency bands which are defined as Industrial, Scientific, Medical (ISM). These bands are exempt from the licensing requirements of UK legislation provided certain conditions are met. These include stipulations that only a certain classification of device can use particular bands and at what power levels. Most devices which operate under the licence exempt clauses of The Wireless Telegraphy Act 2006[65] are classified as Short Range Devices (SRD)[1].

While there are relatively few ISM frequency ranges, there are multiple frequency sub-bands authorised for different purposes; these are documented in UK Interface Requirement 2030 'IR2030'[67]. Despite the fact that multiple SRD bands exist not all are suitable for streaming image & video data and telemetry.

There are five SRD sub-categories that a condition monitoring robot could be classified as:

1. Non-Specific Short Range Devices
2. Industrial/Commercial Telemetry and Telecommand
3. Wideband Data Transmission Applications
4. Wireless Video Cameras - Non Broadcasting
5. Video Distribution for Private Use

While IR2030 provides no specifics as to what constitutes 'Video Distribution for Private Use', classification under this sub-category could be considered stretching its intended meaning. The same could be said of the first item in the list; this classification is only suitable if the device does not clearly fit into any other category.

Sub-categories 1 to 4 listed above allow the use of the 2.400-2.4835 GHz range at a maximum power level of 10 to 100 mW Effective Isotropic Radiated Power

---

[1]Roke Manor Research has produced a wall chart which depicts frequency use in the United Kingdom and indicates the available ISM bands[66]. A textual version detailing various ISM bands is available from The Office of Communications and provides a clearer view of allocations and the associated requirements[67].

(EIRP). This is also the frequency range used in IEEE802.11b/g WLAN. This band additionally has one major benefit: it has worldwide licence exempt status, hence its popularity and wide deployment in WLANs. This global licence exempt status has benefit of allowing systems which use such technology to be deployed worldwide with only minimal modification (if any).

5.1.2 Problems of Unlicensed Frequency Use

A frequency range should not solely be selected on the basis of possible worldwide operability and frequency popularity; the frequency chosen needs to allow the system to operate successfully in the harsh environment surrounding high-voltage lines. In this respect the 2.4 GHz ISM band provides an appropriate starting point.

Frequency popularity often has negative consequences. ISM bands are largely unregulated and their licence exempt status allows any individual or company to make use of them. Even if the conditions of licence exemption are adhered to, it is possible for interference to occur if devices are operated in close proximity. Users of such systems/equipment have no legal recourse to prevent interference, provided the conditions are met; all users need to cooperate, whether it means switching to a different frequency in the range or reducing transmission power.

The intended application should not suffer interference from other transmitters as high-voltage systems are usually situated away from populated areas and hence WLAN, Bluetooth, wireless video senders and other sources of RF transmission in the 2.4 GHz range.

5.2 The Proposed Test System

It has been previously mentioned that not much is known about the effect that high field strength and partial discharges will have on reliable communication at 2.4 GHz. This needed to be investigated before suitability for the intended application could be decided upon. This testing was conducted in The National Grid High Voltage Research Centre here at The University of Manchester.

Before testing could commence, a system for measuring interference with known or measurable background characteristics was needed. A common way of measuring

such interference is through Radio Interference Voltage (RIV) measurements [68]. While this can provide an indication of the interference that might be expected in the field, it will not give an indication of the stability of RF data link, likely transfer rates or the type of data corruption that is likely to occur. A test system which fulfils these criteria needed to be constructed; therefore a system which allowed a maximal length pseudo-random sequence to be transmitted, received and checked for accuracy was designed. The system was to consist of three pieces of hardware listed below and shown in Figure 5.1.

- Transmitter unit (can additionally receive)
- Receiver unit (can additionally transmit)
- Control Room unit
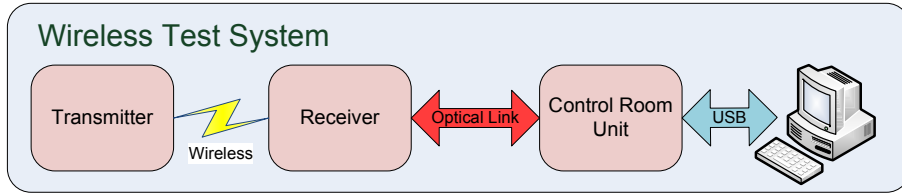
Figure 5.1: Overview of the Wireless Link Test System

The transmitter was to generate a maximal length pseudo-random sequence based on a Feedback Shift Register (FSR) and transmit it at a range of frequencies. The use of a transceiver chip give the unit the ability to receive in addition to just transmitting. This was to allow power saving routines to be implemented should the need arise.

The receiver was to listen at each frequency and process packets received. The payload of the first packet was to be used as the seed for the local PseudoRandom Number Generator (PRNG)/FSR and thereafter for each received packet the PRNG was to be shifted forward one step. A mechanism was required to stop spurious burst errors from potentially skewing the results. The proposed solution was to implement a guard period of approximately ten or 20 correctly matched packets before a synchronisation is declared and statistic collection (re)commences.

Post synchronisation, for every correct comparison between received and local data, a mark character was to be transmitted to a data logging device. To avoid

any unnecessary processing delays or overheads a PC was used for logging duties. When a comparison was not successful both the local and remote values were to be recorded by the logging system.

The transmitted payload would not contain a single instance of the current PseudoRandom Number (PRN), but instead three copies. Prior to the remote and local PRN comparison, in both synchronising and running states, these three copies were to be compared. If there was a single discrepancy all three (remote) values were to be transmitted to the logging system along with local PRN value. This would then allow high level analysis of the corruption to the payload. Through analysis of such data it should be possible to identity if a certain part of the payload structure is more susceptible than others. There are two possible instances when, given that all PRN instances in the payload match, that the latter comparison could fail;

1. A data corruption has occurred in the same manner in all PRN instances, although this is unlikely, it cannot be claimed as impossible.
2. A number of transmissions have been missed.

Both types of error could be identified through analysis of the log; the number of synchronisations that have occurred (at each frequency) would additionally be extracted. This latter information could provide an indication of the quality of the link.

In this model, the need to re-synchronise indicates data corruption at the lowest level. Depending on transceiver options (e.g. forward error correction) or higher level protocols, this data loss may or may not be recoverable. However, these high-level features were not to be implemented for this investigation as they did not allow the type of error to be seen clearly. It is important to note that for the purposes of this assessment, minimal hardware packet handling would be enabled on the transceivers; options such as preamble and Cyclic Redundancy Check (CRC) were additionally to be kept to a minimum.

The data logging stream produced by the receiver unit would not directly be captured by a PC, but instead was to be relayed through a 'Control Room' unit shown in Figure 5.2. The Control Room unit was to link to the receiver via fibre optics. This would allow the transmitter and receiver to have a 'line of sight' and will better represent conditions that the final robot would be expected to

work in, while ensuring safe use of the equipment in the High Voltage (HV) test environment.
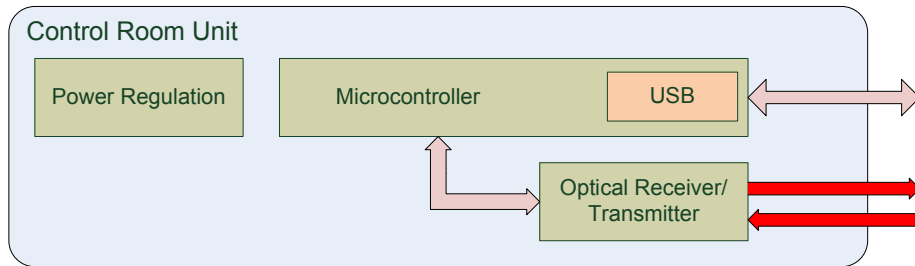
Figure 5.2: Intended Control Room Unit Diagram

The Control Room unit would then interface to a PC using serial emulation over USB; this emulation will permit the using of both HyperTerminal and custom applications. The latter are simplified as most programming languages have built in support to 'attach' to serial communication ports (RS-232): there is no need at present to write custom Universal Serial Bus (USB) drivers.

While this section has detailed the initially envisioned test system, issues experienced during implementation lead to a number of changes being required. These changes will be discussed throughout the next section, which charts the development process of this preliminary communication test platform.

5.3 The Developed Hardware

The transmitter and receiver were the most important devices that needed to be developed for the examination of the radio environment in the region of high voltage overhead lines. The control room unit was designed to minimise interference sources external to the HV test environment and to relay that information in an isolated manner.

Past experience, simplicity and availability lead to the use of PIC Microcontrollers from Microchip for processing duties; this required transceivers with a simple interface but powerful features to be located. Two such devices were located in the

form of the Nordic Semiconductor nRF24L01 and the Texas Instruments CC2500. Beneficially both possess identical core interfaces, namely the Serial Peripheral Interface (SPI), while other pins were dependant on the chip in question.

Importantly, both devices were available in prefabricated modules for quick development; the Nordic device retailing at ¤15 each and the Texas Instruments (TI) at approximately £53 a pair. The difference in the price of these modules is clearly a matter of quality, as will be discussed later.

It should be noted that the TI modules are intended for use with their SmartRF Studio development hardware, which in turn uses their own microcontrollers and Integrated Development Environment (IDE), CodeComposer Studio. Despite the small differences in module construction it was possible to develop hardware platform which would accommodate both chips.

5.3.1   Nordic Semiconductor

The first transceiver explored was the Nordic Semiconductor nRF24L01. This chip could potentially transmit at speeds of 1 Mbps or 2 Mbps, in a bandwidth of 1 MHz or 2 MHz respectively. Frequency use was also restrictive, the user only being allowed to transmit/receive at frequency megahertz boundaries, i.e. 2400, 2401, 2402 MHz etc.

While these speeds would be ideal for the transmission of photographs within a reasonable time-frame, early assessment revealed that they were very poor at reliable and accurate data transmission. A simple test which involved repeatedly transmitting a fixed six byte sequence revealed numerous errors which appeared to be linearly increasing with time. Error types included a single byte being incorrect, one or more bytes being incorrect or the first byte being 'lost', with the remaining bytes shifted down the payload, as seen in Figure 5.3. This last error type will be refered to as the 'Nordic Shift Error'.

No further discussion into the development or use of this chip will be provided. This subsection has been included for the purpose of completeness.
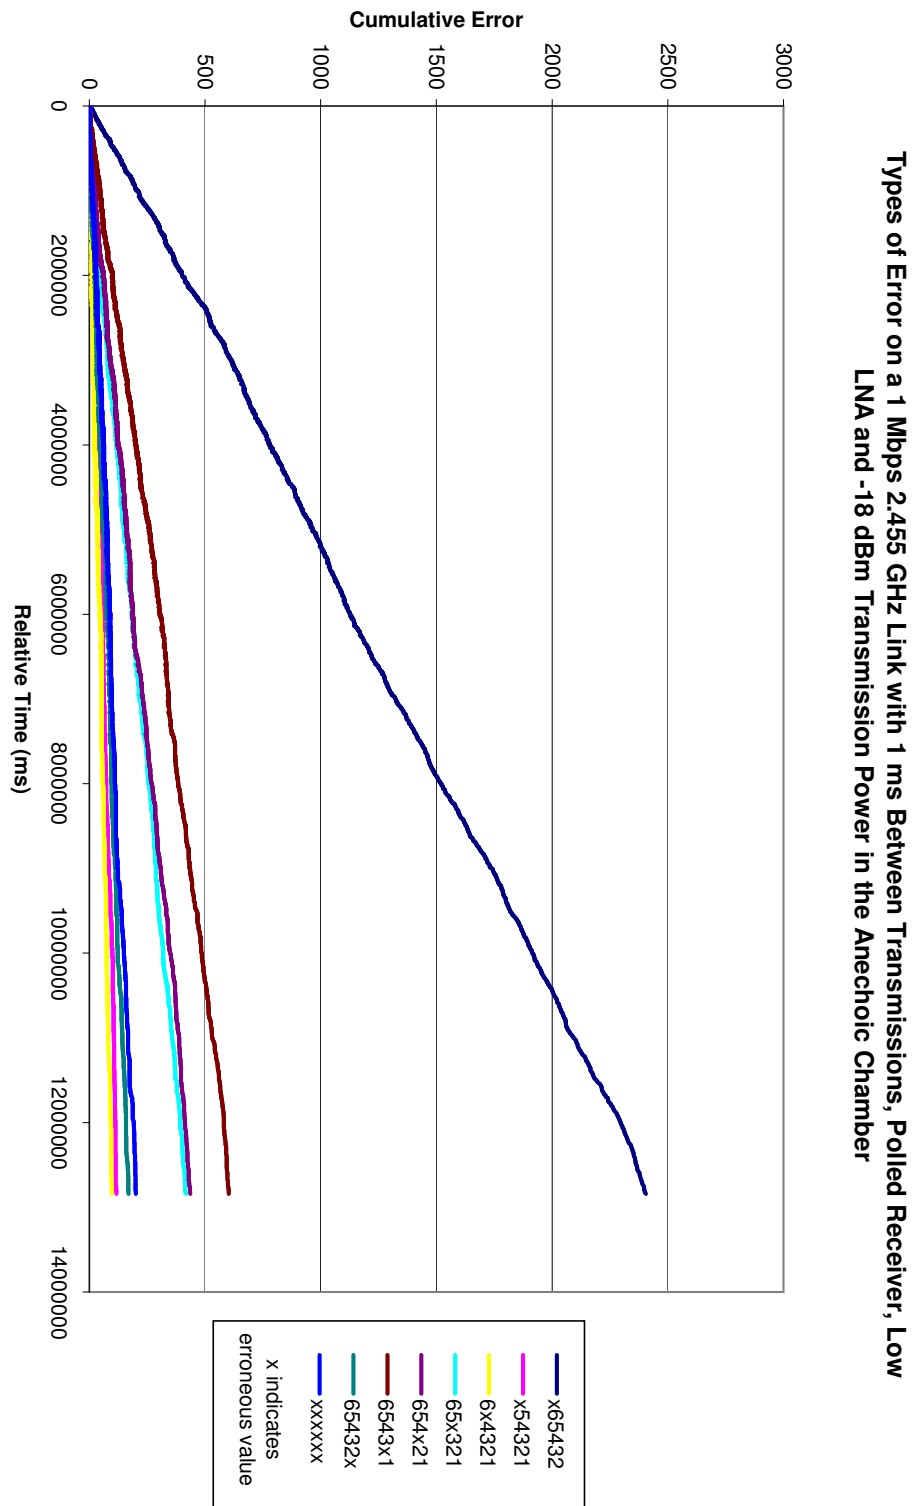
Figure 5.3: Error rates, per type for a Nordic Semiconductor Transceiver

5.3.2 Texas Instruments (ChipCon)

Due to the problems seen in the Nordic Semiconductor nRF24L01 radio chipset, an alternative product was tested; this was the TI CC2500. These devices were purchased as a pre-built modules to aid quick development and integration into built hardware.
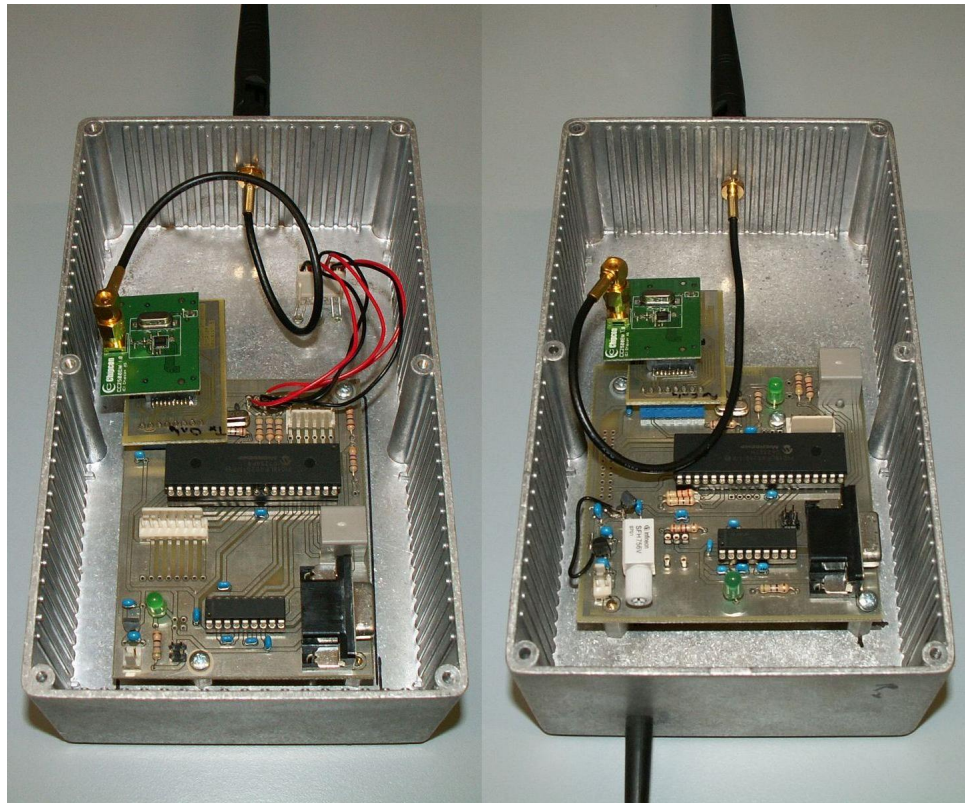
Figure 5.4: Developed Hardware: Receiver (right), Transmitter (left)

Figure 5.4 shows the developed receiver and transmitter with ChipCon devices mounted in their shielded housing. The conversion board can be seen between the ChipCon module and the main PCB. The serial port that can be seen towards the lower right of each image aided development by allowing textual progress information to be sent to a terminal application during debugging.

This chip is highly configurable when compared to the nRF24L01; for example the RX filter bandwidth is user configurable and the transmission frequency has a higher resolution than the nRF24L01. This chip is additionally classed as frequency agile and offers quick turnaround from transmission to reception.

Frequency agility is achieved by allowing the frequency synthesiser calibration registers to be calibrated ahead of time and saved, being written back to the device as required. The frequency synthesiser can also be calibrated after the frequency registers have been set, either automatically after frequency change (after transmit or receive strobe) or manually, by issuing a Frequency Synthesiser Calibration strobe. Depending on the options chosen the blanking interval between each hop is between 90 us and 810 us.

Unlike the nRF24L01, the CC2500 offers customisable channel spacing, a number of modulation techniques and customisable baud rates. These features, together with the Received Signal Strength Indication (RSSI) and Link Quality Indication (LQI) values offered by the chip, make it an ideal candidate to assess the radio environment in a high voltage environment. This chip additionally offers more features that could be used in the field such as Forward Error Correction (FEC), Manchester Coding and Data Whitening.

While Texas Instruments' SmartRF Studio is primary intended to work with their development hardware, the software element can be utilised stand-alone to help configure a range of chips. While a number of preferred combinations exist the user is free to create other configurations. This application was used to create a number of wireless link configurations which were tested in a office environment.

Results from preliminary testing of the chipset indicate that at times the performance is worse than seen with the Nordic Semiconductor chipset, while at others it is supersedes it. The data shown in Figure 5.5 was produced using Minimum Shift Keying (MSK) modulated transmission at 2.433 GHz, with an RX Filter Bandwidth of 812 kHz; it is an example of when the CC2500 is clearly superior; although there are some similarities to the data produced by the nRF2401 (Figure 5.3).

The first resemblance is with rate of error occurrence; errors appear to be occurring at a near constant rate. This appears so until the final stages of the test run, however this stage of the run corresponds to a Monday morning and hence could be caused by external activity. However if the error percentage is considered, there appear to be small periods when no errors occur; this reflected by the sudden decrease in percentage error.

The rapid fluctuation in the first quarter of the figure in the percentage error, could in some ways be a cause for concern, yet may be a beneficial sign. These

fluctuations are caused by sporadic errors; a shallow rise in the cumulative error rate (between two points) indicates that an error has occurred before it was expected, based on the current error rate. A sudden fall in the error rate indicates that an error has occurred substantially long after an error would have normally been expected.

A burst error on the other hand would be expected to produce a curve of decreasing gradient; the more errors that occur in a short space of time, the less influence it has on the percentage error at that point. This is because each plot point presents the occurrence of an error. Conversely, a constant error percentage indicates a near precise periodic error.

Nevertheless there are some marked differences between the results seen here and those obtained for the nRF24L01. Notably the error rate (when assessed over a long period) is less that 0.01%. However as there were some differences between the setups, these results are not directly comparable. The differences are noted next.

The nRF24L01 was configured for Gaussian Frequency Shift Keying (GFSK) 1 Mbps transmission at 2.455 GHz and tested in an anechoic chamber with a -18 dBm transmit power, while the CC2500 uses 500 kBaud MSK modulated transmission at 2.433 GHz with a 0 dBm output power; thus it can be inferred that the CC2500 will survive better because of the high transmission power. Whilst this is an acceptable argument from some aspects, this should have been partly compensated for by the use of an anechoic chamber. While an indication and approximation can be draw from the CC2500 results, direct comparison is not possible and should not be relied upon.

A more important metric would be to assess the types of error seen during the testing of this link. It was previously noted that the nRF24L01 suffered from the so called 'Nordic Shift Error'. Assessing a similar plot from the above test run, Figure 5.6, suggests that no errors of this type occurred; data analysis shows three such errors occurred, however these are not visible due to their low occurrence.

The most common error to occur was where more than one character became corrupted, represented by the error type 'xxxxxx'; with approximately 575 occurrences in almost three days, compared to 2500 occurrences of 'x65432' in the nRF24L01 in three and a half hours.
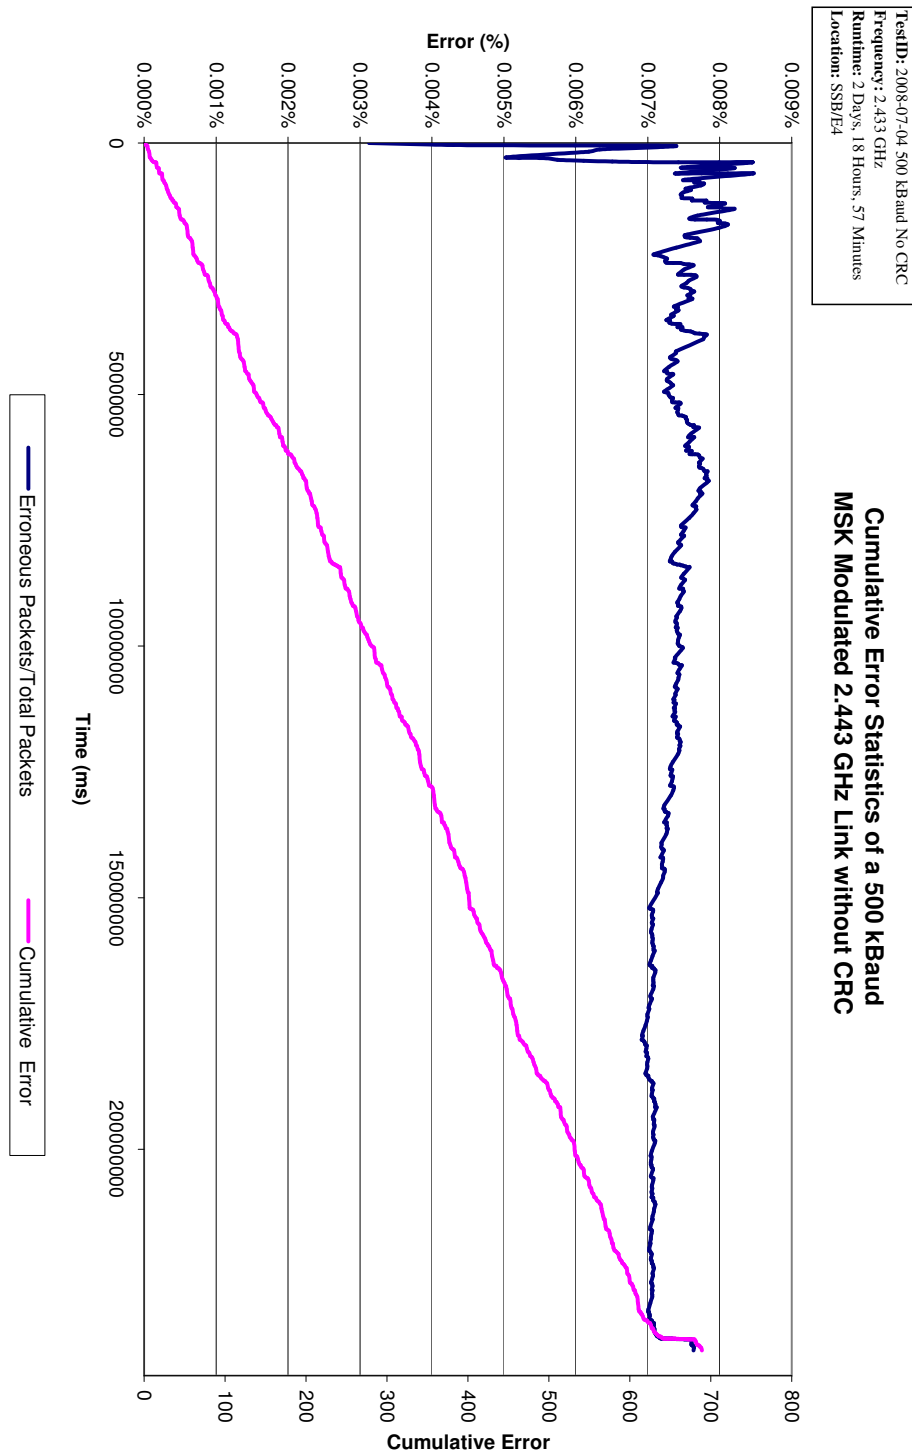
Figure 5.5: Long Term Test of Texas Instruments Unit

This is an early indication that, provided the chipset is configured correctly, the CC2500 can be very stable and robust. Due to this robust nature, the CC2500 was used for all further testing.

5.3.3 Control Room Unit

The largest of the HV laboratories at The University of Manchester is split into two parts; a large test area and a small control room; the two areas are interlocked during HV experimentation. A glass and steel construction separates the two halves. To avoid any potential interference with this structure, it was decided to place both the transmitter and receiver within the isolated HV section of the test laboratory.

This forced the requirement for a mechanism to capture logging data without the use of an electrical connection, i.e. an isolation barrier, and thus facilitate the safe use of the HV test laboratory. The control room unit was created to provide such barrier between the receiver and the data collection system, a desktop computer. The connection between the receiver and the control unit is isolated by means of a polymer optical fibre connection, while a USB connection was used between the control room unit and the desktop computer.

This unit does not perform any processing on the incoming data stream, instead relaying it straight to the USB interface. This has many benefits such as the unit being fairly generic and not having its firmware tied to a particular application. Furthermore, by passing raw data logs to the attached computer for storage allows that raw data to be processed in multiple ways as and when the need arises.

The major part of this unit's firmware is from a reference design provided for developers by Microchip [69]. It has been configured to present itself as a Communication Class Device (CDC) device of the USB specification; it is therefore seen as another serial port and is assigned the next available serial port number e.g. COM3.

The optical communication system has been designed around photo emitters and transistors from Infineon Technologies. To aid firmware implementation these devices were attached to the microcontroller's Universal Synchronous Asynchronous Receive Transmit (USART) pins; therefore byte transmission, start bit, stop bit

Figure 5.6: Error Type in Texas Instruments Test
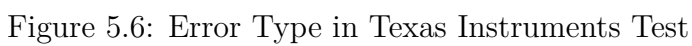
and bit timing were handled in hardware. The alternative was to manually toggle bits on the transmitted end of the link and poll on the receiving end. These operations effectively block the microcontroller from executing other instructions and are therefore best avoided.

However, there is one caveat to this optical link; the operation of the phototransistor. When light is present at the phototransistor's input it produces a logic low output thereby inverting the transmissions. While the data itself could be inverted within the microcontroller, the main problem related to the inversion of the start and stop bits. These two bits are crucial to the correct operation of the hardware USART and hence a solution was needed urgently. One solution was to use an inverter between the phototransistor output and the USART input pin, however the PCBs at this stage had already been manufactured and hence this was not an option for current hardware.

It was discovered that the microcontroller used (the PIC18LF4550) had the ability to treat USART inputs as inversions of the required data stream. However it was found that this functionality was not present in early silicon revisions and only present as of silicon revision B6. When samples using B6 silicon were obtained from Microchip this system became operational.
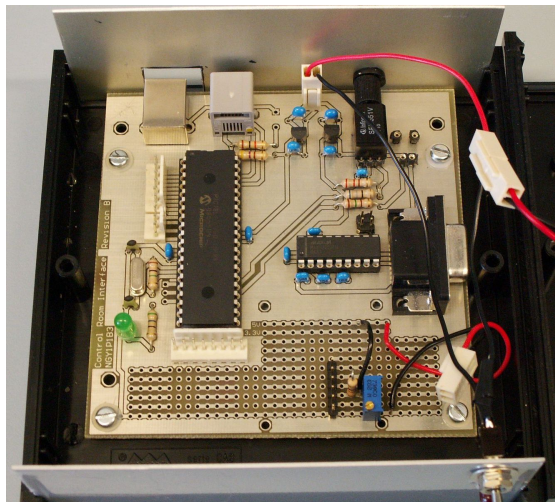
Figure 5.7: Control Room Hardware

The developed Control Room unit can be seen in Figure 5.7 and is shown in its housing. The USB port can be seen in the top left of the image. The back module near the top right is the phototransistor. As it is not envisaged for the unit to be

reprogrammed in the field, the In-Circuit Serial Programmer (ICSP) port (grey socket) is not externally accessible. Unused ports are accessible via headers and provision was made for small scale modifications and additions via a prototyping area (lower part of circuit board).

During the testing of this hardware it became apparent that it was easy to inadvertently block some required USB servicing routines written by Microchip [69], for example having long delays in the 'user' code area of the firmware template. Consequently, a period of extended testing was required.

The PRNG which was initially created for examining the wireless link, discussed in greater detail in the Section 5.4, had already been written at this point. It was therefore the ideal candidate for assessing the stability of the USB connection to the computer and the associated Microchip USB CDC device driver.

A user firmware method was written to send subsequent PRNs[2] over the USB link at fast as the system would allow, i.e. no intentional delays were implemented. This system was subsequently left for a period of 24 hours; HyperTerminal was used to record this data stream to a text file. In this time a file of over 400 MB was created equating to over 200 million PRNs: at least 3200 full cycles of the full 16-bit sequence[3].

A Java application was written which implemented an identical FSR PRNG; this was used to verify that the USB routines on the microcontroller were functioning as expected. This also served as verification that the PRNG routine for the microcontroller worked as expected and matched its Java sibling.

5.4 Spectrum Analysis Methodology

Testing a single frequency and static data will not provide an adequate assessment of any radio environment; therefore it was necessary to test various data combinations at various frequencies within the 2.4 GHz ISM band.

Previous studies have not considered radio communication in this band in significant detail; as such it was essential that fine resolution frequency measurements

---

[2]A software FSR was used to generate these numbers

[3]$2^{16} - 1 = 65536 - 1 = 65535$ pattens

were made. The transmission (and reception) frequency in the CC2500 can be easily changed by altering the channel setting while the frequency 'jump' depended on the currently configured channel spacing.

5.4.1 Pseudo Random Number Generation

While it would have been possible to pre-configure a number of different test patterns for transmission, this would not have sufficiently tested the wireless link. The most taxing method for assessing the link was to transmit every possible pattern of a binary string of a certain length. This was achieved through the use of a maximal length PRNG FSR; this allowed all possible combinations to be tested, while appearing random in nature. The use of a PRNG additionally simplified the synchronisation of the transmitter and receiver.

A 16 bit sequence was chosen for implementation for two reasons;

1. it is a multiple of the 8-bit processing architecture of the PIC Microcontroller and therefore simplifies implementation
2. it provides 65,535 bit patterns, which could be theoretically transmitted across the link in a reasonable time frame

In order to get a maximal length PRNG, the appropriate bit positions in an FSR need to be combined using the eXclusive OR (XOR) function; these positions are referred to as taps. Appropriate tap settings can be calculated, however for speed of development these were obtained from Maxim-IC [70]. This application note suggested that for a 16-bit maximal length FSR, taps 1, 2, 4 and 15 should be used.

The generation code block was written in assembly language rather than the Microchip C language used for the remainder of the system due to the ease with which individual bits can be manipulated, this can be seen in Listing 5.1. The PIC18 series of microcontrollers have an 8-bit architecture; therefore two bytes are required to implement the 16-bit FSR.

The bit relating to each tap was checked sequentially, incrementing a counter if the bit was set. The Least Signification Bit (LSB) of this counter was rotated into a carry flag. This was then rotated into the Most Significant Bit (MSB) of the

Listing 5.1: Pseudo Random Number Generating Code

```
1 movlw 0x00
2 movwf counter
3 btfsc prngLower, 0
4 incf counter
5 btfsc prngHigher, 3
6 incf counter
7 btfsc prngHigher, 5
8 incf counter
9 btfsc prngHigher, 6
10 incf counter
11 rrcf counter      ;move the lsb into carry
12 rrcf prngHigher  ;move carry into the msb of prngHigher
13                  ;the lsb of prngLower -> carry
14 rrcf prngLower   ;carry into msb of prngLower
15                  ;lsb of prngLower is lost
```

higher byte of the simulated shift register. A third rotate instruction moved the LSB of the higher byte into the MSB of the lower byte. The LSB of the lower byte is thus removed from the shift register.

Closer consideration of the above shift register shows that an XOR (Equation 5.1) operation is not being carried out as needed. Instead Equation 5.2 is being used to generate the next bit pattern. However this issue was not discovered until after a number of testing runs, the results of which can be found in Section 5.7 and Chapter 6.

$$A\overline{B}\overline{C}\overline{D} + \overline{A}B\overline{C}\overline{D} + \overline{A}B C\overline{D} + \overline{A}B C D \tag{5.1}$$

$$A\overline{B}\overline{C}\overline{D}+\overline{A}B\overline{C}\overline{D}+\overline{A}B C\overline{D}+\overline{A}B C D+\overline{A}B C D+A\overline{B}C D+A B\overline{C}D+A B C\overline{D} \tag{5.2}$$

Due to this, it was not advisable to correct this error without first assessing the quality of the PRNG that had been designed. This performance analysis was achieved by generating codes until a repetition was spotted, at which point the number of codes between the repetitions would be known: the closer the number

is to the maximum, the higher the quality and the less need there is to rectify the issue at that point in time.

The Java based version of the PRNG that was created to test the USB link of the Control Room unit was used to continually generate bit sequences until the hardcoded seed value was seen. When the generated numbers were placed in ascending order, it was discovered that 65,535 unique numbers were generated, the 65,536th being the seed value.

The number of possibilities is equal to the maximal length sequence; therefore continued use of the generation mechanism was valid without modification of the algorithm. This proved to be of less concern due to the 'bottleneck' and communication timing issues that are discussed in Subsection 5.4.3.

While both the transmitter and receiver can generate identically PRN sequences based on the developed code block, there is no direct synchronisation ability. A description of a possible PRNG synchronisation method was provided previously in Section 5.2. The synchronisation mechanism actually implemented operates as follows:

1. The transmitter starts transmitting PRNs in a periodic manner using the arbitrary seed value 61,455 (1111 0000 0000 1111). A new number is generated for each transmission.

2. The receiver, assuming the frequencies have been synchronised, accepts its first packet and treats it as a seed for the local PRNG.

3. The receiver then accepts the next 20 packets over the link, each time comparing the received PRN with the PRN generated locally based on the seed received. Each time a packet is correctly received in this 'training' window, a period character (.) is sent to the logger. After each comparison operation, the next PRN is generated for use in the next comparison.

4. When 20 packets are successfully matched, an 'S' character is sent to the logger, along with 'carriage return' and 'new line' characters. These final three characters have no real meaning in relation to the quality of received data, only being implemented to enable visual inspection of the data log.

5. Every subsequent packet whose payload matches the locally generated number causes a '-' character to be sent to a logger. This keeps track of the number of error free transmissions in a minimised fashion.

6. If a match is not successful, both the locally generated PRN and the PRN received over the wireless link are sent to the logger. This can then be analysed at a later stage to assess what type of error occurred in the transmission. This then causes the need for another synchronisation; the process restarts from Stage 2.

In Section 5.2 it was noted that the use of three instances of the transmitter generated PRN within the payload would allow detailed assessment of payload susceptibility. This proposal was therefore implemented into the synchronisation system as proposed.

When a packet is received, the three copies are compared and assessed for corruption. This is carried out prior to the remote-local PRN comparison operations noted in the above stages. When corruption is identified, the three copies along with the local PRN are sent to the logger.

5.4.2 Frequency Synchronisation

In a development (office/laboratory) environment it is possible to activate both transmitter and receiver (via power switches) so they will be closely synchronised in frequency. However, for the longer term and during testing in high voltage environments this is not a suitable mechanism to utilise. Furthermore, this requires the start-up sequences of the transmitter and receiver to be closely matched in terms of instruction execution time.

Under the devised channel synchronisation scheme, the transmitter is classified as the master; due to it changing its channel at regular intervals, starting at channel zero, from the moment it is powered. It should be noted that there is no mechanism currently in place for resetting the transmitter remotely. The receiver is classed as the slave and works as follows.

When powered, it starts listening at channel zero. If after a short period of time (~ 3.5 seconds based on a 20 MHz oscillator) it does not receive data, the receiver switches to the next frequency. The receiver listens on each subsequent channel for this period of time until data received. When this occurs it is not known if this data is a consequence of finding the correct transmission frequency or if the data is a result of noise. Therefore the value received is used as a seed value for the

PRNG engine. The next 20 transmissions are checked to verify that the sequence is correctly adhered to. If successful, the system asserts that the correct frequency has been located.

At this point, although the correct frequency has been found, the exact position in the transmission window (i.e. the period for which the transmitter uses each frequency) will not be known. This is resolved by switching the receiver to the next frequency and then listening for transmission. Once a value is detected, this is again used as a seed for the PRNG engine and the next 20 packets are assessed. If correctly matched, the receiver's timer is activated and configured for the same frequency change interval as the transmitter. This mechanism will produce a transmitter-receiver pair which is more closely synchronised than when using the power-base method: equivalent to the time required to transmit 21 packets and activate the timers.

5.4.3   Frequency Dwell Timing

Calculating the exact number of test patterns that can be transmitted in a given time requires knowledge of the packet structure. As this adds to the raw data that is transmitted, the useful data rate achievable will be less than the signalling rate. Additionally, it may not actually be possible to transmit at this rate due to processing requirements, as will be seen shortly.

Given a preamble of two bytes, a two byte sync word, a single byte for payload length information and a six byte payload, a complete packet would contain 11 bytes. Based on this and a 500 kBaud link, it would be theoretically possible to transmit 5,681 packets a second; therefore ~11.5 seconds are required to transfer the entire PRN sequence.

However this is not the case due to limitations with the hardware developed. This timing does not account for processes such as generating the next PRN or 'clocking' the payload into the transmitter and out of the receiver. Furthermore, the initial version of the receiver hardware did not contain on-board flash memory, therefore requiring the use of a computer attached via a simple (RS232) serial connection to collate all data; which is in most cases a fairly slow process. This throttling of data throughput was not a seen during early development as at that point a static

payload and single frequency was used. Conversely in this instance it is intended to test multiple bit patterns, thus the minimisation of bottlenecks is important.

This was compounded by the introduction of the Control Room unit. As previously mentioned, this was attached to the receiver by means of an optical link. The optical emitter and detector units were rated for a maximum data rate of 5 Mbps and this ability was confirmed using a logic analyser. Instead of devising a proprietary protocol for this connection, the microcontroller's onboard USART hardware was used.

This should theoretically allow speeds greater than 250 kBuad to be attained, however during testing it was discovered that little over 10 kBaud was achievable. Furthermore, this was only possible by introducing a 30,000 cycle delay (approximately 6 ms) between transmissions. This meant that less than 166 packets could be transmitted a second. The root cause of this is likely to be the use of Microchip's pre-written USB firmware; whereby it is possible that essential USB servicing is being handled when data is arriving on the optical-serial link.

This subject was not explored further. Over a period of one minute 9960 packets can be transmitted by the developed system. This is approximately equivalent to a full cycle of a 13-bit maximal length pseudo random number generator. This is a sufficiently large number of test patterns with which to assess the link.

5.5 Data Logging and Analysis Algorithms

Listing 5.2 shows an example of the types of logging data that can be generated by the developed system. During a test run this information is logged to a text file for processing at a later time using a Java application. There is a elegance to this raw data file; inspection of this example with the proceeding description of events which caused this output shows that the data file can be assessed in various manners.

A description of the events that cause the data in Listing 5.2 to be generated is presented below;

Lines (1) and (3) indicate a successful synchronisation; it means that a seed packet plus the 20 subsequent packets were successfully sent over the established wireless

Listing 5.2: Pseudo Random Number Generating Code

```
1 ....................S
2 ———————————————————————————————————L12AB  R12AC
3 ...................S
4 ———————————————————————————————R12AB  12AC  12AC  L12AC
5 R12AB  12AC  12AC  L12AC
6 L12AC  R12AB
7 ...L12AC  R12AB
8 00000001
9 L12AC  R12AB
```

link. Twenty-one packets have been sent if this message is seen.

Lines (2) and (4) occur after the occurrence of either line (1) or (3); each '-' (dash) character indicates that a packet has been received as expected (i.e. without corruption). The occurrence of an alphanumeric string indicates an error has occurred. In the case of line (2) the locally generated PRN does not match the value received over the wireless link. In the case of line (4), this error indicates that the three instances of the remotely generated PRN, which were in the packet payload, do not match. In each case the local and remote (prefixed with 'L' and 'R' respectively) are logged for later analysis.

Lines (5) to (7) indicate synchronisation attempts. Lines (5) and (6) show that once the seed value is received, the very next packet is in error. Whilst line (7) shows that a number of packets after the seed value (3 packets in the above example) are successful before the error occurs. The error shown is where the values within the payload match, but there is no match between the remote and local values. It is also possible to have an error of the type where payload packets are in error, similar to the error shown in line (5).

Line (8) shows a change in wireless channel. This is mapped onto frequencies using the following equation: Base Frequency + Channel Separation x Channel Number.

Line (9) indicates an error, however this must be discounted as it is a consequence of the channel change. This is due to the fact that in the devised channel synchronisation scheme, the transmitter acts as the master device. This has the consequence of the transmitter being a finite amount of time ahead of the receiver, i.e. the transmitter will have finished its channel change and started transmitting

the next PRN before the receiver has completed its channel change operation. As the transmission has started prior to receiver being ready to listen, it will at least miss one packet (the actual number being related to the synchronisation delay), thereby causing the two units to be out of step.

Two methods of processing the raw data file have been used; while the statistics they produce are similar, the results produced with the second solution are more meaningful in relation to communication errors. Nevertheless, as both solutions are derived from the same algorithm, albeit with modifications to the final calculations, both statistic collection equations will be discussed.

The Java application works by reading the raw data file a line at a time. Each line can only start with a limited number of printable American Standard Code for Information Interchange (ASCII) characters. Dependent on the character at the start of the line, an appropriate action is undertaken. A number of statistics are collated as the file is processed; these are shown in Table 5.1.

Table 5.1: Java Processing Statistics

Variable Name	Description
trueCnt	The number of good packets after synchronisation
falseCnt	The number of bad packets after a synchronisation
syncCnt	The number of full synchronisations
negSync	A bad packet either at the start of synchronisation or during it
numberOfGoodDuringSync	Number of good packets during synchronisation periods

Referring again to Listing 5.2, the following list indicates when each of the statistics listed in Table 5.1 is incremented:

- Lines (1) and (3) are successful synchronisation attempts, so **syncCnt** is incremented each time. However, for this to happen there would have necessarily been 20 good packets previously. Therefore **numberOfGoodDuringSync** is incremented by 20 each time; increments actually occur individually in a loop as period '.' characters are verified in the line being assessed.

- Lines 2 and 4 show packet reception. As each successful packet match is

shown by a '-' character, each time one is encountered in the line being assessed **trueCnt** is incremented. These two lines end in an error, the system spots this by looking for a 'L' or 'R' character. When this is encountered **falseCnt** is incremented. Firmware design forces a new line after an error and re-synchronisation.

- Lines 5, 6 and 9 show synchronisation attempts which fail straight after the seed character as latched into the PRNG. This event will cause **negSync** to be incremented.

- Line 7 shows a synchronisation attempt which has commenced but failed prior to completing. In this case every period character will increment the **numberOfGoodDuringSync** to increment by one. When the error is encountered, **negSync** will increment.

- Line 8 indicates a change in the transmission frequency. When this is encountered, the above values are displayed in the terminal window and final statistic calculations are performed using the values. Finally they are all reset to zero ready to collect statistics for the next frequency.

A point to note is that only one error can occur after each synchronisation, therefore

$$falseCnt = syncCnt \tag{5.3}$$

The first statistic solution is as follows:

$$temptotalSyncs = syncCnt + negSync - 1 \tag{5.4}$$

Equation 5.4 will calculate the number of synchronisation attempts. It is approximately equal to the number of errors within each frequency; an error after a synchronisation with be accounted for in the **syncCnt** value, while the number of failed synchronisations will be accounted for in the **negSync** value. The minus one is needed because the first received packet after a channel change will be in error due to the small frequency synchronisation delay between the transmitter and receiver.

The alternate statistic solution is more in common with the normal way of assessing communication links. This relies on collating three statistics, termed **n**, **p** and **q**. The total number of packets transmitted is **n**, while **p** and **q** represent the number of correct and corrupt packets respectively. It is then possible to calculate a percentage error, or proportion correct, at each test frequency. The more accurate the calculation of these individual terms the closer the percentage error figure becomes.

It has already been mentioned that the collection of raw receiver data during experimentation enabled assessment in multiple ways. As very little was omitted from the logging mechanism, a very close approximation of **n**, **p** and **q** can be made, as discussed above. In order to achieve this the statistics noted in Table 5.1 need to be combined in a number of manners;

$$negSync + falseCnt - 1 = TotalErrorPackets \tag{5.5}$$

$$negSync + syncCnt = No.of SeedValues \tag{5.6}$$

$$numberOfGoodDuringSync = No.of good packets during sync \tag{5.7}$$

$$falseCnt + 2 \times negSync + numberOfGoodDuringSync$$
$$+ syncCnt + trueCnt = TotalTransmissions \tag{5.8}$$

5.6   Algorithm Discussion

While Equations 5.5, 5.6, 5.7 and 5.8 themselves are accurate, a few assumptions have been made within the processing algorithm which will cause minor errors in actual accuracy. However this is the closest approximation of **n**, **p** and **q** that can be made with the current data format of the raw data file.

Algorithm flaws relate to three areas. Assumption that

- every seeding packet received is free from corruption
- no packets are lost during channel change delays and during logging
- the last packet transmitted on a frequency is not is error

These issues are discussed in detail next.

5.6.1   Assumption of Seed Accuracy

It has been assumed that every seed value has been received correctly, in terms of all three payload instances matching. This problem is of slightly greater concern when multiple errors occur in rapid succession. The seed value can safely be assumed to be correct if synchronisation starts (i.e. if period characters are seen in the log) as this can only happen if the seed was correctly received.

However, problems ensue when the packet after the seed supposedly fails. It cannot be definitely determined whether the failure is due to an incorrect seed being latched into the PRNG or whether the subsequent packet actually failed. There are three possible error combinations that can occur during the initial phase of synchronisation:

1. The seed is correct, the subsequent packet is in error
2. The seed is in error, the subsequent packet is correct
3. The seed is in error, the subsequent packet is in error

There is no differentiation between error types 1 and 2 within the processing algorithm as they increment the number of packets transmitted by two and the number in error by one. Error type 3 is of more concern as the error count needs to be raised by two rather than one.

As error Type 1 and Type 2 lead to identical statistics, the key is to try and uniquely identify error Type 3. As it is unlikely for all three instances to be corrupted in the same way, this issue can partially be resolved by checking the payload of the seed packet prior to loading it into the PRNG. However, some differentiation would be needed between this type of event and when the packet after an uncorrupted seed packet fails in this way. This firmware change then necessitates a change in the processing algorithm to detect the failure of a seed packet.

5.6.2   Assumption of No Packet Loss

It has been assumed that no packets are lost at any point during the experiment. However this is certainly not the case and proved very difficult to resolve with the developed hardware. There are two points at which packets are potentially lost; during a channel change and during the logging of corruption data.

Regardless of how closely the transmitter and receiver are synchronised, the transmitter is fractionally ahead of the receiver when changing channel. Due to this there will be the loss of at least one packet - i.e. the transmitter would have started transmitting its packet before the receiver has stabilised at the new frequency.

Further, packets are lost during the logging of corruption. This occurs because the hardware USART used to control the optical emitter contains a single byte buffer. Therefore when a string of data needs to be transmitted, as in the case of corruption data, blocking operations are used to ensure the string is dispatched as soon as possible over the wired communication channel. This blocking inherently stops further packets from being received.

There is no viable solution to this problem due to the experimentation environment. When assessing a wireless channel for quality, the transmitter and receiver can additionally be attached using a simple or duplex wired channel. This wired channel serves to maintain timing and synchronisation between the two devices; for example, the receiver would be able to inform the transmitter when it was ready to receive data on a frequency, thus effectively closing the channel change delay to zero. Furthermore, the transmitter could be told to stop transmitting while the logging of corruption data was occurring.

However this mechanism cannot be employed in a high voltage environment as the synchronisation connection would provide an additional route to ground (Earth).

5.6.3   Corruption Free Channel Change

This problem is fairly intermittent and is a consequence of the interaction of the channel change mechanism and the synchronisation mechanism. This event affects both the calculation of the number of synchronisation attempts and the percentage error; it only became apparent when after negative percentages were seen in the

Comma Separated Value (CSV) data produced by the processing application.

Negative values were assumed to be solely related to frequencies which contained interference of such severity that no data was successfully transmitted over the link as a consequence of Equation 5.5. However, based on evidence gathered during high voltage experimentation, discussed in Chapter 6, it was discovered that a wider problem exists.

In order to trace the root cause of this issue a detailed examination of the raw data files was required; in particular around the channels affected. This showed that when a negative error percentage was calculated, the associated link data was supremely clean. This was the source of the problem; the application assumes that each channel change causes a single error at the start of the new frequency, leading to the subtraction of one from the number of errors.

Tracing the raw data backwards from the channel in question revealed that an error was received directly before the channel changed. This error would have caused the receiver to be placed in its synchronisation block[4]. This combined with the factional delay between a channel change on transmitter and receiver can produce such a clean data logging instance. The following sequence of events would have occurred:

1. The receiver encounters an error and sends logging data over the optical link
2. The transmitter changes it transmission channel
3. The receiver enters it synchronisation block, using the initial channel
4. The receiver's timer indicates that a new frequency should be used and a channel change is invoked
5. The receiver and transmitter are on the same frequency. As the seed has not been reconfigured, the first packet received is used as a seed.

Resolution of the problem focuses on selectively deciding when to subtract one from the number of errors seen, rather than the implemented blanket policy.

This requires assessment of the types of logging structures that are possible. Emphasis is placed on events at the start and end of the frequency rather than during

---

[4]The firmware was written to allow a channel change while in this block, as without such a facility there is potential for a deadlock situation.

the middle. Refer to Table 5.1 for the purpose of each variable, while Table 5.2 details the various possible structures.

Table 5.2: Possible Logging Error Structures

Error Location (Structure)	Effect on Statistics
Error at start and error at end	negSync = 1, falseCnt = 1, syncCnt = 1
Error at start	negSync = 1, falseCnt = 0, syncCnt = 1
Error at end	negSync = 0, falseCnt = 1, syncCnt = 1
No errors	negSync = 0, falseCnt = 0, syncCnt = 1
Error at start, end and between	negSync = 1+, falseCnt = 1+, syncCnt = 1+ (where falseCnt = syncCnt)
Error at start and between	negSync = 1+, falseCnt = 1+, syncCnt = 1+ (where falseCnt = syncCnt - 1)
Error at end and between	negSync = 0+, falseCnt = 1+, syncCnt = 1+ (where falseCnt = syncCnt)
Error in between	negSync = 0+, falseCnt = 1+, syncCnt = 1+ (where falseCnt = syncCnt - 1)

Two things can be noted from Table 5.2; firstly, the assumption made in Equation 5.4 does not always hold true as initially thought; it is only valid when the final packet received in the previous frequency is not in error. Secondly, there does not appear to be an efficient method for detecting when a literal of '1' needs to be subtracted from the error count as there are no unique set of variable characteristics when this occurs.

This means a modification of the processing algorithm is required, rather than adjustment to the statistics equations. This will involve checking what happens immediately after a channel change. If an error is detected, compensation for the error can be applied at this stage rather during final calculations.

5.6.4 Resolutions

A number of proposals have been made for the improvement of the processing algorithm in order to improve the clarity of the derived statistics, however there need to do so can be debated. Firstly, the value for total number of packets transmitted is already in error and no feasible resolutions have been found. It is known that this number will be higher than the value calculated from the processed log file. The availability of the exact transmitted packet count will lower the error

percentage, assuming that all extra packets are correct; this is yet another unknown value.

Secondly, the assumption of an error at the start and end of a particular frequency, in the worst case, will cause a deviation of two in the error count. As the total number of transmissions increases, this factor becomes less significant. Furthermore, the purpose of these error statistics was to assess the quality of the radio environment in the operational area of high voltage equipment, with a view to discover if broadband communication was viable within the 2.4 GHz ISM band. Small deviations will not significantly affect such analysis in a detrimental way.

What is significant however is when multiple errors are unaccounted for; for example if multiple erroneous seed values are loaded into the PRNG and the subsequent packet additionally fails, these two events will only be detected as a single error. This will cause the error percentage to be calculated as approximately 50% in the majority of such circumstances; whereas a value between 50% and 100% would be more realistic.

Conversely, it could also be suggested that any error percentage of over a fraction of a percent could be treated in the same manner; thereby avoiding the absolute need for precise calculation as suggested above. The theory behind this idea is that a channel over which real-time video needs to be transmitted needs to suffer the minimum of interference. Therefore if a channel is found to suffer from 50% error based on the current algorithm it would not be selected for use. Improving the firmware and algorithm to detect seed errors will only serve to increase the detected error; the choice of channel would not be affected.

These improvements were initially going to be implemented prior to a second round of assessment, however based on the availability of alternative computer platforms the communication development progressed in another direction as discussed in Chapter 8.

5.7 System Stability Testing

It was important to characterise the radio environment using the developed equipment and therefore simultaneously assess the system's stability. This was done

with the following wireless configuration:

- Start frequency: 2.400 GHz
- Channel spacing: 200 kHz (1 minute dwell time, 417 channels)
- Channel spacing: 400 kHz (2 minute dwell time, 208 channels)
- 500 kBaud
- 4 Byte Sync Word
- 8 Byte Preamble
- Minimum Shift Keying

It can be seen that the dwell time and number of channels have been altered. They have been adjusted to allow the entire 2.4 GHz ISM band to be examined in a reasonable time-frame; in this particular case seven hours. The experiments were conducted over a period of 18 and 23 hours for the 200 kHz and 400 kHz cases respectively.

Figure 5.8 and Figure 5.9 show the results of these two experiments. It can be immediately seen that results from the second experiment were more informative; the clarity of the first experiment has been severely affected by the high trace which occurs at around 2.453 GHz; nevertheless conclusions can be made.

Figure 5.8 and Figure 5.9 have some surprising similarities. The number of synchronisation attempts appears to peak at identical frequencies; despite the fact that these experiments were carried out over a 48 hour period, with each cycle lasting seven hours. Furthermore, the interference points appeared at regular intervals. As it is highly unlikely that external interference sources would be found at such intervals, a fundamental problem existed elsewhere in the system.

This problem was traced to a known issue which had been documented in the errata for the CC2500 device [71]. It appears the CC2500 suffers from degraded performance at certain frequencies; specifically multiples of the crystal. This problem exists in all versions of the transceiver and there were no temporary work around solutions. The only advice provided was to avoid the use of the listed frequencies.

The frequencies which are known to suffer signal degradation have been marked on the graphs; seven such frequencies exist: 2405, 2418, 2431, 2444, 2457, 2470 and 2483 MHz. However these alone do not explain all frequencies which have
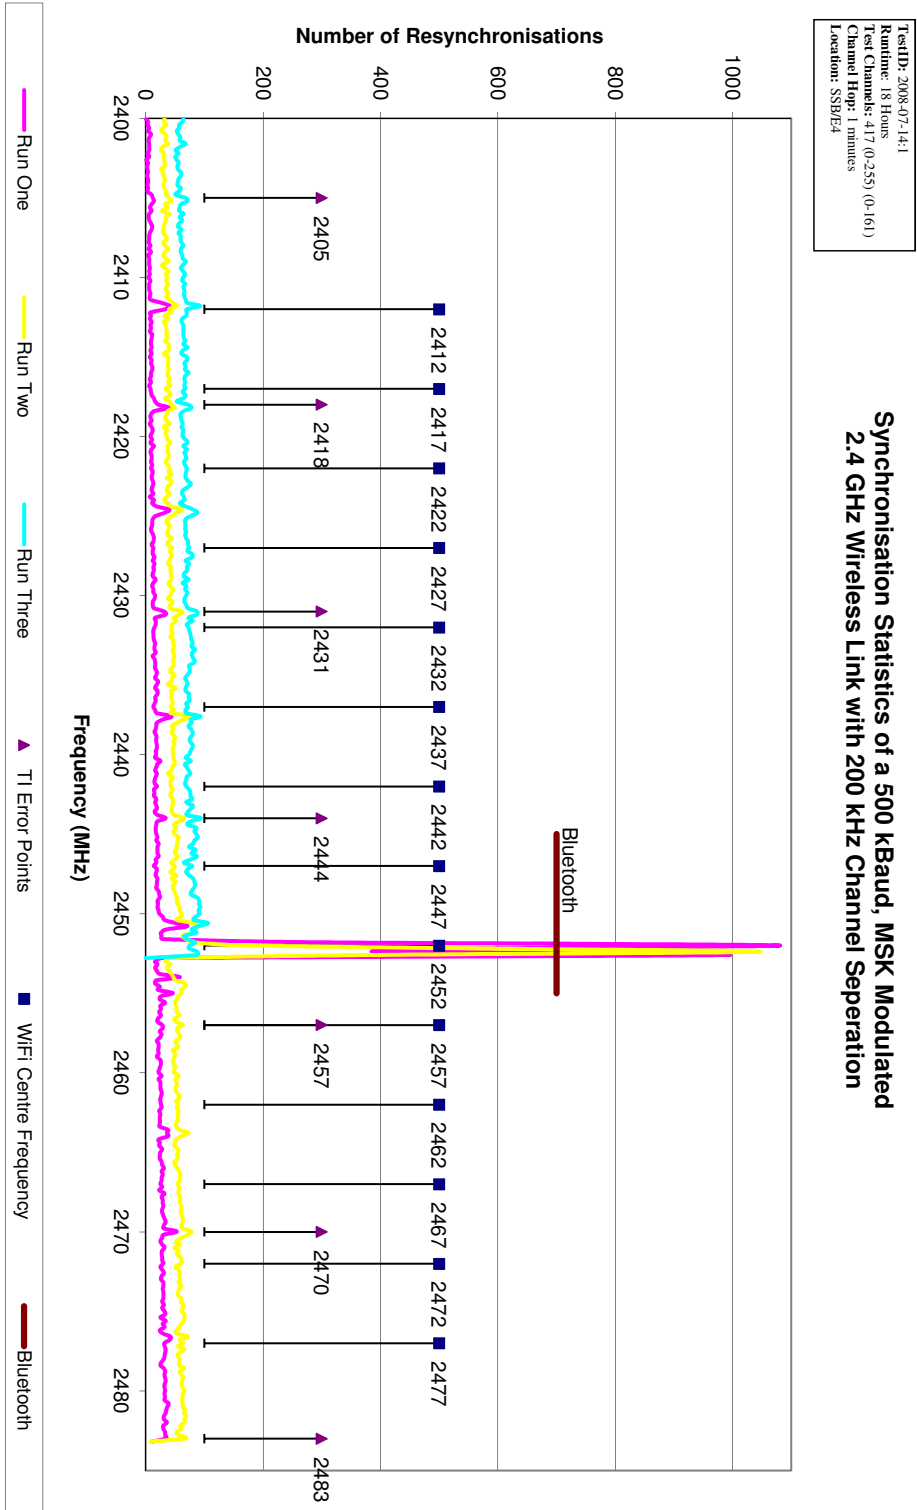
Figure 5.8: 18 Hour 200 kHz Channel Spacing Experiment
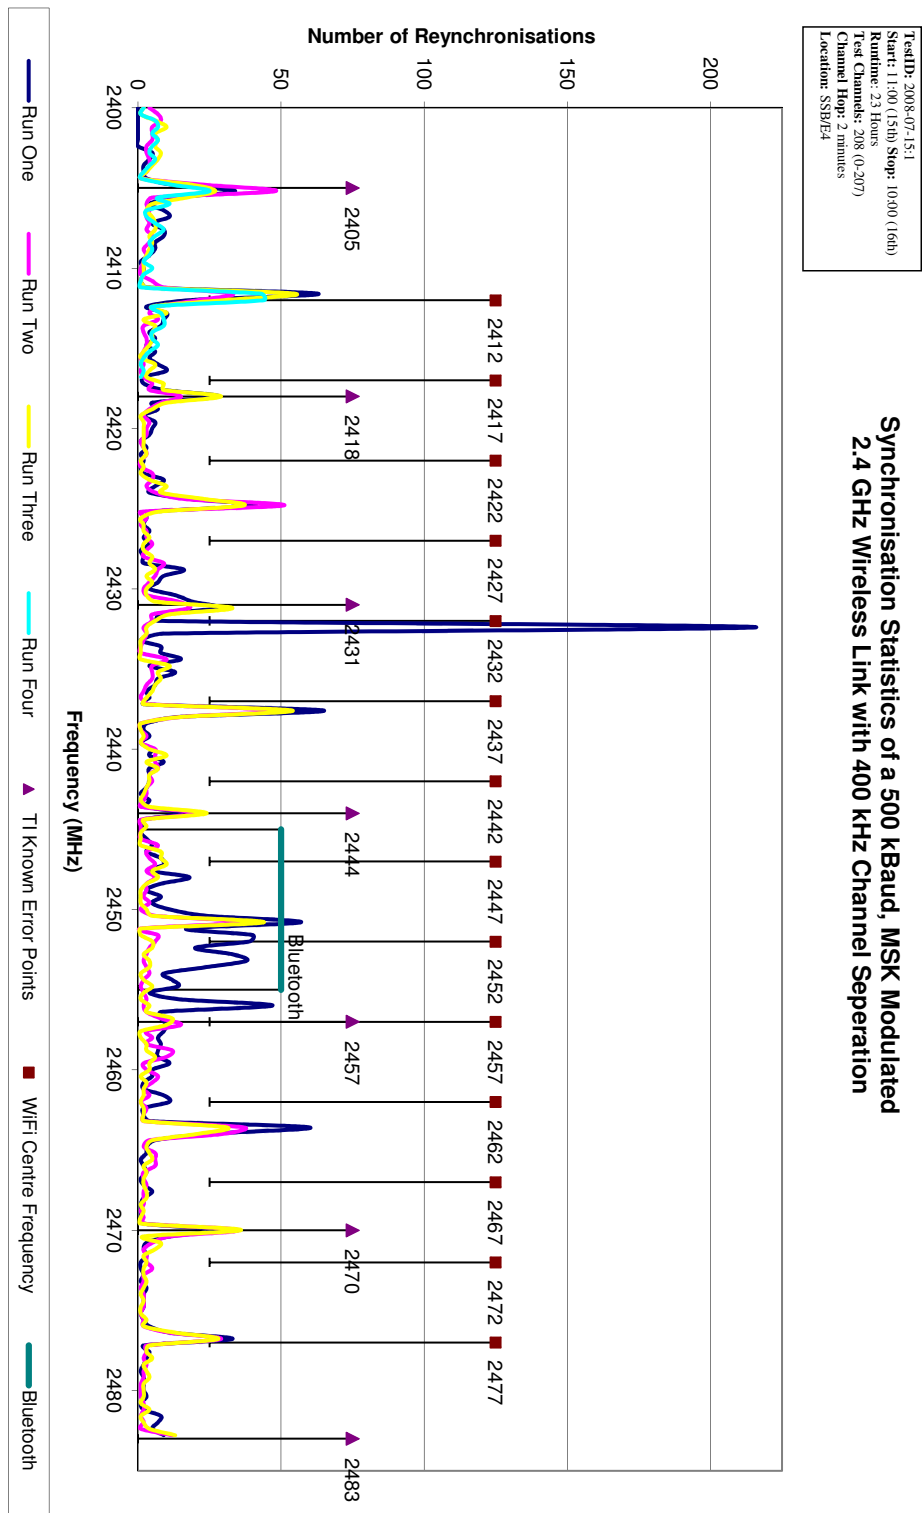
Figure 5.9: 23 Hour 400 kHz Channel Spacing Experiment

been shown to be degraded. Two other popular uses of the 2.4 GHz ISM band are Wi-Fi and Bluetooth; these have additionally been marked on the graphs. Wi-Fi channels have been marked using their centre frequency, while Bluetooth has been added as a 10 MHz band due to it channel hopping scheme.

When these three frequency groups are mapped onto the recorded data, the majority of degraded frequencies either directly coincide or lie near one of them. With reference to Figure 5.9 however, there are two exceptions; around 2.425 GHz and 2.465 GHz. As each Wi-Fi channels occupies a bandwidth of 22 MHz, it is conceivable that these remaining error frequencies are due to wireless networks. Conversely, if this were the case, wideband interference would have been expected around the offending channel.

Although Figure 5.8 and Figure 5.9 share similar characteristics with reference to locations of interference peaks, the former has a number of notable differences. Highly noticeable in Figure 5.8 is the large spike at 2.452 MHz, which then reduces the impact of other zones of interference; although other peaks are visible under close inspection. However, it also noticeable that when utilising a 200 kHz channel spacing, the base error rate is significantly above zero at all frequencies; in contrast to the 400 kHz channel spacing case.

To reiterate, this wireless link utilised MSK as the modulation format. This scheme utilises two frequencies that at 1/2 the data rate apart, i.e. 250 kHz. Therefore, if 200 kHz channel spacing is used and given the fractional delay in the receiver switching to the next channel, interference in the receiver is very likely. This situation exists because of the sequential (as opposed to pseudo random) manner in which frequencies are assessed. As the channel spacing used in Figure 5.9 (400 kHz) is greater than the spacing of frequencies used by the MSK scheme, a subsequent channel does not affect a former channel.

Note the larger than average trace in Figure 5.9 at 2.432 GHz; data relating to this peak was captured during the first cycle of the frequency set. This experiment was started 11 am, meaning the first cycle of the frequency set would have completed at around 6 pm. This seven hour block relates to the period when the university building is mostly occupied. Therefore, assuming that the university uses Wi-Fi channel five, a high Wi-Fi traffic rate would be expected during this time. University occupancy can also account for the increased activity seen in the

Bluetooth band during the same period.

5.8   Summary

In this chapter the Radio Frequency (RF) environment in the United Kingdom has been discussed before a system which would allow the 2.4 GHz RF band to be assessed under high voltage conditions is proposed. The development of a multi-part system is detailed along with the reasoning behind its implementation; to allow flexibility and safe working, while enabling maximum RF coupling. A spectrum analysis methodology is noted along with its implementation across the design hardware devices. The use of a PseudoRandom Number Generator (PRNG) allows a wide range of bit patterns to be assessed with minimal processing and storage requirements. A frequency change synchronisation method has been designed to allow a transmitter and receiver pair to assess the entire band automatically after initialisation.

A flexible data logging mechanism has been designed whereby logging output is proportional to the magnitude of the error seen, which thereby allows errors to be assessed in numerous ways. An analysis algorithm has been specified and a computer application was created to process the log files accordingly. Limitations of the algorithm have been noted but have been concluded to have little bearing on the selection of communication frequency.

Stability testing results are presented for two differing radio chipset configurations; channel separations of 200 and 400 kHz. Results indicate that the latter is more appropriate for long term testing as adjacent channels do not interfere with each other during the channel change operation. Frequencies at regular interval are seen to have higher than average error rates, noted in both cases and over multiple frequency sweeps; these are potentially related to documented chipset errors or other users of the 2.4 GHz band. However as these higher than average errors have been repeatedly seen, they can be accounted for in future test sequences.

With the development of a 2.4 GHz test system complete, the next challenge is to assess the 2.4 GHz ISM band in a high voltage environment; this is discussed in detail in the next chapter. The results from these measurements are crucial, as it will determine what type of communication system would need to be implemented

on the final robotic system. If it is revealed that 2.4 GHz transmissions are no more susceptible to corruption in high voltage environments than in office environments, then a similar office type equipment can be used on the final system. However, if is is shown that RF transmissions are completely unreliable under such circumstances, then the previously noted freespace optical system would need to be investigated.

This page is intentionally blank.

Chapter 6

Preliminary High Voltage Testing

The testing detailed in the previous chapter showed fairly good results; the equipment appeared to be stable and the majority of error points were located at frequencies used by other ISM devices. The next challenge was to test the equipment in a high voltage environment, specifically in the vicinity of a high voltage line.

6.1 Tower Representation

A representation of a single conductor high voltage line was configured in the university's high voltage laboratory, shown in Figure 6.1. The structure left of middle is a representation of a tower, while below centre is a humanoid representation relating to other working being carried out in the department. The transmission line was simulated using 28 mm copper tubing suspended using rope tensioned to the high voltage laboratory's walls. The simulated line was connected to the facility's Phoenix AC Test Set and energised to 231 kV during experimentation.

Figure 6.2 depicts a close view of the transmission tower representation; the transmitter was mounted at the point indicated by red tape, this is approximately 3.6 m above ground and approximately 2 m from the line. The receiver was placed at ground level approximately 3.6 m from the tower. Figure 6.3, 6.4 and 6.5 depict data collated during experimentation with this setup.

Figure 6.3 details the error experienced at frequencies whilst the test line was energised. Frequencies of possible interference sources have been marked as have
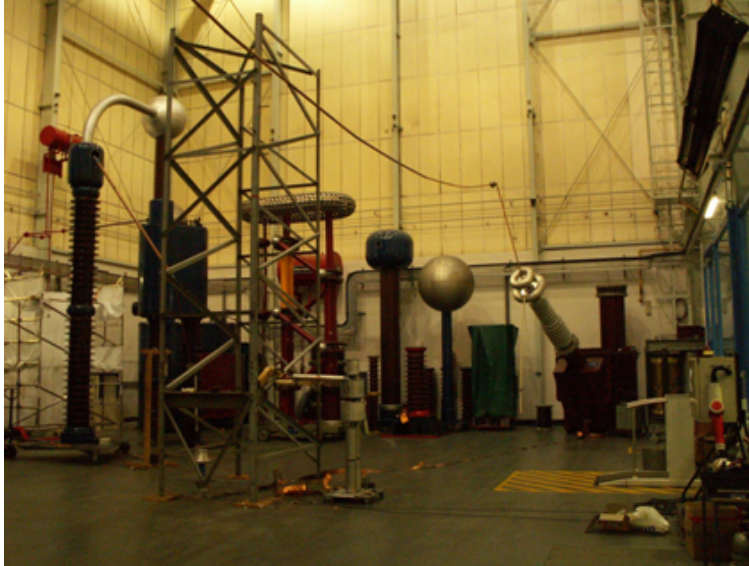
Figure 6.1: Overview of high voltage laboratory configuration
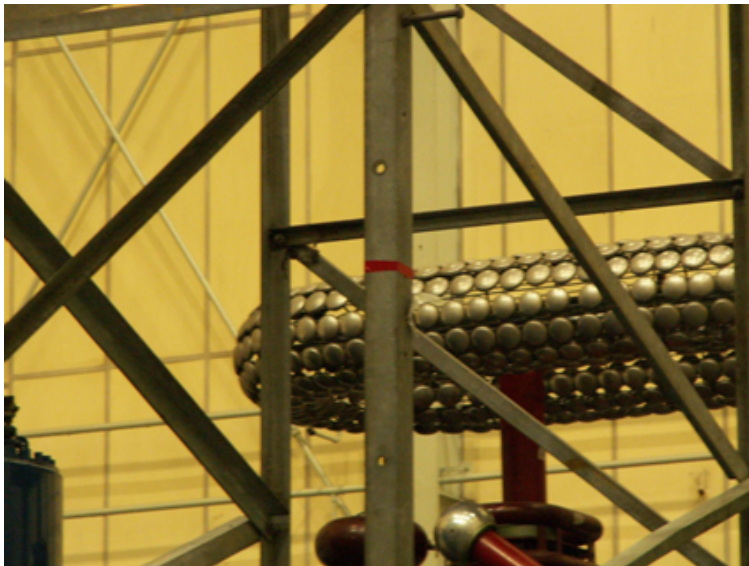
Figure 6.2: View of transmitter mounting point

the frequencies known to cause problems on the TI chipset. It can be seen that although a number of interference points coincide with possible interference locations, a number do not. It is important to ascertain if this interference is due to the energised line or background radiation. Nevertheless, even if it is discovered that the unaccounted for points are due to the energised line, it can be seen that there are enough interference free zones to successfully implement a frequency hopping spread spectrum system.

There is however one factor that needs to be considered; the suitability of the location of the transmitter. While it will experience an electric field and corona sources from the energised line, the device is effectively attached to a grounded structure (the tower). As such it is operating in a substantially different way to how the final robot would operate, i.e. an electrically grounded position rather than electrically floating. This ground structure would provide a level of shielding to the transmitter which would not be available to the robot.

This shielding could be contributing to the low error rates seen. This is because any radio interference generated by the line is more likely to be 'attracted' to the grounded tower than the transmitting antenna. Further evidence of this is visible when these results are compared to those obtained during stability testing of the hardware, Figure 5.9, where it can be see that there are fewer locations (frequencies) of error when the transmitter is grounded. Furthermore, the antenna's ground plane (which is the transmitter's die-cast box) is electrically linked to the tower structure, thereby creating a large ground reference and potentially allowing the antenna to perform better.

Figure 6.4 aims to resolve the debate on the level of background interference. It has been advised that the high voltage laboratory is, in reality, not very well shielded from external radiation. This graph depicts results obtained from an overnight assessment of the high voltage laboratory when the line was not energised: during this time the frequency range was examined five times.

Two things become immediately apparent. Firstly the radiation pattern seen is very similar to that seen in Figure 6.3. Secondly, on the whole it appears that the same frequencies are affected on each cycle of the frequency range. This repetitive interference nature suggests a number of possibilities; the errors seen are due to chipset errors, a constant background radiation, or periodic background

radiation with a periodicity matching the frequency cycle rate used. It is unlikely that the chip is to blame for interference at frequencies not documented as being error prone, thus additional interference is most likely due to varying background radiation.

In order to resolve this discrepancy the radio environment was assessed using a portable spectrum analyser, data for which is shown in Figure 6.5. This graph additionally layers the interference traces from Figure 6.3 and averaged values of data seen in Figure 6.4. It can be seen that while the data sets recorded using developed hardware appear similar, spectrum measurements additionally indicate background energy at other frequencies. It is important to note that spectrum analyser measurements were captured over a short duration at the start of the day.

Based on this analysis, it is clear that there are interleaving interference processes; those that are caused by the TI chipset, those caused by other 2.4 GHz ISM devices (Wi-Fi, Bluetooth) and background radiation. It is the third that is most problematic and also the most complex to trace due to its varying nature.

Detailed study of background radiation is required in order to ease matters. An attempt was made to achieve this by means of a longer period of spectrum analyser assessment; namely overnight, results of which are presented in Figure 6.6. Further resolution was attempted via a 24 hour assessment using the developed hardware, shown in Figure 6.7.

The data presented in Figure 6.6 was collected using the spectrum analyser's 'maximum hold' feature, therefore the values recorded were the peak noise values detected during the measurement period. It can be seen that the peak noise fluctuates in the range 2.400 - 2.465 GHz, remaining relatively noise free thereafter. If the peak levels were constant, more interference would have been expected in previous datasets. This reiterates the fact that background noise levels seen in the high voltage laboratory are not constant.

It is important to note that a few of the spectral peaks are in line with Wi-Fi centre frequencies: 2.412 GHz, 2.432 GHz and 2.442 GHz. Although according to the high voltage laboratories' manager there were no university Wi-Fi access points in that part of the building at the time. Hence the origin of this radiation must be external to the laboratory and adds weight to the premise that the laboratory is
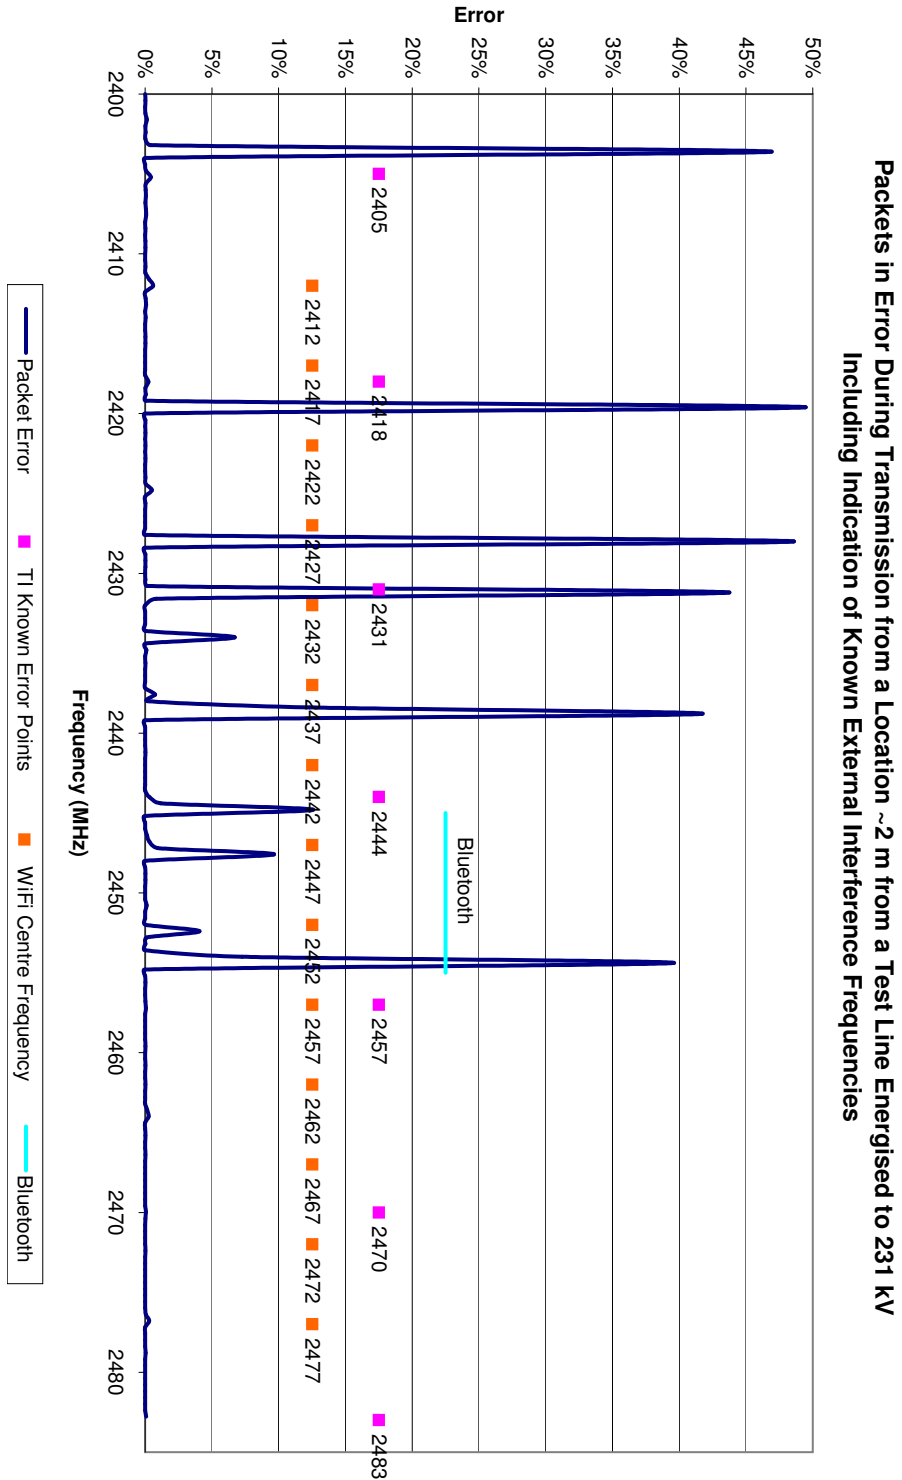
Figure 6.3: Analysis of packet errors under energised conditions
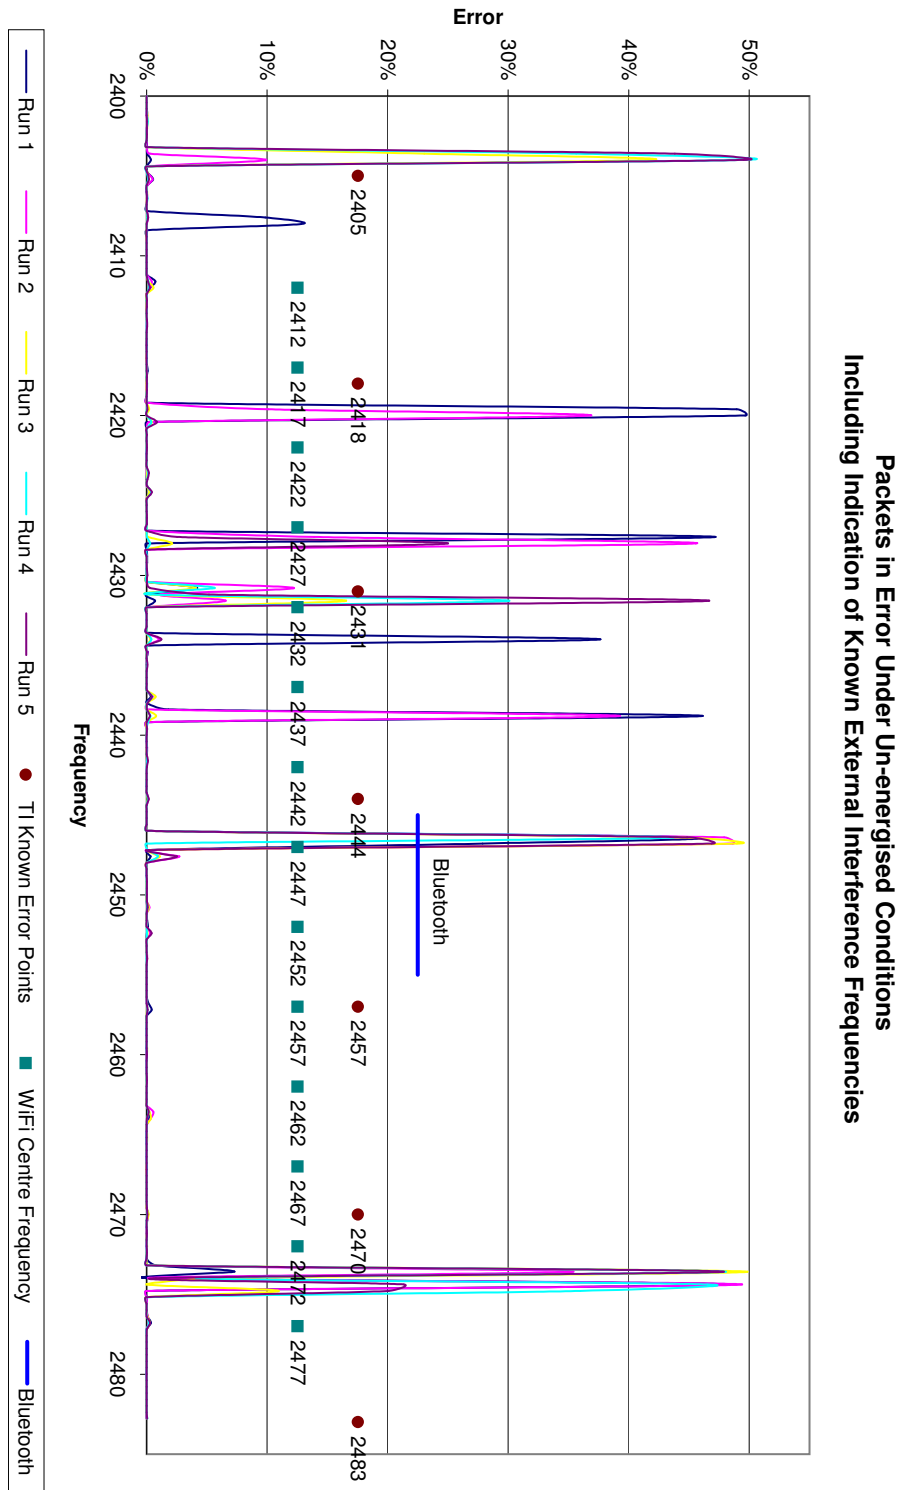
Figure 6.4: Analysis of packet errors due to background interference sources
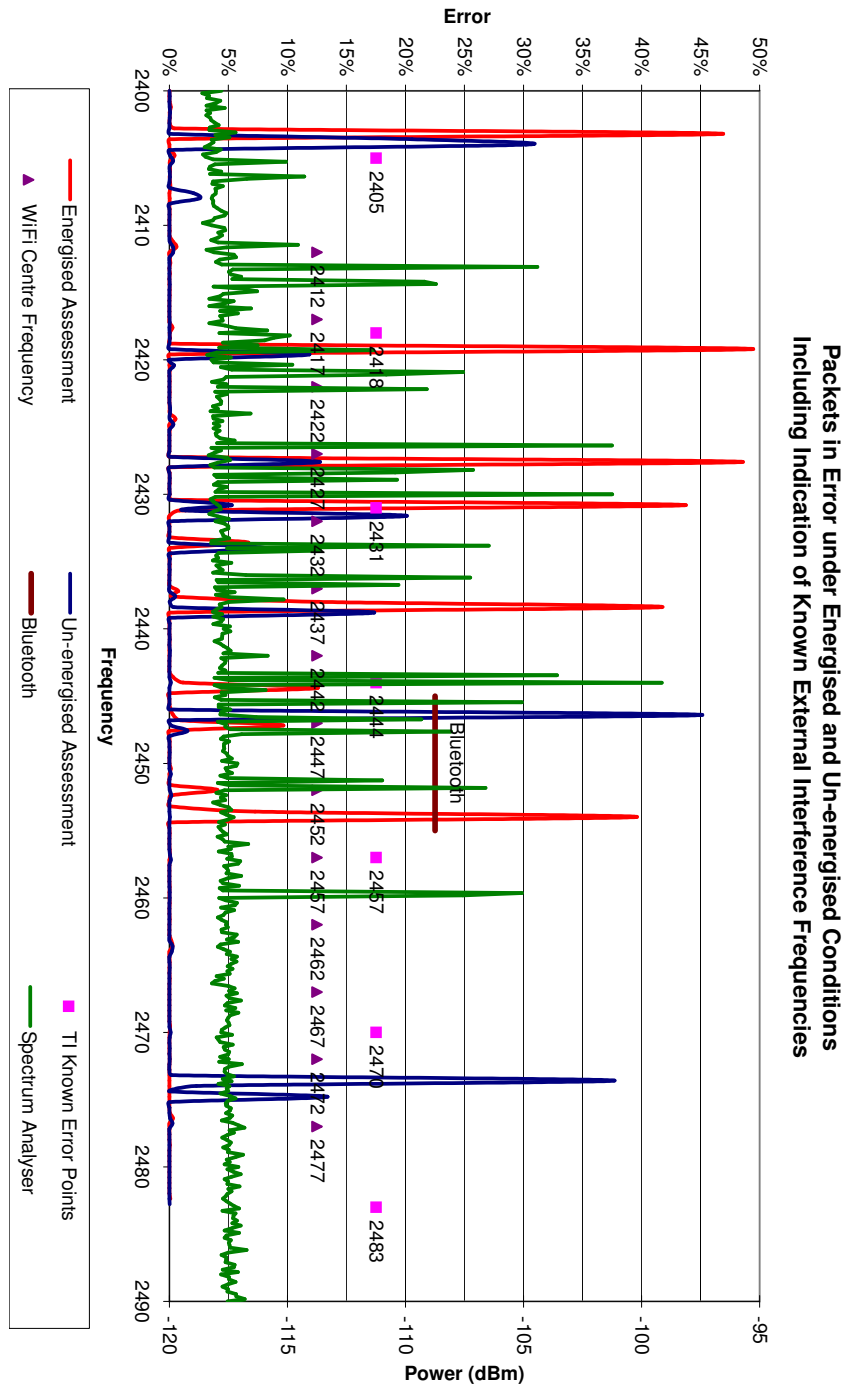
Figure 6.5: Comparison of high voltage analysis and background analysis using developed hardware, with background spectrum analyser measurements
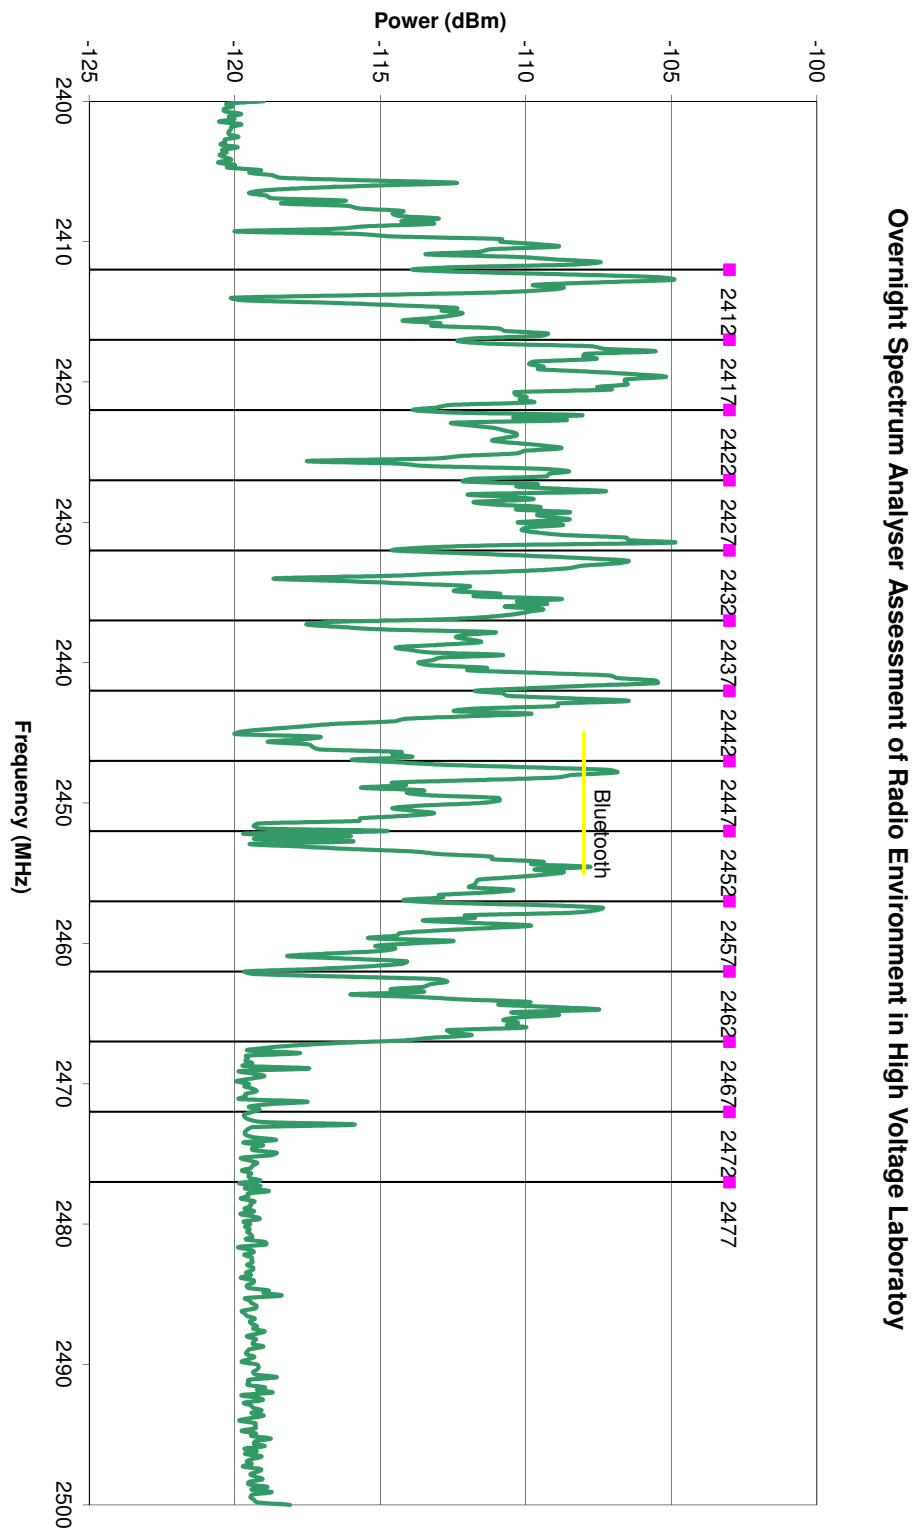
Figure 6.6: Overnight spectrum analysis of radio environment
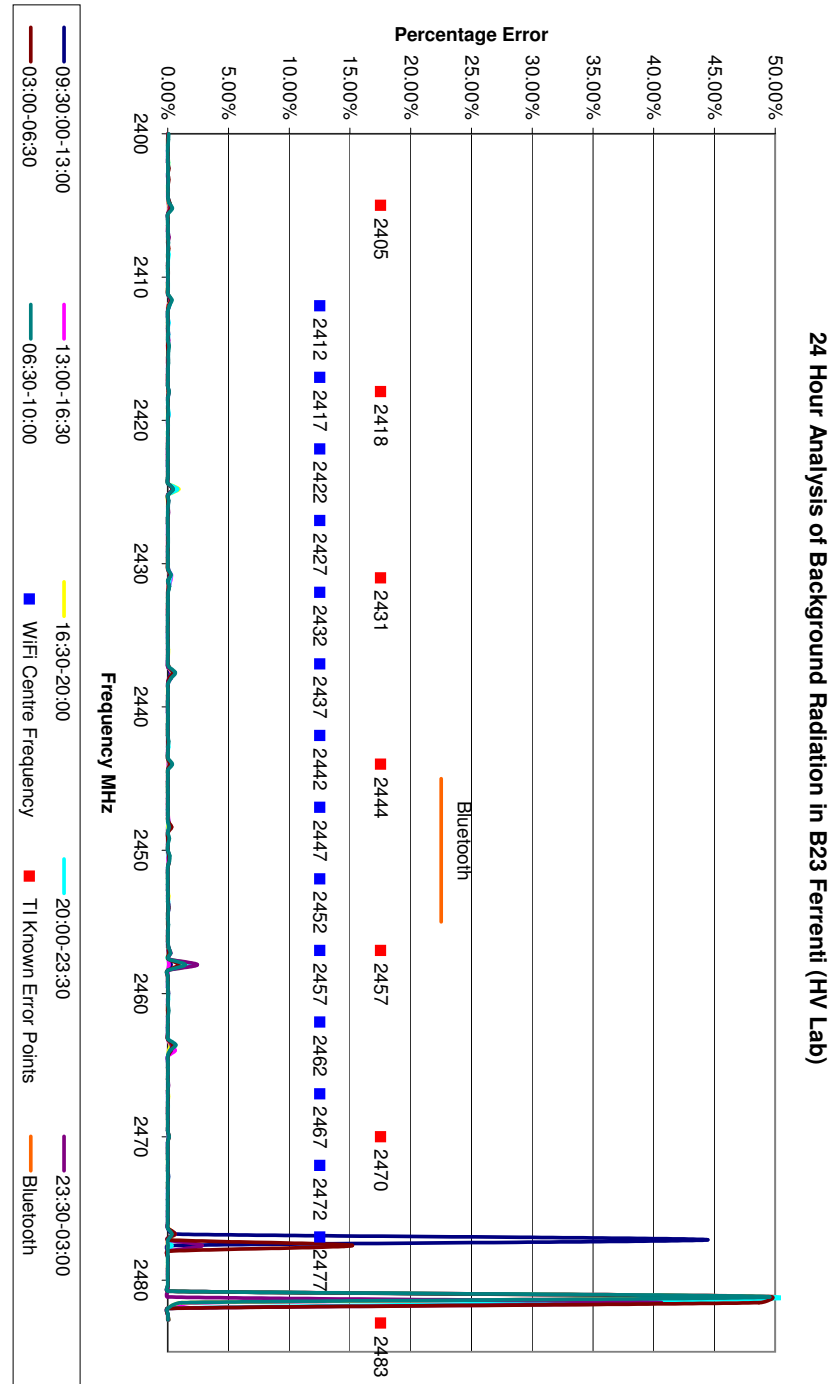
Figure 6.7: 24 hour background analysis in the high voltage laboratory using developed hardware

not well shielded to external radiation.

These facts are seen again in Figure 6.7, which depicts the laboratory to be almost noise free barring two frequencies around 2.48 GHz. Remarkably, even frequencies that should show degradation based on chip errors appear to be fairly clean; although detailed inspection reveals that low levels of interference were recorded.

In order to confirm the variable nature of the background noise environment over time, a further assessment sequence was undertaken. While the first spectrum measurement (Figure 6.5) was a single snapshot of the spectrum and the next (Figure 6.6) captured the maximum noise values during overnight hours, this assessment was performed during the day and recorded complete range measurements approximately every four minutes[1] over a period of 6.5 hours. This allowed noise measurements to be compared against both time and frequency simultaneously. Measurements were made with a 2 dBi whip antenna in the range 2.4-2.5 GHz.

Figure 6.8 shows the result of this assessment. In this graph colour represents the noise level with reds and oranges representing higher relative noise and hence potentially more interference. Due to the discrete nature of the datasets recorded there has been some interpolation in the creation of the graph.

Regardless of what has been indicated by the previous results, Figure 6.8 shows two zones of near continuous wideband interference. These coincide with WLAN Channels 6 and 11. As it was known that there are no wireless access points in the vicinity, this unwanted spectral noise must have emanated from a source external to the laboratory. Closer inspection reveals that there are additionally spots of low level interference in other parts of the spectrum. The possible sources of these spurious emissions are numerous, for example cordless phones, Bluetooth, ad-hoc WLAN and microwave ovens to name a few. Some of these types of equipment are know to be present nearby but not directly in the test laboratory. Due to working methods in the test facility, their complete removal from the test site was not possible.

The datasets presented above describes the environment to be both noisy and noise free over the measurement period, consequently it is quite difficult to accurately comment on this state of affairs. In order to properly assess the background and thereby arrive at a better average, longer duration tests will be required. British

---

[1]This was the spectrum analyser 'sweep time'.
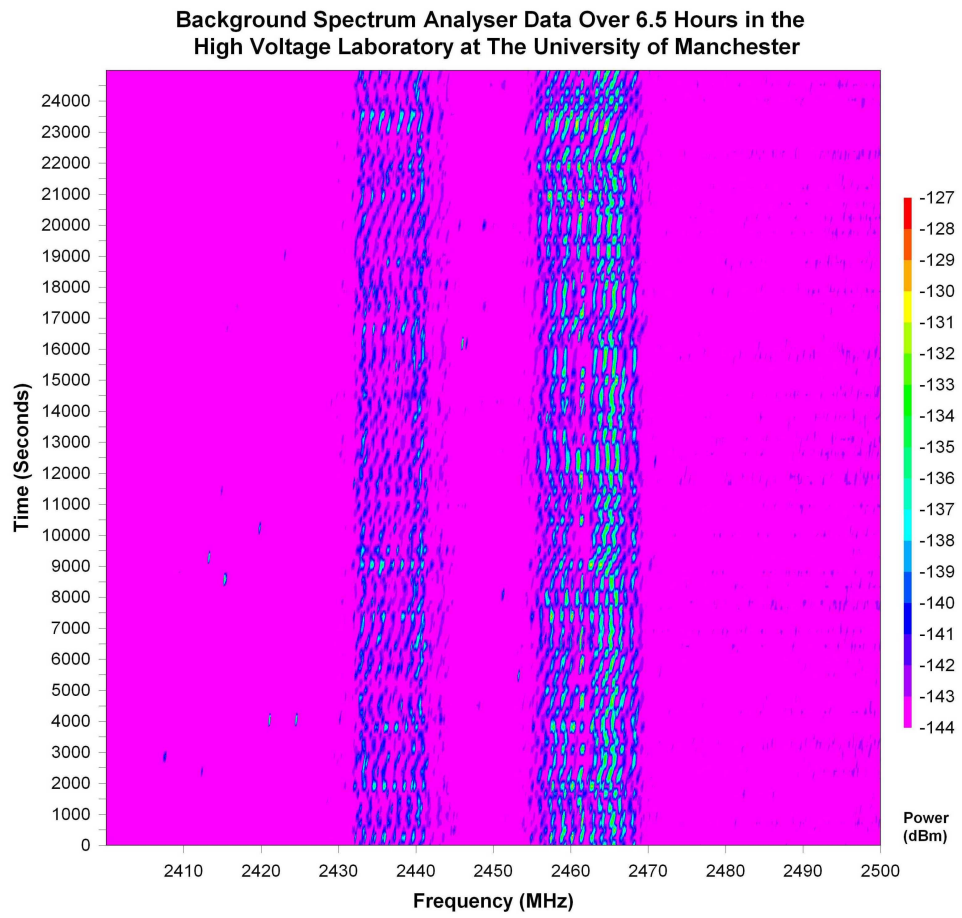
Figure 6.8: Day time RF analysis of high voltage laboratory over 6.5 hours

Standard 5049[63] actually indicates that measurements of such interference in the field should be conducted over a period of a year, to suitably account for widely varying atmospheric conditions.

While this may or may not be a cause for the varied background radiation recorded, assessment over such a duration would allow interference caused by external equipment to be adequately averaged. Although as other users will need access to the laboratory, it is not practical to undertake a year long study. While a duration of one week is more feasible, it still may not be appropriate to 'block' the use of the high voltage laboratory in this way; a solution to this was not realisable. Therefore, while background noise statistics have only been collated over a short time-frame, the detail they contain should be adequate enough for comparison to high voltage test results.

6.2 Interference Assessment

Chapter 5 discussed the use license free wireless communication & its associated problems, the development & testing of spectrum analysis hardware and the creation of error analysis algorithms. The previous section provided a preliminary assessment of the background noise present in the high voltage test facility at The University of Manchester, together with a single, but basic, high voltage test scenario.

While the latter provided an early insight into possible radio interference sources, the experiment was fairly arbitrary and hence may not be easy to reproduce. What was required was a number of easily reproducible test scenarios; two of which were developed.

The first scenario saw the developed transmitter suspended halfway between a set of parallel plates, with the antenna parallel to the ground. The base plate was connected to Earth, while the top plate was connected to an AC Resonance Test Set. The plates were spaced approximately 1 m apart and the receiver was situated at ground level approximately 5.8 m from the transmitter and not in contact with the ground plane. The potential on the top plate was increased from 0 kV/m to 100 kV/m in 5 kV/m steps. This experiment's aim was to deduce the effect of field strength on interference in the 2.4 GHz ISM band. Figure 6.9 shows the
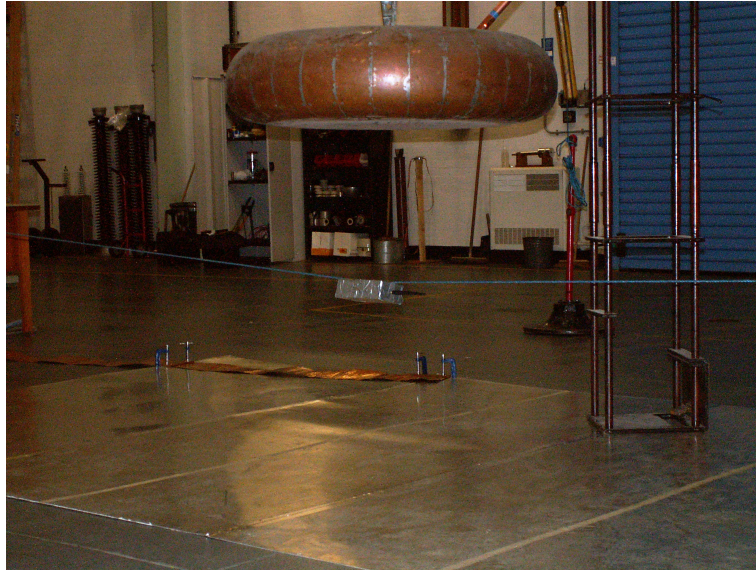
Figure 6.9: Parallel late field strength experimental set-up

experimental set-up.
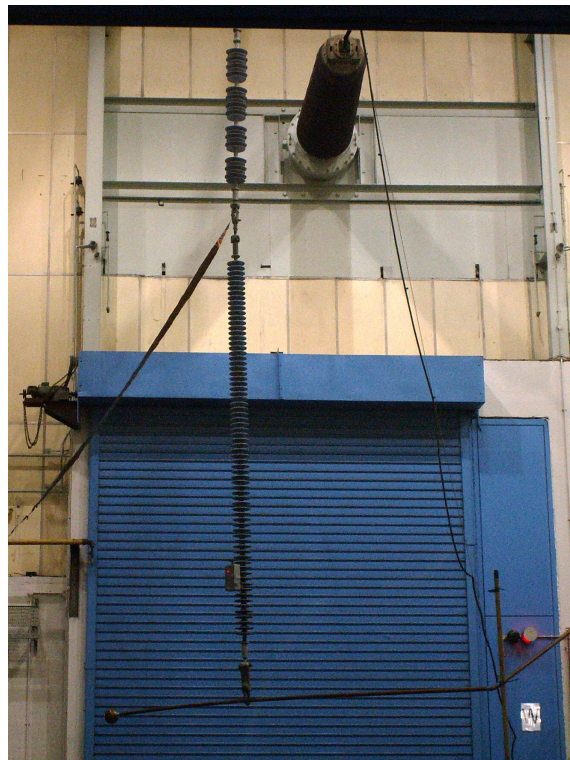
Figure 6.10: Experimental set-up for representation of single conductor high voltage line

The second scenario consisted of a representation of a single conductor high voltage line, it's aim was to remove the abstract nature of the test scenario. In addition to high field strengths, the robot will be exposed to partial discharges and corona around fittings. Here the transmitter was physically attached to a high voltage insulator at varying distances from the energised (hot) end. The transmitter was taped on the insulator with the antenna perpendicular to the insulator and its tip pointing away. The first measuring point was approximate 56 cm from the earth end. The receiver was placed on the ground approximately 8.7 m away. After each test the transmitter was moved approximately 25 cm closer to the live end; equivalent to six insulator sheds. The same high voltage source was used and the test potential was ~231 kV[2]. Figure 6.10 shows the experimental set-up.

In both cases data was recorded to a Secure Digital (SD) card for subsequent analysis using the algorithms detailed in Section 5.5. The data transmission format and collection & real-time checking strategy is reiterated next.

Each transmission was 19 bytes in length, consisting of a 4 byte sync word, 8 byte preamble, 6 byte payload and 1 byte address. The payload consisted of three copies of a 16 bit PRN generated on the transmitter. Once transmitter-receiver frequency synchronisation is achieved, one instance of the PRN from the payload contents of the first packet is set as the seed for the local random number generator. Thereafter, each packet has its three instances cross-checked before being compared to a value generated on the receiver itself. Successful matches result in an 'OK' symbol being written to memory, while cross-check or transmitter-receiver match failures result in the entire payload and the receiver generated value being written to memory.

The system developed used the MSK modulation scheme, which uses a bandwidth equal to half the data rate. The data rate used was 500 kbps, leading to a bandwidth of 250 kHz. In order to reduce inter-channel noise, channels were separated by 400 kHz; this could be reduced to 300 kHz (or less) to maximise frequency utilisation, while still maintaining adequate separation.

---

[2]231 kV line-to-earth is equivalent to 400 kV line-to-line as used on high voltage pylons.

6.2.1 Experimental Results

The primary information extracted during data analysis was a detailed account of error rates at each of the tested frequencies for the different test cases. Due to the volume of information extracted from the collated datasets, it is impractical to display them here. A table listing the types of data extracted from the raw datasets[3] can be found in Appendix A, while the data files themselves can be found on the attached DVD.
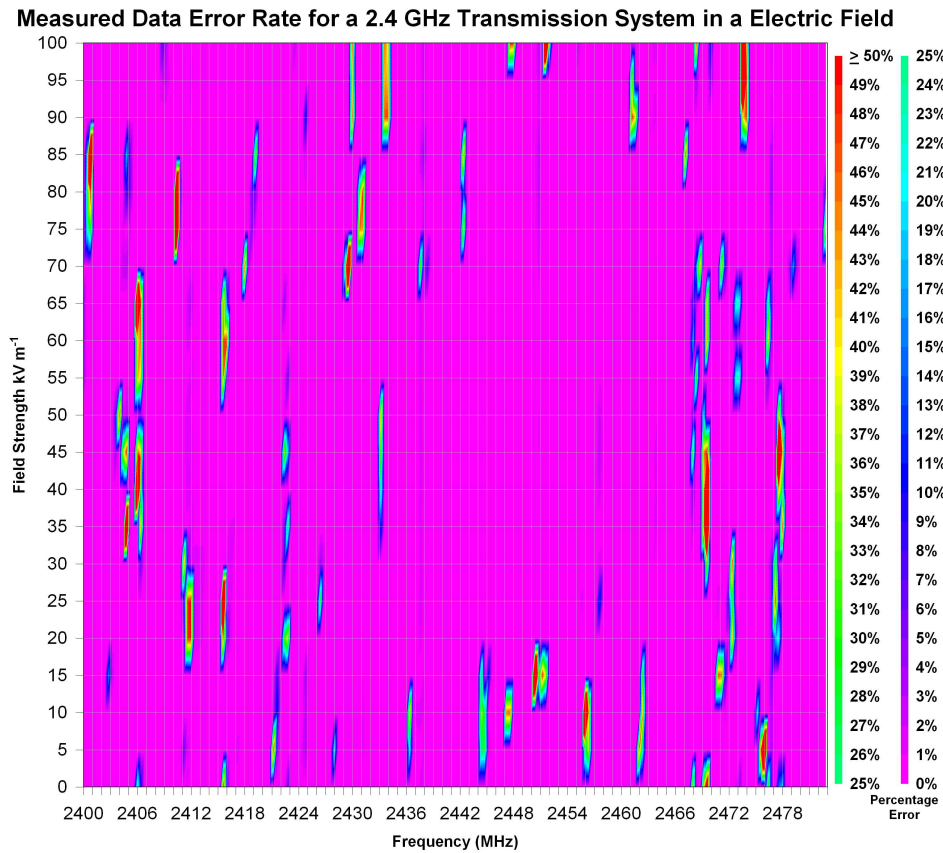
Figure 6.11: Error rate in each RF channel versus field strength

Each dataset produced a comparison of error rate against frequency in the absence of audible corona from the high voltage test rig. While possible to analyse in isolation, better conclusions can be drawn if all collected datasets from a single experiment series are compared simultaneously. Figure 6.11 shows such a comparison for the first scenario, presented as a surface plot. Here the colour represents

---

[3]This refers to the data that was recorded to the SD data card.

the error at a particular frequency and field strength. A gird at 1 MHz intervals is shown for clarity during analysis.

It should be noted that the field strengths presented in Figure 6.11 are idealised by neglecting the distorting effect the transmitter would have on the localised field. It should also be noted that there is some level of interpolation between the plot points due to the discrete nature of collated datasets. Due to the nature of the analysis algorithm (Subsection 5.6.4) the maximum error rate is capped at 50%[4]. However, as channels with error rates this high would not be selected for use under most circumstances, not being able to differentiate higher errors rates is acceptable.

Field strength dependence would result in horizontal bands of similar colours across all frequencies, the colour only vertically varying based on the field. Frequency dependence would result in vertical bands, independent of field strength, possibly varying horizontally in colour based on transmission frequency. Figure 6.11 clearly shows that the errors were not field strength dependant, but were in part frequency dependant.

Partial frequency dependence was actually expected given the licence free nature of the 2.4 GHz ISM band, the known error frequencies on the transceiver used and from background noise analysis (Figure 6.8). However this cannot explain why errors are not clearly visible at known error frequencies[5]. Initial experiments with this transmitter/receiver pair indicated that there was a marked increase in the number of re-synchronisations[6] at or near most problem frequencies. However the majority of these experiments were conducted under unenergised conditions for either hardware testing or the assessment of background noise. The cause of this intermittent error was not probed further at this point as frequencies with significant interference should be avoided to maintain reliable communication.

As noted before, the grid superimposed on to the data presented in Figure 6.11 is horizontally spaced at 1 MHz intervals. The figure clearly shows that there are numerous zones where data can be transmitted with less than 2% error; for example between 2401-2402 MHz, 2439-2442 MHz and 2480-2482 MHz. Given that 1 MHz can support up to three of these 250 kHz channels when separated at

---

[4]errors at and above 50% are not differentiable from one-another.
[5]2405, 2418, 2431, 2444, 2457, 2470 and 2483 MHz
[6]In early experiments error percentage was not calculated.

300 kHz, 18 channels can be established within the previously specified zones.

While 500 kbps is not sufficient for real-time video transmission, 9 Mbps[7] is more than adequate. Therefore there is sufficient scope to develop a system which aggregates multiple low data rate channels into a single high data rate channel. This has the benefit that the selected channels do not have to be adjacent and can be spread throughout the 2.4 GHz ISM band. Adding frequency hopping to this, while complicating the firmware design, will improve the ability to avoid interference.

Another interesting statistic that can be drawn from the captured data is the average bit error for packets in error, shown in Figure 6.12. This average is calculated by summing the total number of bits in error in the channel and dividing by the number of packets in error, rather than the total number of packets received. This provides an initial indication as to whether the errors are random or 'bursty' in nature, with a higher average indicating the later.

When analysing this graph it is important to note that only errors are shown; if the bits in error are considered over the total number of packets received the average number of bits in error per packet would be lower. This information allows an informed choice to be made as to whether forward error correction techniques can be used with confidence.

The next stage in identifying if errors are random in nature would be the assessment of the number of error bits within each packet. This needs to be displayed separately for each field in turn as there is no adequate way to convey four dimension data (field strength, transmission frequency, number of bits in error and frequency of occurrence). Figure 6.13 shows such a plot for the data gathered in assessed 65 kV/m field.

This figure shows major errors in about five frequency bands: 2411-2412 MHz, 2415-2416 MHz, 2469-2470 MHz, 2472-2474 MHz and 2476-2477 MHz. Comparing these frequency bands to those on Figures 6.11 and 6.12 shows fair to good correlation to points of error.

A burst error is signified by a number of adjacent bits being in error. Therefore, while this data cannot provide any concrete conclusions, it can be indicative of the

---

[7]18 x 500 kbps
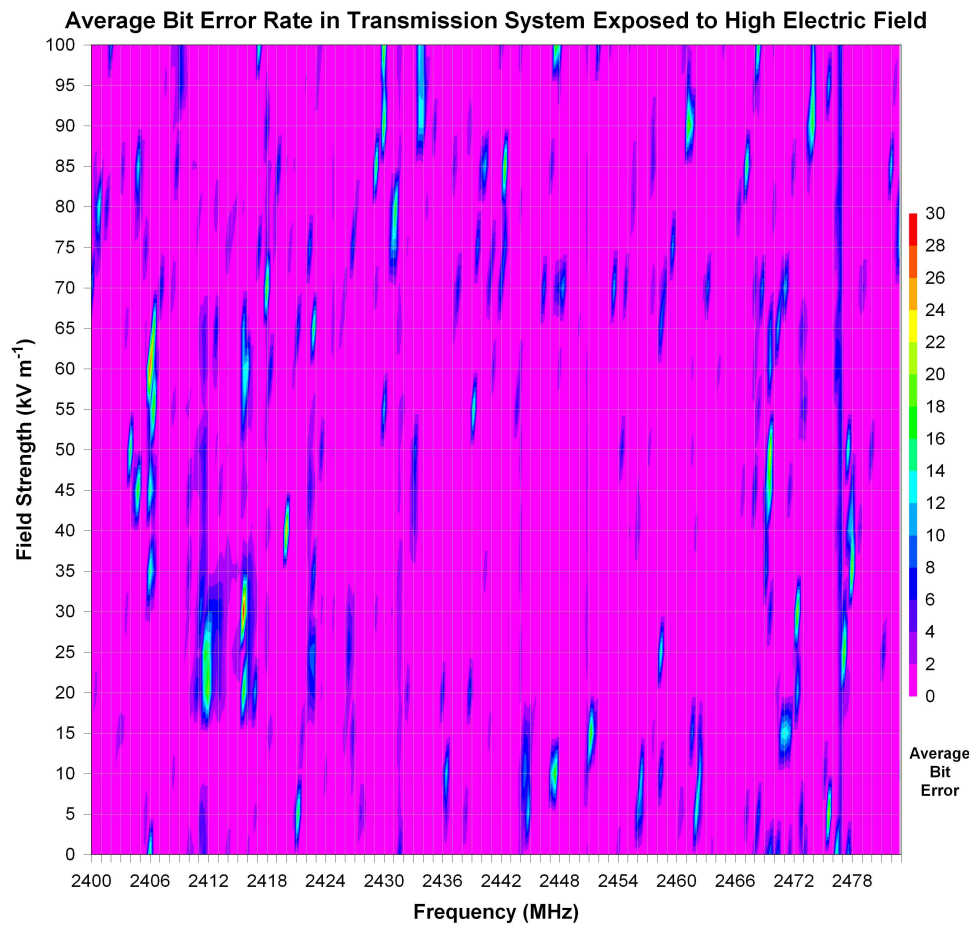
Figure 6.12: Average number of bits in error per packet in error for a transmiter located within various electric field strengths
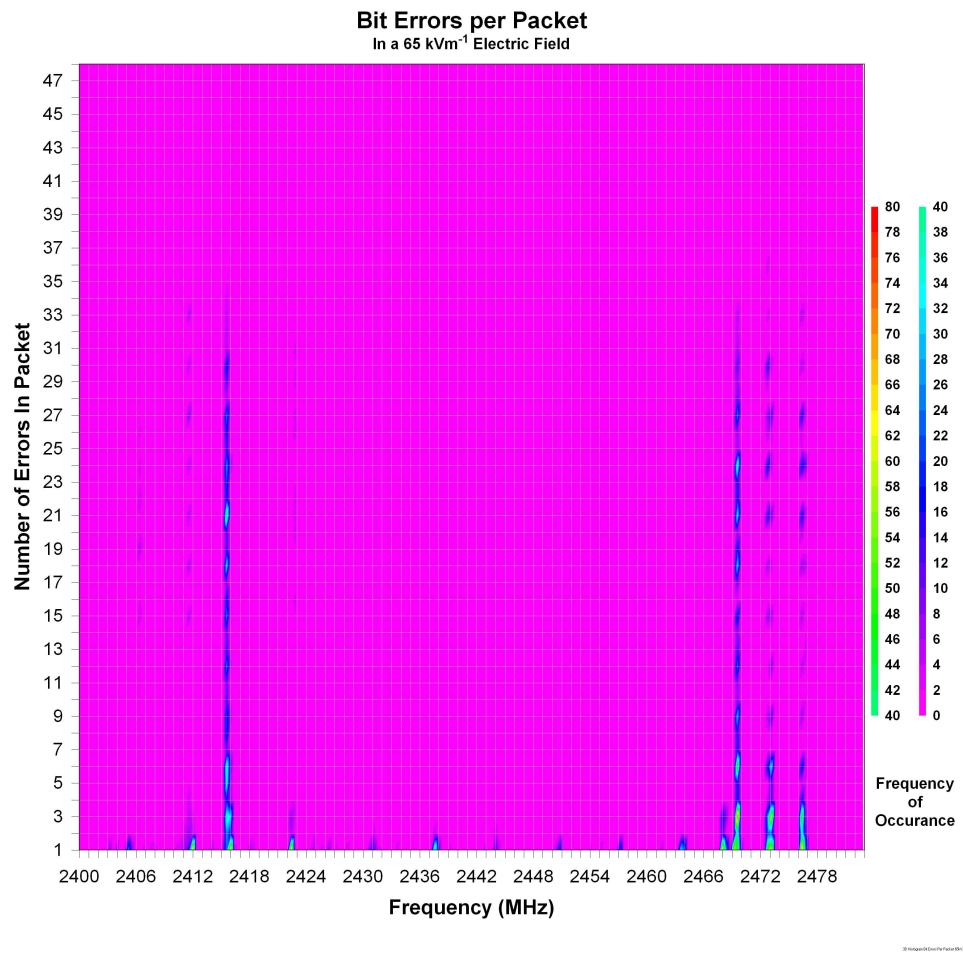
Figure 6.13: Bit errors per packet for a transmiter located within a 65 kV/m electric field

likelihood of burst errors. The packet payload length is 48 bits, thus the greater the number of bits in error, the higher the likelihood of the error being 'bursty' in nature. For example if 24 bits were in error, it is possible that the error could occur on alternate bits, a burst length of one, however this is not very likely. A better conclusion would be that the errors are grouped; a certainty if the number of bits in error is over 24.

The datasets from scenario two can also be presented in isolation however more information can be gleaned if presented simultaneously. Figure 6.15 shows a surface plot depicting the error rate at various frequencies at various distances from the Earth end of the experimental set-up. While the previous experiment experienced no audible corona during assessment, the high voltage rig used here had visible corona at the experimental voltage of 231 kV; shown in the composite picture in Figure 6.14. In an similar vein to Figure 6.11, the visualisation does possess some level of interpolation between data points; which is most profound between 'distance from the Earth end' measurements.

Figure 6.14: Composite image depicting corona near energised end of experimental rig.
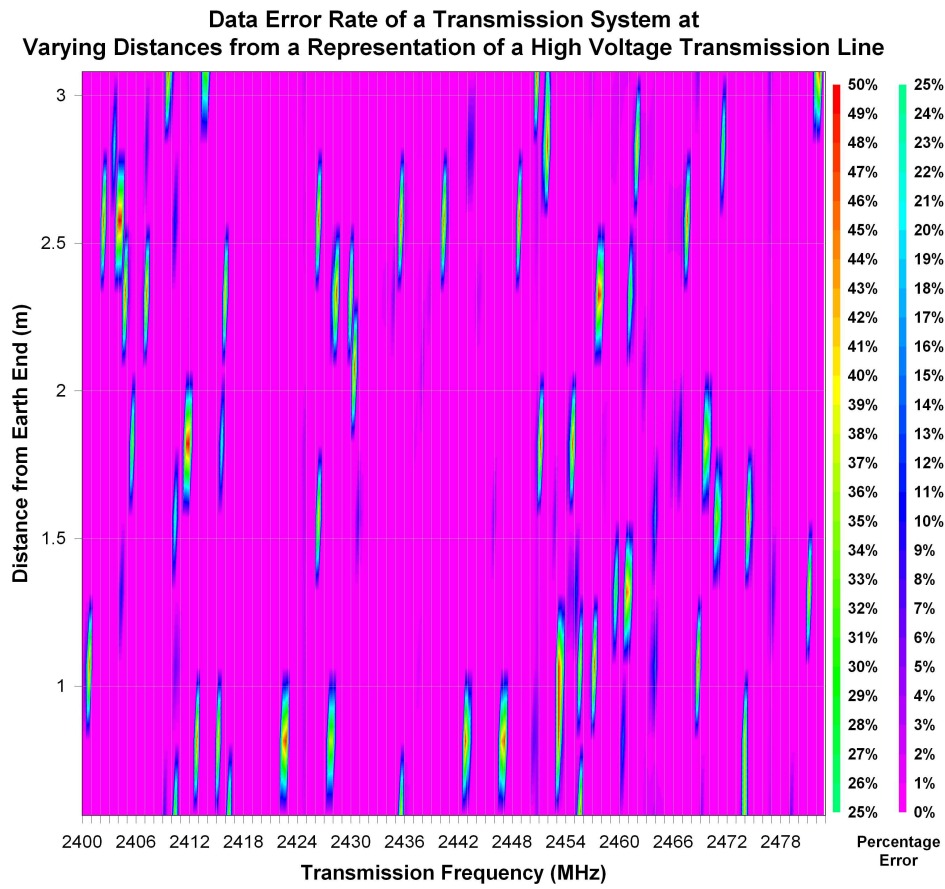
Figure 6.15: Measured data error rate of a 2.4 GHz transmitter at various distances from a representation of a high voltage transmission line

While this plot suggests that the environment is more error prone than with a high field alone, which can possibly be attributed to the presence of corona, closer inspection reveals that the error rate is actually less severe. This is indicated by fewer areas highlighted in red. In addition there appears to be larger bands with less than 2% error: 2417-2422 MHz, 2431-2434 MHz and 2437-2441 MHz. This availability bodes well for single channel broadband communication.

An important point to note with these representations is the granularity at which the data points exit. The frequency was the automatically controlled variable, which had a fixed spacing of 400 kHz between plot points. The distance from the Earth end was the manually controlled variable and was limited by the design of the insulator being used. A six shed separation was used between data sets. The error rate was the measured variable. Without this context it could be assumed that some interference phenomenon was occuring at regular intervals along the insulator string.

Interestingly, distances between 0.56-1 m from the Earth end appear to have the most consistent error readings in certain channels: 2412-2413, 2415, 2422-2423, 2427-2428, 2443, 2446-2447, 2454, 2460 and 2474 MHz. The other 50 cm distance bands do not show as greater a consistency in single channels. Additionally it was assumed that the areas close to the live end would suffer greater interference due to corona and partial discharge activity than areas near the Earth end.

However, judging by Figure 6.15, areas less than 1 m or greater than 1.5 m away from the earth end suffer from multiple zones of greater than 25% error. Between 1-1.5 m errors seems to be very small or non existent (for the most part) at frequencies below 2450 MHz.

One particular distance of interest is 81.3 cm from the Earth end. Inspection of Figure 6.15 shows that this distance produces red zones at a number of frequencies. Taking the error data from this test and extracting the bit error information produces the plot seen in Figure 6.16, which depicts the number of error bits within each packet.

In comparison to Figure 6.13, which is related to the first experiment, Figure 6.16 indicates that more frequencies appear to suffer burst type errors. This is despite the fact this test had the transmitter relatively close to the Earth end. These figures compare cases where the transmitter is exposed and not-exposed to a source
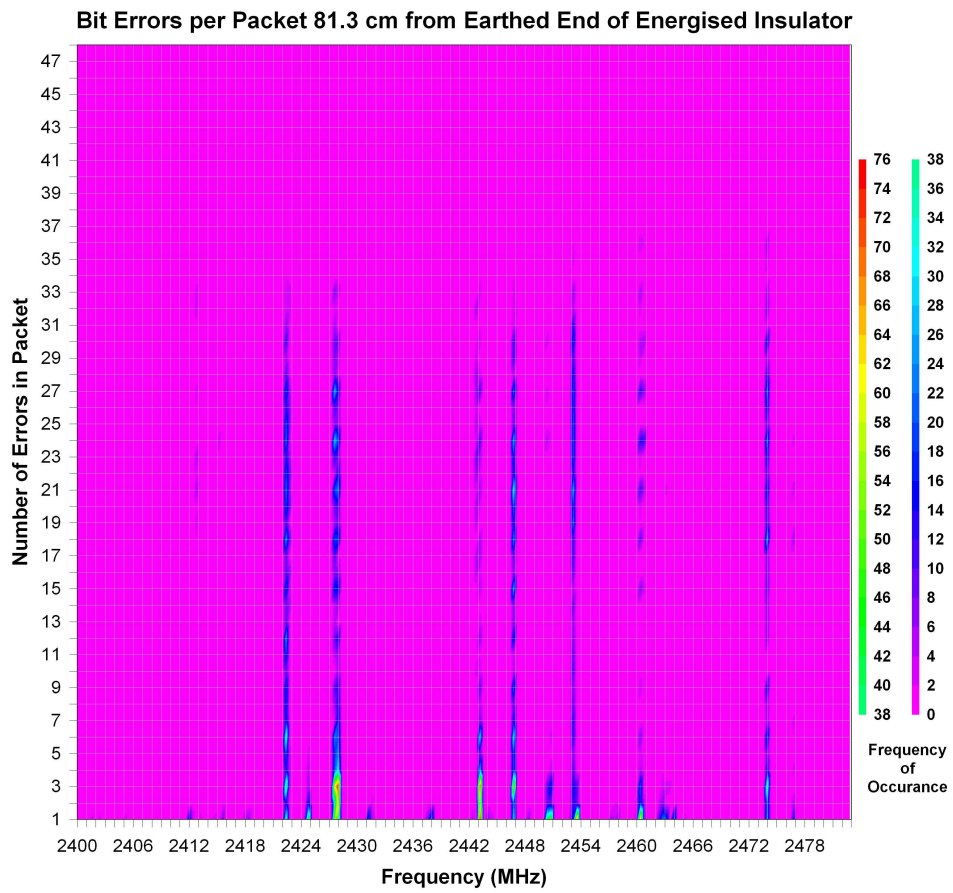
Figure 6.16: Bit errors per packet for a transmitter 81.3 cm from the earth end of a line energised to 231 kV

of corona discharge. This provides evidence that corona is potentially more detrimental to RF transmission than high field strengths alone, in that corona presence seems to increase the range of frequencies effected by transmission error.

6.3 Summary

Taking the equipment discussed in Chapter 5, a number of experiments were conducted under high voltage conditions. These included transmission from a location between parallel plates and from various locations on a composite insulator connected to a test line energised to 231 kV.

Results indicated that the error rate was more dependent on frequency than on either field strength or distance from the Earth end of an energised insulator. However in each case sufficient bandwidth existed for interference free communication.

Although other researchers have investigated the use of 2.4 GHz systems in HV environments, notably substations[43, 45, 46], results from either field trials or representative laboratory experiments have not been published. Therefore the extensive research presented here provides a good basis for further exploration of the practical uses of 2.4 GHz transmissions systems in the vicinity of HV systems.

Some determination can also be made for the effect of corona on the RF link. The presence of corona appears to increase the number of channels that suffer transmission error, however in comparison to the isolated field strength assessments, the error appears to be of less severity. However, simultaneously, those reduced errors appear to be 'bursty' in nature, which from a communications system's point of view presents a more complex operational challenge.

Elevated error levels were seenin both experiments throughout the frequency range and varied with each frequency sweep and additionally occured at frequencies outside those known to be problematic. Similar results were seen in long duration unenergised assessment and via a spectrum analyser; which also showed strong signs of two locally operating Wireless Access Point (WAP) units. This suggested that external interference sources were present and had to be carefully considered in any analysis.

While these experiments have proceeded to practically assess the level of radio

frequency interference in high field strengths and in the presence of corona, there are certain shortcomings. It has been noted that a single channel link possesses a bandwidth which is too small for the transmission of high quality real-time video. Multiple channels could be aggregated together to create one high data rate link which would then be sufficient. However this complicates both the hardware and firmware design, needing some level of parallel processing to control the multiple transmitters simultaneously.

Therefore, while this research forms a base of knowledge, with a level of detail not see in published literature, it needs to be extended with data captured directly on a high data rate link. This link can be achieved with WLAN technology; possible because WLAN also makes use of the 2.4 GHz ISM band. Development of a WLAN test system to meet these needs is discussed in the Chapter 8.

This page is intentionally blank.

Chapter 7

Vision System Platforms

There are three routes by which the image capture interface could have been realised: a pre-built module, a custom solution around a readily available imaging element, or a USB (computer) webcam. Each solution has different merits and failings which are discussed below.

Typical pre-built modules have the benefit of being supplied with matching lenses. Whilst most modules are small they are still larger than the 1 cm height restriction inherent in the target insulator's physical design. Furthermore, as their size reduces so does the available resolution. Additionally, most devices output either composite or 'super-video' signals. These signals would have to be digitised prior to compression and transmission. Based on the systems discussed in this chapter this may not be a complicated task (see Section 7.2). However as the problems outweighed the benefits, the idea of using pre-built modules was quickly dropped.

The customised camera design route seemed the most promising. This solution was based on a imaging element which provided a parallel digital interface for imaging data. This interface could then be attached to one of a number of processing systems for image creation[1] and compression prior to transmission. This was the initial solution selected for development. The imaging element selected was a Complementary Metal Oxide Semiconductor (CMOS) device from Aptina Imaging with a Inter-Integrated Circuit (I^2C) interface and 3 Megapixel resolution. However due to implementation complexities, development on this solution

---

[1]The imaging element selected used the Red, Green, Blue (RGB) Bayer data format. This needs converting into full colour data (per pixel) prior to further processing.

was halted despite its superiority.

The webcam option was initially rejected as it would require a device with USB Host functionality and a device driver; of which the latter was considered more problematic. However with the discovery of the advanced video processing platforms this solution became a viable alternative. It was this option that was ultimately chosen in combination with an Open Multimedia Application Processor (OMAP) based platform. Although this was only decided upon at a late stage after problems with the initial solution.

While the image capture and processing solution that was selected is noted above, it is prudent to explain the other solutions that were investigated. Discussions additionally depict PCBs that were designed for imaging element to processing system integration and explains why the solutions were not developed.

7.1 Standard Microcontrollers

Due to the simple nature of the imaging element's interface, it would have been theoretically possible to use either a 16 or 32 bit microcontroller to handle the management of the imaging element and processing of image data. This would also be an ideal solution if the low data rate modules, used in preliminary communications testing, were to be used for the final system due to experience gained during the earlier development process.

There was one possible problem though, the question of whether a microcontroller would have enough processing power. Aptina state that for a resolution of 1,310,720 pixels (1280x1024), a frame rate of 27^2 frames per second is possible. This equates to 35,389,440 pixels that need to be captured per second. Given that data can only be captured on either a rising or falling clock edge, a clock speed of approximately 36 MHz is required to latch pixel data into local memory. Further instruction cycles would then be needed for

- the blending of the single colour pixels into full colour representations
- compressing the resultant image
- packeting the compressed data for wireless transmission

---

[2]25 to 30 frames per second usually provides 'smooth' video.

- handling the wireless interface

Given these concerns, although technically possible to implement, it would only be suitable for low resolution low frame rate situations. If only photographs[3] were required with no stipulation on their update frequency, this would have been an ideal solution, additionally minimising cost and possibly development time. A dedicated video processing system was therefore required: two equally capable devices were discovered in the form of the DaVinci Media Processor and the Open Multimedia Application Processor (OMAP), discussed next.

7.2 DaVinci Digital Video

The DaVinci Media Processor, sometimes refered to as DaVinci Digital Video Technology, is a family of processors from Texas Instruments designed specifically for video processing duties. Devices come in three general varieties: DSP only, ARM processor with dedicated video hardware, or with both DSP and ARM cores. Additionally some devices in the family include a digital capture interface for connection to image sensors, such as the one previously selected.

These processors include all the essential features that were required to capture & compress images, prepare them for transmission over the selected channel implementation and to control the robot. The best implementation strategy would have been to use one of the multi-core devices (DSP + ARM) which contained a capture interface. However, there was one potential flaw in this strategy: programming the ARM processor. While there were multiple choices for the operating system, the most prevalent seemed to be 'MonaVista' Linux. While the source code is open-source, MonaVista's version is highly customised for each supported hardware platform and development tools & parts of the platform seem to only be available via subscription. As such this solution did not appear to be appealing.

The next option would be to use a DSP only device with capture interface. Texas Instruments DSPs are programmed using their Digital Signal Processor/Basic Input Output System (DSP/BIOS) platform, while development libraries and tools (CodeComposer) are included with the purchase of a hardware development kit. A

---

[3]Any resolution up to that available on the selected imaging sensor.

153

kit comprising of the DM6437 processor based development board was purchased and the associated training course was attended. In reality, the latter did not prove to be useful because the course focused on the modification of existing firmware applications to see the effect rather than the development of video application from the 'ground up'. At this course it was suggested that custom boards based on this type of processor should be avoided due to difficulties in designing, fabricating and populating boards containing high density Ball Grid Array (BGA) devices.
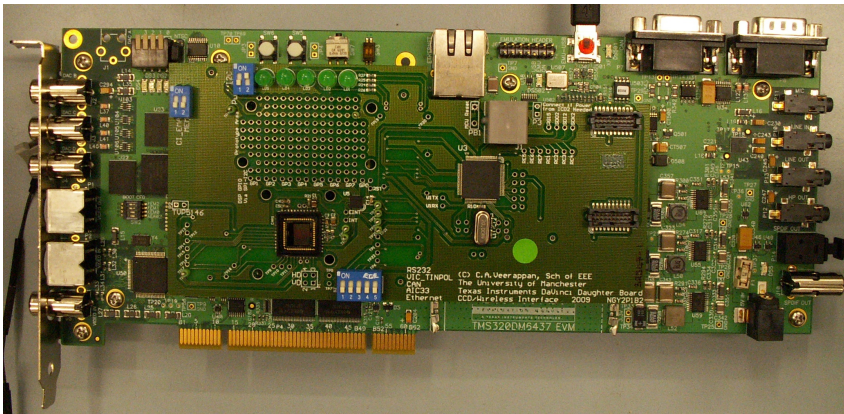
Figure 7.1: Interface board mounted on development kit

A PCB (daughterboard) was then developed which interfaced the selected imaging sensor to expansion sockets on the development kit. A 32-bit PIC Microcontroller was also included on the PCB and interfaced to the DSP's 'host port'. Image data was to be acquired via the capture interface with the DSP converting[4] and compressing the raw data prior to passing it to the microcontroller. The microcontroller would then interface to the wireless transmission system. A further benefit of this implementation is the fact that the DSP can be 'booted' from the host processor, therefore allowing the algorithms running on the DSP to be upgraded easily on a final solution if necessary. Figure 7.1 shows the interface board developed mounted on the development kit.

There was however one caveat with the daughterboard's design. The sensor uses a control interface based on I^{2}C; a two-wire, multi-device, command-response, address based, communications protocol. The address is the key to enabling multiple devices to share the same two lines, compared to the SPI where individual 'slave select' lines are used to direct data. The problem here related to the fact that,

---

[4]Photographs to JPEG images and video to MPEG streams.

while the majority of I^{2}C devices allow the address to be changed within a small range, the imaging sensor selected used a fixed address.

This fixed address unfortunately clashed with a video decoder chip included on the development board. While this video decoder chip had an 'output enable' option, this did not necessary mean that the device would also stop 'listening to' and 'responding on' the I^{2}C data interface. This lead to the possibility that two devices may simultaneously drive the data line, which would then cause possible malfunctions. This situation had to be avoided and a solution came in the form of an SPI-to-I^{2}C protocol bridge. It was ultimately this protocol bridge which caused this development route to be abandoned.

Firstly, the DSP selected does not have a true SPI module, rather, the included serial port can be forced to act as one. Furthermore, the Digital Video Software Development Kit (DVSDK) did not include a DSP/BIOS compliant[5] SPI Device Driver, therefore one had to be written. Due to the lack of clear guidance or documentation on creating DSP/BIOS[6] compliment device drivers, a driver was created directly using the device's Chip Support Library (CSL), which is an Application Programming Interface (API) for the configuration and control of on-board peripherals. However this would not necessarily be DSP/Basic Input Output System (BIOS) compliment and as such its interaction with the driver for the image sensor and with the readily available codecs could not be guaranteed.

Programming using the CSL was not the issue, but rather the lack of a concrete standard for SPI implementation. The datasheet for the protocol bridge was inaccurate in places, such as the maximum operating speed and the state the slave select line should take between packets of data. This meant that the on-board serial port peripheral could not be used as it could not be slowed down to adequate levels and 'insisted' on toggling the slave select line between packets.

This then meant that the protocol had to be 'bit bashed', a procedure which is not problematic in principle. The issue was now due to the DSP's pseudo multi-cored architecture. Split into two sets of four functional units, this allows up to eight

---

[5]This essentially entails that multiple peripheral drivers would be able to co-exist and operate with causing each other interference.

[6]The term DSP/BIOS is used in multiple situations. It is refered to as the real-time kernel (minimum operating system) that runs on the DSP. However documents suggest it can also be used during development to 'graphically' configure peripheral devices, whose implementation would be Chip Support Library (CSL) based.

instructions to be executed at once under certain conditions. As 'bit bashing' typically uses manually counted delays and bit toggles, it was entirely possible for the delays to be executed by one unit while the toggles are handled by another. It was exactly this that was seen during development and testing with a logic analyser revealed the toggles would occur in adjacent clock cycles with the total bit delay following.

Although this was eventually resolved by careful placement of instructions (it is possible that the compiler/optimiser had a role in the problems) it was decided that the time taken to make the protocol bridge functional was too high. Furthermore, part way through this development cycle, the Open Multimedia Application Processor (OMAP) was discovered. The only reason for the continued development after its discovery was the fact that the code developed for the DSP, specifically conversion and compression algorithms, could be ported to the new hardware. However due to the projected time required to fully develop the image capture firmware application, this route was abandoned.

7.3 The Open Multimedia Application Processor

The OMAP processor family from Texas Instruments shares its roots with the DaVinci processor line. While the latter is very video centric in nature[7], the OMAP line focuses on more generalised (but multimedia orientated) cases where both general purpose and digital signal processing duties are required and where power consumption is of concern. While high end sub-families include additional communication subsystems and are aimed mobile multimedia devices, the OMAP35xx line is possibly the most generic range and is of most interest in this situation. The OMAP35xx always includes an ARM Cortex-A8 GPP and can include a C64+ series DSP and a graphics accelerator dependant on the chip variant. Like the DaVinci line, these processors also include an image capture interface.

The main benefit of this processor line is the availability of a sub $150 development board, directly from Texas Instruments, aimed clearly at open source development, opposed to the DaVinci methodology. The board in question is known

---

[7]Product literature states products like digital video recorders and in-flight entertainment systems.
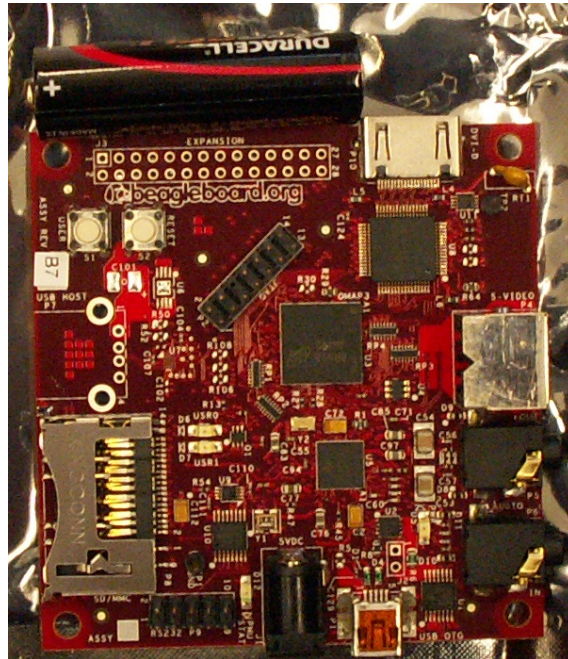
Figure 7.2: BeagleBoard Platform

as a BeagleBoard and reportedly consumes 2 W (5 V @ 400 mA) when running and is shown in Figure 7.2. While this board is primarily targeted as a development platform, other versions for use on a commercial basis are also available, in particular from Gumstix Inc.
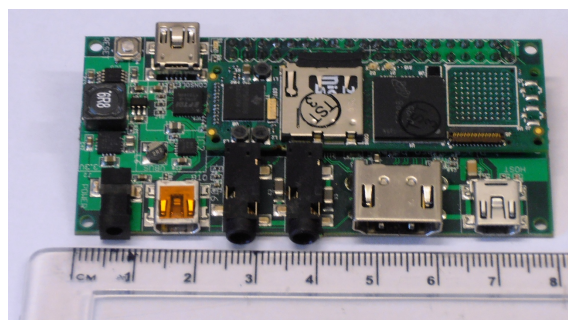
Figure 7.3: Gumstix Overo Earth COM with Summit Baseboard

This company's boards are almost identical to the BeagleBoard and have been designed in modular fashion: a main computer module and an expansion module. This then allows customised expansion boards to be developed which only include the features that are required for the task in question. Importantly, the computer

modules include access to the processor's image capture interface; something the BeagleBoard does not offer. The combination chosen for the robot was the Overo Earth COM and Summit expansion board[8] and is shown in Figure 7.3; this is approximately 2/3 of the size of the BeagleBoard.

Both platforms can run multiple operating systems however the most popular seems to be Angström Linux, a Linux distribution tailored for embedded systems. Both systems described above were purchased for development and testing purposes. The former because it seemed to have the widest community support, the latter because of the image capture interface and its smaller footprint.

The image capture interface exposed on the Gumstix device is in the form of a flat-flexible cable socket on the main computer board. In order to use this interface a second[9] expansion board was required. The board also needed to provide 'level-shifting' functionality as the processor operated at 1.8 V while the imaging sensor operated at 3.3 V. This board's PCB design can be seen in Figure 7.4. Prior to commissioning fabrication and assembly, the programming requirements and source code availability for the image capture hardware was assessed.
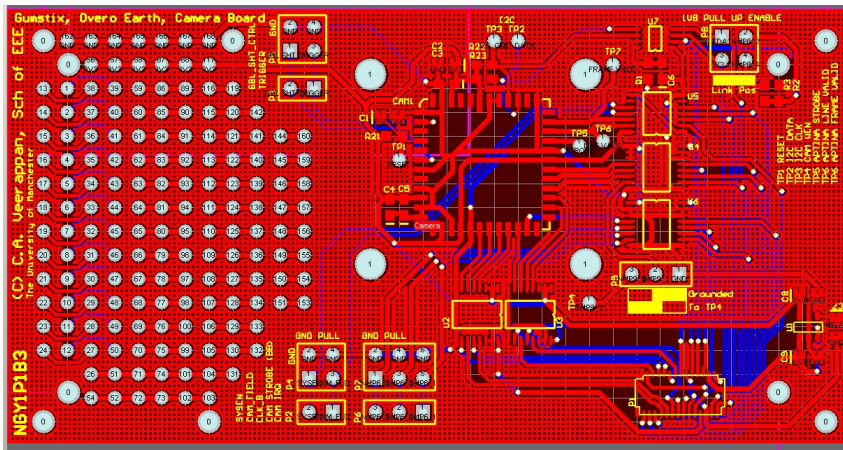
Figure 7.4: Designed Gumstix Camera Board

This revealed that the capture hardware is not well documented, supported or known in the wider developer community. While there is a project to develop

---

[8]These were actually the only boards in the Overo range at the time of purchase. The main computer contains the ARM only processor variant. New COM modules now include ARM+DSP variants and variants with WiFi and Bluetooth interfaces.

[9]The Summit expansion board being number one.

camera support for the OMAP3 processor, it did not seem to be very active when first investigated[10]. Although this recently seems to have changed, the documentation still appears to be an issue. Furthermore, there is no clear guidance about how to integrate these components into the Linux system. On this basis, despite the fact that a customised camera provides a superior investigation system, development on this potential vision system was temporarily suspended. While further development would be viable over a longer period, in the time remaining there was a danger that it would detract from the main research objective; its continued development would also require substantial low level programming; essentially a software engineering exercise.

While this apparently leaves the robot without a viable vision system, a proposal that was initially rejected can now be used due to the availability of a Linux operating system: a USB webcam. This was initially rejected due to need to write drivers for the hardware platform controlling the robot, at one point potentially a PIC Microcontroller. As the control platform is now Linux based, and with the up take of Linux in the wider community, the availability of suitable drivers has increased; although it should be noted that not all webcams operate within Linux. Therefore a number of different cameras were tested with capture software to assess which provided the best solution; this assessment is in Sections 9.1 and 9.2, which also discuss the implementation of live video and still image capture systems respectively.

The use of Linux also enables the use of a 'wireless USB dongle' as the radio frequency interface. This would then allow the use of a wireless enabled laptop (or similar) at the operator end of the link. Captured video could then be converted into a streaming media format and then sent in a similar manner to streaming video services on the Internet. The realisation of this link is documented in Section 8.1.

Furthermore, as briefly noted in Section 4.3, such a platform can be used to control the robot's motors in addition to running the image capture, format conversion and compression duties. This multi-purpose ability is discussed in further detail in Chapters 8 and 9.

The combination of a small, powerful and 'open' platform combined with an 'open' operating system and associated compilation tools provides the ultimate flexible

---

[10]http://www.gitorious.org/omap3camera/

embedded solution. Software, tools and utilities exist for the majority of tasks that users would normally want to accomplish. In the event one cannot be located, it can be written and compiled by the user themselves. Combined with an active developer communities for both the BeagleBoard and Gumstix Overo, help, advice and support (in both directions) is always possible.

7.4   Summary

The vision system is a major component of the monitoring system being developed. This chapter has detailed three ways in which such a system could be built and controlled. A customised solution was noted to provide the highest quality and flexibility but at the expense of cost and development time, with both proposed solutions (DSP and OMAP based) experiencing a combination of hardware and software issues. A system based on consumer webcams and wireless interfaces connected to a community supported Linux based hardware platform built around an OMAP chip was ultimately selected. The implementation of this system is described in Chapter 9; this chapter additionally includes an initial investigation into the physical optics required for the vision system.

If remote real-time video streaming and photo capture shown to be viable from an energised insulator, this system could be adapted to implement either upgraded or customised imaging solutions. It is additionally important to remember that a visual light camera is just one type of sensor that can be implemented.

If it can be shown that a complex sensor such as a camera can survive high field strengths and the presence of partial discharges, then the integration of other devices will also be possible in principle. Other devices include thermal imaging sensors for the detection of insulator heating[15, 16] and ultra-violet for the detection of organic growth[6] or corona[17, 53].

Other possibilities include acoustic detection as noted by Lang, Allen & Zhou[15], however it is possible that the robot's mechanical systems may interfere with reliable corona detection. Measurement of the electric field distribution is also a possibility however the presence of the robot itself will alter the field distribution.

Chapter 8

Communication System Development

After the successful results from the preliminary communication experiments reported in Chapter 6, the initial aim was to develop a communications system based on the aggregation of multiple 500 kbps channels to form a single broadband link capable of sustaining live video. However after investigating possible vision systems, the availability of low power Linux computers opened another possibility: the use of WLAN technology[7].

The 'G' variant of this technology uses the same license free 2.4 GHz ISM band as used during preliminary investigations. This specification has a maximum theoretical data rate of 54 Mbps which is sufficient for real-time video streaming. This technology is widely used in consumer and business environments so has a proven track record of reliability. Furthermore, the use of a Linux computer allows the use of a consumer webcam as the capture element, as detailed in the previous chapter.

This chapter details the method by which the high speed wireless link was examined. Details of the high voltage test scenario and results are additionally presented. Results are then compared to a simulation of the electric field distribution on an energised composite insulator.

Results from extensive RF spectrum measurements made in un-energised conditions are presented along with comparative measurements made in an office environment. Spectrum results are then compared to data transfer results obtained

during high voltage experimentation and potential interference sources present within the high voltage test facility are debated.

The antennas used for these tests were an 18 dBi Yagi and a 2 dBi whip. Additionally circular waveguide antennas were examined to see the difference in performance. The Yagi and circular waveguide antennas were characterised and results compared to the results obtained during high voltage testing.

Due to stability problems the developed system was additionally tested in an office environment to assess operational performance and to verify that the proposed test methods were suitable. Results were analysed to assess the types or performance issues that would typically occur in the intended operating conditions.

8.1 High Speed Wireless Communication

The initial test strategy involved the creation of an 'ad-hock' link and streaming a video file located on the mini-computer, as depicted in Figure 8.1. When placed in a high voltage environment, the degradation of the video stream would visually indicate the level interference. This would be a qualitative assessment rather than a numerical assessment; furthermore, when performing visual interference assessment across multiple test scenarios it may prove difficult to ensure the same benchmarking is used throughout.
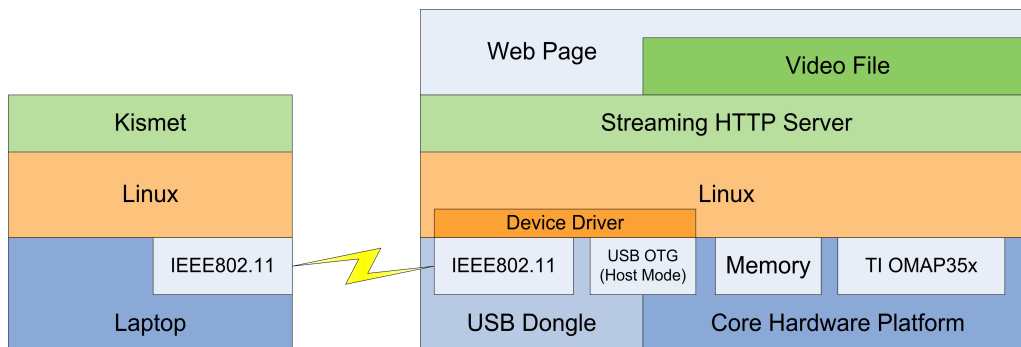
Figure 8.1: Initial Broadband Testing Architecture

Due to this possible uncertainty in the performance assessment, it was decided to transfer data files across the wireless link and then use an application such as WireShark to monitor the over-the-air traffic. This monitoring data would enable

the wireless-packet loss rate to be calculated and then would enable the production of distance verse interference graphs similar to that seen in Section 6.2.

The creation of a high speed wireless link based on WLAN technology should be a fairly trivial proposition, if the devices operate as expected. The main problem was poor wireless network support within Linux. Furthermore, it appears that hardware vendors sometimes change a device's internal wireless chip-set between versions of the same product. This issue is compounded by the fact the majority of devices have been developed specifically for Windows and the fact that many manufacturers do not specify what chip-sets are used in their devices.

The chip-set used is an important bit of information for a Linux user as it enables them to make an informed decision on how likely the wireless adaptor is to work with their system. Out of four wireless chip-set's tested, only two reliably functioned on the Linux platform being used; the Realtek rt73 and ZyNas zd1211. The wireless adaptor selected additionally needed to have an external antenna connector. This was to allow the body of the adaptor to be shielded within a Faraday cage, only leaving the antenna exposed to the high voltage environment.

Another issue which presented itself was the fact that an ad-hoc wireless connection could not be created between the Linux BeagleBoard and a desktop computer running either Windows or Ubuntu Linux. This was even the case when the desktop masqueraded as a master device[1]. This was resolved by sourcing a WAP and using the 'infrastructure' mode[2]. It was possible to configure the Linux computer to automatically establish the wireless connection during the boot process by modifying the 'interfaces' configuration file.

At its heart this link uses protocols from the Transmission Control Protocol/Internet Protocol (TCP/IP) software/protocol stack to enable the transfer of data. The most important layers of this stack, for purposes of assessing link quality, are the transport and application layers; which for the purposes of the performed data collection are highly linked. Of particular interest are Transmission Control Protocol (TCP) and User Datagram Protocol (UDP), which belong to the transport layer of the TCP/IP software stack. The former is used in File Transfer Protocol (FTP) while the latter is used by the Trivial File Transfer Protocol (TFTP), both

---

[1]This is when a computer acts like a wireless access point.
[2]In Linux, this is also known as managed mode.

residing in the application layer.

As the names suggest FTP and TFTP allow files to be transfered from one location to another. FTP is intended for use with large files to and from, local and distant servers. TFTP was originally intended for small file transfers to local devices, for example managed network hardware. While still mainly intended for local transfers, newer specifications have removed file size limitations. Two testing methods were proposed to assess the broadband link; transfering a predefined data file across the wireless link using both FTP and TFTP.

Under TCP, connections and data transfers are assured; if packets are lost they are automatically re-transmitted. Therefore, by monitoring the wireless network traffic it is possible to assess how many 'bad' or missing TCP packets were received and how many re-transmit requests were made. Therefore, as FTP uses TCP as its underlying transport mechanism, received data should be correctly ordered and identical to that transmitted. By repeating the transfer a number of times for each test case, an average link quality assessment can be calculated. Furthermore, the intended packet sniffing application, WireShark, can also provide RSSI values on a packet-by-packet basis. This can also be used to provide an indication on signal and link quality.

Under UDP, connections and data transfers are not assured; if a packet 'gets lost' or becomes corrupted, it is ignored or used 'as is' respectively. Therefore, the use of TFTP should allow a different sort of comparison to be made. By comparing the received file with the original, a calculation can be made as to the raw data loss in the transfer. Out-of-sequence packets are somewhat more complicated to detect and undo within the received file, but can be revealed by WireShark monitoring.

Another possibility would be to compare the average number of packets transmitted for a FTP transfer and TFTP transfer of a particular file. This would allow an estimation of the additional overhead needed to make communication more reliable.

8.2 Test Methodology

Basic WLAN connectivity checks between a BeagleBoard system and a WAP were undertaken prior to high voltage environment testing. Shell scripts were written to automate the process of transfering 100 copies of a 1.2 MB PRN file located on the BeagleBoard to a desktop computer, via both FTP and TFTP. The BeagleBoard was loaded with FTP, TFTP and streaming media servers and a Secure SHell (SSH) service for remote control. The desktop computer was attached to the WAP via a cross-over Ethernet cable.

The same computer had a Belkin F5D7050edV3 IEEE802.11g USB wireless networking adaptor fitted. This adaptor when operating under a Linux operating system has the ability to be placed in 'monitor' mode. This causes the 'air-interface' to pass all packets detected on the current wireless channel to the next layer of the network protocol stack regardless of the packet's destination[3].

The shell script incorporated information on current channel number and number of times the file had been transfered within the local filename. This allowed all transfered files to be identified uniquely; important in the case of TFTP transfers, where the files were to be compared to the original file for comparison and data loss calculation purposes.

This configuration appeared to remain stable for a number of hours, the wireless link only failing when the battery's charge fell to below the level required to maintain the required voltage output. This occured after approximately three hours; based on the calculations this was the expected runtime given the system's estimated current draw and the battery's ideal capacity.

The device was configured to automatically connect to a wireless access point during its boot-up sequence. Basic network utilities were used to detect when the wireless link had been established. After confirming communication between the logging computer & BeagleBoard and configuring the wireless adaptor for monitoring, the proposed test procedure was as follows:

1. Configure the wireless adaptor channel
2. SSH into the BeagleBoard via the established wireless connection

---

[3]Under normal operation only packets marked with the adaptor's MAC address are passed on the upper layers of the TCP/IP protocol stack.

3. Start capturing wireless packets with WireShark

4. Run the FTP shell script redirecting the output to a text file

5. Reboot the BeagleBoard, instructing it to switch to a new channel on restart

6. Change to next channel on WAP

7. Save data logs for subsequent processing

8. Back to 1

WireShark was used to capture wireless packets on a channel basis. Transfer statistics can be obtained from these capture file including detailed information on the TCP protocol as used for FTP transfers. Additionally, by capturing packets via a wireless interface, RSSI values were available on a packet basis. This could then be used to calculate the 'RSSI energy spread' of packets involved in the transfer.

The FTP process generated a number of statistics for each file transfer including data rate, transfer time and port number used. This information was written to a log file and subsequently used to produce an average data rate for each channel at each location. Port numbers could also be cross-referenced with the packet captures recorded with WireShark if required.

8.3 The Meaning of WireShark Statistics

A TCP packet header contains a number of items; among these are the sequence number, the acknowledgement number and packet length. At the start of a transmission session the starting sequence number is randomly generated. The next sequence number is the current value plus the previous packet's length. The packet's length is also used within an acknowledgement packet by being added to the current acknowledgement number. Sequence numbers are updated prior to the transmission of a new data packet. Acknowledgement numbers are updated prior to sending an acknowledgement for a received packet. This mechanism thereby allows the sender to know what packets have been successfully delivered and allows the receiver to make a better judgement on potentially missing packets[72].

As previously noted WireShark can calculate a number of statistics from the TCP packet follow. A description of each statistic follows:

Bad Checksum: This can be caused in different ways. Firstly it may be due to a mechanism called Checksum Offloading, whereby the hardware calculates the checksum rather than the TCP/IP stack. As no packets are generated by the wireless monitoring station errors due to Checksum Offloading are unlikely.

Out-of-Order Segments: This is tracked by WireShark rather than by the TCP protocol itself. By using the sequence and packet length values, WireShark is able to calculate the expected sequence number of the next packet. If the sequence number is lower than expected, the WireShark knows the packet has arrived out-of-order.

Previous Segment Lost: As the receiver knows the sequence number and packet length of the previous packet, it can calculate the next expect sequence number. If the sequence number of the next received packet is higher than expected, the previous segment can be assumed to be lost. The protocol then issues a subsequent 3 ACKs using the expected acknowledgement number (i.e. a duplicate ACK)[73]. This is a somewhat dubious statistic to use when using a wireless card in monitoring mode as there is no guarantee that all transmission are being captured. For this value to be used with confidence, the wire-line exiting the wireless access point would have to be monitored instead. This value is associated with Retransmissions and Fast Retransmissions.

ACKed Lost Segment: This indicates that WireShark detected an acknowledgement to a packet that it did not see itself. As the acknowledgement would not have been transmitted if the packet was not seen, this value is less relevant. When capturing wireless network data this type of event is more likely as the monitoring station will not necessarily see all data being communicated between sender and receiver.

Fast Retransmission: This mechanism is one of four TCP protocol extensions used for TCP Congestion Control. First devised in the late 1980s, it has been re-documented in Internet Engineering Task Force (IETF) RFC2581[4] [73]. If the receiver does not receive the packet (sequence number) it is expecting, repeated ACKs with the expected sequence number are generated until the situation is resolved. If the sender receives three identical ACKs (one acknowledgement and two duplicates), the sender needs to immediately retransmit the requested packet,

---

[4]Request for Comments (RFC)

regardless of the state of the retransmission timer.

Retransmission: This is the 'standard' method of retransmission. When a packet is transmitted, the sender maintains a timer. If this timer expires before an acknowledgement for the packet transmitted is received, the sender can assume that the packet has been 'lost in' or 'dropped by' the network. The missing packet can then be transmitted.

Duplicate ACK: Based on the previously noted statistics, duplicate ACKs can be generated for multiple reasons; Out-of-Order packets and Missing Packets/Fast Retransmission mechanism. According to RFC2581[73], network issues can also be the cause of certain duplicate ACKs; for example duplication of packets in either direction, or packets being re-ordered as they pass through the network. However, as the test network set-up is close to a point-to-point link as possible, network duplication or re-ordering is unlikely.

8.4 Electric Field Strength Distribution

It was suggested that the electric field distribution on a insulator was not uniform. If true, this may lead to field strengths at certain parts of the insulator exceeding those tested in Section 6.2 and be at such a level as to cause corona. While corona was noted in previous experiments, it was thought this was limited to the lower (energised) end of the insulator and the copper tube representing the line.

To verify this suggestion a representative composite insulator exposed to 231 kV was simulated in electric field modeling software. The lower end of the insulator to exposed to 231 kV while the alternative end was set as 0 kV. The insulator was simulated in air consisting of a 0.36 m^2 area centred perpendicular to the central shaft and extending along the insulator's length. Simulations were performed with the guidance of a colleague, Dr Vidyadhar Peesapati, to demonstrate the use of the 'Opera' modelling software and advise on its limitations. Figures 8.2, 8.4 and 8.3 show the results of the simulation.

Figure 8.2 depicts the complete insulator and initially suggests that the field is uniform. However closer inspection shows that the ends of the insulator experience higher fields and it is only in relation to this that the reminder of the insulator has
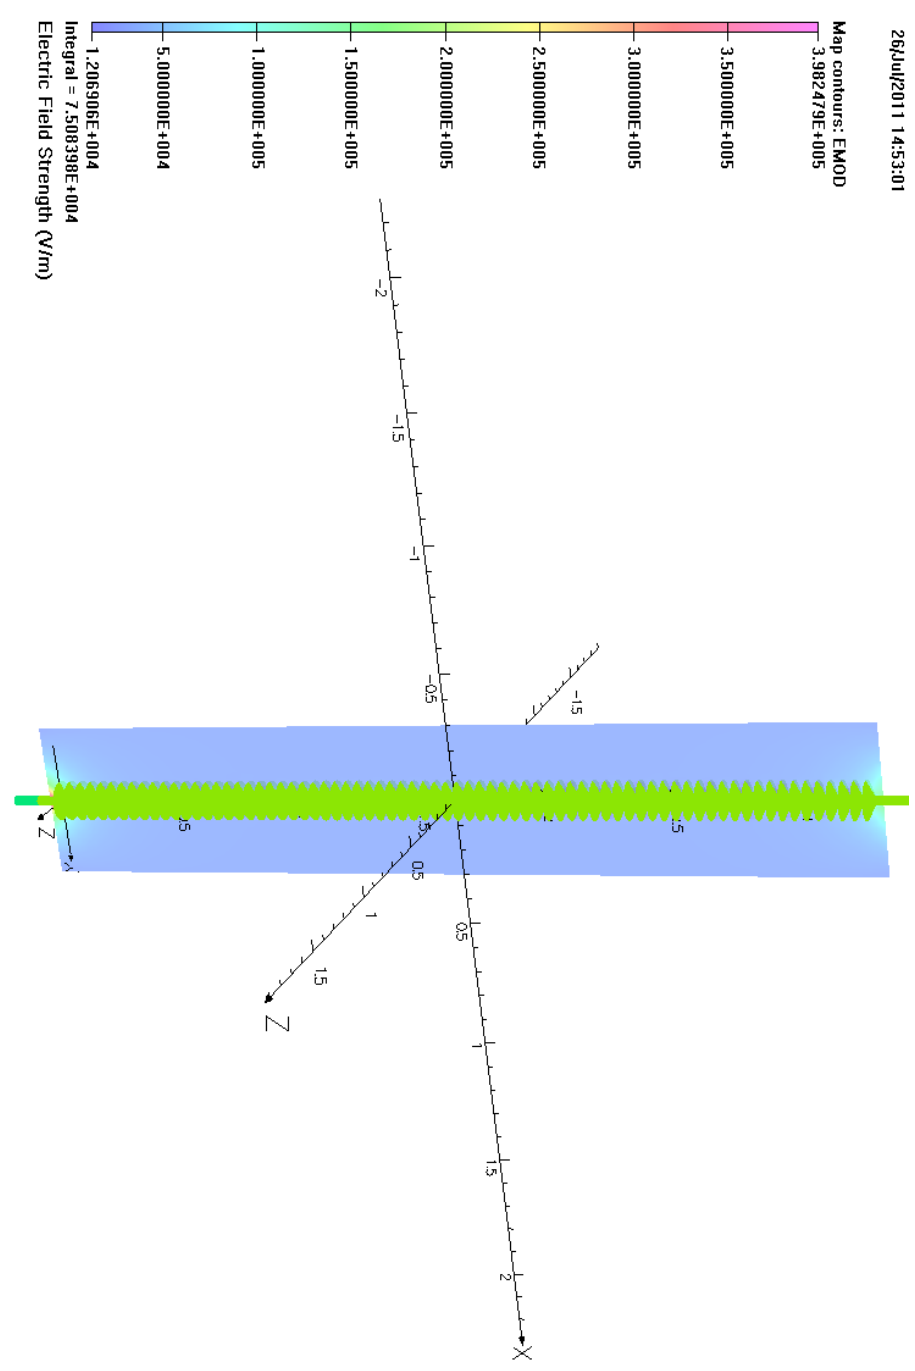
Figure 8.2: Overview field model of 231 kV silicone composite insulator
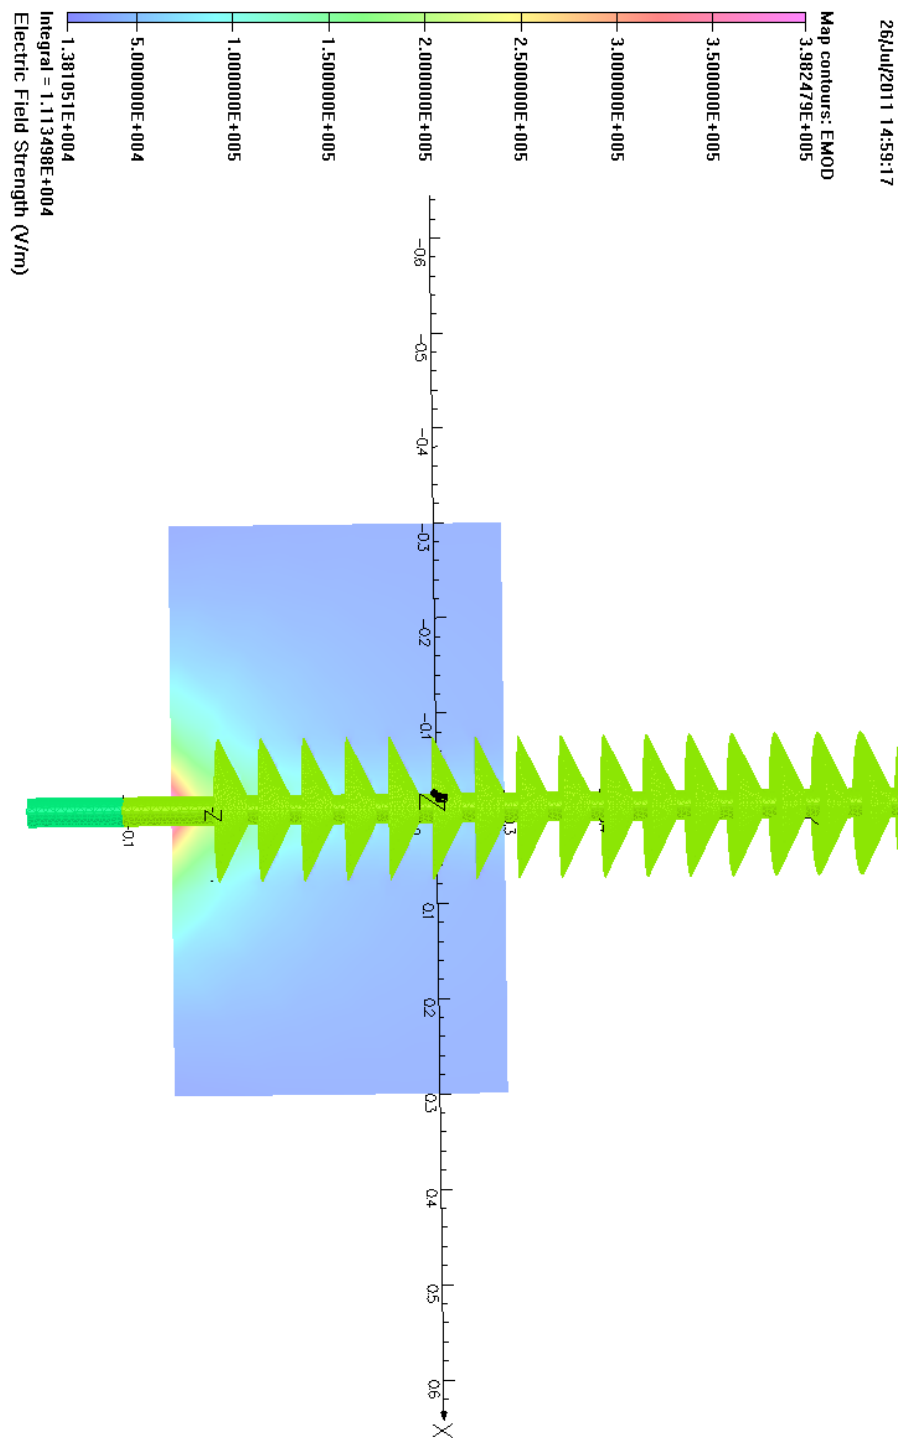
Figure 8.3: Field model of energised end of silicone composite insulator
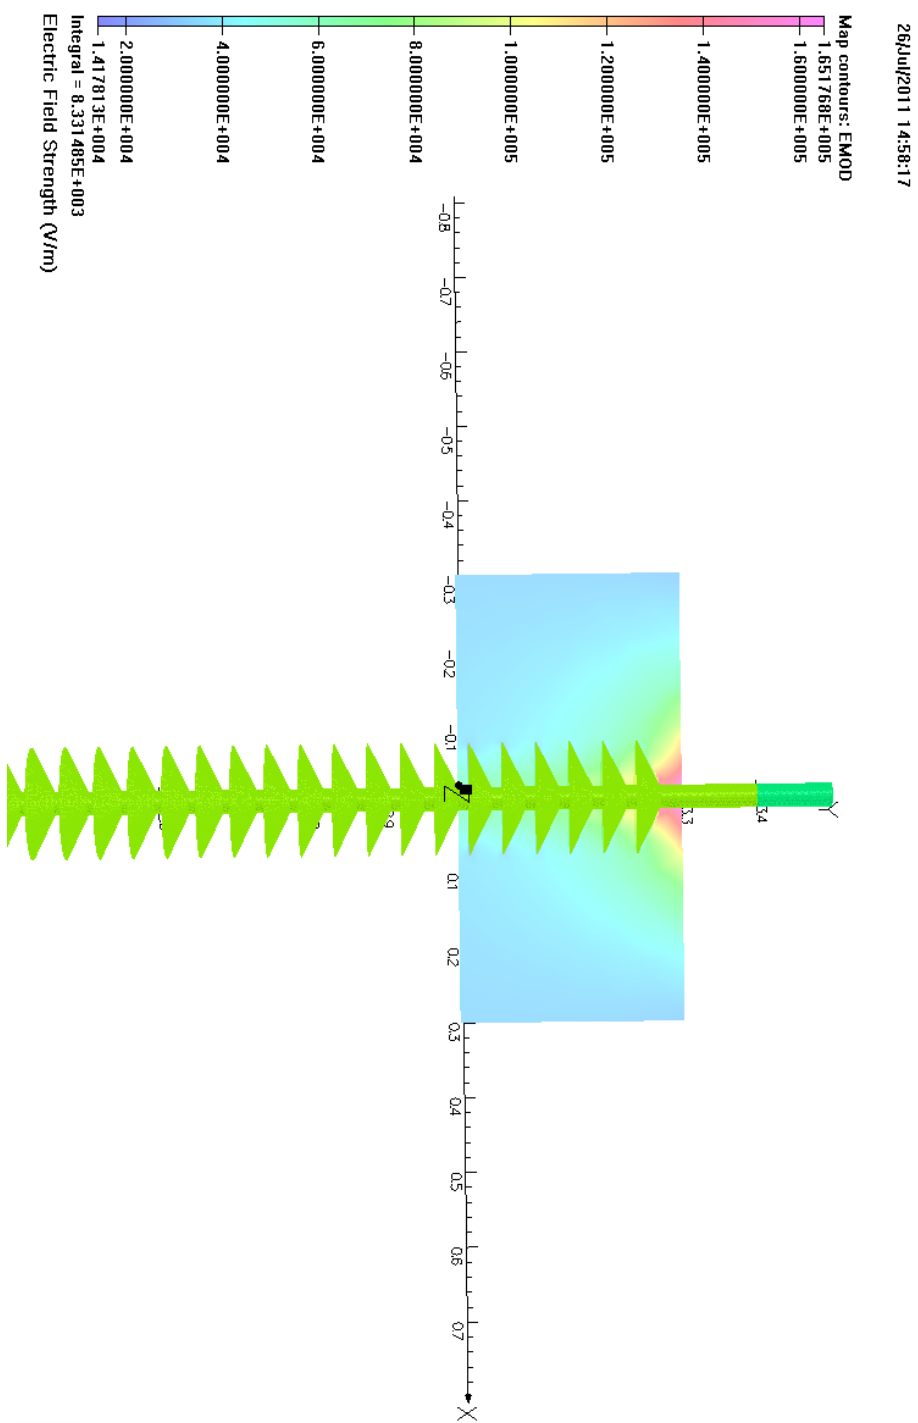
Figure 8.4: Field model of Earth end of silicone composite insulator

a uniform field. This higher field and the non-linearity around the closest sheds to the Earth and energised points are shown in Figures 8.4 and 8.3 respectively. Further depiction of the non-linearity is only limited by the presentation mechanism, whereby most of the variation occurs within a small dimensional range in between two extreme values and is therefore expressed as a single colour.

Subsequent domain research conducted through simulation[74, 75, 76, 77] and experimentation[78, 79] additionally verifies the fact. However it was also noted that simulations need to take into account the fittings, the line and the tower itself as these effect the field present. Regardless of this omission from the simulations noted above, the general conclusion stands: both high voltage and Earthed ends experience a higher field strength than the centre portion of an insulator.

While the presence of corona did not seem detrimental to the experiments discussed in Section 6.2, WLAN uses a larger bandwidth and hence may be more susceptible. Therefore given that both the earthed and energised ends of an insulator experience higher electric field strengths and that it is known that the lower end directly experiences corona discharges (Figure 6.14), it can be theorised that the earthed and energised ends may see a lower relative data rate than the centre portion of the insulator. This theory can be tested through a number of experiments as will be discussed next.

8.5   High Voltage Environment Test Scenario

Testing was to be undertaken on a representation of a single conductor line at ~231 kV; identical to the test scenario used for the latter part of preliminary communication testing. To reiterate this consisted of two composite insulators suspended from the high voltage laboratory's crane. The first insulator (approximately 1 m in length) served to isolate the crane from the experimental rig. A Earth connection was attached to the connection between the first insulator and the insulator being used for investigation. The lower end of the second insulator supported a 28 mm copper tube attached at one end to the laboratory's AC Resonance Test Set. A metallic sphere was attached to the alternative end of the copper tube to reduce the electrical stress that would otherwise exist at the copper tube's exposed end. A representative set-up can be seen in Figure 8.5.
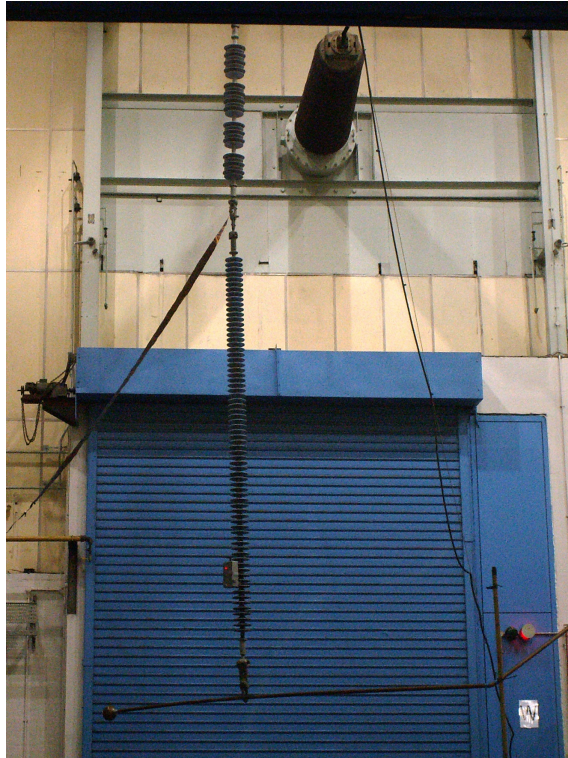
Figure 8.5: Representation of single conductor high voltage line

A BeagleBoard, a high capacity battery, a wireless adaptor and associated support electronics were mounted into a die-cast enclosure. This unit was then attached to the insulator string supporting the energised line. When tested in the high voltage laboratory, briefly under energised conditions and extensively un-energised, a stable WLAN connection between the access point and BeagleBoard could not be maintained regardless of WLAN channel. Typically the link was maintained for only a few minutes before the BeagleBoard's association with the WAP was lost. While less than two minutes would be required to transfer 100 1.2 MB files with transfer rates far below the theoretical maximums offered by WLAN, laboratory safety procedures meant that a longer period of stability was required between booting the BeagleBoard (a manual process) and starting the test sequence.

Due to the basic software set installed onto the BeagleBoard, this wireless link would not re-establish itself without user intervention. The link could be re-established by issuing a number of commands to the BeagleBoard via it's console

interface[5]. However, with the device in a high voltage interlocked environment and the only console access via a 10-pin header on the PCB, this could not be achieved while the experimental set-up was energised.

The possible causes of this link failure are discussed and investigated in Section 8.7 which examines background radio energy levels and Section 8.8 which investigates the device's operation and stability in an office environment.

8.5.1 Interference Assessment: WLAN Data Transfer Results

Although is it known that the devised system could be stable under continuous operation, based on previous problems it was decided to modify the wireless link methodology. Rather than create the link using a client architecture the decision was taken to implement a bridging architecture using two WAPs. This then allowed the link (including link re-establishment) to be handled solely by the WAPs themselves and hence avoids the problems associated with not having a connection manager running on the remote computer. A USB to Ethernet adaptor was used to provide the BeagleBoard with a physical network socket. A 18 dBi Yagi antenna was used at the ground station to improve directivity and signal gain over a typical whip antenna.

Figure 8.6 shows the collective results of the data rate experimentsas a function of both shed and WLAN channel number. Figures 8.78.88.98.10 depicts the average data rates on a channel basis as a function of shed number. Figure 8.11 shows the average data rate over all channels as a function of shed number..

There were 72 sheds on the insulator used for the experiment and measurements, in most cases, were conducted at five shed intervals. Shed one is defined as the Earth end. Each test sequence (assessment of 13 WLAN Channels) took a maximum of 150 minutes to complete. This experimental dataset took approximately 10 days to acquire over a period of about a month.

While there is no outright correlation between channel number, shed number and data rate as initially theorised, a few observations can be made. Firstly locations near to the energised line show poor data rates in certain channels, as do locations

---

[5]The commands are 'ifconfig wlan1 down' and 'ifconfig wlan1 up'.
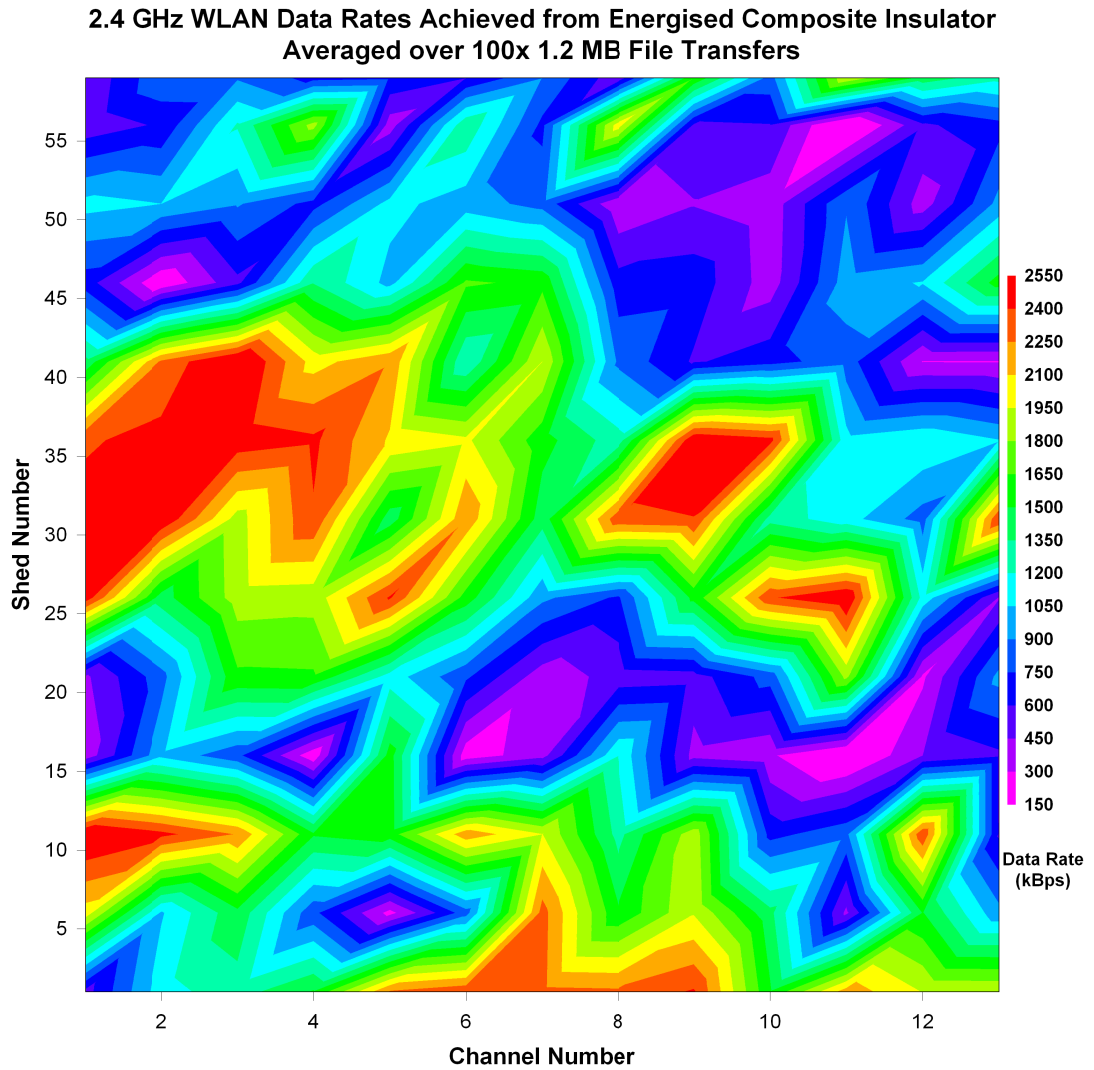
Figure 8.6: Average date rates achieved on all 2.4 GHz WLAN channels from various positions on an energised composite insulator
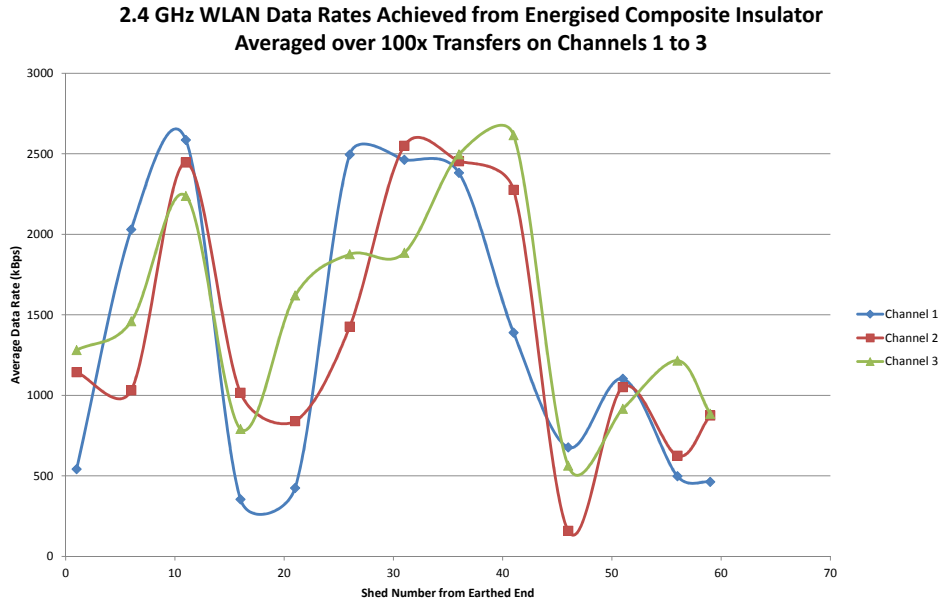
Figure 8.7: Average date rate as a function of shed number. WLAN channels 1 to 3
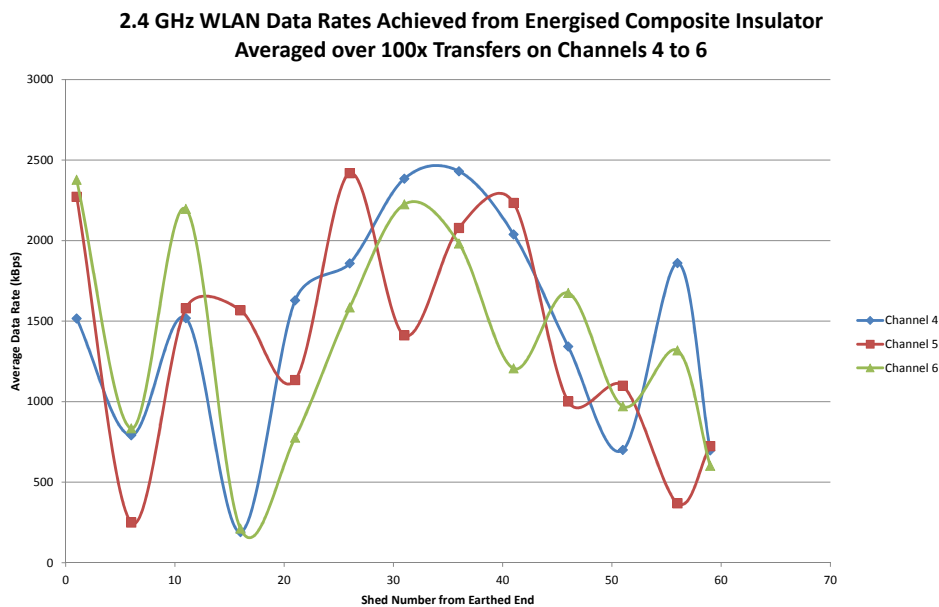
Figure 8.8: Data rate as a function of shed number. WLAN channels 3 to 6
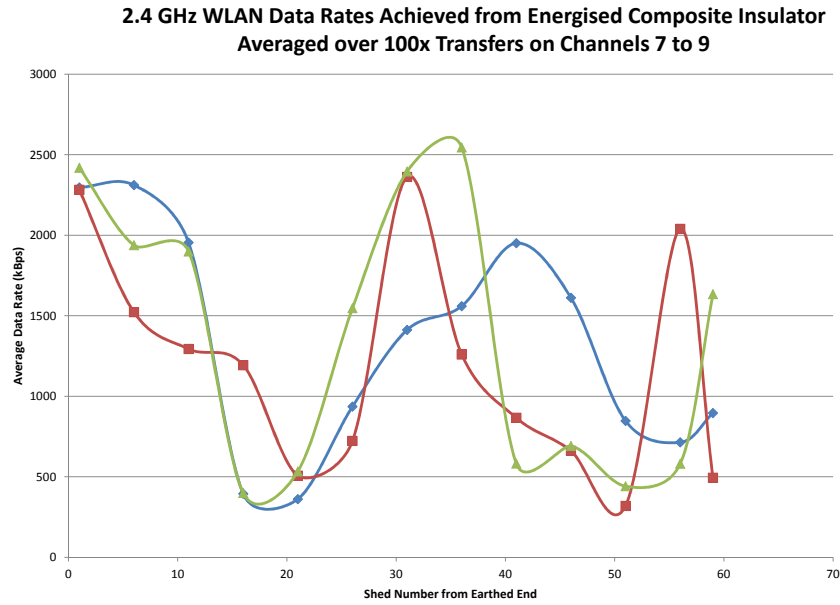
Figure 8.9: Average date rate as a function of shed number. WLAN channels 7 to 9
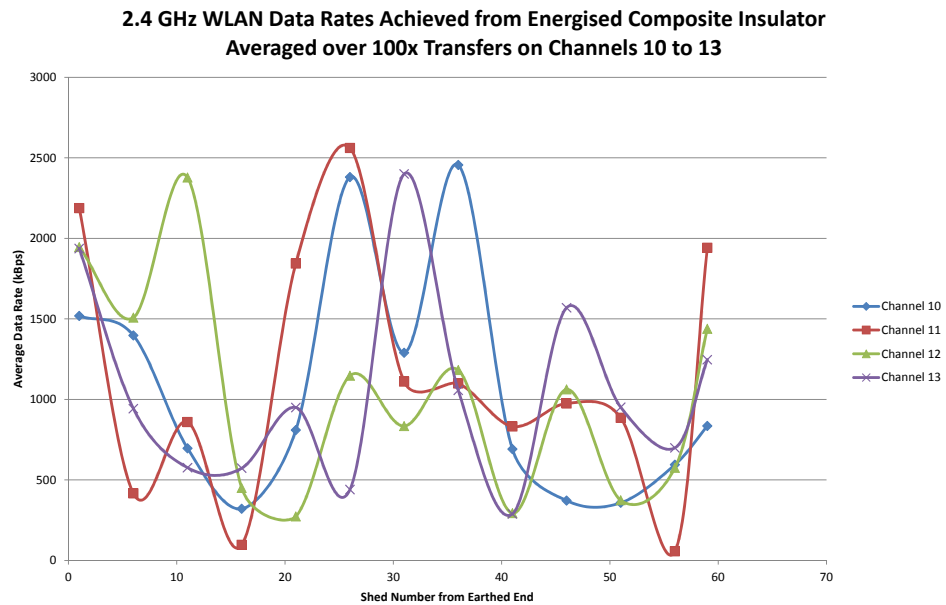
Figure 8.10: Average date rate as a function of shed number. WLAN channels 10 to 13
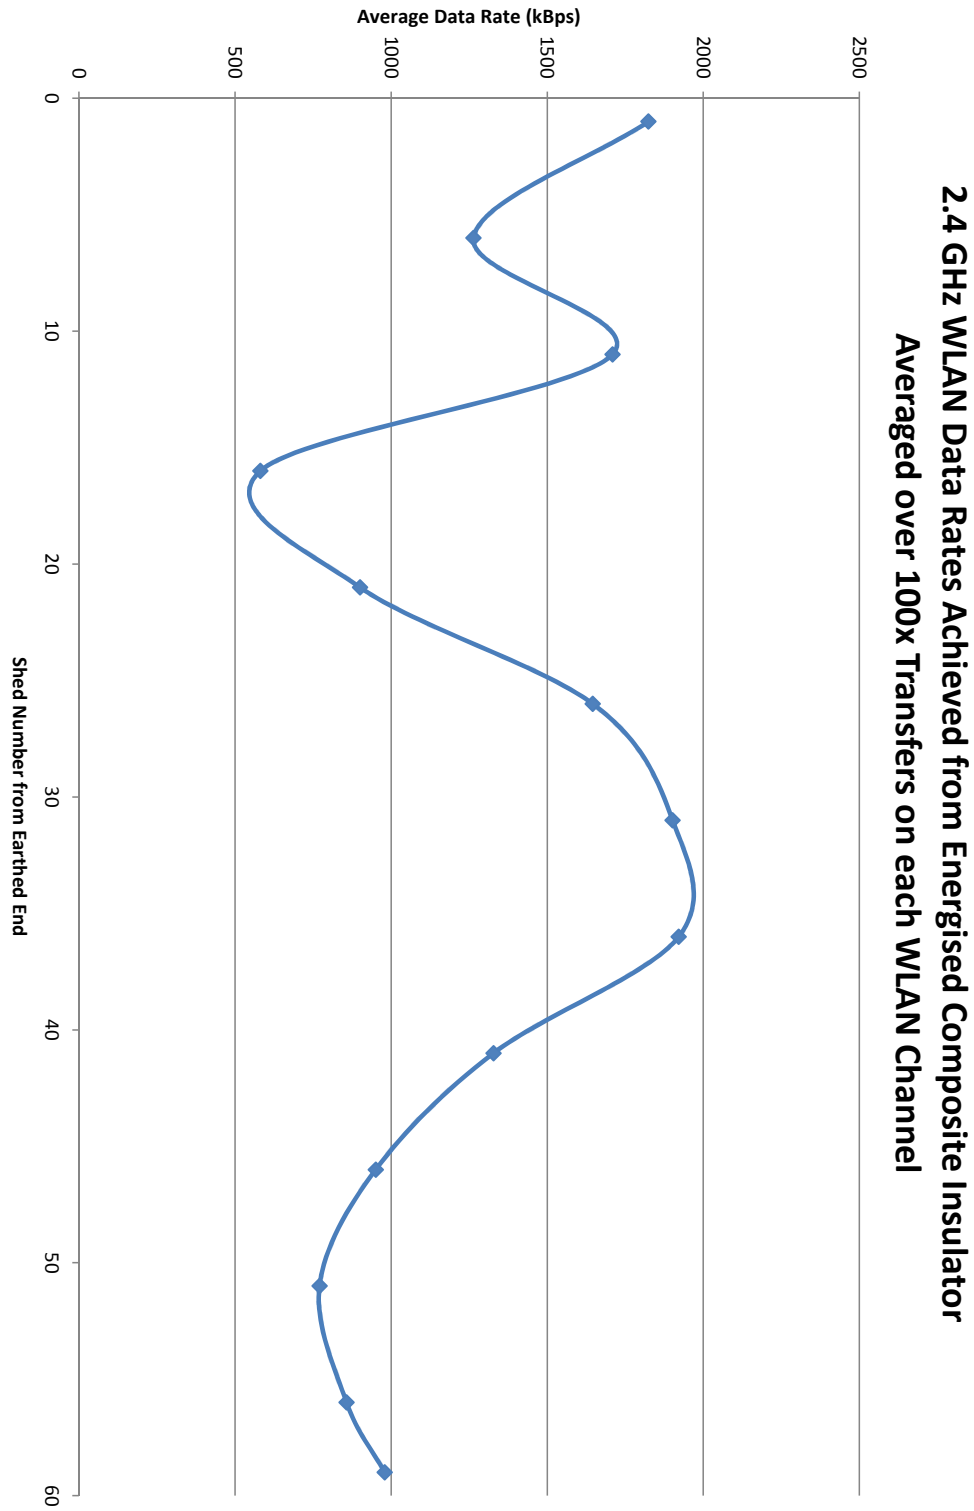
Figure 8.11: Average date over all WLAN channels as a function of shed number.

between sheds 15 & 25 in most channels. This leaves a large band between sheds 25 & 40 that experienced good rates, repeated in certain channels between shed 1 & 15. Shed 1 itself reveals interesting results, whereby data rates are above 1500 kBps for all channels except channels 1-3, higher than what the proposed theory suggests.

When overall average data rate is examined (Figure 8.11), it can be confirmed that the centre of the insulator has the highest data rates, which quickly reduce as the energised end is approached. Interestingly, when refering to this figure, it can be seen that the data rate appears to higher at the very extreme of measurements, i.e. sheds 1 and 59. Refering back to Figures 8.78.88.98.10 it can be seen that these two locations experience a wide range of data rates on difference channels and it is this that appears to be causing an increase in average data rate than when compared to adjacent sheds. Similar analysis on channel data rates for sheds 5 and 11 additionally reveals that a wide range in data rate exists for these locations, hence the reason for the trough and peak at sheds 5 and 11 respectively.

The overall average data rate was 1280 kBps, the maximum and minimum seen was 2615 kBps and 57 kBps respectively, with a theoretical maximum of a of 6750 kBps. These findings add weight to the notion that some external source is contributing to the results recorded.

Figure 8.12 shows the negative result of this external interference; here data rates from a transmitter located at shed 57 are shown for a number of channels recorded at different times over two days. While Channel 1 appears fairly consistent, Channels 4, 8 and 13 show dramatic changes which in certain cases are for experiments hours apart.

It was additionally important to assess if the Yagi antenna's gain or directivity was more beneficial in the data rates achieved. Figure 8.13 shows the results from an attenuated on-line and off-line test in comparison to the original test. In approximately half the channels the two former configurations lead to a higher data rate. A data rate of almost 2000 kBps was recorded in Channel 6 in the attenuated configuration, while Channel 11 showed a very poor link quality across all scenarios.

Figure 8.14 reveals a clear picture of the use of a directional antenna and also of the nature of the external interference. This graph shows the average of all file
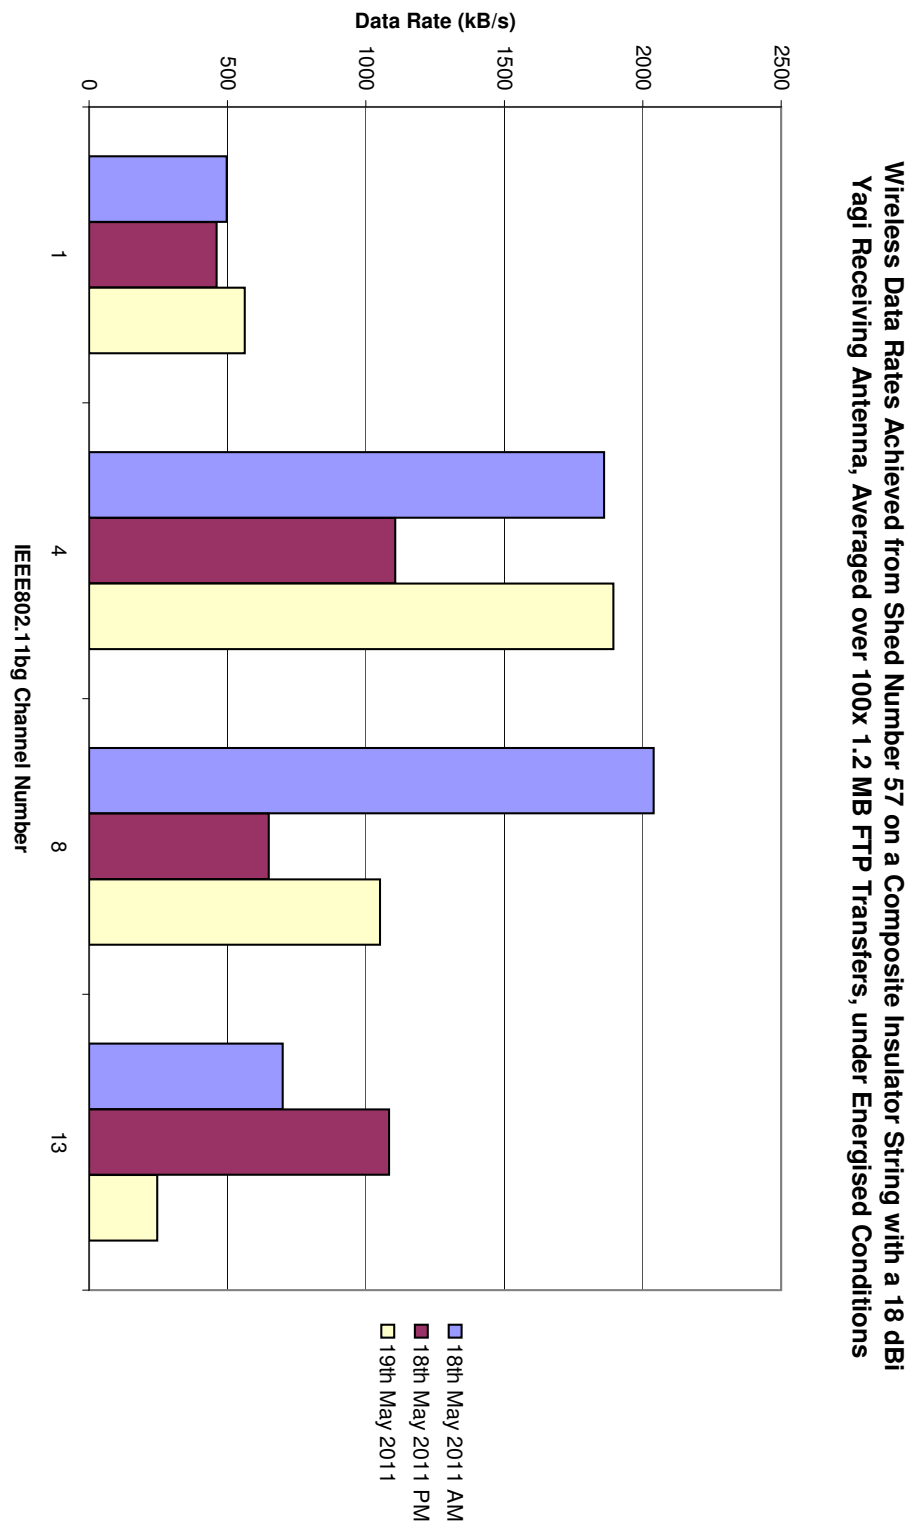
Figure 8.12: Date rates achieved from shed 57 of energised composite insulator
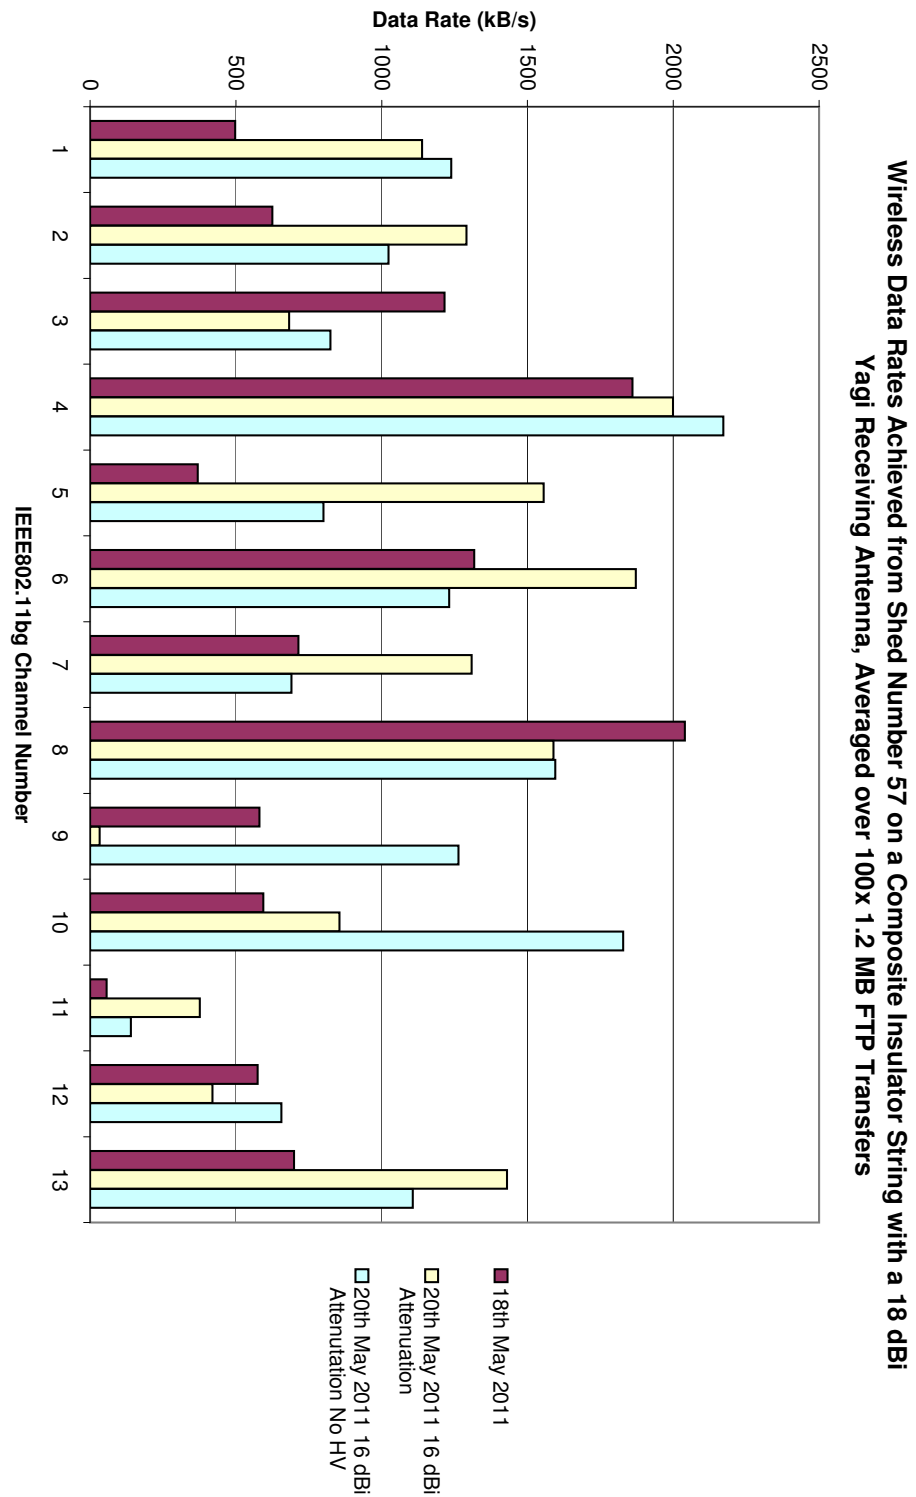
Figure 8.13: Data rates achieved from shed 57 of a energised composite insulator with 16 dB attenuation at the ground station
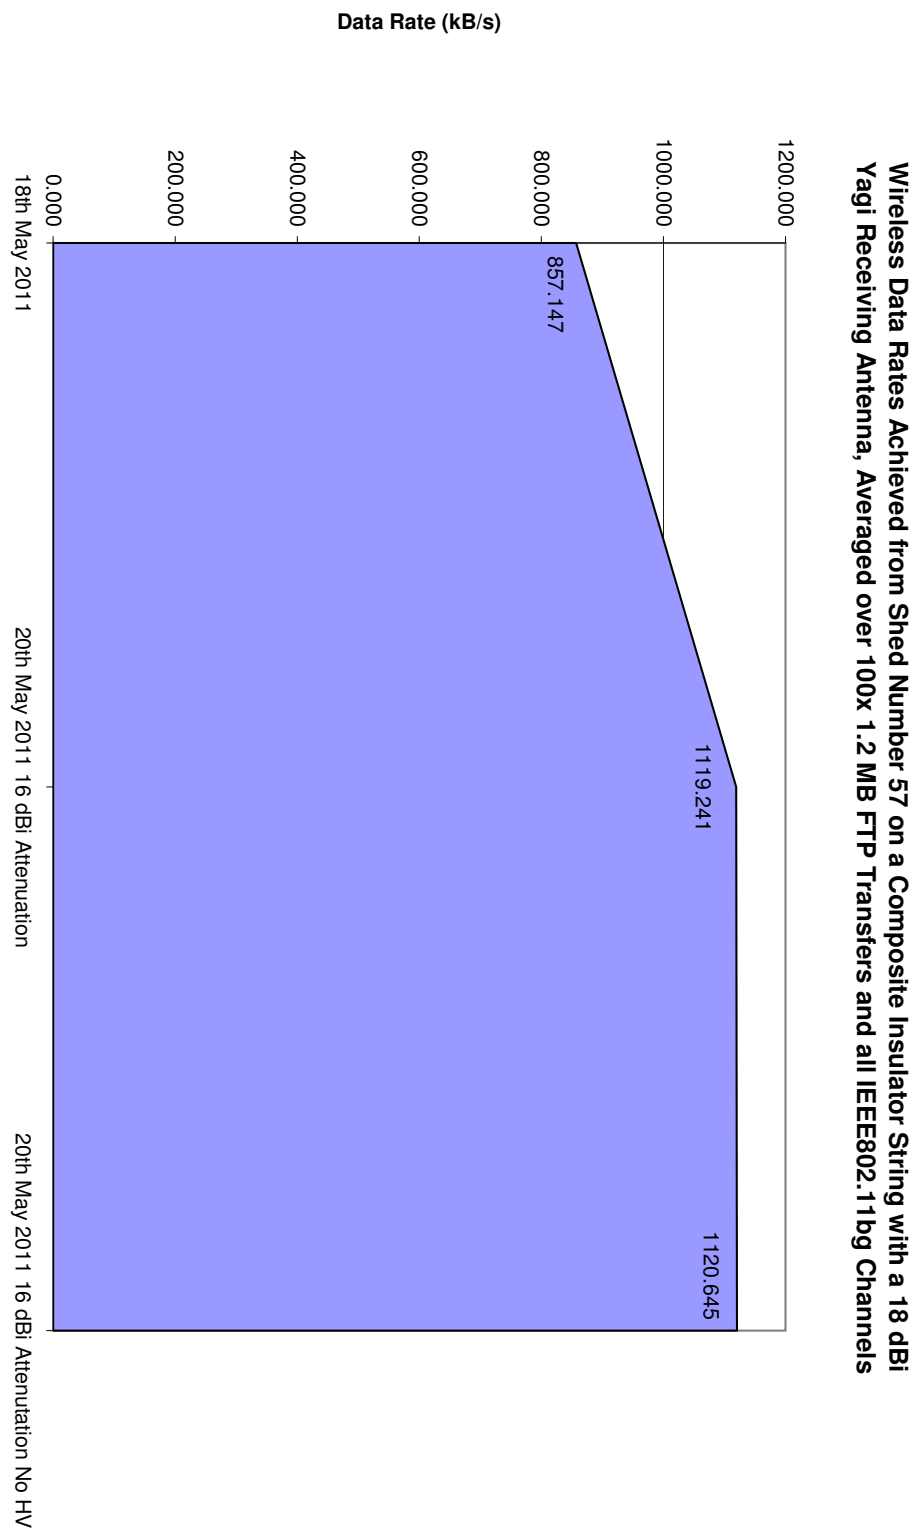
Figure 8.14: Average 2.4 GHz WLAN data rate achieved based on position on energised composite insulator

transfers from a particular location (i.e. across all channels). While the graph suggests that the use of 16 dB of attenuation will cause the data rate to increase, this has most likely been caused by a decrease in external interference. More interesting is the comparison between the energised and non-energised data rates, where only a 1.4 kBps increase was recorded.

These results confirm that the presence of corona or high field strengths is not enough to severely degrade the performance of a WLAN link, with the bigger danger being external interference. The results additionally suggest that a Yagi antenna's directivity plays a greater role in maintaining link quality than its high gain factor.

The final reliable measurement was made at shed 59, beyond which the remote WAP experienced repeated non-fatal failures. Two experimental runs were conducted with the transmitter at shed 61 and were successful until Channel 6, after which the remote hardware failed to respond.

This was initially thought to be related to diminished battery capacity on both occasions, however further analysis revealed otherwise. This came to light when further attempts with a known 'good' battery failed at Channel 1 itself. A examination of the hardware revealed that the BeagleBoard was still functional and it was the WAP that had entered an unknown state; in particular the Ethernet 'link light' was illuminated but neither the power nor WLAN indicators were active. A reboot of the device restored the WLAN Bridge.

Actively using the ping diagnostic tool while increasing the system voltage revealed that connectivity was lost between 215-220 kV. Reducing the voltage in some cases restored the connectivity. A further experiment was conducted at shed 71 where the system lost connectivity at approximately 150 kV.

It is suspected that the WAP's radio circuit is overloading in some manner and is entering a fail-safe mode of some description. This is the only explanation that provides for the fact that a power cycle of the WAP restores functionality.

In order to protect the WAP's radio circuitry a number of cylindrical waveguide transition antennas were constructed. Highly directional in nature, they can be constructed using a food can of the sort baked beans are bought in. The can's diameter and the position & length of the active element determine the optimised transmission frequency. A 10 cm diameter is best for transmission at 2.4 GHz and

small adjustments of the position and length can tune the antenna for a particular WLAN channel[80]. Information on the tuning of these antennas can be found in Section 8.6.

Instead of attaching the cylindrical waveguide transition antenna to the die-cast box arrangement used previously, it was attached to the complete robot arrangement. When this arrangement was mounted at shed 70 of the composite insulator and the set-up energised, discharges to the robot's Faraday cage were clearly seen when the line voltage crossed 100 kV. A video of this event can be found in Appendix B. At this point communication with the robot was lost, with a visual examination of the powered circuit boards suggesting that the WAP had completely failed.

However, in a similar manner to when the WAP seemingly failed previously, a number of power cycles was all that was required to bring the device back into a functioning state. The only casualty of this discharge activity was the powered USB hub attached to the Gumstix mini computer.

Once repaired the robot was then attached to shed 64: the half-way point between the energised line and shed 59. While previously the link failed at 210 kV and 150 kV for shed positions 61 and 71 respectively, in this configuration the link was successfully maintained beyond a line voltage of 231 kV.

Results can be seen in Figure 8.15; while initially looking poor, a comparison to Figure 8.6 shows that this is roughly what would have been expected from the previous configuration at this location.

Therefore these results in part confirm the theory that was proposed in the previous section, whereby the energised end of the insulator experienced a lower data rate than the centre part of the insulator. However there is some uncertainty over the higher than expected data rates at the earthed end which cannot be explained fully at this moment. It should also reiterated that corona may not be the only contributory factor to the reduction in data rates as revealed by experiments in which 16 dB attenuation was used.
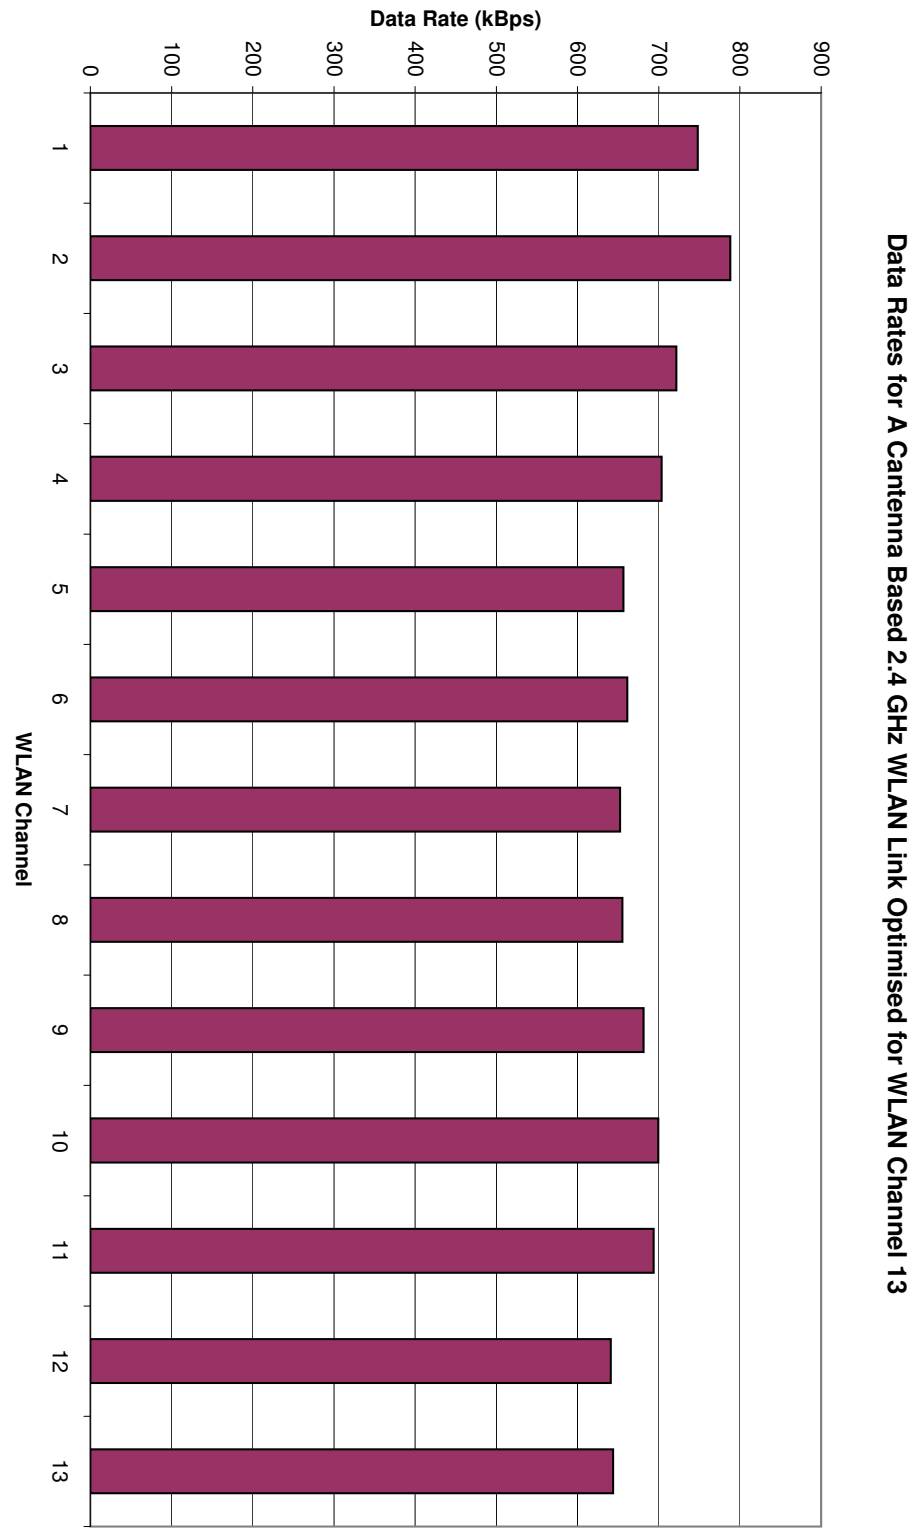
Figure 8.15: Data rate results for a cylindrical waveguide transition antenna link from shed 64

8.6 Antenna and Link Dynamics

Given the varied results across the WLAN channels, it is prudent to assess the Yagi antenna's properties and the quality of the link formed for the initial high speed trials. Figure 8.16 shows the results for the S11 Reflected Energy measurements made for the 18 dBi Yagi antenna used during the investigations. Measurements were made with and without the extension cable used during the investigations.

The antenna itself shows a small gradual loss of performance as the frequency increases from 2.4 GHz to 2.5 GHz. Comparison to Figure 8.6 shows minor correlation in that the average data rate appears to decrease as centre frequency increases.

Figure 8.17 show the energy transfered between whip and Yagi antennas, when the former is located at opposite ends of a suspended composite insulator. The plots on this figure are based on averages calculated from three samples taken at each position. The figure suggests that greater energy transfer would occur when the transmitter is located at shed 1, the Earthed end of the insulator, compared to when located at shed 57, the energised end. The cause of the latter's apparently poorer performance may possibly be attributed to the anechoic chamber situated close to the test area, which could be interfering with radio transmissions as the test device is brought closer to the energised end.

Similar measurements were also preformed for the cylindrical waveguide transition antenna link test scenario that was investigated. S11 Reflected Energy measurements were initially used to tune the antennas to the frequency of interest. As the test scenario proposed to use WLAN Channel 13 for the video link and remote computer control, the antenna was tuned to approximately 2.472 GHz, the centre frequency of WLAN Channel 13.

The results of the tuning can be seen in Figure 8.18 along with S21 Energy Transfered Measurements and the data rates achieved at each channel. The figure indicates that the ground-based cylindrical waveguide transition antenna was tuned to a higher precision than was attained for the antenna attached to the robot. However both tuning results are acceptable given the crudeness of the antennas' design and construction. It can also be seen that the link quality increases around the region of the tuning match, falling slightly on either side.
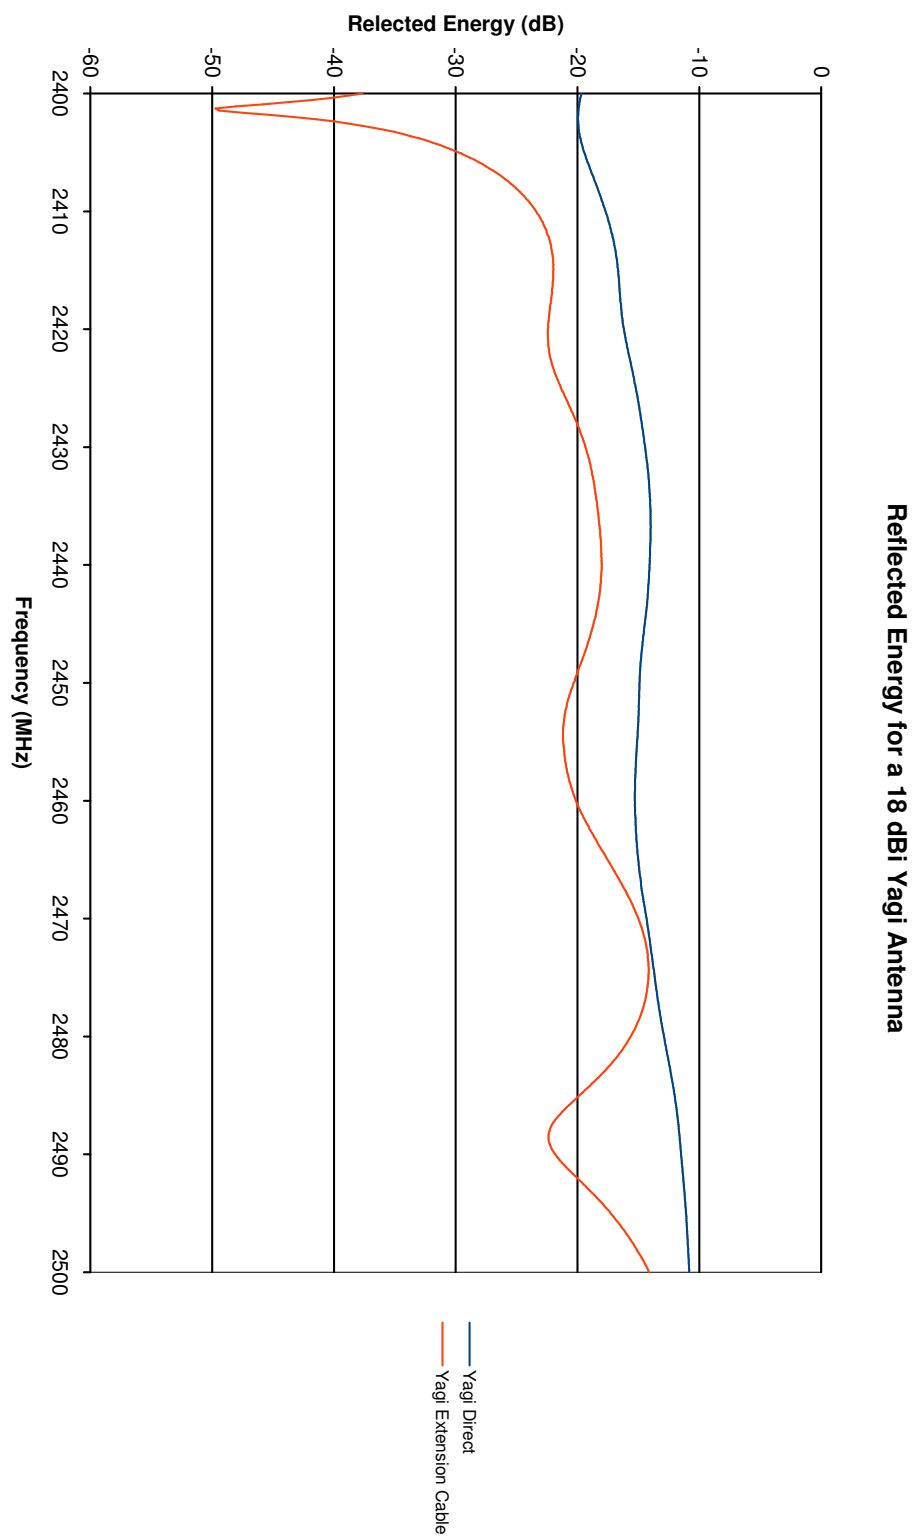
Figure 8.16: S11 measurements for 18 dBi Yagi antenna
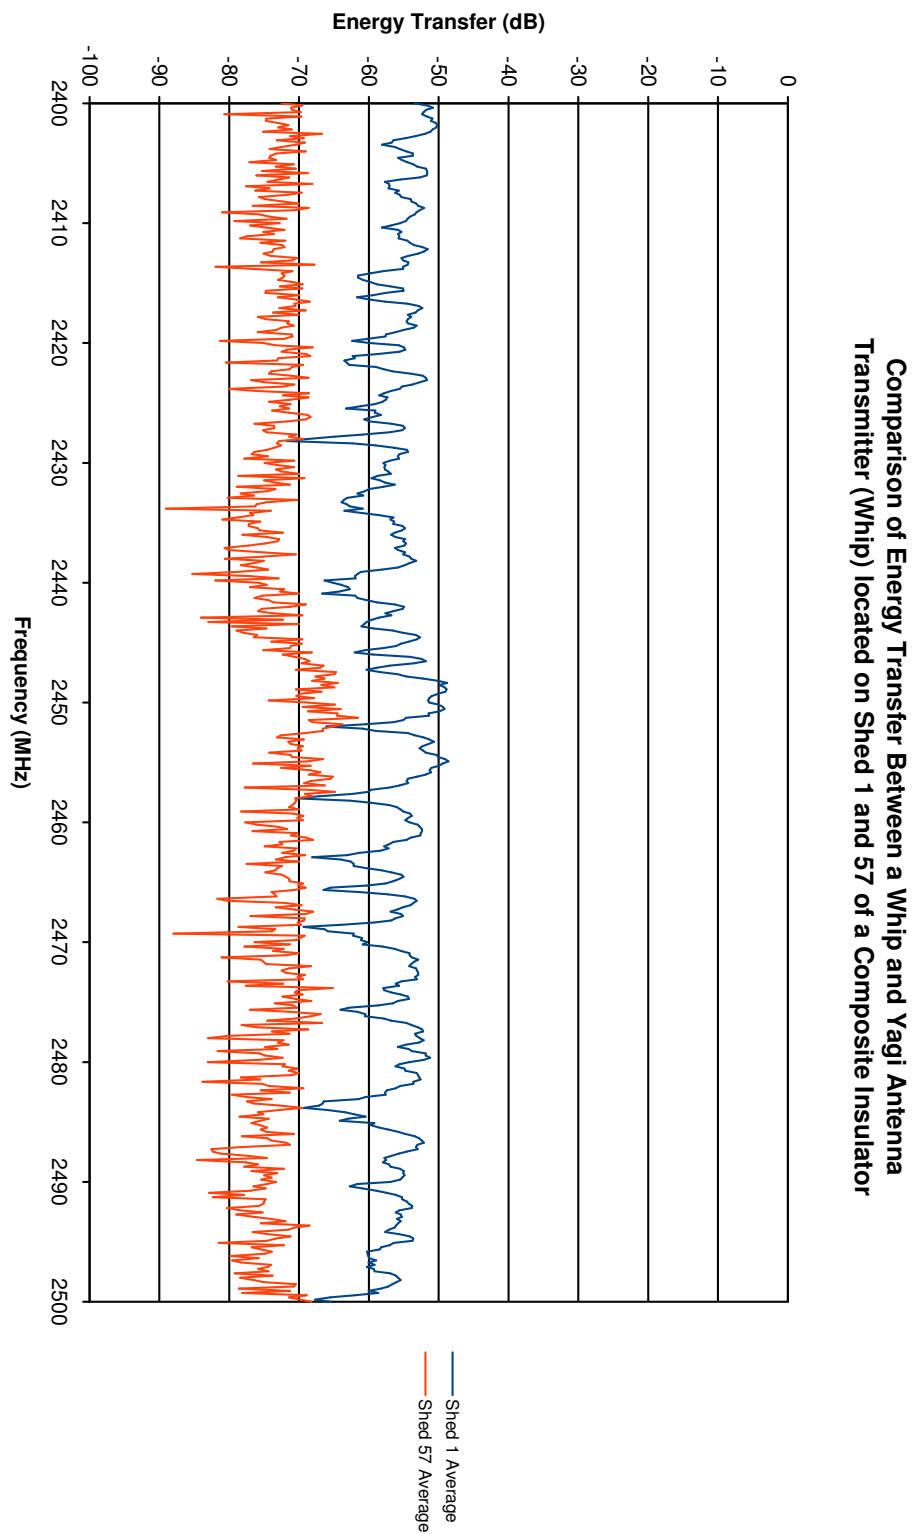
Figure 8.17: Transfered energy measurements between Yagi and whip antennas
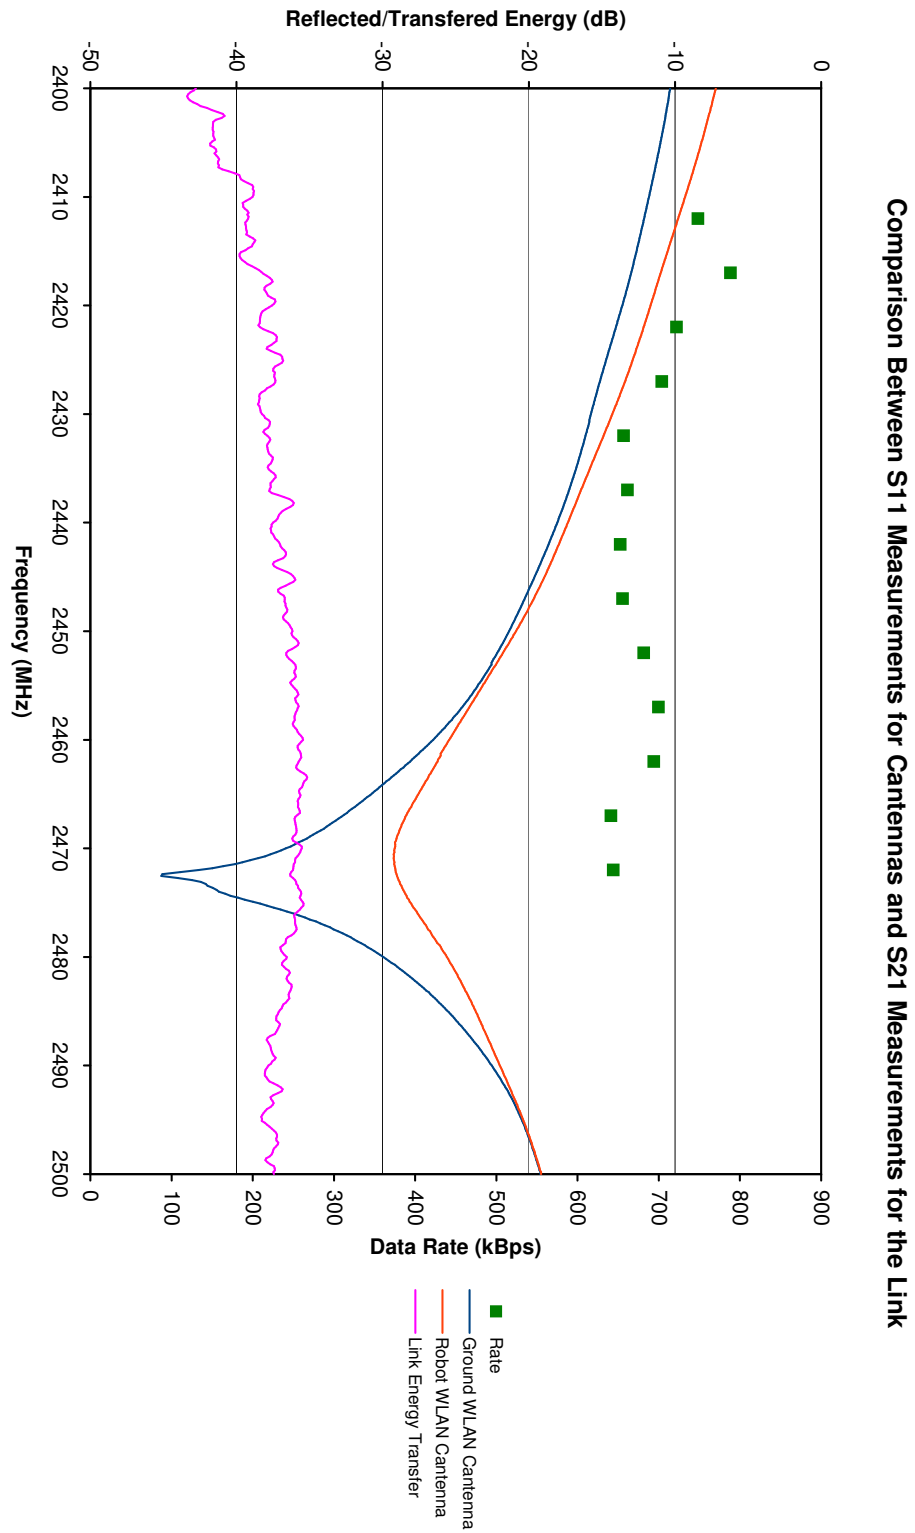
Figure 8.18: S11 and S21 measurements for a cylindrical waveguide transition antenna radio link

As mentioned previously the data rates seen when the robot was mounted at shed 64 were towards the lower end of that seen at other locations and to what is known to be possible over WLAN. However the data rates' lack of correlation to either the antennas' tuning or link profile is of greater interest. It was expected that the data rate at channel 13 would be greater than that seen for other channels, however it was one of the lowest seen during this test sequence. Transfer rate peaks are seen at Channel 2 and between Channels 9-11, the former at a frequency for which neither antenna was particularly well tuned to.

Comparison to Figure 8.16 actually suggests that the Yagi antenna is better tuned at this frequency[6]. However comparison to the S21 Energy Transfered diagrams in Figure 8.17 shows that the cylindrical waveguide transition antenna link performs better. However, it should be remembered that this link was shorter than the Yagi-Whip link and thus may account for some of the signal improvement.

Figure 8.19 shows the predicted electric field strength surrounding a cylindrical waveguide transition antenna when viewing its radiating face. Externally a higher than average field is present which quickly reduces before slowly increasing. While the can apparently acts as a Faraday cage Figure 8.20 suggests otherwise. Focusing on the radiating element, the figure shows that a high field strength is actually present within the void; highest at the element's tip.

This simulation was for the can located at the centre of a 1 m^3 volume of air, with source and ground applied to upper and lower surfaces of the cubic volume. An field of 100 kV/m was applied; higher than the average estimated along the centre section of an insulator (Figure 8.2). This then infers that the field surrounding the can under operating conditions should be less than estimated here.

8.7 Detailed RF Spectrum Analysis

While it would be ideal if the test environment remained stable throughout a 3-year research programme, this is not a realistic expectation. Since RF measurements were first taken to the point at which high data rate measurements were made, a number of things changed in the wider vicinity of the HV laboratory. Of primary significance is the deployment of two university WAPs which had the potential to

---

[6]2417 MHz WLAN Channel 2
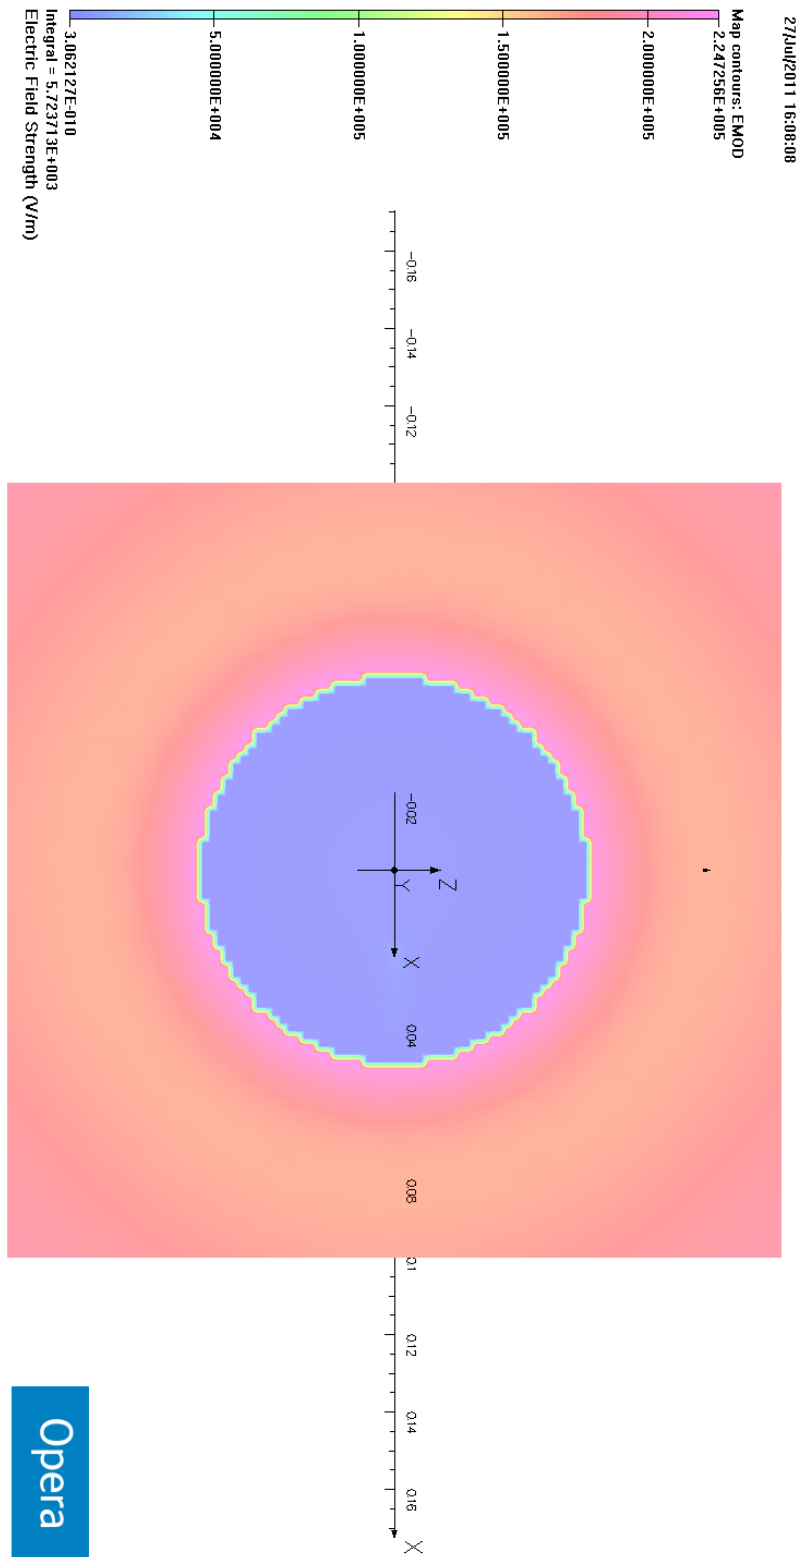
Figure 8.19: Model of the constructed cylindrical waveguide transition antenna in 100 kV/m electric field
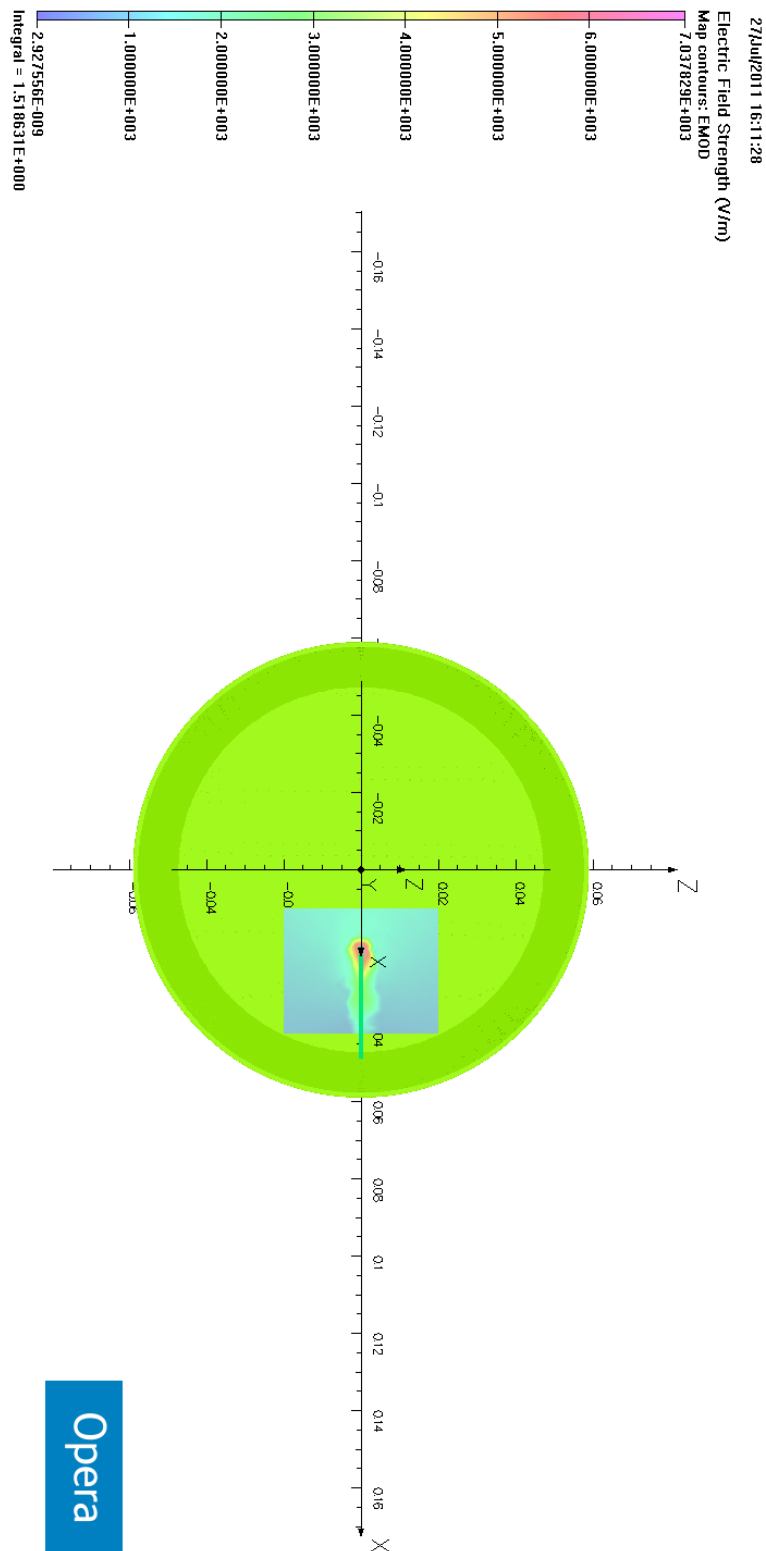
Figure 8.20: Active element of the cylindrical waveguide transition antenna in 100 kV/m electric field

conflict with measurements being made.

However a more immediate concern was the instability of the wireless link during high data rate energised trials, where the link did not appear to be stable for more than 60 seconds. This was in comparison to the development environment where hours of stability were seen, with only a diminishing battery charge causing link failure. All 13 WLAN channels showed similar link stability problems.

This had many consequences, the primary one being the loss of communication with the remote device. As a high level 'connection manager' was not used, this link could only be re-established via the device's physical console port, which meant retrieving the device from its position on the insulator. While a secondary wireless communication link joining this console port to additional hardware on the ground was created for the Gumstix variant of the platform, a similar interface was not developed for the BeagleBoard. It was therefore essential to identify the causes of this link failure.

8.7.1   Background RF Spectrum

Early assessments of radio frequency interference can be gleaned from initial HV experiments, as presented in Chapters 5 and 6. During this time the developed equipment was left running for a number of hours in unenergised conditions. While this primarily revealed that the interference experienced during energised trials was not significantly greater, it also showed high packet reception accuracy for the majority of the spectrum. Further insight can be taken from previous long term spectrum measurements, Figures 6.6 and 6.8 replicated as Figures 8.21 and 8.22.

Figure 8.21 shows the noise ceiling measured overnight in August 2008 using a 2 dBi omnidirectional antenna. The variation in energy during the daytime over approximately six hours recorded using an 8 dBi omnidirectional antenna[7] in April 2010 is depicted in Figure 8.22, clearly showing two bands of higher than background levels of radio energy.

The centre frequencies of these bands closely resembled the centre frequencies of WLAN Channels 6 and 11. These were initially thought to be from wireless networks external to the university, but were later discovered to be internal. Regard-

---

[7]Doradus Model: 26-1360
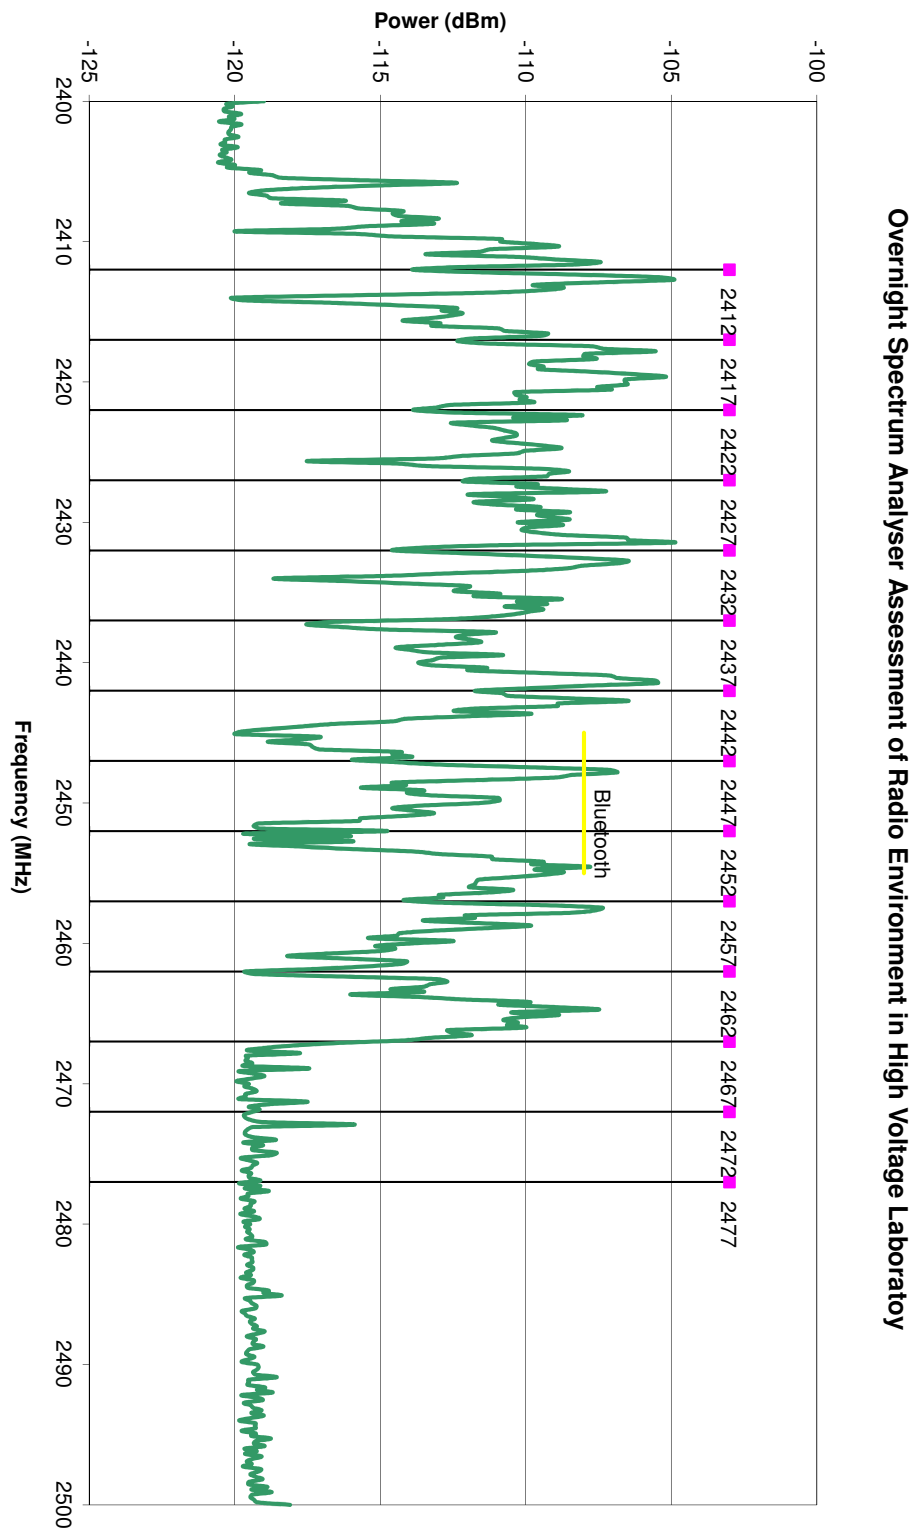
Figure 8.21: Overnight Spectrum Analysis of Radio Environment
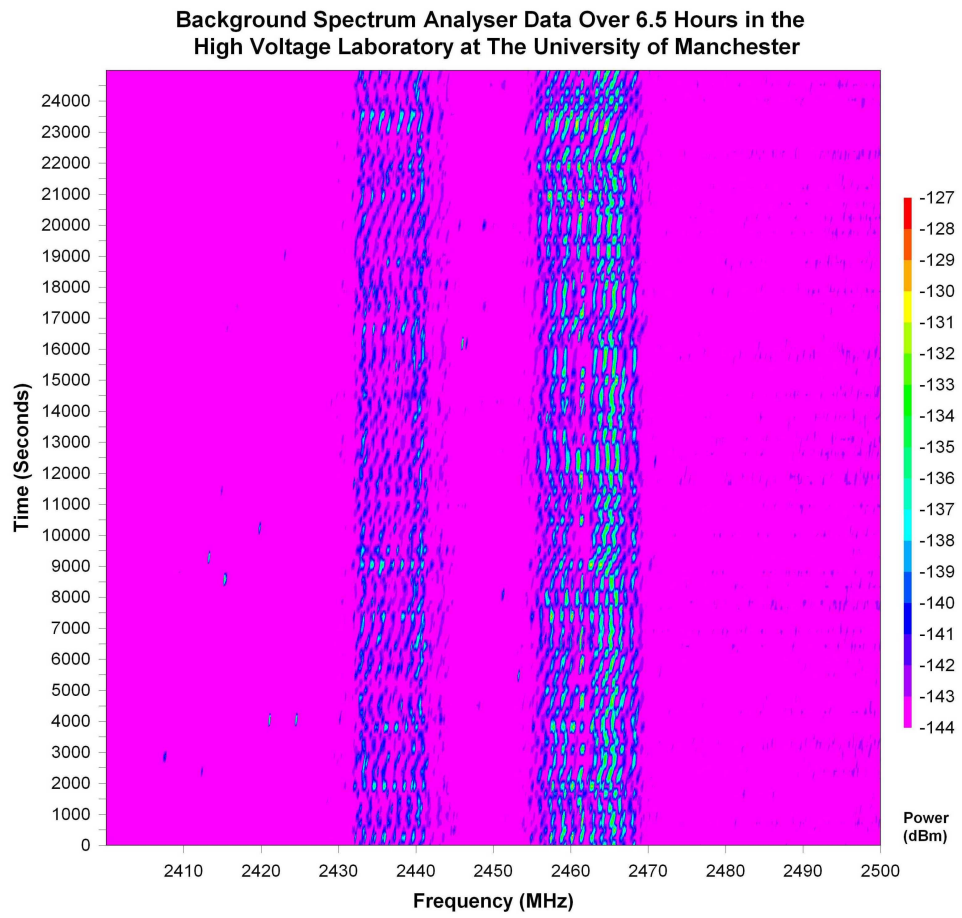
Figure 8.22: Day Time Spectrum Analysis of High Voltage Laboratory

less of the presence of these channels there was an expectation for other channels to operate successfully, although possibly at a reduced capacity.

This then suggests that either the low bandwidth and modulation scheme used previous were not susceptible to interference present or the high levels of external interference visible now were not there at the time. Without similar test sequences under current conditions, a conclusion was not possible. The former suggestion does not seem likely as spectrum analyser measurements made when the stability problem was discovered suggested the interference was broadband in nature, covering the majority of the 2.4 GHz band.

Measurements were taken with the spectrum analyser configured to hold the maximum value seen while using a 18 dBi Yagi antenna. Figure 8.23 shows the results for various directions in the HV laboratory. When compared to Figures 8.21 and 8.22 the noise ceiling appeared to be 20 to 30 dB higher.

While the change in antenna may have lead to increased sensitivity, the increases seen cannot be purely attributed to antenna gain. Therefore, if similar external interference was present during initial investigations, it is unlikely that high accuracy data transfers would have been seen; particularly when no forward data correction algorithms were used. This reiterated the need for more detailed investigation into the source of this increased interference.

The first task was to repeat the long term cycled spectrum measurements using the same 8 dBi omnidirectional antenna. Measurements were made over an 11 day period at 14 minute intervals; similar measurements were made in an office environment over 30 hours for comparison. Results can be seen if Figures 8.24 and 8.25 and were captured in December 2010 and January 2011.

The office based results are clearly noisier, showing increased activity in WLAN Channels 1, 6 and 11. Energy in Channel 11 cannot be seen in the former chart, while energy for Channels 1 and 6 are significantly lower. The office in question is the same environment where the system was developed and initially tested. It is unlikely that the energy seen here was not also present when the system was developed, as such an explanation for system instability in 'quieter' environment cannot be currently given.

However a difference in the RF bandwidth used for these measurements meant direct comparison between them and Figure 8.22 was not possible. However, as-
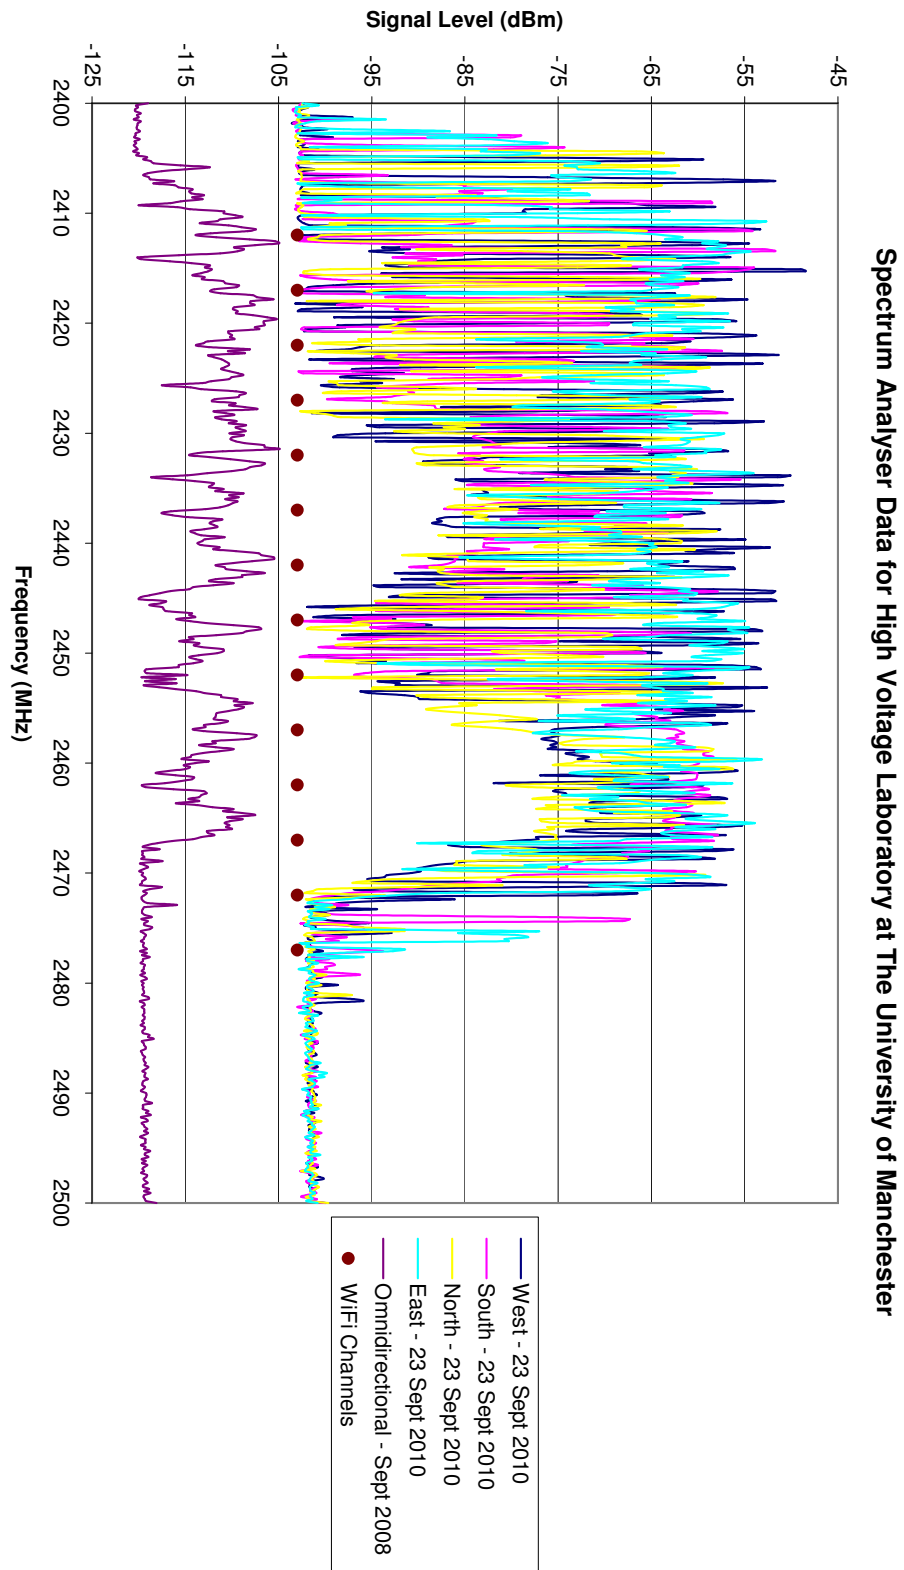
Figure 8.23: High voltage laboratory snapshot spectrum measurements
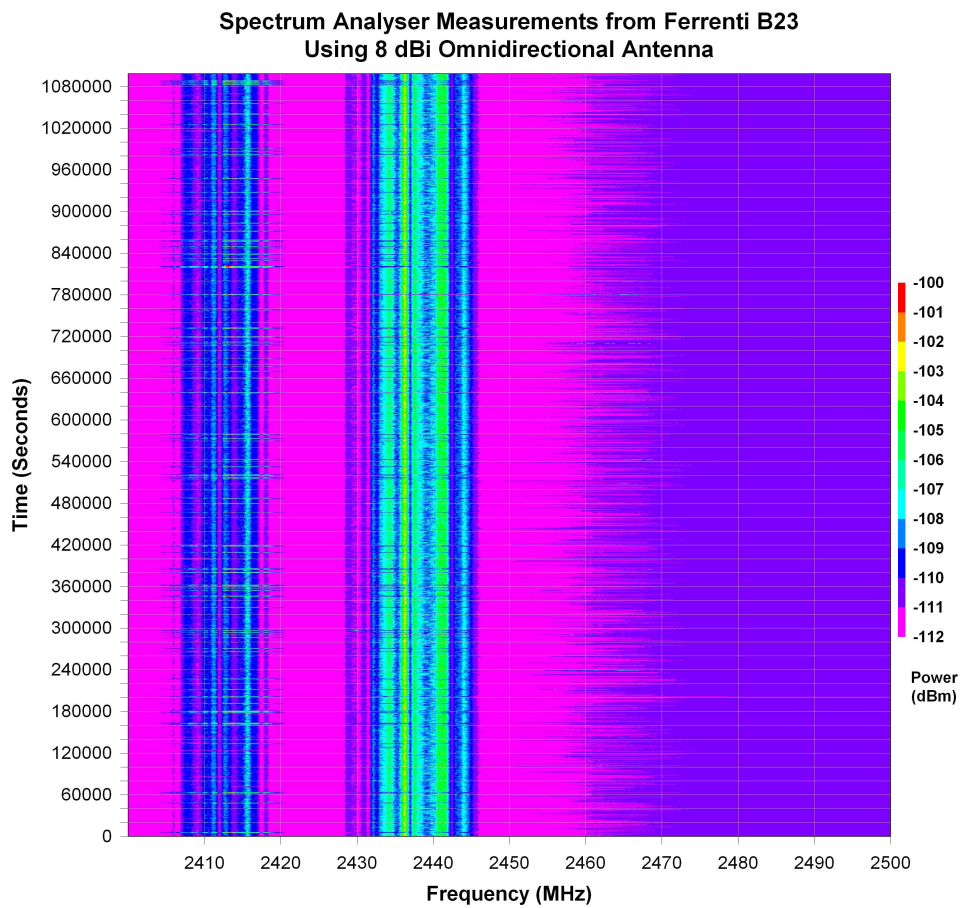
Figure 8.24: High voltage laboratory spectrum measurements (11 days)
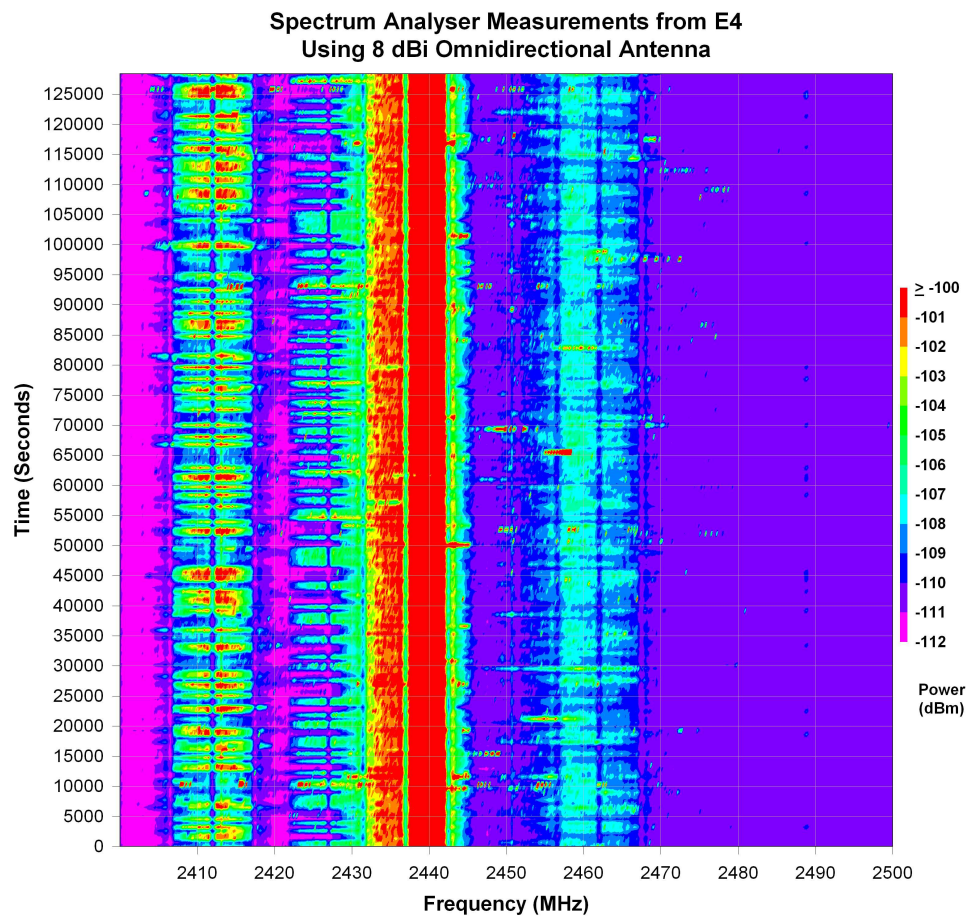
Figure 8.25: Office based spectrum measurements (30 hours)

sessing these results independently in line with the increased RF Bandwidth figure, again suggests no reason as to why the WLAN link failed so frequently.

8.7.2   Directional Spectrum Analysis and Reflected Signals

Previous results indicated a possible relationship between field strength and possible data rates & signal strengths, however anomalies prevented a more confident statement from being made. It has long considered interference within the test environment to be random in nature and it has been proposed that it may be a contributing factor to the anomalies.

In order to identify the interference source, the radio energy radiating from different directions within the laboratory was assessed in May 2011. As before a two minute 'maximum hold' window was used in conjunction with the same 18 dBi Yagi antenna. A second dataset was also collected but with the addition of 16 dB of in-line attenuation, with the aim of simulating an antenna with the gain of a typical whip antenna (2 dBi) but with the directivity of a Yagi.

Figure 8.26 and 8.27 show the collated spectrum results for the non-attenuated and attenuated datasets respectively. Of primary interest is that only WLAN Channel 6 is clearly visible in both datasets, with small energy traces from Channel 11 noticeable when the antenna was not attenuated. This data suggests that the WAP operating on WLAN Channel 6 is located directly behind the laboratory's southern wall, with Channel 11 being generated towards the west. However the laboratory's western side is an external wall, this evidence is complicated by the fact the spectrum measurements from the control room with the antenna facing in a south-westerly direction gave a strong energy trace for Channel 11 (see Figure 8.30).

There are two possibilities for this turn of events; firstly, the partition between the control room and the laboratory itself is enhancing signals in certain bands. This suggestion however is very unlikely and is not supported by other results. Another suggestion is reflected signals which have circumstantially caused results to suggest that the WAP is located in a false position. This could be verified by repeating this directional spectrum analysis in multiple locations within the laboratory and producing similar charts; if truly beyond the westerly wall all further datasets
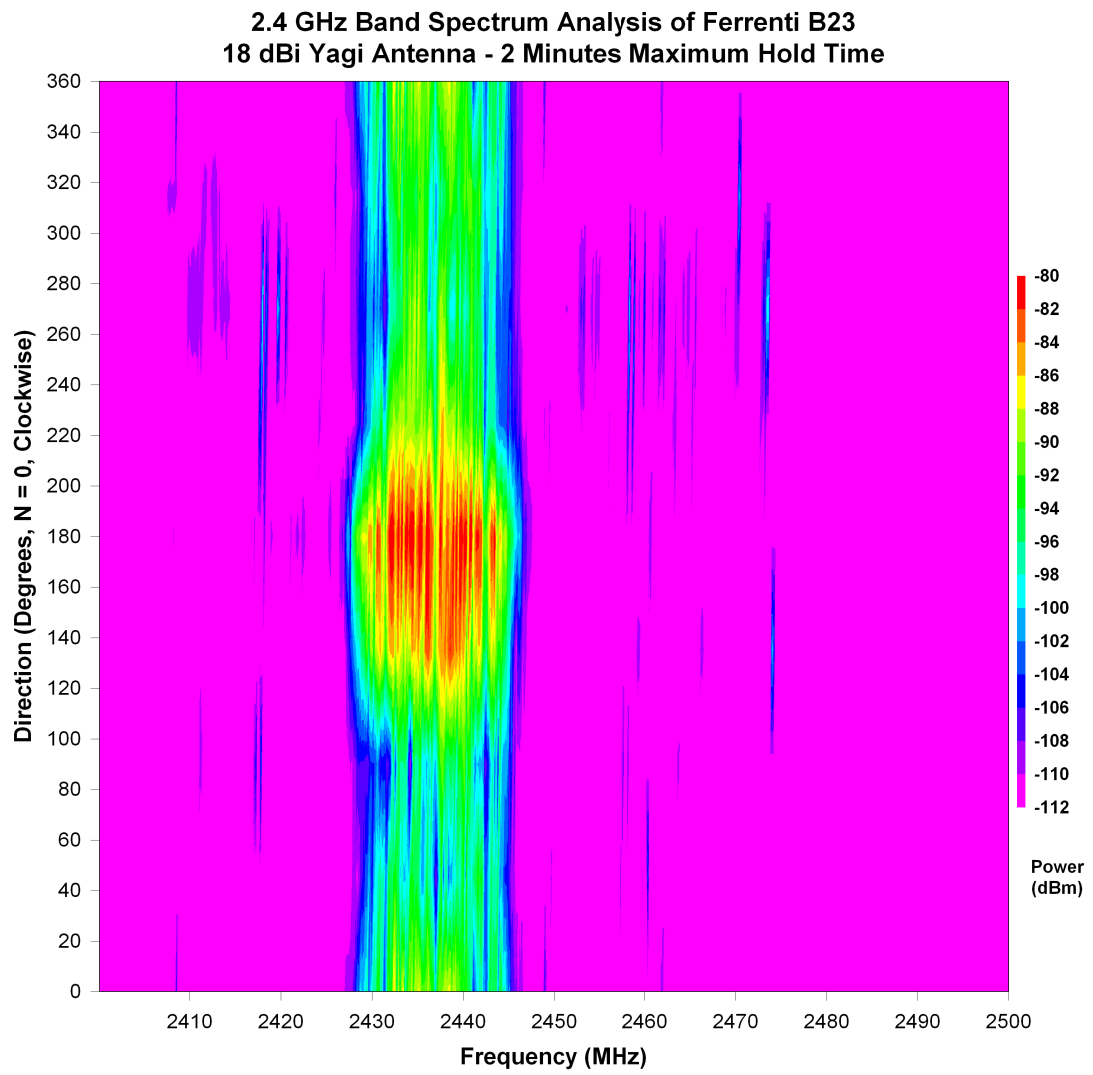
Figure 8.26: Directional spectrum analysis of high voltage laboratory
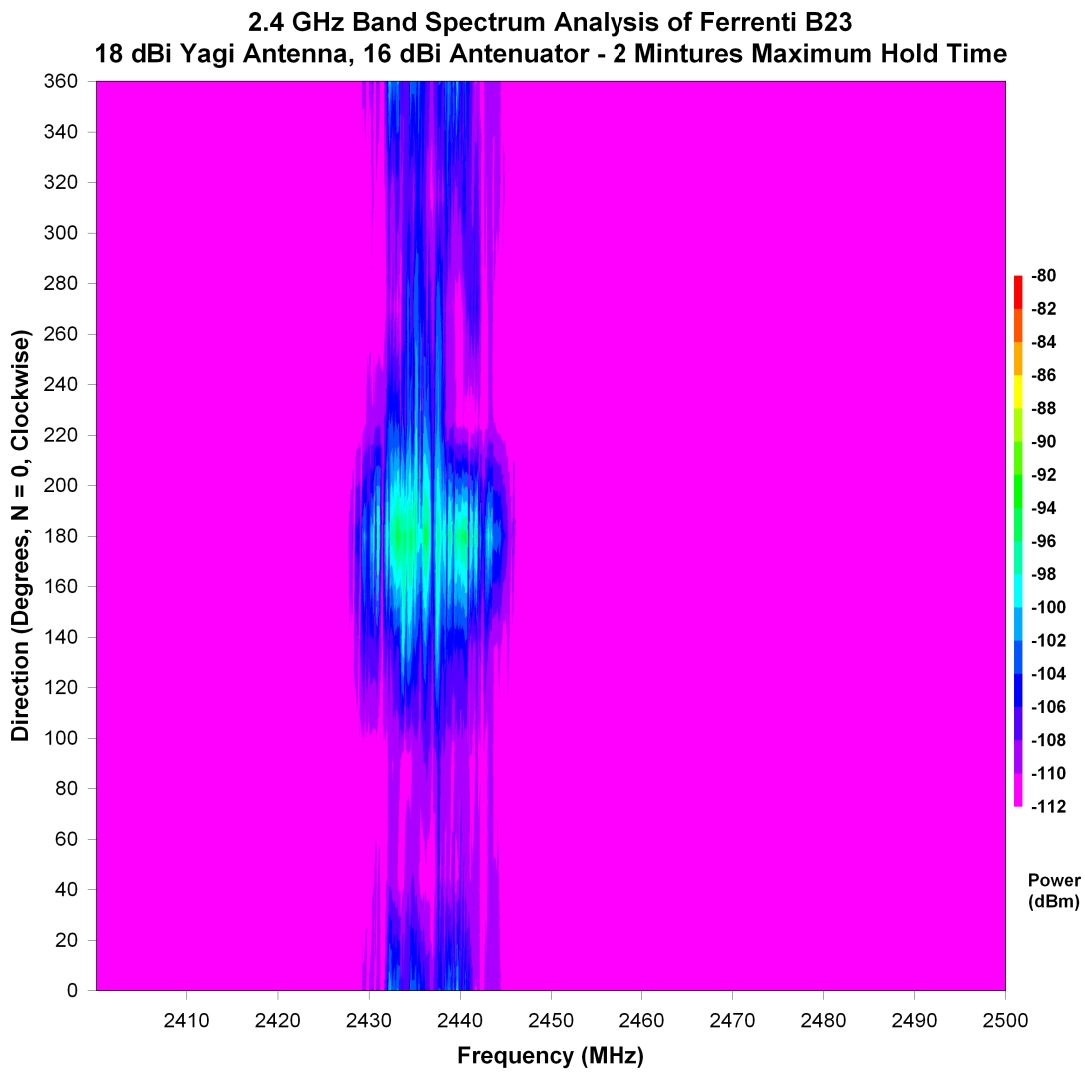
Figure 8.27: Directional spectrum analysis of high voltage laboratory using 16 dB attenuation

should corroborate this suggestion. However, as the actual location of the WAP is irrelevant to the wider research, further effort was not put into radio-metrically locating the device, although it provides a interesting concept.

The theory of reflected energy also manifests itself in the Channel 6. As noted previously, the highest energy source appears to be generated from the south. However there is also energy detected at other compass directions; notably the north experiences a higher energy level than either the east or west. If reflections were not present, energy would normally only be expected in the general direction of the source. This multi-directional energy can also be seen on the attenuated dataset.

There is however one other possibility for some of this energy, namely clients of the WAP in question. While this may account for the low energy levels seen in some directions, it is unlikely to fully account for the increased energy seen directly opposite the source.

Further evidence of reflected signals is visible from the TCP statistics obtained from Wireshark shown in Table 8.1. Wireshark recordings were made via both wired and wireless interfaces for Channel 13 when the transmitter was located at shed 57[8]. These figures are from the test case where the signal was attenuated under energised and non-energised conditions. Refering to the meaning of these

	Energised Ethernet Capture	Energised Wireless Capture	Un-energised Ethernet Capture	Un-energised Wireless Capture
Bad Checksum	46794	1	46632	0
Out-of-Order Segments	0	1001	0	1006
Previous Segment Lost	1	15781	5	8597
ACKed Lost Segment	1	14739	1	31881
Fast Retransmission	2	26	5	17
Retransmission	0	182	31	151
Duplicate ACK	32	34408	36	28224
Average Data Rate (kBps)	1322.68	1430.14	1433.32	1106.56

Table 8.1: TCP Transfer Statistics for Transmitter at Shed 57, Captured via Wired and Wireless Methods

---

[8]Due the processing speed of the computer these recordings were made over consecutive file transfer operations.

statistics as noted in Section 8.3 a few conclusions can be reached. Firstly Previous Segment Lost and ACKed Lost Segment values can be neglected as it cannot be guaranteed that the wireless listening station receives all transmitted packets, hence detections of this type will be common in wireless data captures. The similar is true of Bad Checksum as they may be caused for many reasons.

The more useful indicators are Out-of-Order Segments, Duplicate ACKs, Fast Retransmissions and Retransmissions, where in most cases WLAN captures shown a higher number of packets marked as such. Considering the receiving host re-orders packets rather than any intermediary network device, if Out-of-Order Segments were truly an issue, it should be visible regardless of the capture interface. Again, Duplicate ACK values are approximately 1000x greater for WLAN captures as opposed to wired captures; this further reiterates the possibility of reflected signals. Also note the decrease in data rate between the wired and wireless transfers for the non-energised trial. These trials were conducted within 5 minutes of each other yet possess 200 kBps difference in transfer rate, which additionally suggests an rapidly varying external source of radio interference.

In the comparison between non-attenuated and attenuated evaluations, it appears that the antenna gain is of greater benefit than the directivity. Although it should be noted that the measurements could not have been achieved without the directivity itself. This is in comparison to the results from Section 8.5.1 which revealed the opposite was true when a communication link is considered.

For comparative purposes Figure 8.28 and 8.29 show similar spectrum measurements obtained eighth months earlier. Figure 8.28 has had its scaling adjusted to match the those used in Figures 8.26 and 8.27, while Figure 8.29 uses a more appropriate measure.

Analysis of these figures suggests that the laboratory has lost a number of sources of radio energy over the intervening months, indicated by the increased energy levels seen in September. However it should be noted that in these earlier datasets the spectrum analyser was found to have been configured for a filter bandwidth of 300 kHz rather than 100 kHz as used in the latter more structured experiments. Therefore, while the results are not directly comparable, a number of things can be noted from these earlier results. There are clear energy signatures for WLAN Channels 6 (2437 MHz) & 11 (2462 MHz) and hints of a WAP operating on either
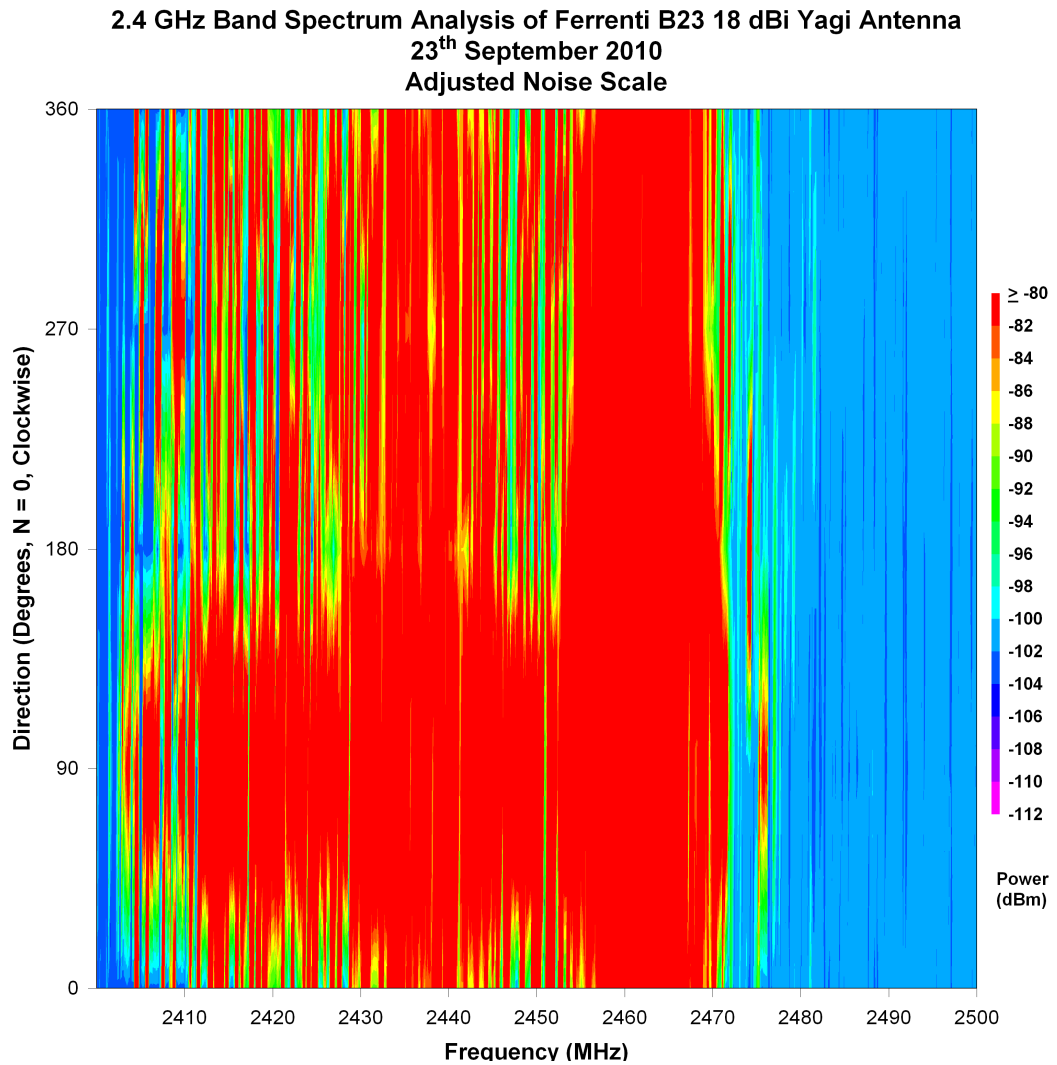
Figure 8.28: Spectrum measurements in September 2010, scale adjusted to match Figure 8.26
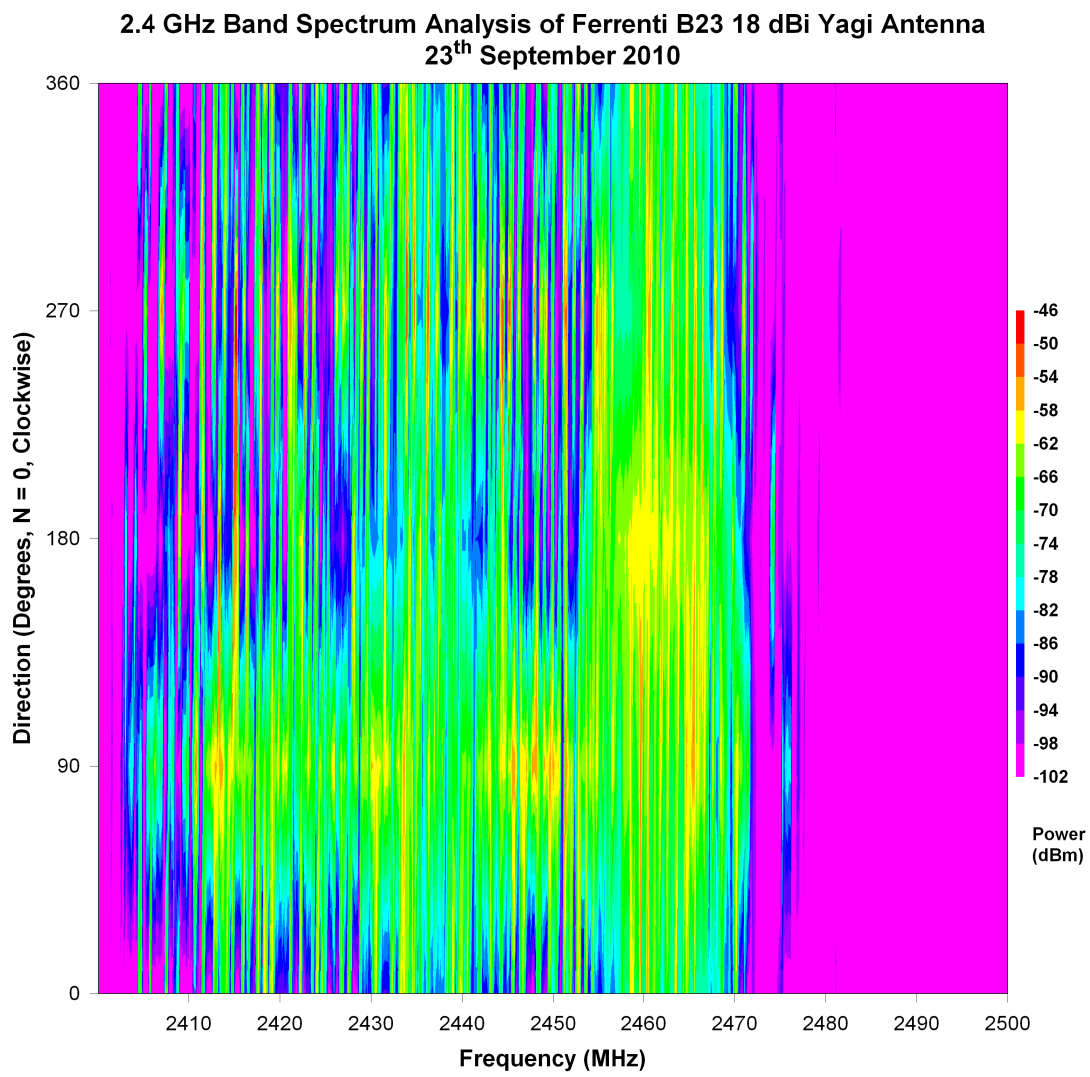
Figure 8.29: Spectrum measurements in September 2010

Channel 2 (2417 MHz) or 3 (2422 MHz). Furthermore the WAP(s) generating these signals appears to be located towards the east beyond the laboratory's control room, a change from the current apparent operating position.

8.7.3 RF Spectrum Generated by Energised Insulator

In order to assess if the insulator itself was radiating radio energy, the Yagi antenna was directed at the previous experimental set-up and the insulator energised to approximately 175 kV, with subsequent 5 kV increments. At each stage the spectrum analyser was used in 'maximum hold' mode for approximately two minutes after which the data captured was stored. The trace was cleared prior to the next 'maximum hold' operation.

Figure 8.30 shows the result of this analysis. It can be seen that there is no substantial radio energy besides two bands corresponding to WLAN Channels 6 and 11. Note how energy in Channel 1 is no longer present in comparison to measurements made in January 2011; similarly missing in May 2011's directional analysis.

The cause of multiple energy 'streaks' at certain line voltages, while appearing indicative of some kind of correlative phenomena, can be simply explained as a consequence of limited datasets, the garph's inherent interpolation and the fact that datasets cannot be simultaneously captured. If the latter were possible then external energy would be identically captured within all datasets, something which cannot be guaranteed here.

8.8 Baseline Testing of WLAN Transfer Rates

In order to further confirm the stability and quality of the developed system a number of baseline tests were carried out. This used the client-WAP method of communication. The system was tested in office environments over distances of 3.5 m, 8.5 m and 20.8 m; during which only two issues were encountered. Firstly after transmission had completed on Channel 8 of the 3.5 m test, where the BeagleBoard's WLAN interface appeared to be associated with WAP[9] but

---

[9]Based on the information provided by the iwconfig command

2.4 GHz Spectrum Analysis of Energised Composite Insulator
18 dBi Yagi Antenna - 2 Minute Maximum Hold Time
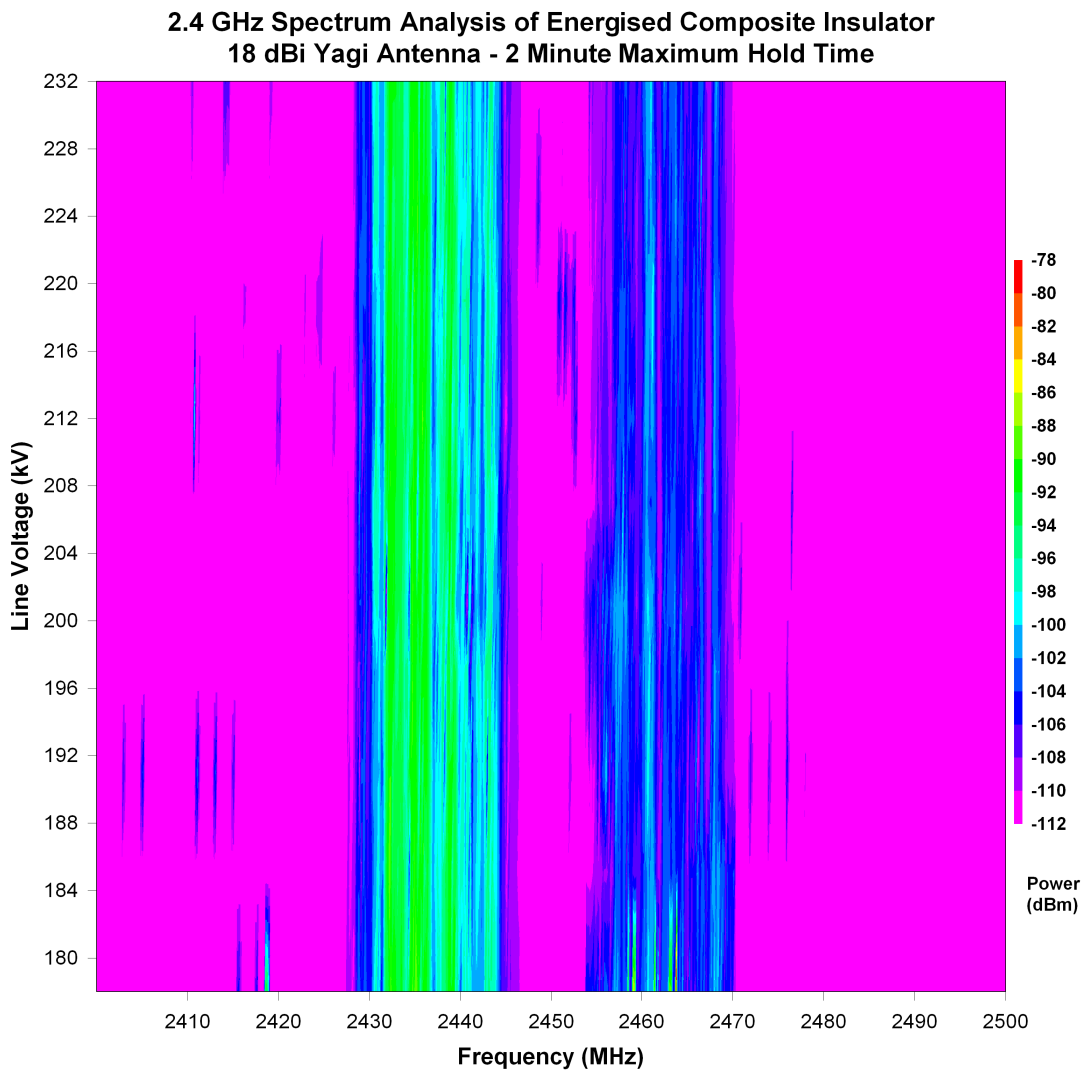
Figure 8.30: Spectrum analysis in direction of energised composite insulator

communication over this link was not possible. As this issue did not appear again in any channel at either of the other two test distances, it was concluded that this was either down to a external burst of radio energy within Channel 8 or by a software glitch in either the WAP or BeagleBoard.

The second issue related to file transfers using TFTP; attempted a number of times at Channel 1 at 3.5 m separation. Here the link appeared to fail after approximately 14 files were transfered. When the transfered files were examined, the files were revealed to be identical. While this is what is needed for a file transfer, it prompted the need for the TFTP mechanism to be re-examined.

It was known that TFTP typically uses UDP for data transfer and it was assumed that TFTP itself adds very little in terms of integrity checking, merely serving to reconstruct the file as UDP packets are received. Therefore, based on UDP's connectionless nature, it was assumed that received files would not necessarily be identical to the original. However, an in-depth look at TFTP's design[81], revealed that TFTP implements a lock-step mechanism whereby every received packet is acknowledged. If the sender does not receive this acknowledgement, the previous packet is retransmitted; the next packet is not sent until an acknowledgement for the previous packet is received.

This lock-step architecture greatly slows down the transfer process and this was observed during the initial TFTP transfer attempts. Furthermore, as only a single packet is in flight at any given time, packets will always arrive in order, although they may be duplicated. Duplications are identified via the 'block number' field in the TFTP header.

TFTP was initially considered as mechanism for assessing how degraded a UDP based video stream would become under high voltage conditions, on the basis that no 'high level' integrity checking is performed at the UDP level itself. Subsequently, if the TFTP protocol lacked the lock-step mechanism, then by comparing the transfered file size to that of the original, the average data loss through either packet loss or corruption could be calculated.

However due to the lock-step mechanism little additional information above that discoverable via FTP can be gleaned from TFTP. On this basis, it was debatable whether to continue to test with TFTP. This little additional benefit combined with the wireless link stalling after approximately 14 TFTP transfers lead to the

decision to stop TFTP transfer testing. One possibility given adequate time would have been to a write a small application to transfer and receive data over a UDP connection without implementation of higher level integrity checking functionality.

8.8.1 Results from Baseline Testing

As before, data rates were averaged on a channel and test-case basis, this is shown in Figure 8.31. In all channels except 8, transmission at 20.8 m separation shows at least marginal improvement over transmission at 3.5 m. Transmission at 8.5 m appears to be at least marginally better in half the available channels than at 20.8 m. Of particular interest are Channels 10 and 11, where there is a distinct improvement at 20.8 m compared to either 3.5 m or 8.5 m; however an average transmission rate of >1500 kBps is reasonable. At a separation of 20.8 m the average data rate was >2250 kbps. Based on these results, lack of data capacity does not appear to be a cause of the link failures experienced in the high voltage laboratory.

Wireshark packet captures were subsequently analysed to see if anything was failing at the packet level; the TCP statistics and the packet signal strength, reported as RSSI values in dBm, were of most interest.

The wireless capture interface was placed in the transmission path close to one of the WAP's antennas; while the interface should see all the traffic flowing between the WAP and BeagleBoard, it could not be guaranteed. Therefore analysis may indicate a higher figure for missing TCP Packets and Lost Acknowledgement than was actually experienced. An alternative mechanism to capture all the communication between the two devices would have been to capture on the wired network interface. However this had the disadvantage of losing the 'radio' information which would have been removed by the WAP.

Figures 8.32, 8.33 and 8.34 show the TCP transmission statistics collated by WireShark for transmission distances of 3.5 m, 8.5 m and 20.8 m respectively; Section 8.3 details the meaning of these statistics. These are acceptable values given that a wireless packet capture was used and that not every packet transmitted is guaranteed to be captured. At all channels and test scenarios values for 'Previous Segment Lost' or 'ACKed Lost Packet' appear to the consistent, with
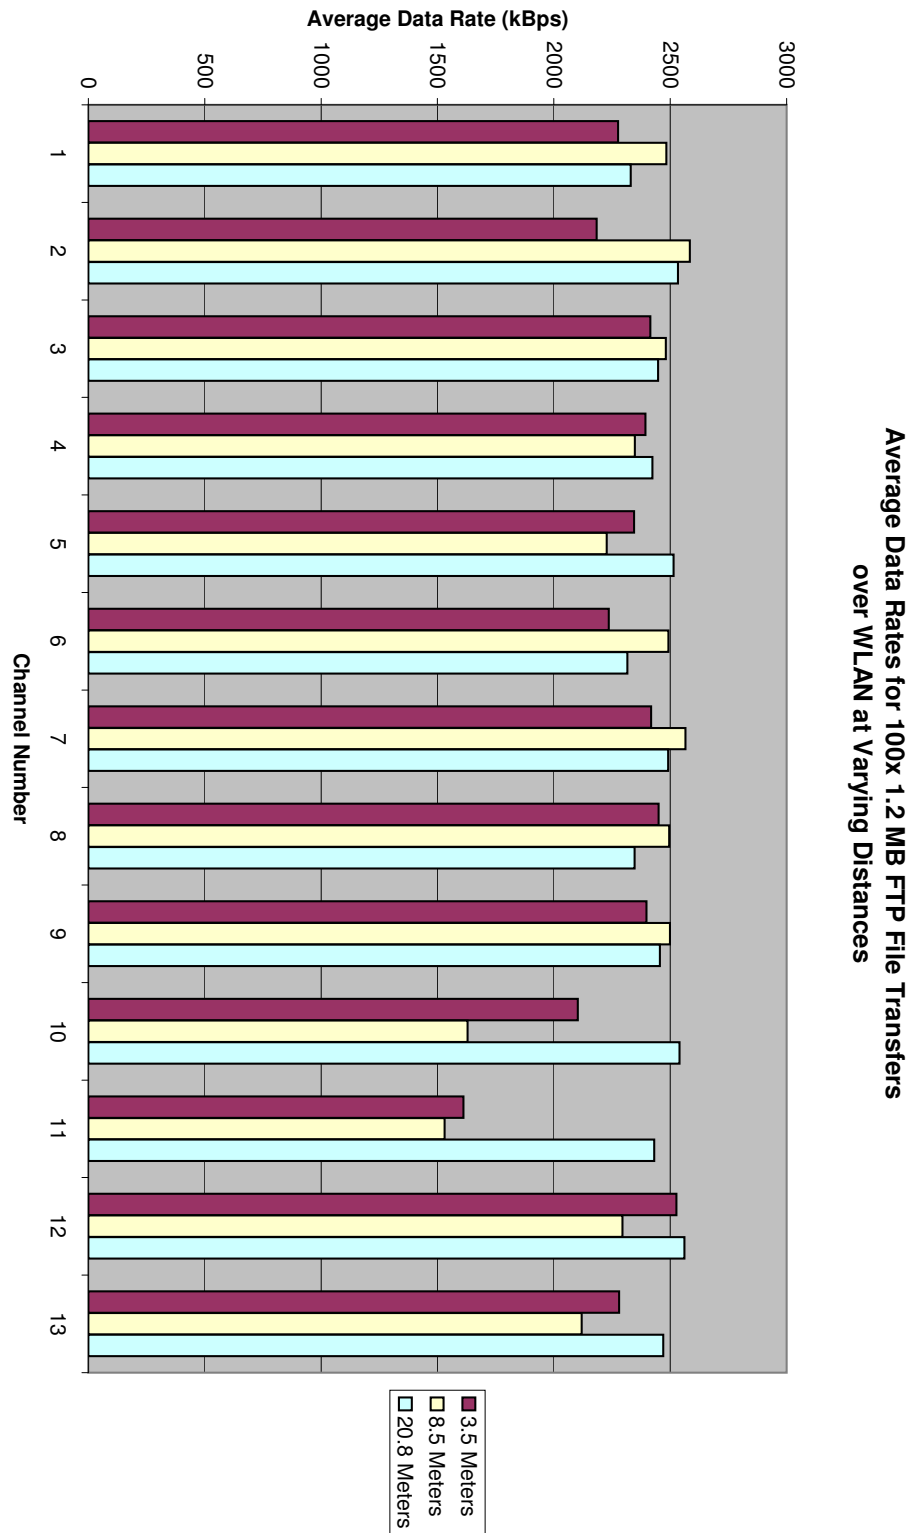
Figure 8.31: Average data rate between WAP and BeagleBoard over Different Distances
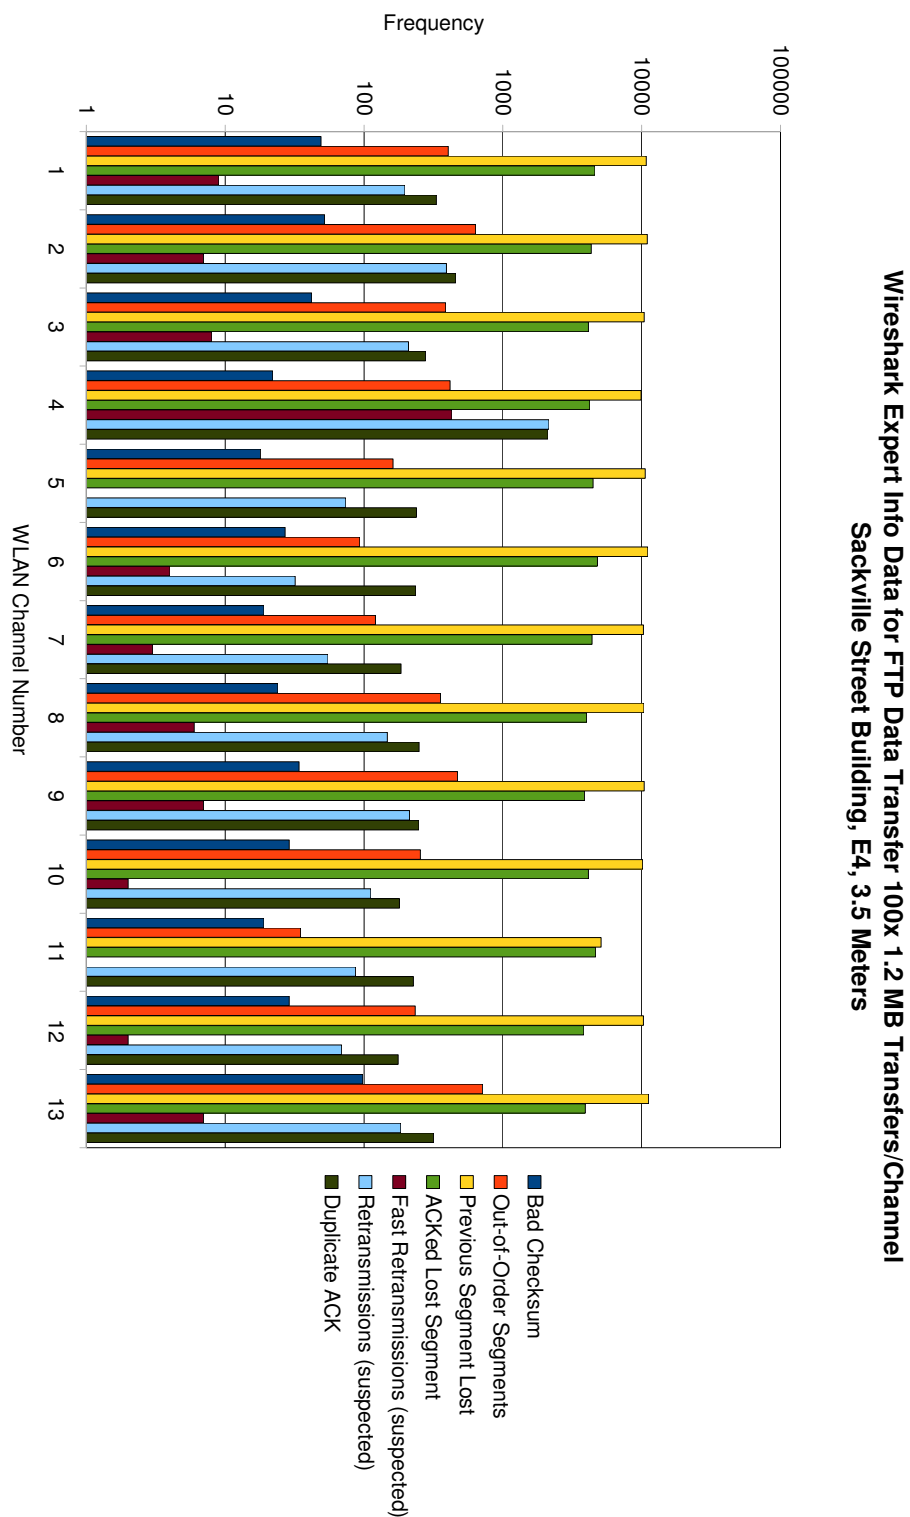
Figure 8.32: Wireshark Expert Info Data for FTP at a Transmitter-Receiver Separation of 3.5 Meters
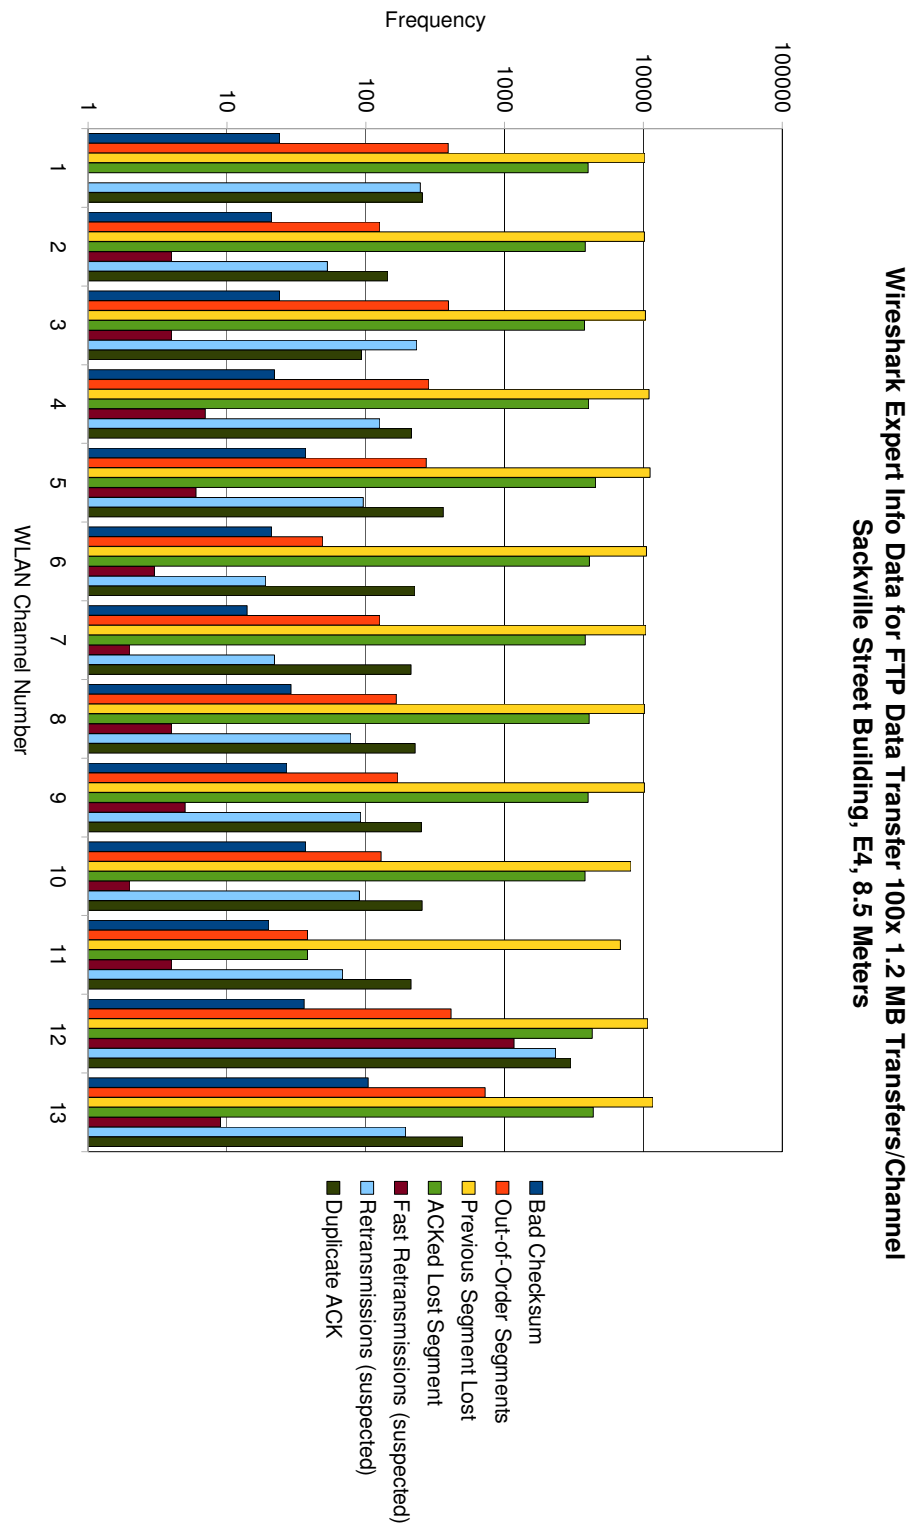
Figure 8.33: Wireshark Expert Info Data for FTP at a Transmitter-Receiver Separation of 8.5 Meters
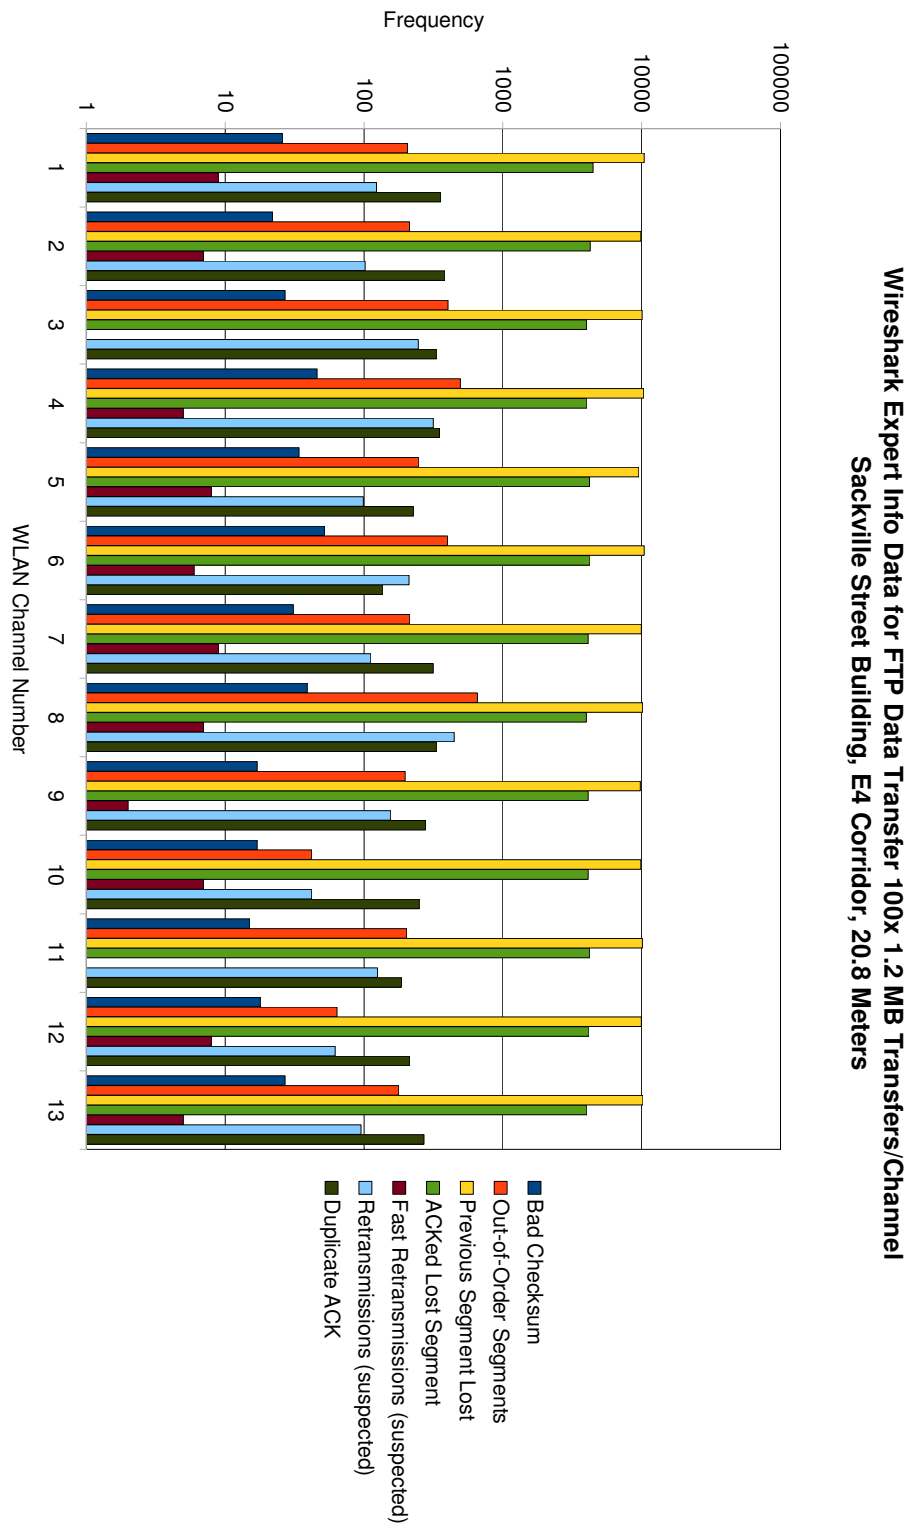
Figure 8.34: Wireshark Expert Info Data for FTP at a Transmitter-Receiver Separation of 20.8 Meters

the exception of Channel 11 at 3.5 m and 8.5 m.

More revealing statistics relate to Out-of-Order Segments, Duplicated ACKs, Suspected Fast Retransmissions and Suspected Retransmissions. On a large network, TCP packets may take differing routes between source and destination, thereby possibly causing packets to arrive out of order. In the configuration used here there is only one route between source and destination, so one could assume that packets should always arrive in the correct order.

However this single route concept is only valid at the network-level and above. At the physical-level, multiple routes may be traversed due to multi-path propagation of the radio transmission; debated in more detail in Section 8.7.2. The frequency of Out-of-Order Segments is similar across all ranges on Channel 3, while Channels 7 and 13 show similarities at separations of 3.5 m and 8.5 m. On other channels there appears to be a wide spread in measured values.

Fast Retransmissions are broadly similar across all channels and separations, numbering less that 10 in the majority of cases. Notable exceptions are at Channel 11 at 3.5 m and Channel 1 at 8.5 m where no Suspected Fast Retransmissions were recorded and for Channel 12 at 8.5 m where Suspected Fast Retransmissions exceeded 1000. However, this channel also experience higher that average Suspected Retransmissions and Duplicate ACKs, as did Channel 4 at 3.5 m. In other cases Suspected Retransmissions and Duplicate ACK frequencies are broadly similar.

Prior to assessing the wireless signal strengths, Wireshark's filtering mechanism was used to isolate packets which were destined to, send from or broadcast from the BeagleBoard. This then removed packets sent from other wireless networks operating on the same channel in the vicinity as well as broadcast, Address Resolution Protocol (ARP) and Radio Beacon packets from the 'home' WAP.

When the filtered packets were combined on a test-distance basis, each data set contained approximately 500,000 data-points. The most efficient way to visualise this information was to first calculate the frequency distribution in 5 dBm bins from -100 dBm to 0 dBm. As the number of packets recorded in each channel differed, it was more appropriate to graph percentage occurrence rather than raw number of packets in the bin. The results are shown in Figure 8.35, 8.36, 8.37.

As expected, in general, as the transmission distance increased, more packets with a lower RSSI were seen, extending to the -60 to -65 dBm bin in Figure 8.37. In
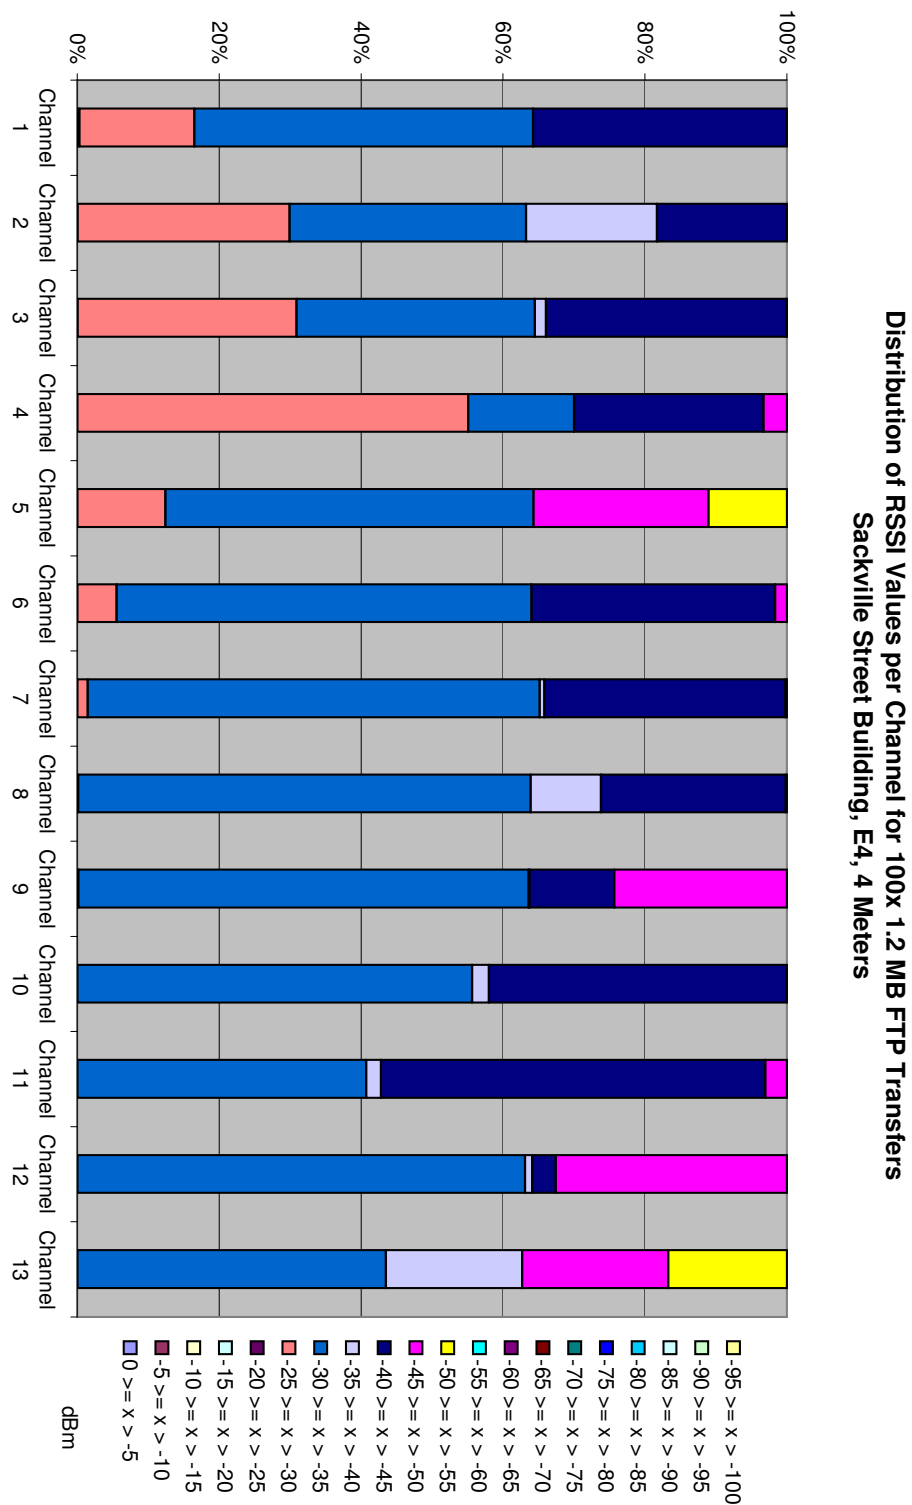
Figure 8.35: RSSI Frequency per WLAN Channel for a Transmitter-Receiver Separation of 3.5 Meters
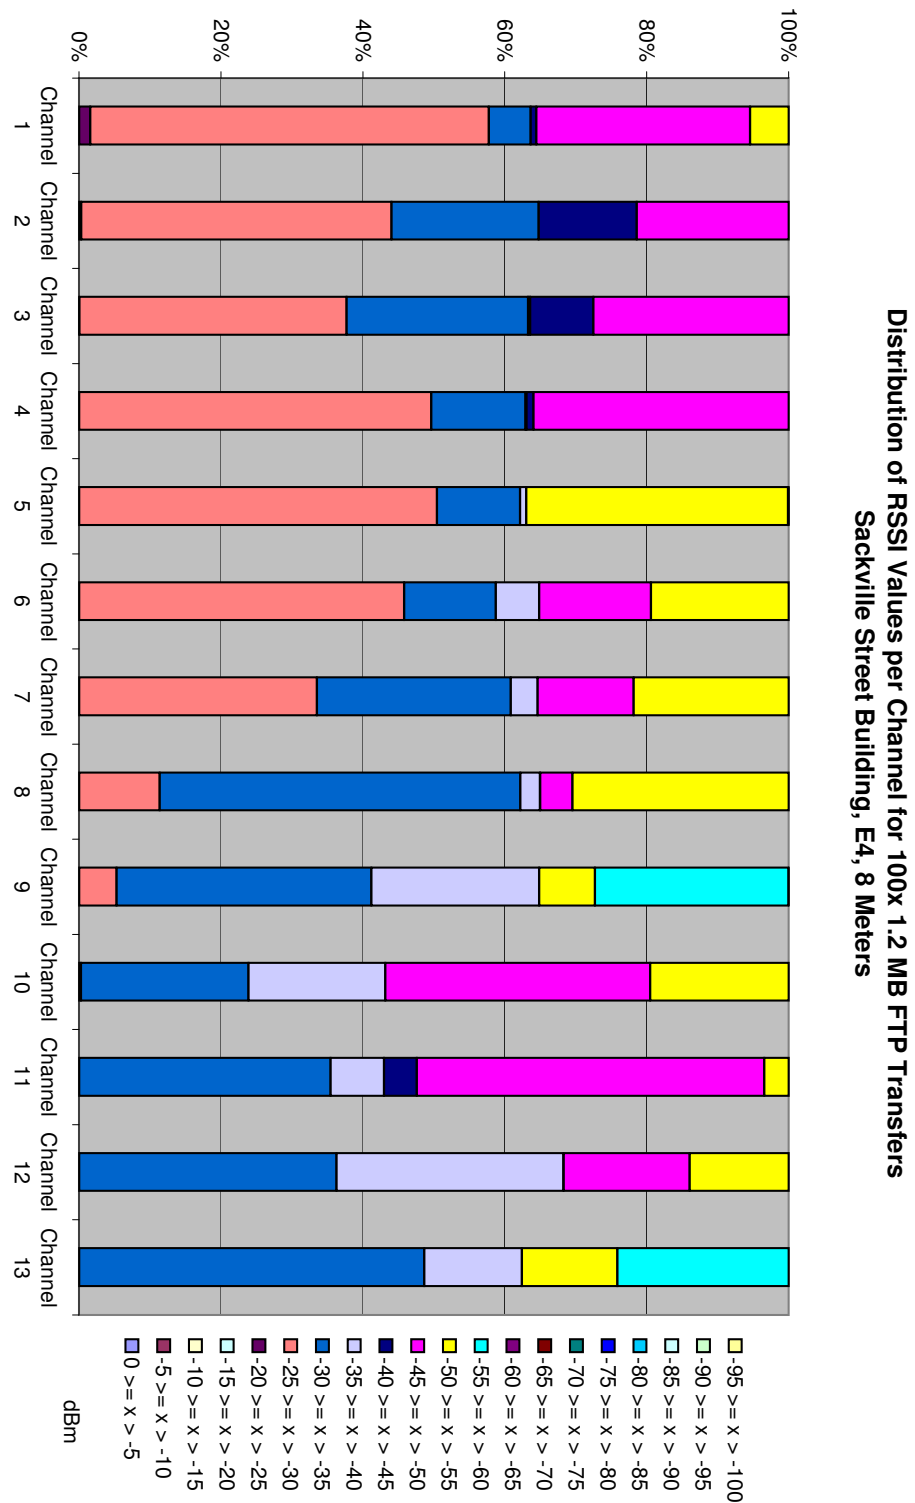
Figure 8.36: RSSI Frequency per WLAN Channel for a Transmitter-Receiver Separation of 8.5 Meters
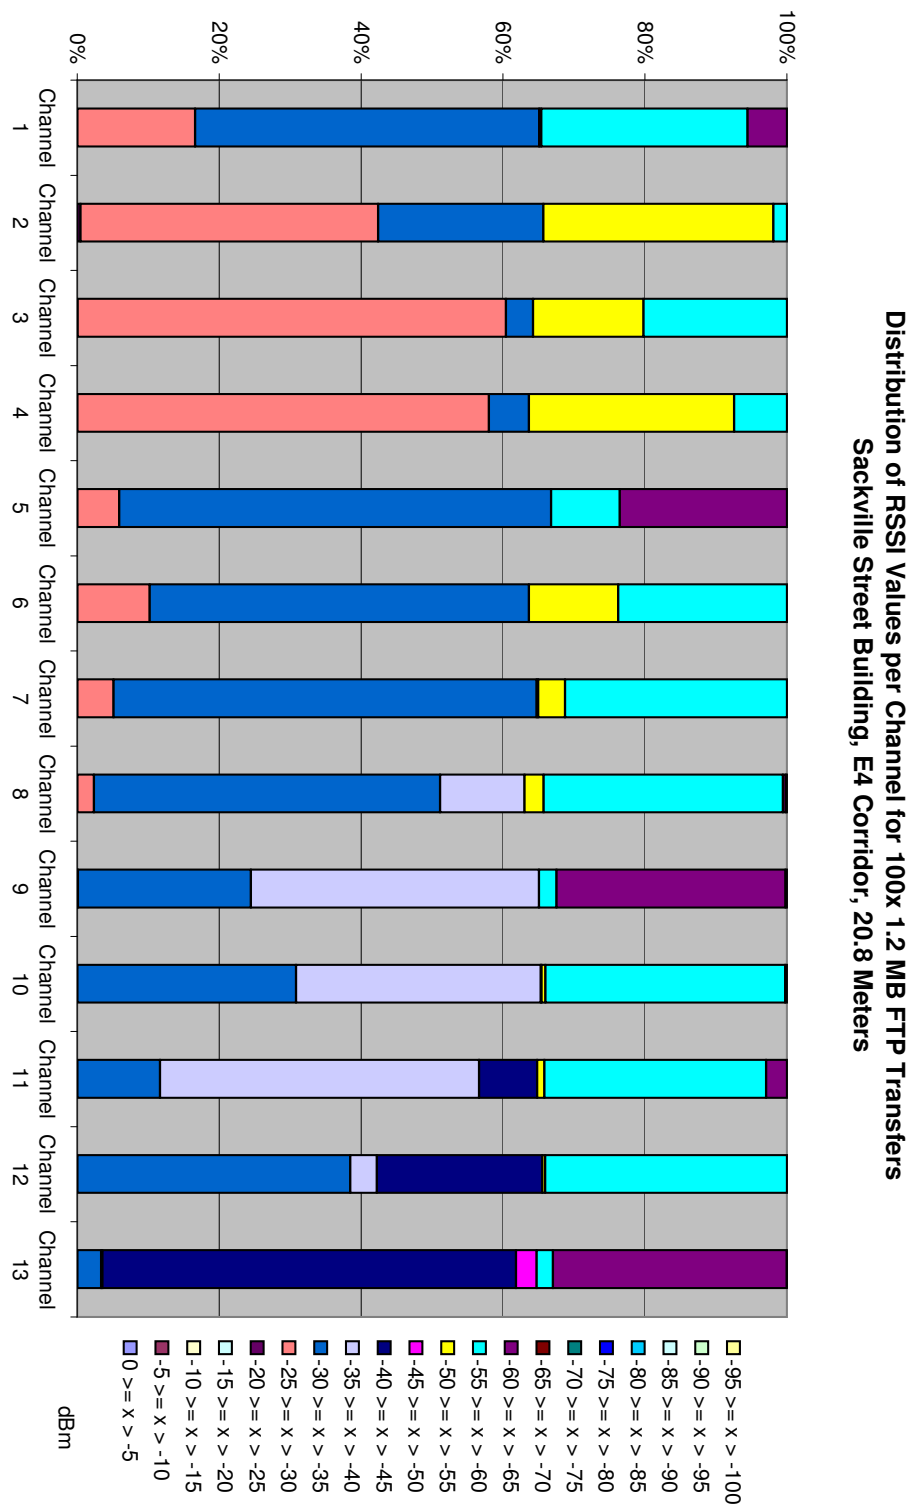
Figure 8.37: RSSI Frequency per WLAN Channel for a Transmitter-Receiver Separation of 20.8 Meters

the 20.8 m experiment, no packet had a RSSI greater than -50 dBm across all channels, compared to only Channels 5, 9 and 13 at 8.5 m, while the minimum RSSI was -55 dBm at 3.5 m in Channels 5 and 13. Furthermore in the majority of cases RSSI values are consistent within each channel at each distance.

8.8.2   Conclusions on Baseline Results

Based on this analysis, no definite conclusion can be given as to why the WLAN link was unstable within the high voltage laboratory. Given that the theoretical maximum data rate of a IEEE802.11g link is 6912 kBps and the minimum achieved was 1500 kBps, a link utilisation of at least 21% was seen. Transfer rates were consistent at each distance thereby negating any concerns about reduced RSSI.

It was on this basis that work detailed in Section 8.7 was undertaken.

8.9   Summary

This chapter extends work presented in Chapter 6 by discussing the development of a transmission system based around WLAN technology. Results from energised testing on a representation of a high voltage transmission line shows that data rates above 1000 kBps are achievable along the majority of a composite insulator string, exceptions being between sheds 10 & 20 and above shed 40: when considering a combined average data rate over all WLAN channels.

Consideration on an individual channel basis suggests that locations close to the earthed (up to shed 16) and energised (above shed 40) ends appear have the wide ranging recorded data rates. At shed 1, Channels 1-3 have data rates below 1500 kBps, while Channels 5-9 and 11 experienced data rates in excess of 2000 kBps and Channels 4 and 10-12 had data rates in between. Above shed 40, the most experienced data rates were below 1500 kBps, with the majority actually be less than 1000 kBps. However a few cases were seen where the channel data rate reached between 1500-2000 kBps and hence lead to a increase in the overall average data rate across all WLAN channels at sheds 51-59.

It was shown that whip antennas have limits of usability within the representation tested which can then be exceeded through the use of cylindrical waveguide transi-

tion antennas which can provide additional protectionfrom corona, to the radiating element. While the simple antennas constructed only raise the limits marginally, dedicated design and fabrication should allow for further improvements. Based on this and evidence seen in Chapter 6 further suggests that corona may be the cause of the data transmission errors seen, especially at the energised end.

Corona may even be able to explain the low overall average data rate seen at shed 16. Surface damage or contamination during handling and use may have lead to increased corona activity at or around shed 16, which thereby caused the data rate recorded to fall in comparison to adjacent measurement locations.

The results presented in this chapter additionally hints at the possibility of a link between field strength and achievable data rates, especially for sheds above number 26 from the earthed end. Through verification and refinementthis could lead to the development of a new field strength measurement method.

In light of early stability problems encountered extensive experiments in a different environment were conducted to ascertain equipment stability. This showed that the equipment could indeed be stable over long durations and revealed no clues for the cause of repeated failures seen in the wireless link while in the test laboratory.

This then lead to a detailed study of the radio environment in the HV laboratory. Through detailed spectrum analyser measurements, there is a suggestion that multipath effects and external sources of radio energy are being seen during investigations, as such corona cannot be treated as the sole contributory factor to the reduction in WLAN data rates. The varying levels of unwanted radio energy and their potential interference with gathered results have been noted.

Chapter 9

Vision System Development

In Chapter 7 it was discussed that a consumer webcam attached to a Linux based OMAP platform would provide the quickest development route to a functional system. This chapter documents the way in which real-time video and image capture on the Linux platform was acheived, the problems encountered and their resolutions. The camera's mounting position and the optical manipulation of images to improve image size are also noted.

9.1   Live Video Streaming

The BeagleBoard was indented to be used as a proving platform for the high data rate wireless link, prior to moving to the Gumstix Overo to develop the camera interface. While searching the OpenEmbedded[1] application repository for FTP and TFTP servers, a number of webcam utilities were discovered. These tools allowed either photographs or streaming video to be captured from a webcam; the most interesting was `w3cam`, which captures still images from the webcam and presents it as a web page via a Common Gateway Interface (CGI) script.

There was one major problem however, the `w3cam` service did not seem to be capable of decoding the video captured by the webcam; the result being images as shown in Figure 9.1. This was the case for both a Creative Labs Live NX

---

[1]OpenEmbedded is a open source meta-data repository containing information on building/-compiling software for embedded systems.

Figure 9.1: Static Images from W3Cam

and a Logitech QuickCam E2500, while a Microsoft LifeCam NX-6000 failed to be recognised by the Linux installation itself. Research indicated that, as with wireless networking adapters, support for web cameras is notoriously patchy. The NX-6000 is supposedly a `uvcvideo` device and is shown as fully compatible on the uvc-driver development page[2], while the Live NX is supposedly supported by the `gspca` driver. It is possible that the drivers could have been out of date, however attempts to build the latest version of the `gspca` driver from source failed via BitBake[3] as did a native compile of the `uvcvideo` driver.

Incidentally, the Creative Labs NX Live did function correctly via the `mplayer` media player application on the Angström Demo Image[4]. This indicates that a small configuration error may be causing the problem; regardless, as a solution which functions over a wireless link was needed no further development with `w3cam` was undertaken. Forum messages suggest that an alternative solution for viewing webcam video was `FFmpeg`. This application includes an additional component, `FFserver`, which uses `FFmpeg` to capture a video feed then make it available via

---

[2]http://Linux-uvc.berlios.de/

[3]BitBake is the build tool used as part of the OpenEmbedded metadata repository/system.

[4]This is a pre-compiled kernel and file system available online at http://www.angstrom-distribution.org/demo/beagleboard/

a web server in a number of 'streams' as defined by a configuration script. Each video format available to the client is termed a 'stream'.

`FFmpeg` has multiple input and output options for file-formats and video codecs; while it tries to make logical guess as to the correct options, it is not always successful. In this case the options had to be manually specified. As detailed information about the actual output generated by the webcams is not available, this entailed testing a number of file-formats and assessing how the output appeared. Flash video appeared to be the most stable video output option and was selected while assessing input file-format options. Output was captured to file rather than streamed via the server during this assessment period. Assessment details can be found in Appendix C and captured video files on the attached DVD.

The command that worked with both webcam's, for streaming operation, was found to be

```
ffmpeg -s vga -f mjpeg -i /dev/video1
        http://192.168.1.42:8090/feed2.ffm
```

However when capturing video in this way frames of video were continually dropped. This caused motion to appear very sluggish as can be seen from captured videos included on the DVD. Flash video is not the only output solution, although it may prove to be the best for embedding into a webpage; alternatives include MPEG, Multi-Part-JPEG (MPJPEG), AVI and WMV.

Using `FFmpeg` to output as MPEG was successful, however when combined with `FFserver` the browser/operating system insisted on attempting to download the entire file before playing it. MPJPEG did function, however is only supported by browsers similar to FireFox. AVI files appeared to suffer interlacing issues while WMV and ASF/ASX (Microsoft's streaming file types) additionally failed to work.

9.2 Still Image Capture

It was also envisaged that still images could be captured from the system. This task was somewhat problematic; when a still image was requested through the server, the file continually downloaded. This is in fact correct operation from the

browsers point of view; it needs the entire file before attempting to display it. However `FFserver` doesn't support single JPEG streams, they are instead treated as a Multi-Part-JPEG. Hence when requested, based on the file extension, the browser expects an end-of-file 'marker' which is never sent by the server[82].

The Debian Bug Report Logs[82] which noted this inability to handle single JPEG images additionally posted two patch files to remedy the situation. These were download and FFMpeg's BitBake script[5] was altered to include this change. When the build operation was invoked it did not successfully complete. Build messages indicated that one of the two patch files could not be applied to its target. Further investigation showed that the patch file and the target differed greatly in terms of the position at which the patch file had to be applied. This in turn suggested that the patch applied to a different version of `FFmpeg` than included in the repository.

This was confirmed by examining the source code of the application; `FFmpeg` appears to have had a major rewrite between recent versions. Tracing similarities in the patch to locations in the re-organised code, along with the aid of compiler error messages, allowed a new set of patch files to be created. These patch files then lead to a successful compilation of `FFserver` with support for single JPEG image streams; an example of which can be seen in Figure: 9.2, a 640x480 pixel JPEG from a Logitech QuickCam E2500.
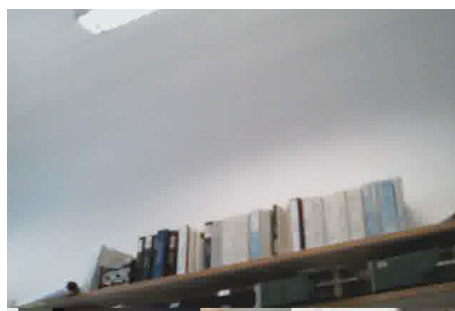
Figure 9.2: Single JPEG Image Capture

---

[5]Termed a 'recipe'

9.3   Processor Utilisation for FFserver

`FFserver` provides a helpful status page when streaming video. This page details items such as the available streams, the number of times each stream has been requested and the amount of data transferred within a stream; another useful feature is the current CPU usage. The level of CPU utilisation while capturing, converting and streaming video feeds was a source of concern. A Video Graphics Array (VGA) sized capture with a 320x240 pixel Flash video stream conversion at 25 frames per second causes a processor utilisation of ~80%, while a single JPEG stream used marginally less at 70%.

It is important to note that these values were noted with problematic video capture devices; therefore the driver compatibility and frame drops issues noted above could be contributing to the high CPU utilisation. This would certainly be true if the video encoding was similar to MPEG, where differences from one frame to the next are encoded rather that the entire frame of video. In this case, if the next frame is 'dropped' by the video capture mechanism, the whole scene changes rather than a small minority. This could then lead to a large overhead to encode the difference.

There is no direct route to solving the video quality issue via a webcam; it is a case of finding one that 'works well' with the available drivers and `FFmpeg`. However as the creation of a webcam system is to assess if video transmission is possible in a high voltage environment, albeit in a crude manner, high quality video may not be absolutely required.

High quality video can be achieved via the camera interface of the Gumstix device as discussed in Section: 7.3. However due to time restrictions the webcam system was used, in which case the CPU utilisation had to be considered. One possible route that may reduce the computational burden of capturing, converting and serving live video is the use two devices.

Currently, video is captured by `FFmpeg` and converted into an 'ffm' video stream[6], `FFserver` then takes this feed and converts it into a number of outbound streams as specified by the configuration file. This operation is shown in Figure: 9.3.

Alternatively, it should be possible to use one device to capture the video and for

---

[6]'ffm' is an 'FFServer Live Feed' stream
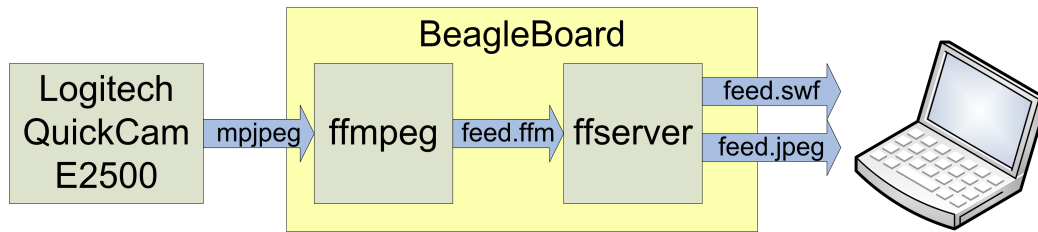
Figure 9.3: Single BeagleBoard Implementation

FFserver to present only the 'ffm' stream to the outside world. A second device would also run an instance of **FFserver** whose input feed would be the 'ffm' stream from the first device, as shown in Figure: 9.4. As the first device is only performing one conversion (from Motion JPEG (MJPEG) to FFM), the computational load should be reduced. However, due to time restrictions this mechanism could not be explored.
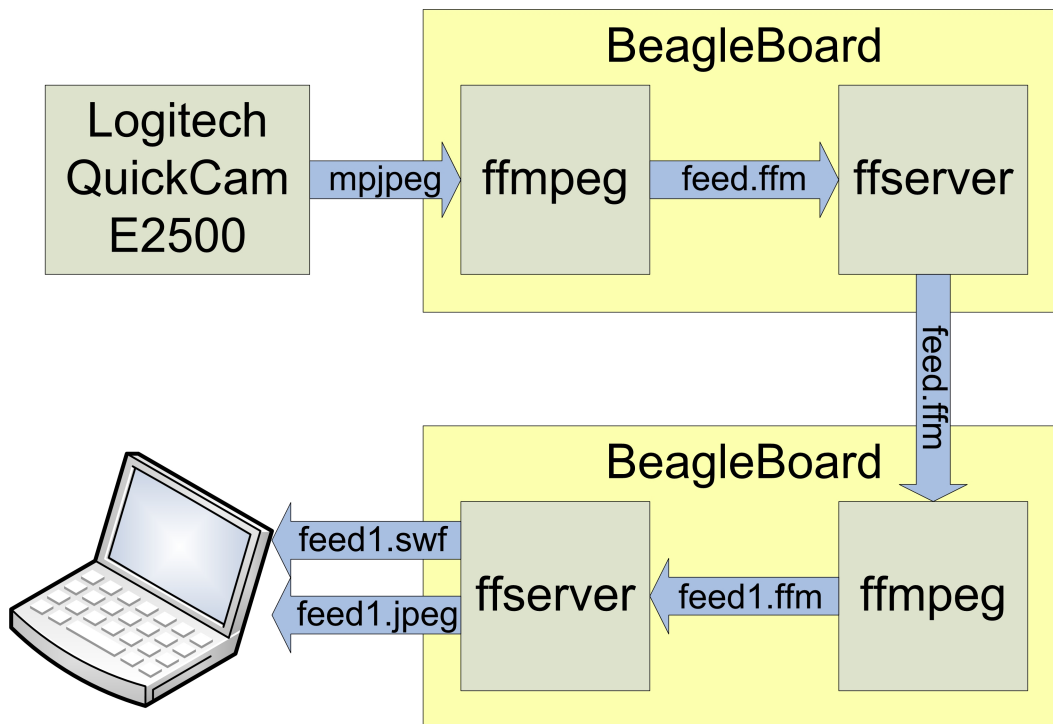
Figure 9.4: Duel BeagleBoard Implementation

9.4 Camera Mounting and Image Clarity

When mounted to the carriage, the camera's field of view will be perpendicular to the orientation of the insulator shed. A prism was to be used to capture images of the surface, by directing light through 90° and into the camera's optics. As noted previously, given the clearances at the insulator's central column, the prism's entry and exit surface could each only be 1 cm^2. The prism was to be held at a fixed distance from the camera with both pivoting around the camera's base to image a swept section of insulator as shown in Figure 9.5.
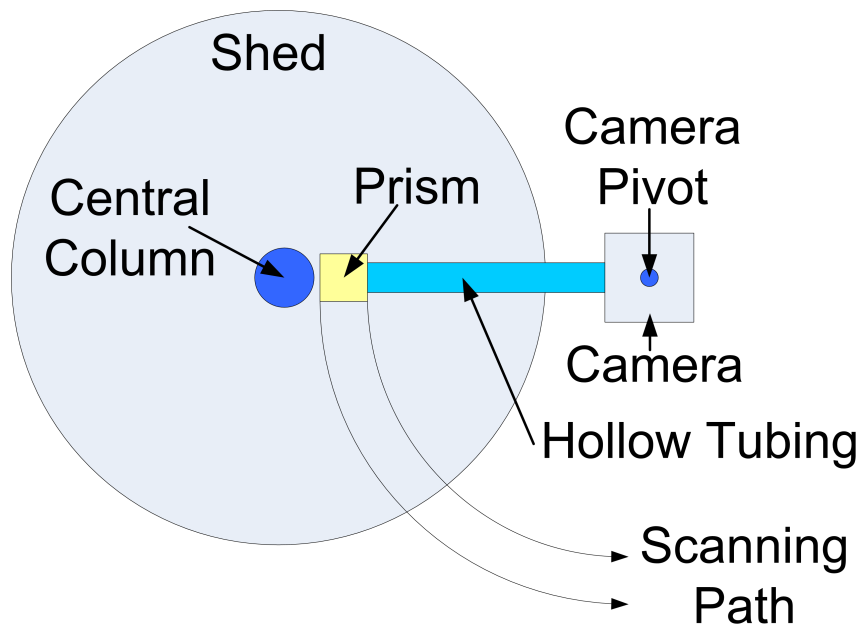
Figure 9.5: Camera Scanning Path

A light pipe was used to hold the prism at a fixed distance; this also ensured that only light directed through the prism was captured by the camera. Blackboard paint was used to provide a non-reflective surface on the inner wall of the light pipe. Without this coating, light directed though the pipe followed multiple paths to the imaging sensor and thereby distorted the image.

The webcam used for this initial investigation had a resolution of 640x480 pixels, covering wide field of view. When the light pipe was attached, the field of view naturally reduced. Figure 9.6 shows a image taken from the webcam using the

Figure 9.6: Image captured using coated light tube, without prism, approximately 4 cm away from edge of insulator shed

provided capture software[7]. It can be seen that while only the area of interest is captured, it only occupies a small area of the imaging sensor.

A better proposition would be if it filled a larger proportion of the 640x480 sensor area. This meant that an additional lens was required to alter the manner in which light coming though the pipe was directed to the webcam's own lens. A number of lenses with differing focal lengths and lens apertures were sourced and examined in conjunction with the webcam and light pipe.

The lenses sourced were as follows: an acrylic Fresnel lens, two planoconvex acrylic lenses and one biconvex polystyrene lens. The Fresnel lens had a focal length of 10 mm. One planoconvex lens had a focal length of 32.7 mm and an aperture of 9 mm, while the second had a focal length of 40.8 mm and an aperture of 12.8 mm. The biconvex lens had a focal length of 20 mm and an aperture of 14 mm. The lenses additionally had varying physical sizes; this meant that each lens could not be tested in similar conditions.

---

[7]The provided viewing and capture software was used for the assessment of the light pipe and prism because it provided a clearer view than through the embeded platform, which required more in-depth exploration into Linux compatible webcams.

Figure 9.7: Image captured without prism, showing opposite side of robot. Captured from intended camera mount position

Figure 9.8: Image captured without prism using 9 mm lens, showing opposite side of robot. Captured from intended camera mount position

Figure 9.9: Image captured without prism using 12.8 mm lens, showing opposite side of robot. Captured from intended camera mount position

Figure 9.10: Image captured without prism using light tube, showing opposite side of robot. Captured from intended camera mount position

Figure 9.11: Image captured without prism using light tube and 12.8 mm lens, showing opposite side of robot. Captured from intended camera mount position

All lenses except the Fresnel could be fitted within the webcam's own light guide, at distances dependent on their aperture. Images captured with the use of these lenses are shown in Figures 9.7, 9.8, 9.9, 9.10, 9.11.

Figure 9.7 shows what the webcam sees using only its own optics. The use of the 9 mm lens clearly magnifies the image as shown in Figure 9.8; while the 12.8 mm lens further magnifies the image as shown in Figure 9.9. Figures 9.10 and 9.11 show the camera's view when using the light tube with and without the 12.8 mm lens; a similar magnification effect can be seen.

While a sharp focus could be obtained with each lens by adjusting the webcam's own focus wheel, the captured images were not any larger. This could be resolved by having the lens further than the focal length away, whereby the scene directed through the lens would focus to a point (at the focal-length) and then start enlarging again.

The only lens that could fit within the light pipe was the planoconvex lens with a 9 mm aperture and 32.7 mm focal length. By positioning this lens within the light pipe, at a distance further than 32.7 mm from the webcam's own lens the size of

captured image could be increased.

Figure 9.12: Capture through light pipe, with 9 mm lens within light pipe. Noticeable lack of focus.
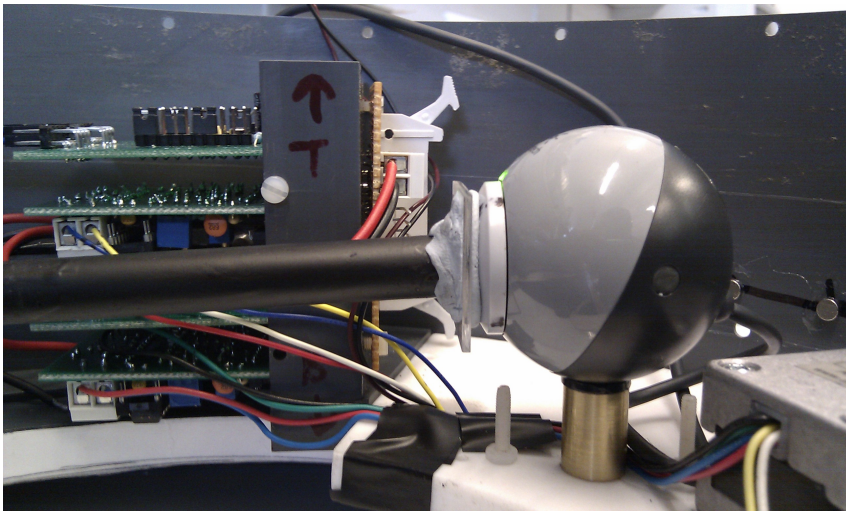
Figure 9.13: Depiction of Fresnel mounting location in respect to light pipe and webcam

However, while the image was substantially larger, it could not be brought into focus as seen in Figure 9.12. Placing the Fresnel lens between the webcam and

Figure 9.14: Capture through light pipe, with 9 mm lens within light pipe and Fresnel lens externally. Distinct improvement over Figure 9.12

the light pipe, depicted in Figure 9.13, provided a distinct improvement as can be seen from Figure 9.14. While the current image is adequate for proving purposes, any commercial system would have to implement additional lenses in the chain to improve the focus, magnification and portion of viewport occupied by the usable image.

The Fresnel lens however had to be removed as it interfered with the insulator shed as shown in Figure 9.15.

This is an ideal juncture in which to compare an external view of contamination on an insulator shed with that seen through the developed optical solution. Figure 9.16 shows a external view of some surface contamination. A number of distinct marks can be seen towards the centre of the image. Figure 9.17 shows the mark to the right as captured by the webcam.

Due to the lower levels of light around the inner parts of the insulator sheds, images appeared too dark. This was resolved by attaching LEDs to the light pipe to illuminate the area being inspected. Further work may have to be done in either constructing a imaging device (whereby the image processing can be

Figure 9.15: Fresnel lens interfering with the insulator shed
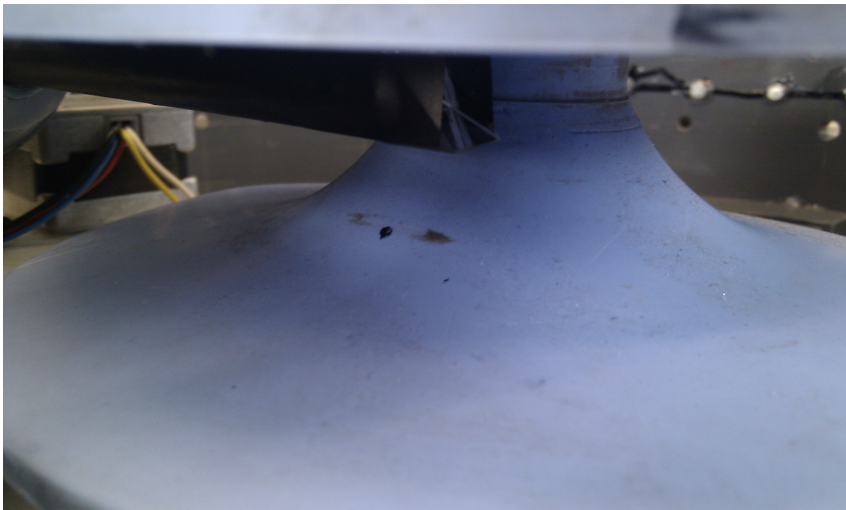
Figure 9.16: External image of surface contamination

Figure 9.17: Surface contamination captured by webcam

intricately controlled) or using more controllable webcam[8]. However as the current implementation can be used to prove the concept of real-time video transmission, further investigations were not undertaken.

9.5 Energised Trials

Given that a transmission system has been created and shown to function adequately under energised conditions, the next step was to integrate the live video streaming functionality discussed above. This would verify if complex electronic sensors function reliably when exposed to high field strengths and partial discharges.

The developed video system was placed into a die-cast enclosure with the previously assessed wireless transmission hardware along with a battery and voltage regulator. It was subsequently tested on a representation of a transmission line, equivalent to the set-up used in Chapters 6 and 8.

---

[8]It is known that certain Logitech webcams have a high degree of controlability under Linux, while others have none or fail to work at all.

The enclosure was attached to sheds 36, 54 and 63, as numbered from the Earthed end. In each case the line voltage was increased from 0 kV to 231 kV and then decreased to 0 kV. Real-time video was viewed on a laptop screen during each ramp-up and ramp-down. The video stream was additionally recorded on a video camera. These files are listed in Appendix B and are available on the attached DVD.

These captures show that it is possible to transmit real-time video from an energised insulator. The video degradation seen in the captures is not due to the energised line. Note how the top half of the video is clear with only the bottom suffering. This is due to a Linux webcam driver issue which was previously noted earlier in this chapter.

Tests at shed 36 and 54 passed reliably, whereas live transmission ceased at shed 63 when the line voltage reached approximately 200 kV. First thought to be a power issue, it was traced to a failure of the USB hub. It was this component that failed during data rate assessment as noted in Section 8.5.1. At present it is unknown as to why this component should fail rather than the more vulnerable wireless transmission system.

9.6 Summary

Based on decisions made in Chapter 7, this chapter discussed the implementation of video streaming and image capture software on the computer platform previously selected. `FFServer` was found to provide the best solution, however compatibility of webcams with Linux itself was a cause for concern. It is currently thought that these compatibility issues are related to high processor utilisations seen on the control platform.

The system has been tested under high voltage conditions and shown in principle to function appropriately, however there is some concern over the repeated failure of the USB hub. However, the majority of the electronics included in this experiment survived and hence shows that complex electronics can survive in areas of high field strengths and partial discharges. This then allows for other sensor packages to be deployed within the robot being developed and paves the way for the development of robots for other high voltage needs.

It has also been noted that use of a light-tube to restrict the viewport also causes the image to occupy a small area of the captured frame. It has been shown that in principle lenses can be used to manipulate light directed through the light-tube in a manner which causes the image to occupy a large portion of the webcam's sensor. However further development is required in this area to simultaneously enlarge and re-focus the image within the limited physical space.

This page is intentionally blank.

Chapter 10

Mechanical System Development

The robot being developed required three different movement actions; to clamp onto the insulator, to move along the insulator and to scan the insulator surface. Although consideration was needed during the design of these systems to ensure that mechanical 'interference' between systems did not occur, it was possible to treat each task separately; both mechanically and from a control point of view. This is possible because it was envisaged that these systems would not need to operate in parallel.

This then allowed the robot's mechanics to be slowly expanded as time permitted; thereby allowing for delays caused by problem mitigation in other areas. This option proved useful as part way through the project's time-frame it was concluded that an attempt to design and implement all the required mechanics would not succeed if a high standard was required. As such it was decided that the target would be to design and implement the surface scanning mechanics. This would then be integrated with the communication and camera subsystems already under development; thereby sufficiently assessing the ability to transmit live video from an area of high electrical field strength and partial discharges.

The idea of separating control duties on a task basis can also be extended to a 'per motor' basis, should the task require more than one motor. This allows motor driving, control and problem resolution activities to occur close to the site of the motor in question. This then avoids the use of excessively long wiring looms for power distribution and communication. Despite the fact that not all the mechanical components will be constructed, implementation possibilities and

initial designs are presented here for completeness.

10.1 Solution Investigation

10.1.1 Surface Scanning

Due to the insulator's form, there were severe limitations on the possible surface scanning mechanisms; the main issue related to the central column to which insulator sheds are attached. Typically if a flat surface needs to be scanned there are two basic options. Firstly, if a strip sensor was available, then it would be possible to replicate the scanning methodology of a computer scanner or photocopier, thereby producing a full image in a single pass. The second option would be to replicate a 'XY Pen-Plotter', where the camera replaces the pen. This would then allow an image to be built up over a small time-frame. However both these methodologies are inhibited by the central column which blocks a 'clean' path over the surface.

The 'XY Pen Plotter' solution can however be modified to account for the central column and in doing so is transformed into a system using polar co-ordinates. Figure 10.1 shows such as system. In this depiction, the camera would be fixed to the end of a threaded rod, whose other end is linked to a linear actuator mounted within the main carriage. This carriage is mounted on a circular track and a stepper motor used to provide circular movement. A tubular can would exist at the outer diameter of the circular track[1] to provide some level of shielding and to mount power, control & communication electronics.

This solution will however not be possible in the shown form; the threaded rod which moves the camera from the circumference to the centre, when in the former position, would breach the outer surface of the robot. Not only would this be problematic to fabricate, the rod would also potentially be a source of discharge as it will enhance the local electric field due to its pointed shape. A possible solution was to increase the diameter of the robot such that this rod would be fully contained within its chassis; however as one of the requirements was for the robot to be as physically small as possible, this was not viable.

---

[1]This is omited from the diagram from clarity.
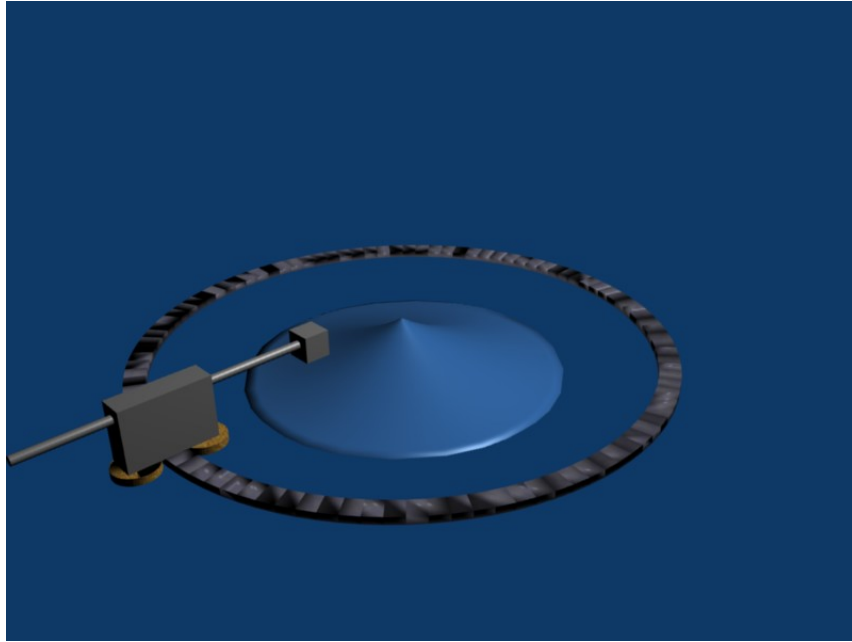
Figure 10.1: Polar Coordinate Camera Movement

A small modification was suggested in the form of a multi-element rod-within-rod arrangement; thereby avoiding the need for any appendage to extend far beyond the robot's circumference and allowing the robot's diameter to be kept small. Furthermore, to relax the camera's dimensional restrictions, it was decided to mount the camera on the carriage and use a prism-mirror to acquire surface images. Rotating this prism-mirror through 180° would then allow the second surface to be examined. Figure 10.2 shows the initial movement path of the prism-mirror in relation to the camera. Engineering drawings of this design can be found in Appendix D.

However fabricating a motor controlled retracting rod-within-rod arrangement with the specified tolerances was not a trivial undertaking. Furthermore, this may have implications for the camera's focus. A solution was found whereby the prism could be held at a fixed distance away from the camera, with both pivoting around the camera's mounting point, shown in Figure 10.3.

For completeness it should be noted that, while capturing a large portion of the shed in a single frame is possible, it has not been considered for a number of reasons. Figure 10.4 depicts two sheds of a typical insulator. There are then two possible mounting positions the camera could take, shown on the left and right
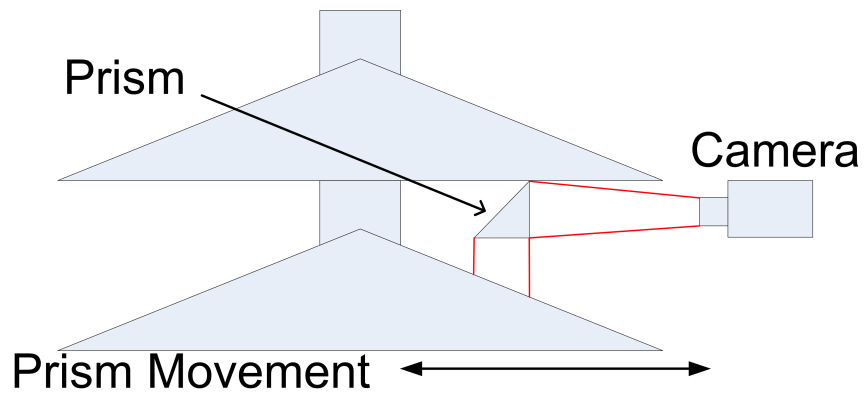
Figure 10.2: Moverment of Prism Mirror
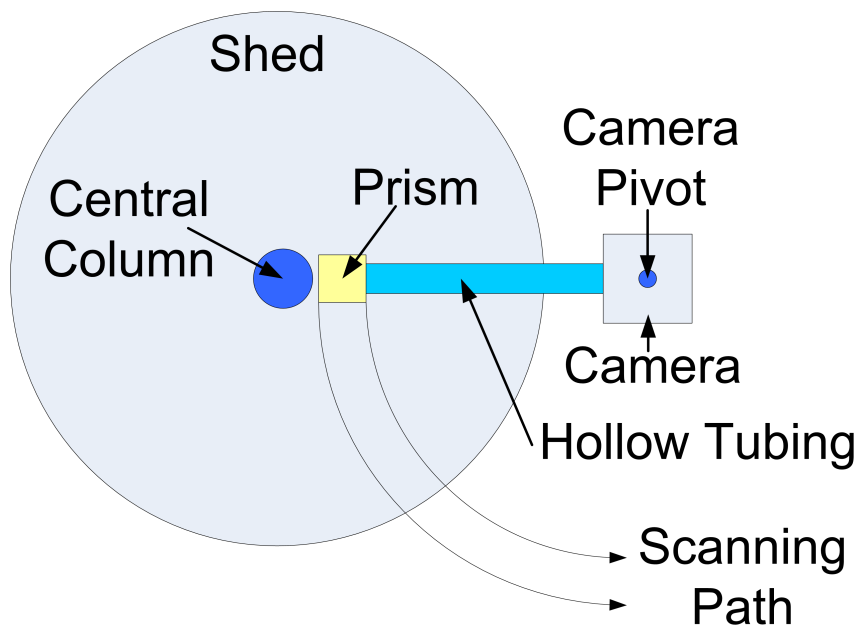
Figure 10.3: Camera Scanning Path

of the diagram. The camera shown on the left focuses on one surface at a time, while the camera on the right aims at both surfaces simultaneously. The camera on the left would additionally need to be capable of changing its orientation and position, as indicated by the arrow. Both these solutions would not be able to see the strip of insulator hidden by the central shaft; hence images from 180° apart would be necessary. Further still, as indicated by the red 'viewing field' lines on the diagram, the far field may show considerable optical distortion which would not be acceptable.
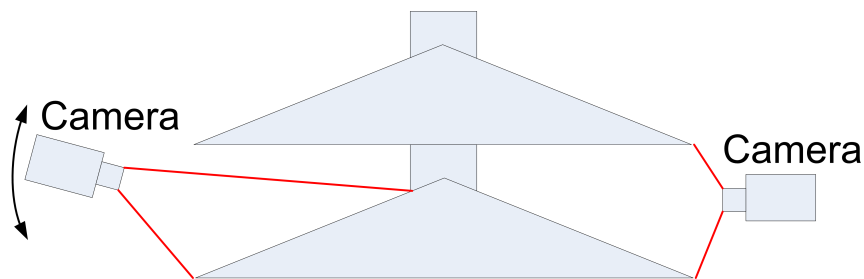
Figure 10.4: Wide Area Photography

10.1.2 Lateral Movement

Lateral movement was probably the most complicated mechanical task and due to time constraints these systems were not physically developed. While a number of possible mechanisms have been theoretically explored, they may not be the only options nor may they be optimal solutions. The core issue is that the mechanism needs to provide controlled descent and accent while providing adequate grip to support the robot's entire weight both when stationary and when moving. Furthermore, this has to be achieved without damaging any part of the insulator's surface. Additionally, due to the requirement that the robot not adversely affect the path length, this has to be achieved in limited vertical space.

There were two broad strategies for the support of lateral movement: use the sheds or use the central column. A cross-section of a more developed idea which uses the sheds for support is shown in Figures 10.5, 10.6 and 10.7. The black blocks in these depictions represent support flaps which are hinged at the robot's case; only two of six equally spaced hinge points are shown for clarity. Rather than being directly hinged onto the casing, they are hinged onto motorised vertically aligned

runners; which are then attached to the robot's case. These motorised runners, which are also not shown for clarity, allow the support flaps to move along the robot's vertical length.
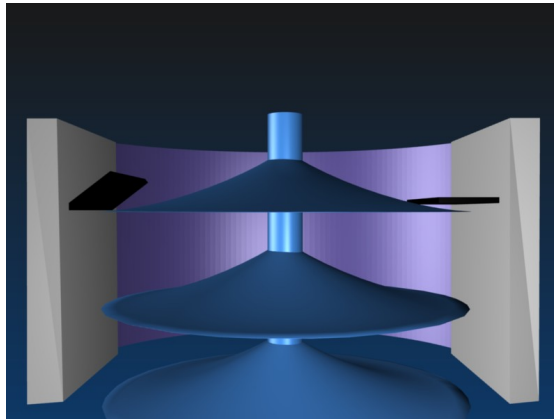
Figure 10.5: Motorised Pivot: Hinge Set One Opening

The support flaps are operated in two sets of three. Within each set, the flaps are spaced at 120° intervals and each set is offset by 60°. As there is one tab every 60°, the robots weight is evenly distributed over the shed. Figure 10.5 shows flap set one being opened; at this point the robot's entire weight is supported by the flap set two. Once fully retracted, all three flaps move along their respective runners until aligned with the second shed, where upon the flaps open and part of the robot's weight is transfered to the second shed, as shown in Figure 10.6.
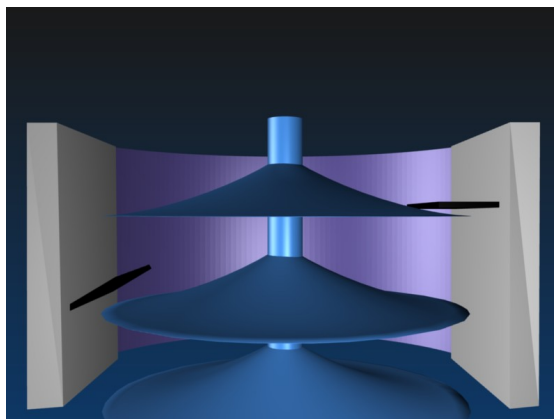
Figure 10.6: Motorised Pivot: Hinge Set One Moved to Next Shed and Opening

At this point flap set two begins to retract, where upon the robot's weight is

supported solely by the second shed. Now, rather than flap set two using their respective runners, those of flap set one are put into reverse, thereby causing the robot as a whole to move down to the second shed. Once aligned, flap set two engage with the shed, thereby creating six points of contact. This is shown in Figure 10.7. An animation of this motion can found in Appendix B.
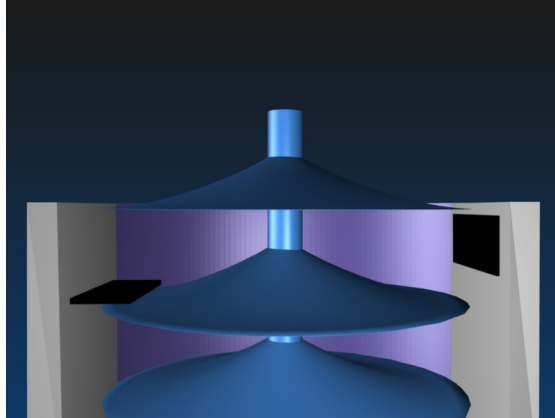
Figure 10.7: Motorised Pivot: Hinge Set Two Opened and Robot Moving to Next Shed

There is however one problem with the above scheme; the amount of torque required to hold the hinges when supporting the robot's entire weight. A slight modification of this design is possible which avoids the need for motorised hinges/pivots. The modification replaces the pivot motors with two preformed rings, one for each pivot set; one such ring is depicted in Figure 10.8.

Figure 10.8: Mechanical Pivot: Pivot Set Unlatched

The ring has three equally spaced dips into which the pivots gradually falls as the

ring rotates. Once in the dips, the ring (now connected to the previously described vertical runners) can move down along the robots length. Once in position, if the ring rotates further, the pivots will latch, Figure 10.9, and allows the second ring unlatch and move as needed. An animation of this rotating motion can found in Appendix B.
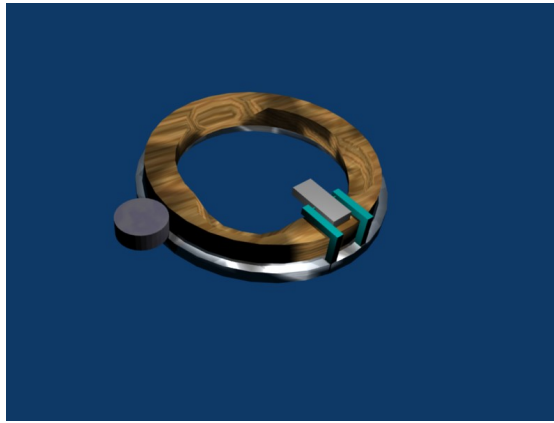
Figure 10.9: Mechanical Pivot: Pivot Set Latched

The 'shuffling' motion envisaged for the motorised pivot based solution cannot be implemented in this case. Previously, each pivot was technically independent even though they were to be operated in sets; as such all six pivots could exist in the same horizontal plain simultaneously. In this particular solution a physical ring connects a set of three pivots and therefore both rings cannot be aligned at the same insulator shed.

Figure 10.10 aims to clarify this position by way of a solution; it is clear that both 'Ring 1' and 'Ring 2' cannot share the same physical space, at best 'Ring 1' can rest on 'Ring 2'. The solution involves off-setting the runner 'sets' as shown[2]. Starting as shown in the diagram, 'Ring 2' would first have to move to the third shed. Once in place, 'Ring 1' can then be aligned with the second shed. Reversing the direction of the runners after all pivots are in position will allow the robot body to be positioned appropriately.[3]

---

[2]Only one runner (black bars) from each set is shown for clarity.

[3]If the robot's mechanics are implemented as shown, the described sequence is just one of many possible.
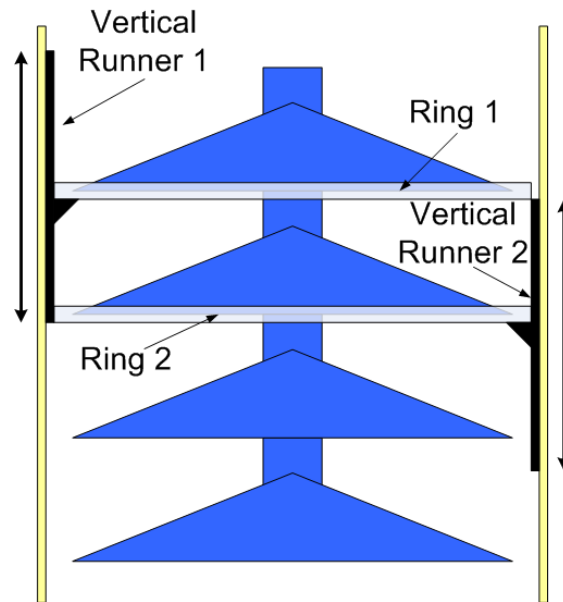
Figure 10.10: Mechanical Pivot: Vertical Runner Movement

10.1.3 Insertion and Removal

This was initially considered to be one of the more trivial aspects of the mechanical components, however in retrospect this proved not be the case and heavily depends on the solution chosen for the vertical movement. In general terms there is only one broad solution to meet the requirements of insertion and removal; the robot will need to split into two or more parts to enable it to wrap around the insulator. This is required because the insulator is fixed to the energised line at one end and the tower at the other. The key task here is deciding on the best mechanism for this split.

Firstly, the best solutions are those which can be fully inserted from one side, thereby allowing insertion to occur from the furthest possible distance from the transmission line. It also avoids the need for two operators to mount the device or for a single operator to overstretch themselves. This effectively enforces a hinge-based solution and rules out those which fully split into parts.

Furthermore, failure is not an option; this means that the solution has to be as simple as possible to minimise the probability of failure. Therefore, it would be ideal if the solution could be manually operated rather than electronically oper-

ated[4]. Additionally, simplicity means that the maximum number of parts could be three; any more potentially requiring complex mechanical linkages.
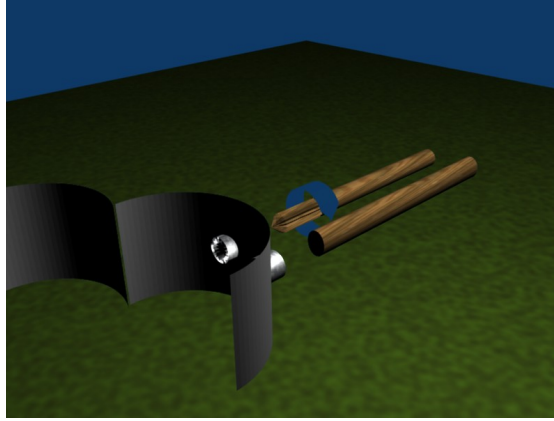
Figure 10.11: Two Part Robot

The first solution considered is shown in Figure 10.11; here the robot is split into two parts and is controlled from one side. This solution requires the use of two hot-sticks; one to engage with the mechanics of the hinge mechanism and the second to support the device when robot's 'shell' is being closed. An animation of this closing motion can found in Appendix B.

Assuming that the robot's weight will be evenly distributed between its two parts, it will be considerably unstable (side heavy) when fully open. Secondly, a mechanism for achieving this motion which sits flush against the inner surface of the robot's case could not be easily conceived. If three parts are used however, something similar to the layout in Figure 10.12 could be implemented.

Importantly, it should be noted that the opening mechanism here was inspired by a 20 year old patent application [55, 56, 54]. The left and right parts of the casing behave like a set of 'double doors', a threaded rod is placed in a slot located near door's hinge with the stationary section, in grey in the figure. The two rods are then attached via a gear arrangement to a third shaft which exits through stationary section. By rotating this third shaft, it is possible to rotate the two threaded rods in different directions. Based on maintaining a tight fit between the threaded rod and it's locating slot in the door, it should be possible to manually open both doors simultaneously from a single point on the robot's chassis. A

---

[4]Also allowing the robot to be retrieved fairly simply if the electronics fail.
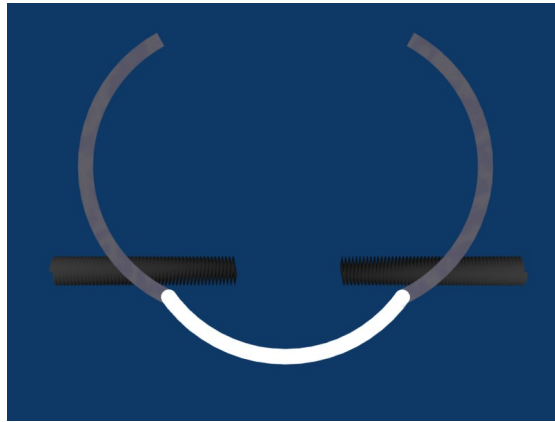
Figure 10.12: Three Part Robot

similar mechanism is used on a pair of precision compasses.

It is known that the layout depicted in Figure 10.12 suggests that the rods may interfere with the sheds, particularly when moving from one shed to the next. However this figure only intends to demonstrate the principle mechanism; further refinement of the design would be required if it were to be fabricated. However, a more serious issue is the protruding rods. These may experience electric field enhancement and thereby cause corona discharge; something which needs to be avoided. No solution to this issue had been envisaged at the time of writing.

Regardless of the mounting mechanism chosen, consideration was needed for the requirements of the surface scanning and vertical movement mechanisms. As previously noted, the camera moves around a circular track; this would need to be split to match the multiple case parts. While not necessarily a problem, it does add another possible point of failure whereby the camera mounting may get 'stuck' while moving between the sections. However there is no simple resolution to this and the camera track was split in the fabricated design.

Complexity in implementing the vertical movement however depends on the chosen solution. If the motorised hinge solution discussed previously is selected, due to the physical independence of each hinge and vertical runner pairing, no additional complexity is caused by the split casing. The ring-base hinge solution would suffer the same pitfalls as the surface scanning mechanism and can therefore solutions follow the same reasoning.

10.2   Solution Selection

As previously noted, time became a serious issue with this project; as such certain areas were scaled back in order to achieve the main objective. The main casualties in this scale-back were the mechanical components, on the basis that if one motion could be demonstrated to function properly in a high voltage environment, others should also function given sound mechanical construction. Therefore, as one of the main intentions was visual monitoring, the surface scanning mechanism was implemented.

As vertical movement and insertion & removal were not considered, different mechanisms could be employed to mount the device onto the insulator during testing. As this particular implementation was only going to be tested within the High Voltage Laboratory at the University of Manchester, it theoretically could have been constructed as a single part (i.e. without hinge). The insulator would then be inserted during the construction of the experimental setup. This scenario however leaves the issue of how to support the device against each shed.

A solution was to use two semi-circular perspex plates with another semi-circle cut-out at the centre of the straight edge; the cut-out matching the circular dimensions of the central column of the insulator. These plates are then attached to the robot, also constructed in two halves. These two halves are connected using a hinge on one edge and a magnetic clamp on the other, shown in Figure 10.13. One semi-circular plate is attached to each half; therefore, when closed the robot is appropriately aligned with one shed. A second set of semi-circular plates is also attached to the robot at a lower position to provide additional support and weight distribution.

The camera and electronics carriage was mounted onto a ring shaped platform attached vertically between the two clamping plates. This platform has a guiding groove cut into it; a similar (but inverse) shaped guide was attached to be bottom of the carriage. This arrangement allows the orientation of the carriage to be maintained and stops the carriage colliding with robot's casing. The inner edge of the ring platform contains a toothed track which is then matched to a gear mounted underneath the carriage. This gear is then connected to a stepper motor mounted on the upper surface of the carriage. A 'pancake' stepper motor is attached to the underside of the carriage and its shaft is connected to a spindle which passes through the carriage. The selected webcam is then mounted onto this spindle. The
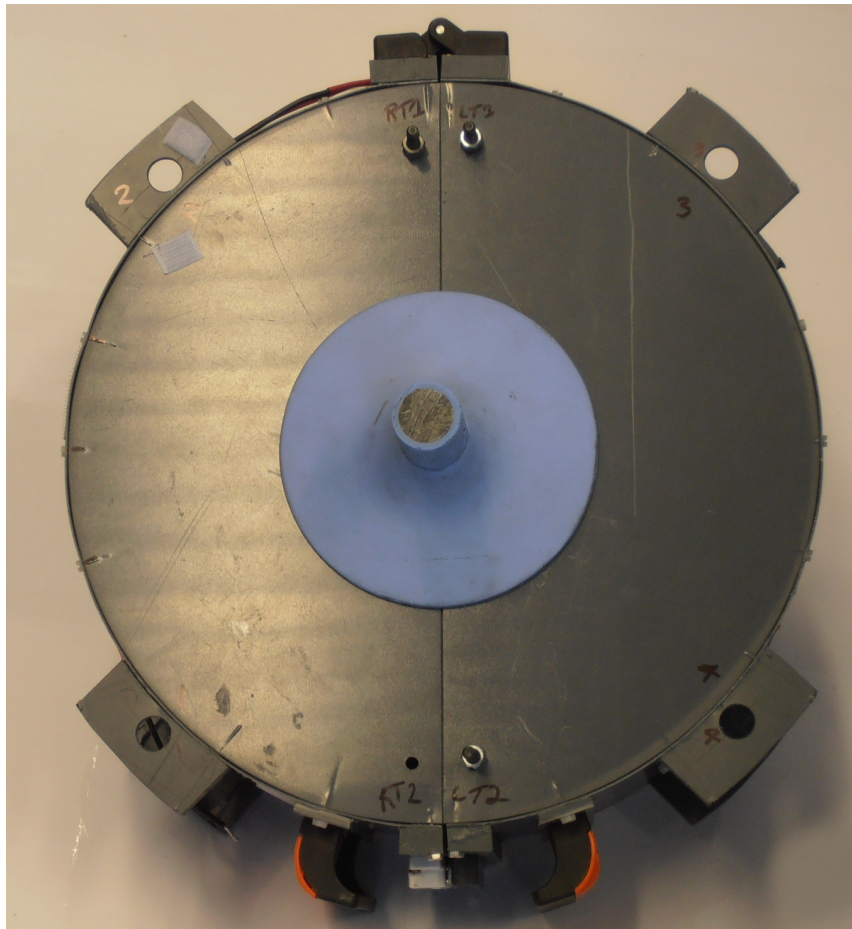
Figure 10.13: Developed robot clamped to insulator

carriage, complete with camera, motors and PCBs can be seen in Figure 10.14.
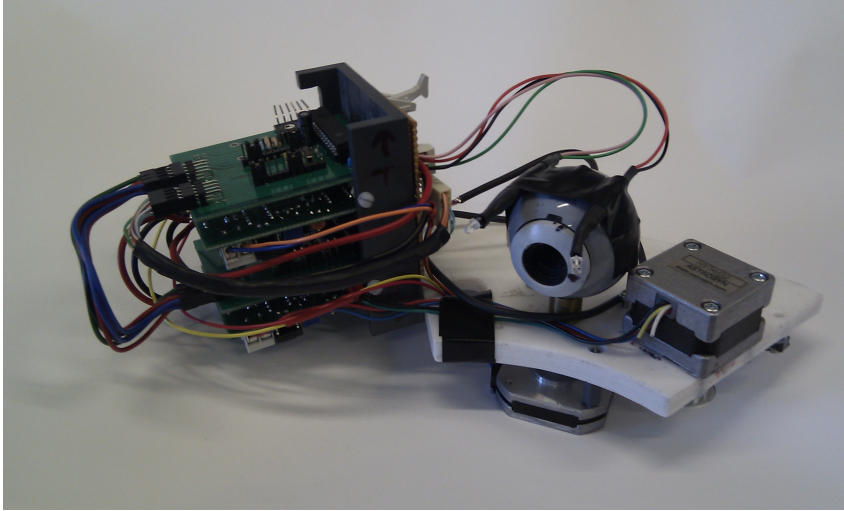
Figure 10.14: Mounting carriage with camera, motors and motor controller boards

Although the fabricated robot is largely similar to the suggested design, shown in Appendix D, a number of modifications were made during construction by The University's mechanical workshop; for example the carriage platform was extended and the aforementioned clamping method designed.

10.3 Electronic Control System

As discussed earlier, only the surface scanning mechanics were to be constructed. This required the use of two motors; one to rotate the carriage around the shed and a second to pivot the camera & prism arrangement. For the purposes of modularity, it was decided that making a single tunable motor controller module would be beneficial while also reducing complexity. This platform contains three key elements: a motor driver, a sensor input port and a microcontroller. The microcontroller provided localised control over the motor driver and monitoring of the sensor input. Furthermore, the motor driver and the microcontroller were optically isolated from each other to protect major electronic systems.

In order to minimise its footprint, this module was constructed in two parts: motor driver & optical isolators and microcontroller & communications link, shown in Figure 10.15. These two elements were then connected via a 10-pin con-

nector mounted at a right angle to the PCB surface. Two of these modules then stacked and subsequently mounted onto the motorised carriage as can be seen in Figure 10.14.
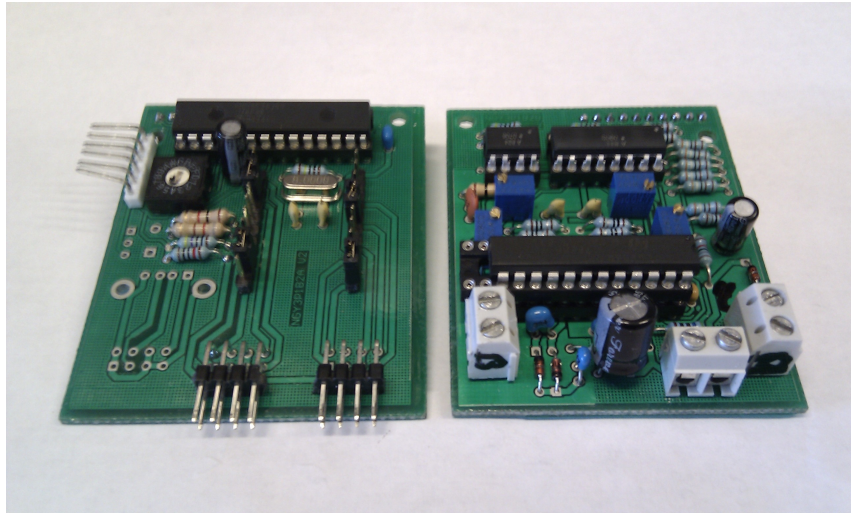
Figure 10.15: Motor driver module

The microcontroller on the module accepts requests from the central controller over a serial link. The request is then performed and a response returned to the central controller. The response takes the form of either a 'request successful' flag or an error condition. The sensor port has been implemented to allow any sensor with a singular output to easily interfaced; regardless of whether that output was analogue or digital. Versatility was added by allowing the driver off-time and winding current limit to be selected via trimmer resistors. Stepper motors with windings rated up to 2.6 A could be used with the designed circuits.

In order to track the current location of the carriage and rotational position of the camera, a position encoding system was required. The mechanism for tracking the position of the carriage uses alternating magnetic discs placed at regular intervals on inner surface of the robot's shell and a hall effect sensor connected to the module's sensor port.

The sensor selected provides a single-bit digital output to denote the detection of either a north or south pole. The detected state is held indefinitely until a change in field is subsequently detected. Therefore, by recording the number of state transitions and knowing the number of transitions along the entire track, it

will be possible to know to a reasonable degree where the carriage is located.

While this location tracking system was implemented, small variations in the fabrication of the platform on which the carriage resides meant that the gear attached to the carriage was not always correctly aligned with the toothed track on the platform. This took the extremities of either causing the motor to stick or rotate freely. This subsequently meant that the position tracking system became unreliable in its initial incarnation.

The camera system additionally needs a positioning system to ensure the prism arrangement does not collide with the central column or the carriage. The ideal solution for this is the use of an infra-red ranging module. While a solution was assessed separately, it could not be implemented within the robot due to size restrictions. Implementation can only be achieved through design and fabrication of a custom ranging module, with dimensions of a few millimetres.

10.4   Summary

This chapter has proposed and analysed various options for the implementation of the robot's mechanical system. The splitting of the mechanical system into three components has allowed each motion to be considered individually in the first instance. Implementation drawbacks with certain designs were noted and potential resolutions were given. It was subsequently noted how closely the vertical movement and insertion/removal mechanisms were related in practice, whereby decisions taken for one mechanism can substantially affect the design of the other.

The development of a tunable motor control module was proposed which could be used throughout the system to provide localised control and management of each motor individually. This increases system modularity and reduces the complexity of the communication channel to the master controller. The modular nature also allows the system to grow as needed; beneficial in relation to the decision to limit construction to only the insulator surface scan mechanism. Furthermore, the developed motor driver module is not limited to use on the developed robot; with a redesign to switch to surface mount components, board dimensions can be reduced with the result being a compact and versatile unit capable of supporting stepper motors with winding capacities of up to 2.6 A.

Two position tracking mechanisms have been proposed, one based on the detection of magnetism and the other on infra-red ranging. While these have been seen to operate in principle, either mechanical tolerances in fabrication or minimum clearance distances have prevented them from being fully implemented within the robot.

This page is intentionally blank.

Chapter 11

Systemwide Communication and Control

Chapter 10 briefly discussed the motor controller system and how it has been designed to accept movement requests from a central controller. This chapter discusses the methodology of the central controller and the design & implementation of the flexible communication system. Various solutions for remote control of the robot are detailed, including the concept of website based hardware control.

11.1 Central Control Platform

As documented previously in Section 7.3 the central control platform was based around a Gumstix Overo mini-computer. During early experimentation with a similar system, a BeagleBoard[1], concerns were raised with the stability of the USB subsystem and hence WLAN connectivity. These issues were most likely caused by attempting to draw more current than could be supplied by the regulator used, which led to a collapse in the supply voltage. This then would have caused the USB subsystem to experience errors.

However experience gained while developing this system revealed that in this event the computer remains semi-functional due to the processors's 1.8 V core. Only

---

[1]The BeagleBoard and Gumstix Overo are virtual identical systems in terms of functionality and schematic design, mainly differing in physical size.

peripherals using a 5 V source, i.e. the USB subsystem, failed. Although, as WLAN connectivity is provided via USB, any failure also resulted in a loss of direct communication via either telnet or SSH.

In order to track this stability issue it was decided that it would be prudent to monitor current usage in the final system in real-time. At this point it was also decided that development should shift to the Gumstix Overo platform that was to be used as part of the final system. Furthermore, in order to mitigate for the possible loss in wireless LAN connectivity to the robot, it was thought advisable to implement a backup control system in addition to monitoring functionality. A customised expansion board was designed for the Gumstix Overo to meet these needs.

11.1.1 Gumstix Expansion Board Design

Dual processor control would require both the Gumstix computer and backup processor to have access to the wired communication link to the motor driver modules. As all Gumstix signalling was at 1.8 V rather than the 3.3 V level that was to be used on the majority of the expansion board meant that both level translators and multiplexers had to be used to combine the various control requirements. The multiplexors are controlled by the backup processor to allow for switch-over if the main system fails. A basic USB 1.0 hub was also implemented on the PCB to avoid the need for an external device.

The backup and monitoring system was based on a PIC32 microcontroller and the TI/ChipCon wireless transceiver that was thoroughly tested during the preliminary communication analysis discussed in Chapter 5 and 6. Three forms of monitoring were implemented; current monitoring of the various subsystems on the designed PCB, local temperature monitoring and monitoring the 'console messages' from the Linux Gumstix Overo. By default, console access is via a Serial-to-USB translator chip on the Gumstix PCB, however it was possible to intercept the raw serial signals via the provided expansion connector and relay them to the PIC32.

The PIC32 microcontroller line offers many features such as Dynamic Memory Access (DMA) and automated Analogue to Digital Converter (ADC) sampling, thereby allowing for more efficient program operation. The DMA module allows

data to be transfered from one memory location to another without intervention from the CPU. Memory locations not only include the data memory but also the memory mapped peripherals. In the event that a memory location is specified, there is the facility to automatically increment the address, thereby allowing one block of memory to to be transfered to an another block.

The DMA module has been used to allow Linux console messages to be transfered automatically from the Universal Asynchronous Receive Transmit (UART) to one of a number of memory buffers. The transfer is also automatically triggered as characters are received by the UART and terminated when a ASCII 'carriage return' character is detected. As the memory buffers are filled they are dispatched for transmission over the ChipCon wireless link. DMA cannot be used for this transfer process due to the need to monitor the status of the transceiver during transmission.

The buffers are configured to be used in a circular fashion to allow for delays in transmission or a number of serial transmissions in rapid succession; although this should not be a problem due to relative transfer speeds. The Gumstix serial port operates at 115,200 bps while the SPI link to the ChipCon transceiver operates at 5.25 Mbps. Therefore, even with the overhead needed for status monitoring, SPI transfers should always complete prior to serial transfers.

The other end of the wireless link uses the hardware developed for the receiver unit of the preliminary communication link. The firmware was altered to allow the unit to transmit & receive data, use an RS232 link to the computer rather than the optical link, have two modes of operation (supervisor and controller) and to output a text-based menu system when in controller mode. A screen capture of the text-based control system can be seen in Figure 11.1.

When in supervisor mode, console messages were then available independent of any remote terminal connection (e.g. telnet or ssh) running over the WLAN. When in controller mode, current and temperature data could be requested in addition to movement requests. This platform provides a low power and stable monitoring and backup-control system for the robot investigator. The independent terminal link could also provide a mechanism to shut down the main computer in a graceful manner in the event of a supply voltage collapse in a production system if required.
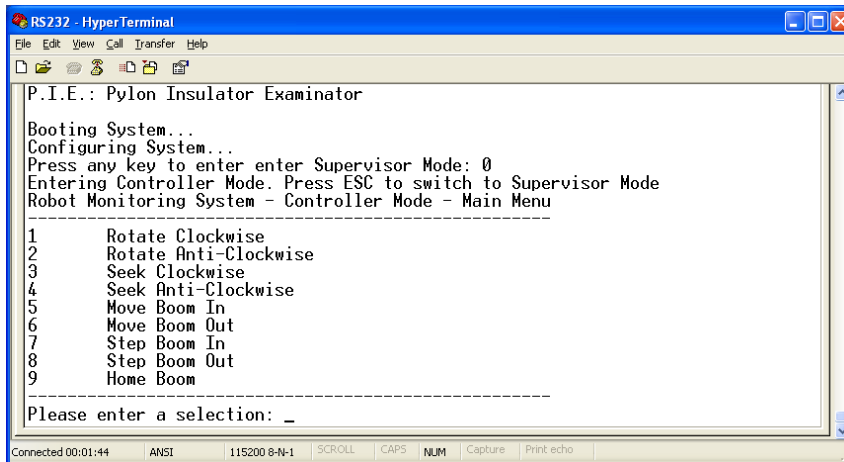
Figure 11.1: Console interface for robot mechanical control

11.2 Robot Wide Communication Network

The communication link needed to allow the central control platform to communicate with the microcontroller on the driver modules in a simple yet robust manner. Furthermore, the central controller was to be mounted on the outer surface of the inner shell in a fixed position. This meant that the data cable had to pass through the inner shell to reach the driver modules. In order to simplify cable routing the number of cables used needed to be kept to a minimum.

There are a number of different network topologies that can be implemented to connect the central controller to the driver modules. The star and daisy-chain topologies were considered for implementation, however the formerF was found to be disadvantageous. While the star topology has the benefit that a broken node does not effect operational ability, it also requires more physical space on the controller PCB (i.e. each remote node requires an individual port to attach to). This solution is therefore not expandable beyond its initially designed specification and additionally breaks the initial requirement of minimising cable usage.

The daisy-chain topology's expandability is only limited by the program memory and processing power of the central controller. Although it is true that if one node fails communication is also lost with all nodes further down the device chain.

The link was designed to use serial signalling, either UART or I^{2}C. It was additionally designed to allow a differential driver module to be installed without the need

for a redesign; this can be used in the event that non-differential signals proved inadequate; Figure 11.2 shows the developed hardware

Each device on this link is assigned an identification number. This identification number then forms part of data packet that is sent along the link. The data packet then includes the data item indicating either motor action or operation response code. Each packet is additionally framed with a start and stop byte to allow for accurate detection.
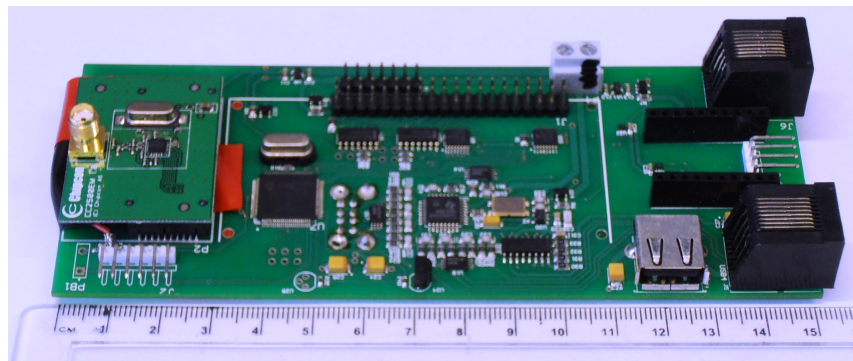
Figure 11.2: Microcontroller Board with ChipCon Transceiver

When a packet is received by a driver module its address byte is checked. If the byte matches the module's address, the instruction within the packet is carried out and a response code is sent to the controller. If the address match fails, the packet is transmitted to the next device in the chain; this onwards transmission occurs until the packet arrives at the correct device. The response is then sent in the opposite direction where it is passed from one device to the next until it arrives at the central controller.

Each driver module was programmed for a particular stepper motor and had its identification number hard coded into its firmware. The central controller was statically programmed to control two motors, i.e. those needed for surface scanning of the insulator shed. The former is due to a limitation in driver module's hardware; in particular having no mechanism of selecting ID or function in hardware (e.g via rotary or DIP switch). The latter is due to simplified firmware coding during the development phase and can be modified to intelligently detect modules on the link and their function.

Updating the design to include switches would allow the module ID and motor

type to be selected in the field. Combined with a 'super' firmware package which includes the functionality for all the required stepper motors would allow the module to truly become universal and replaced as needed.

11.2.1 Wiring Harness

Twisted pair cable was initially considered for the provision of power and communication to the microcontroller. The USB signal for the camera was also to be routed over a twisted pair cable. It was thought that the twisted nature of the conductors would enhance signal immunity to the high electric field around the insulators.

However, there are a number of disadvantages of using such cable which were not identified until the component parts of the robot were integrated. The main issue is with the cable's lack of flexibility in terms of bend radius - the cable would need move fluidly with the carriage around the track while not interfering with the insulator or the carriage itself. To enable this the cable would need to be located between the carriage and the robot's shell, shown in Figure 11.3. The cable would also need to 'double back' on itself to allow the carriage to move to the furthest extreme from the cable's entry point.

A typical 8-core twisted pair cable has a diameter of approximately 5-10 mm. As can be see in the photograph, there is insufficient space for a single diameter of type of cable, let alone the four that is necessary when two lengths of cable are doubled back on themselves. Furthermore, this cabling system would require the cable to be bend through 90° to enable it to pass through to the exterior of the robot, where the major control electronics were to be located.

While the two twisted pair cables cater for the low-power electronics, the motor drivers themselves are optically isolated and powered from a separate source. This requires two further conductors which face the same clearance constraints as noted above. A solution which minimised the number of individual cables required.

Inspiration can be drawn from an ink-jet printer. Such devices require data and power to be sent to a moving carriage (the print head and cartridge) from the control unit. This is achieved with the use of a Flat-Flexible Connect (FFC). Similar to ribbon cable, they are sub-millimetre thick and are not composed of
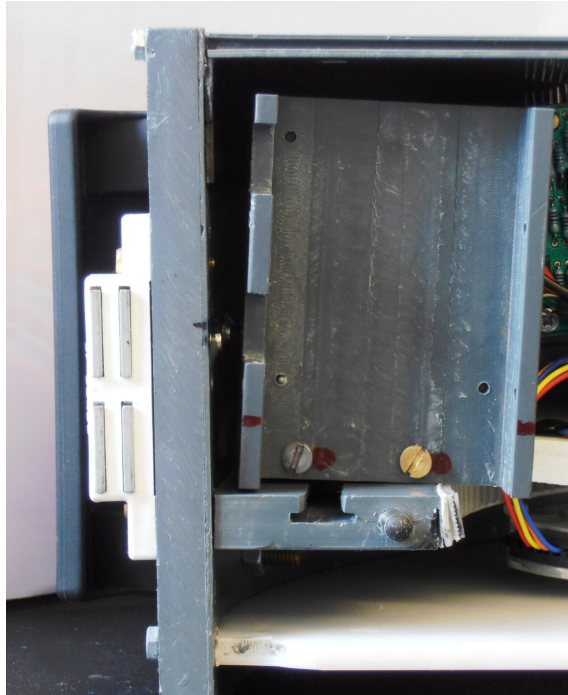
Figure 11.3: Clearance between robot body and carriage

individually insulated conductors; they are similar to PCBs whereby copper tracks are placed onto a thin plastic film.

Due to their thin nature, the 'doubling back' issue would not be a problem and if a FFC with sufficiently enough circuits could be located, then a single cable could replace the many initially envisaged. However a commercially available FFC of sufficient length with the requisite number of circuits could not be located. Those that were located (having insufficient circuits) were additionally prohibitively expensive.

It was therefore decided to use standard ribbon cable. As the carriage positioning system uses magnets located half way up the side wall of the robot, it is necessary for the ribbon cable to be positioned away from this area. While it could be routed above or below this row of magnets, the former position was chosen to avoid potential conflict with the carriage's base. When this was initially tested, the ribbon cable was found to be too flexible and moved with its own weight; additionally blocking the magnets and obstructing the carriage.

Although a number of mechanisms were tested in order to resolve this issue a

reliable solution could not be found. While brush contacts could be used for the transfer of power, it is not a viable solution for the communication of data. Further investigation needs to be carried out into multi-circuit cable reeling mechanisms to successfully resolve this problem.

11.3 Website Based Remote Control

It was initially intended for the robot to be controlled and video viewed via a combination of a custom hardware device and desktop application. It was also the view that a custom wireless link would be established by aggregating multiple CC2500 modules. However with the discovery of the Linux capable OMAP platform, control via a web interface became possible.

Most modern network infrastructure equipment (routers, switches, access points) have a web interface to aid configuration. These web application interfaces, by necessity, will have to make changes in the configuration of the underlying hardware. The question arose as to how this was actually achieved in practice? Internet based web applications usually write/modify information held in databases or data files on the web server. It is certainly possible that a similar methodology is used whereby the application invokes changes in a data file and the data file periodically read by a process on the device and changes are acted upon accordingly; however a substantial overhead is used in polling the data file.

The WAP purchased for the high data rate communications testing used open source firmware. This provided an ideal opportunity to investigate how this device implements interaction between the web interface and the underlying hardware. Browsing the firmware source code revealed that the system utilised Active Server Page (ASP) scripting technology. The web pages that form the complete web interface, when called invoke some action as appropriate to the request/change made; this included directly accessing memory locations and changing values. While this approach as been shown to work, it requires an in-depth knowledge of memory usage on the target system and may still involve an element of polling to check for changes. A better solution is a server-side scripting language which could interact with installed applications.

Before continuing with the process required to achieve web interface control, it is

important to detail the General Purpose Input/Output (GPIO) structure of the OMAP chip and the Angström operating system[2].

The Angström distribution abstracts GPIO pins into the file system, thereby allowing interaction by means of reading and writing to 'files'. An example follows:

```
root@/sys/class/gpio>echo 168 > export
root@/sys/class/gpio>echo ''out'' > gpio168/direction
root@/sys/class/gpio>echo 1 > gpio168/value
```

The above sequence will add pin 168 to the active pin list, creating a folder for the pin with 'files' for modifying the 'direction' and 'value'. Writing appropriate values to these files enables the pin to act as an input or output and to read or set the value respectively.

As pin configurations can be changed by writing to 'files', an application or script can be created invoke the shown instructions in sequence, as opposed to manually invoking each on a command line. If this application (or script) is then called by the web server, in response to a page request, the process of remote access is achieved. One possible server side scripting language with which this can be implemented is the Common Gateway Interface (CGI).

CGI scripts can be either written in C or Perl, both of which have the ability to write to files in the Root File System (RFS). While the former is compiled mechanism, the latter is interpreted. Each method has its benefits, such as C programs being quicker to execute but were string manipulation is harder, while in Perl strings are easy to handle but are slower to run. In this situation the speed difference is marginal due to the limited application length.

Regardless of the implementation method chosen, the first thing that should be returned to the web-browser is the content type; after which can follow the content of the page. It was decided that C would offer better facilities regardless of the difficulty of handling strings; the difficulty being mitigated by that fact that not much textual information would need to be passed to the application.

A C library exists to aid in the development of C based CGI applications: `cgilib`. This library aids in the dispatching of the 'content-type header' and other textual

---

[2]As both the BeagleBoard and Gumstix computers are essentially identical, the same operation sequence is applicable to both systems.

information to the browser. The source code for this library had to be cross-compiled for operation on the ARM core. This could be either done via BitBake or natively on the BeagleBoard. While the former would generally be regarded as the simpler option, it required the creation of a BitBake 'recipe' and information on doing so was sparse.

Compiling natively required the installation of a set of compilation tools including the `gcc` compiler. These tools already have BitBake 'recipes', hence a new console-image could be created with ease. As a cross-reference to the library is required when compiling CGI applications, this also had to be performed natively on the BeagleBoard. As the CGI applications developed were only small-scale and compilation was only during the development cycle, this extra complexity was acceptable. For larger applications and longer development time-scales, it would be beneficial to create a BitBake recipe to handle the cross-compilation and installation to the RFS.

Once successful library compilation was achieved, a number of applications were written and compiled:

- to activate the pin 168
- to output logic 1 on pin 168
- to output logic 0 on pin 168

A LED was interfaced to pin 168 and the applications invoked via the command line; the result being the successful toggling of the LED. The next stage was to invoke the same operations via a web interface. At the basic level this meant placing the application files in the cg-bin directory and requesting those files directly via a browser's address bar; bar a few issues with the browser's caching mechanisms, this test also succeeded. This demonstrates that it is possible to control physical output operations on the BeagleBoard via a web browser; an extension of these applications would then allow the control of the various motors required by the robot.

There is one caveat that should be mentioned due to potential security implications. In order to configure hardware on the Linux based BeagleBoard `root` privileges were needed. By default, the Apache web server on the BeagleBoard starts under the `daemon` user. This then caused CGI applications to be run with

`daemon` user privileges. As the application try to modify port states, the Linux security model blocked the attempt. Additionally, a default build of the Apache web server could not be started under the root user due to the security risks of allowing potentially harmful scripts/applications to be executed remotely[3].

To circumvent this issue the web server had to be recompiled with a specific flag to enable the server to be run with `root` privileges. Due to this potential risk, if this remote access mechanism is taken forward to a working robot, it cannot be connected to a public network i.e. it should not be controlled over the public Internet.

While the web-server control route works well for network devices and has been shown to be possible via the OMAP platform, it is not the only remote control mechanism; alternatives should be mentioned for completeness. Two of the alternatives are command line based and are similar to accessing the device over a direct serial connection; these methods are telnet and SSH.

While it is possible to invoke control commands in this way, the user will need to be comfortable using a text based interface and a second system would be needed to monitor live video and JPEG images; thus introducing further complexity to the ground station device. A third solution would be to run a graphical environment on the BeagleBoard and then 'transport' that graphical environment over the wireless connection to a Linux computer on the ground; it is not known what effect this method will have on CPU utilisation. However, as this route was not exploited, further analysis into CPU utilisation was not undertaken.

11.4 Summary

This chapter has detailed how the design of the central controller incorporates dual processors to create a system with built-in redundancy. It is detailed how a PIC32 microcontroller's functionality is harnessed to create a monitoring and backup system to the primary processor. The design of a flexible and extensible communication system for use on the developed robot has additionally been discussed in conjunction with issues relating to it physical implementation.

---

[3]This risk does not exist when running under the 'daemon' user as only limited access is granted to applications executed.

While the central controller and individual modules are currently configured partly manually through firmware changes, flexibility exists for the creation of a self managing system; whereby new motor control modules can be added and used immediately without firmware/software modification. Secondly, for development purposes control has been limited to the PIC32 microcontroller; however the long term intention is for the robot to be controlled remotely via the primary processor, with the PIC32 purely acting as a backup and power monitoring device.

Finally a web based control strategy has been explored and demonstrated in an effort to provide a 'user'friendly' interface for operators of the system. While known to be technically viable, its implementation has computer security implications and hence should not be used in its current form.

Chapter 12

Case Study: The Composite Cross Arm

The worlds ever increasing demand for electrical power has previously been noted, as have some of the problems utility face when attempting to reliably satisfy it. This research has considered the problems utilities face when trying to assess their transmission network assets, in particular the insulator strings, while the circuits are energised. However this is not the only issue that utilities face, another is system capacity.

This issue is becoming increasingly difficult for utilities due to planning restrictions surrounding the installation of new transmission routes and difficulties with right-of-way for the uprating of existing ones. The Composite Cross Arm (CCA) is intended to be a solution to this problem.

12.1   The Technology

The CCA is an innovative technology that has been developed by The University of Manchester in conjunction with an industrial partner, EPL Composite Solutions Ltd. Traditionally transmission towers are of steel construction and an insulator of some form is used to isolate the steel tower from the energised line. The CCA innovation proposes the construction of the cross arm from composite material

rather than steel[1]. The use of a composite material for the cross arm itself provides isolation from the energised line which would normally be provided by a suspension insulator string.

Providing the insulating structures are correctly designed the need for an additional suspension insulator is negated, therefore the energised line can be directly connected to the nose end of the cross arm. Furthermore, by not installing the suspension insulator, it is possible to raise the height of the line by a number of meters without raising the height of the tower itself. Raising the height of the line thereby increases clearance distances, which in turn allows for an increased system capacity on an existing line.

This product is therefore intended to be retrofitted onto existing tower structures[2], therefore allowing system capacity to be increased without the need to acquire new right-of-way permissions. This ability is a major advantage when a country has increasing energy demands yet obtaining planning permission for additional transmission routes proves difficult or impossible.

A complete cross arm consists of four members; two horizontal and two diagonal. The diagonal members are similar to typical composite insulators while the horizontal members are of novel design. Figure 12.1[3] shows the developed CCA.

Members are constructed from a Glass Reinforced Plastic (GRP) core with silicon sheds. The GRP core provides the mechanical strength with the silicon sheds providing the appropriate creepage path between the line and tower mounting points.

As with all insulation systems it is important that the damage does not occur during installation, surveying and maintenance operations. This poses additional difficulties for the CCA solution.

Through an un-energised field trial in the Scottish Highlands, The University of Manchester, their design partner and power companies have gained experience in the installation of such technology to existing transmission tower structures. In this trial four CCA structures were installed as replacements for steel counterparts

---

[1]The tower itself would maintain its steel construction.

[2]Although it could also be conceivably be used on new designs for either size reduction or capacity increasing purposes.

[3]Courtesy Arago Technology

Figure 12.1: Installed Composite Cross Arm

on a transmission circuit that was being decommissioned. One centre phase on each of four adjacent towers were replaced.

Installation knowledge and expertise is just one support area that this technology needs if it is brought into the commercial arena. Other areas include the surveying and maintenance of the structures post installation. This will not only allow the developers to ascertain real-world performance, it will allow utilities to properly manage their assets.

While this un-energised trial has made use of wired sensors to monitor mechanical aspects of the newly designed structure, this is not viable on energised transmission routes. Furthermore, with a steel cross arm a worker is able to easily reach the cross arm's nose if required without fear of causing damage; the same is not true of the designed system.

Although the CCA could mechanically withstand the additional weight of a worker, having been designed to withstand ice and wind loading, it cannot be easily achieved without the possibility of causing damage to the silicon sheds. This problem was encountered during the field trials when monitoring devices at the nose had to be replaced after initial cross arm installation.

In this situation, the devices were accessed by suspending a ladder from the steel cross arm above the installed CCA. However, if each cross arm is of composite construction, this solution could no longer be used. This is where the research presented here can be applied.

12.2   Application of Research

While this research has focused on developing a robot to assist utilities with the monitoring of existing insulation structures in order to detect damage, the technology has wider applicability. One area is the ongoing monitoring requirements of the CCA that is being developed at The University of Manchester.

The issues of accessing parts of the composite structure, in particular the nose, has been previously mentioned. This research has demonstrated that high data rate communication is possible between an energised structure (a suspension insulator) and a base station, using both omnidirectional (i.e. whip) and cylindrical

waveguide transition antennas.

Whereas the WLAN system may have a high power requirement, the 500 kbps link that was additionally presented has low power requirements, which could possibly be achieved through power scavenging. Therefore it is conceivable that a small self-contained sensor unit could be designed and permanently placed at the nose.

Example sensors include accelerometers and gyroscopes, which will allow utilities and the developers to assess cross arm movement during use. Theoretically any low power sensor with relatively low data rate requirements could be integrated into the module. In this case the module would transmit its data to a collection station on the tower itself where a more substantial power scavenging system could be implemented. Data could then be transfered over either the mobile telephone or public WLAN networks.

Another possibility, although not directly related to the CCA, would be to attach the module to the line itself in order to detect galloping lines. Again the unit would transmit its data to the nearest tower. In this situation the module could gather its power from the transmission line itself.

A more dramatic concept would be the construction of a robot to 'walk' along the horizontal members of the CCA. The robot could then photograph surface conditions, make electric field measurements[4] and possibly retrieve samples of surface contamination for later analysis.

Compared to the limiting mounting mechanisms available for the monitoring of suspension insulators, the facets of the horizontal members of the CCA provide more distributed and ridged areas of contact and hence are less likely to damage the insulation surface. The horizontal nature additionally means that split design for the robot would not be required, however the creation of a 'walking' motion will take careful and considered design.

---

[4]It is known that the presence of the robot will effect the field distribution, however most electrical methods of field estimation will suffer similar problems.

12.3 Summary

This chapter has considered another scenario where robotic and static monitor systems would be of benefit to utilities as an asset management tool. The same systems would also be of benefit to researchers developing the technology to gather real-world performance figures.

It is important to note that the CCA is just one area in which the research presented can be applied. The wireless data transfer and image capture systems can equally be applied in other high voltage environments for example substations and in situations where it is undesirable to have a physical connection between the monitoring unit and the data collection & processing device.

A more detailed account of early unenergised monitoring work undertaken for the CCA can be found in Appendix E.

Chapter 13

Conclusions and Future Work

It has been shown in principle that it is possible to communicate from a location on an energised high voltage line representation to a ground station. Collated results indicate that achievable WLAN data rates vary quite dramatically in the environment of The University of Manchester's high voltage laboratory. The highest individual result recorded was 2615 kBps while the lowest was a mere 57 kBps. Combining all individual cases from all channels resulted in an overall average throughput of 1280 kBps, with the theoretical maximum data rate of WLAN being 6750 kBps: an average utilisation of ~19%. This then indicates that it is feasible to use WLAN technology for the creation of a data link between a high voltage condition monitoring device and a ground station. With the introduction of a periodic link assessment mechanism which detects poor link performance and initiates a WLAN channel change, it should be possible to reliably communicate with the monitoring system in most eventualities.

Within this domain it has been revealed that standard omni-directional 'whip' antennas only have limited scope for use in such a communication system. It has been shown in principle that a cylindrical waveguide transition antenna link provides shielding for the radiating element and the attached radio circuitry. It was further discovered that antenna directivity is of greater benefit than antenna gain when considering link quality, particularly at the ground station end, where a stable link could not be maintained when a omnidirectional (whip) antenna was used. These are important results as communication and interference from such an environment and proximity has not previously been considered or discussed in

literature.

However there are some limitations to the presented results. The problems that were encountered with background interference in the test facility have been noted. Furthermore due to the random nature of this interference combined with the duration over which measurements were gathered meant that the external sources could not be easily removed from results. Therefore while results have to be viewed with some caution, a general potential trend can be identified; namely interference may be directly linked to field strength, particularly towards the energised end. This was identified when it was discovered that the field around an energised insulator is not as uniform as first thought. This knowledge then allows for the development of a new tool for the estimation of field strength as will be discussed in Section 13.3.

In relation to the camera system, it was shown that it is possible to capture video and still images at a remote location and transfer them wire-free to a receiving station. Furthermore, by using appropriate lenses images can be enlarged and brought into focus. However compatibility issues between consumer webcams and embedded Linux distributions affects the quality of available imagery. This may also be the cause of the high processor use seen during development and experimentation. Image quality of the webcams used appear adequate however when integrated into the condition monitoring robot, image quality deteriorates. This suggests that a fully customised and hence intricately controlable, camera system may be more appropriate for any future system.

A robot capable of examining surfaces of composite insulators has been designed and partially built with the assistance of the Chris Kirkham from The School of Electrical and Electronic Engineering's Mechanical Workshop. The component constructed consisted of the surface scanning mechanism. While seen to operate in certain situations, when programmed to operate in a manner as required by a final system, mechanical tolerances in the construction of the mounting carriage's movement mechanism caused the system to be unreliable. As such an estimate for scan-time of each insulator surface could be assessed. Possible designs for vertical movement along an insulator string and mechanisms for insertion and removal from a installed insulator have additionally been designed and presented.

A dual control strategy has been designed consisting of a primary link operating

over WLAN and a secondary low data rate narrow band non-protocol link intended for backup and monitoring purposes. Such an implementation is strongly recommend in any final system, as it provides a mechanism to communicate with the remote device in the event of protocol-based communication becoming unreliable.

Finally, this research has not considered a power provisioning system in great detail, instead choosing to use a number of battery cells. While this is acceptable during development, assessment of the power requirements of the complete robot is required and needs to be related to both power scavenging and available battery capacity choices prior to integration into a final system.

If batteries are used in the final system, it is essential to ensure that enough capacity exists to avoid failure at an inopportune moment; for example part way along an insulator string, which in turn may prove difficult to retrieve. As noted in Section 2.5 there are two possible battery strategies that could be implemented: for one day's operation of for operation over one assessment sequence[1].

If one day's operation is targeted, then the batteries used can be highly integrated into the device, thereby reducing mechanical complexity. However the batteries would need to have a quick charge cycle, such that operators waste no time waiting for the device to charge. Alternatively, if a reduced operation time system is targeted, the batteries would have to be easily changeable. While such a system would potentially increase mechanical complexity (by requiring a battery compartment and access hatch) it may actually reduce the weight of any final system.

13.1 Design Recommendations: System Trade-offs

This research has considered the subsystems required for practical condition monitoring robots. This has lead to a number of interesting operational and design questions. It is clear that while an all encompassing system would be the ultimate device, it is not without its own drawbacks; in this respect a number of design trade-offs are required.

The time required to scan each insulator shed is related to the number of cameras

---

[1]Assessment of one insulator string.

used and the image quality. Two cameras would potentially half the time required to examine an insulator string, but would then in turn increase the complexity of the mechanical design and potentially reduce reliability. Furthermore, the use of two cameras would increase the bandwidth requirements for live transmission, while a 'store and review' approach to assessment would require a doubling of storage media write speeds. This would also be the case if the required image quality increases.

The increase in image quality and the use of a second camera would additionally increase the processor power required, which would in turn potentially increase the power consumption of the system. However it should be noted that the power requirements of the mechanical components will still outweigh that of the remainder of the system. This would subsequently either reduce the operation time of the system or increase the device's weight through the need for larger capacity battery packs.

All these areas could in turn potentially reduce the overall operational life of the system through the increase in complexity and its associated maintenance requirements. As such reliability appears key in any system developed which in turn means reduced complexity. This essentially would suggest a single camera system with either a live transmission or 'store and review' approach and a variable image quality to maintain a level of flexibility. Such a system would potentially be less complex and hence more reliable, but at the detriment of longer insulator assessment period.

13.2 Future Work

While the various basic ideas needed for an insulator monitoring robot have been demonstrated, a number of components have been crudely designed in order to further testing. Other components have only been developed to an extent where demonstration is possible and a concept can be proved. This therefore leaves a margin for improvement and refinement of most components.

It has been demonstrated that a cylindrical waveguide transition antenna link operates at closer proximities to the energised line. However the antennas constructed were of simplistic design and crudely tuned to the WLAN channel of interest.

Furthermore, tuning required sub-millimetre accuracy in both the length of the radiating element and it positioning from the back wall of the tube. Due to these concerns, future work will involve the design and fabrication of a sub-millimetre accurate cylindrical waveguide transition antenna, tuned to a central frequency in WLAN band.

While the imaging system functions in principle, there are currently a number of shortcomings. The major issue relates to compatibly with Linux; with the current webcam causing the capture mechanism to continually generate errors. These errors are thought to be related to frames being dropped, subsequently thought to be causing the processor utilisation to run at over 80%. Therefore, another future task involves assessing a large number of commercially available webcams and cross-comparing resolution, Linux compatibility, resultant CPU use, physical size and cost.

An alternative solution is to re-explore the idea of creating a fully customised image capture solution using the Gumstix's camera interface. At the time the support (source code repository) was minimal and it was considered inefficient to spend time developing low-level software. However as the hardware which interfaces the camera to the image capture interface already exists, further potential work includes a re-assessment of current status of the source code repository and developing code if required.

Both solutions also require detailed assessment of the necessary focusing optics. When initially implemented the light guide restricted the view-port to a small area of the captured image. It was then shown in principle that the image could be enlarged with careful placement of a lens within the light guide itself. However this then caused the image to be blurred and required a secondary lens to remove the distortion. Although this functioned as expected outside the robot, a clear focus could not be obtained at the required working distances.

Full details of the distances between target and imaging device will be needed from which focusing and magnification optics can be selected or designed. This analysis will be easier for the fully customised solution as the webcam's own optics will not effect any calculations. Ideally these optics will fabricated at fixed locations within the light guide itself, thereby making the system more robust.

Although the various mechanical motions required for a full system have been

designed, only the surface scanning mechanism was implemented. This motion itself is not without its own flaws, mainly related to the mechanical tolerances to which the components have been made. It was discovered that while the carriage moved cleanly over certain sections, at others the carriage's drive gear lost contact with the drive track. This then made the implemented carriage positioning system unreliable.

Primarily the carriage and its mounting platform will need to be re-manufactured to a high degree of precision with the drive track possibly being formed with the mounting platform as a single unit. Secondarily the shed-to-shed movement and clamping systems will need to be implemented. The latter is the greater issue as during investigations with the developed system it became apparent that it was too heavy, weighing ~9 kg. Firstly, the weight was causing the insulator shed to become damaged, something that needed to be avoided. Secondly, at ~9 kg it is was challenging to fit the robot to the insulator single-handedly when situated directly next to the insulator. The end device needs to be much lighter if it is to be attached via hot-stick or some other mechanism. This needs to be investigated with further assessment of the Faraday cage used to protect the electronics; it was seen that when the device is at its lowest point, the protection offered by the current design is insufficient to protect some of the systems.

Currently a dual control system is used, where the imaging data is transmitted over a WLAN channel and robotic control is achieved over a secondly non-protocol based 2.4 GHz link. However, circuits were designed with single link communication and control in mind. Motor control programs can be written for main processor and exposed via CGI scripts and directly merged with the video stream in a single interface. Furthermore, the backup hardware (secondary link receiver) can be miniaturised into an 'USB Stick' and its control system additionally merged into the main application user interface.

At present a simple battery system is used for power provisioning. These batteries, while replaceable, requires a significant amount of disassembly of the robot to achieve. This is something that cannot be tolerated in a final system. Different battery technologies need to be explored and associated capacity and physical size & weight need to be compared with the power requirements of the complete system.

13.3 Ancillary Outcomes

During the assessment of the radio environment surrounding a composite insulator numerous measurements were made. Taking into account various deficiencies, these measurements indicate that higher data rates are seen when the transmitter is located towards the centre of the insulator string. Comparing these with simulated field distributions, there is a suggestion that higher field strengths lead to reduced data rates.

Discussion with other academics revealed that presently a 'clean' method to assess the electric field distribution around insulators is not available. Therefore by miniaturising the transmitter system, basing it either on WLAN or non-protocol 2.4 GHz transmission and using a non-intrusive patch antenna, a system to indirectly measure field strengths could be created. These devices would need to be placed at regular intervals on the insulator, either permanently or when assessment is required.

However before this can be taken forward further assessment of the radio environment is required. This involves preforming similar tests to those preformed within this research but in a confirmed radio quite environment (an anechoic chamber) to ensure external sources of 2.4 GHz radio energy are not contaminating gathered results.

13.4 Closing Remarks

This research has developed and shown the constituent parts required for the construction of a complete insulation monitoring robot. As noted in the Section 13.2 there is still research and refinement to be conducted in order to produce a fully working system. Some of this work is outside the domain of experience, for example optical imaging and high precision manufacture.

Obscure obstacles were additionally encountered, whereby sub-systems function as expected and remained stable for extended durations during the development process, but repeatedly failed during planed testing. A specific reason for this was not ascertained. A change in the implemented WLAN methodology to use the commonly available 'Bridging Mode' on WAPs allowed the 2.4 GHz band to be

assessed successfully.

It was subsequently shown that sufficient data bandwidth exists for the transmission of data from areas of high field strength and partial discharge. One form of data is real-time video which has been successfully demonstrated on an energised representation of a high voltage line.

As complex electronics have been shown to work successfully under energised conditions, this research can provide a base from which other complex robotic high voltage inspection robots can be built. It is also important to note that visual light cameras are just one of many sensors than could be implemented, others include infra-red and UV imaging elements, acoustic transducers and electric field probes.

Regardless of the issues encountered valid research outcomes have been achieved, particularly in the domain of radio frequency and data rate knowledge, which has been gathered through practical means rather than simulation.

This knowledge can now lead to the development of new monitoring tools to aid utilities in managing their networks and to the creation of systems to aid research in high voltage engineering, including new mechanisms of electric field distribution analysis.

References

[1] National grid: Appendix ii main features of a transmission line. Online. National Grid plc. [Online]. Available: http://www.nationalgrid.com/uk/ LandandDevelopment/SC/devnearohl_final/appendix2/

[2] National grid emf: Geometries for calculating fields from power lines. Online. National Grid plc. [Online]. Available: http://www.emfs.info/Sources+of+ EMFs/Overhead+power+lines/Calculating/geometries/

[3] S. Montambault and N. Pouliot, "Field experience with linescout technology for live-line robotic inspection and maintenance of overhead transmission networks," in *Applied Robotics for the Power Industry (CARPI), 2010 1st International Conference on*, oct. 2010, pp. 1 –2.

[4] J. Toth, N. Pouliot, and S. Montambault, "Field experiences using linescout technology on large bc transmission crossings," in *Applied Robotics for the Power Industry (CARPI), 2010 1st International Conference on*, oct. 2010, pp. 1 –6.

[5] W. Hongguang, J. Yong, L. Aihua, F. Lijin, and L. Lie, "Research of power transmission line maintenance robots in siacas," in *Applied Robotics for the Power Industry (CARPI), 2010 1st International Conference on*, oct. 2010, pp. 1 –7.

[6] J. Looms, "Live working on high-voltage lines," *Physical Science, Measurement and Instrumentation, Management and Education - Reviews, IEE Proceedings A*, vol. 128, no. 2, p. 89, march 1981.

[7] S. Wallstrom, A. Dernfalk, M. Bengtsson, S. Kroll, S. Gubanski, and S. Karlsson, "Image analysis and laser induced fluorescence combined

to determine biological growth on silicone rubber insulators," *Polymer Degradation and Stability*, vol. 88, no. 3, pp. 394 – 400, 2005. [Online]. Available: http://www.sciencedirect.com/science/article/B6TXS-4FHJGC6-1/2/2c85023f30f8fd8e357ea829754c4ee8

[8] IEEE, "Ieee standard for information technology-telecommunications and information exchange between systems-local and metropolitan area networks-specific requirements - part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications," *IEEE Std 802.11-2007 (Revision of IEEE Std 802.11-1999)*, pp. C1 –1184, 12 2007.

[9] J. H. Dunlap, J. M. Van Name, and J. A. Henkener, "Robotic maintenance of overhead transmission lines," *Power Delivery, IEEE Transactions on*, vol. 1, no. 3, pp. 280 –284, july 1986.

[10] A. Santamaria, R. Aracil, A. Tuduri, P. Martinez, F. Val, L. Penin, M. Ferre, E. Pinto, and A. Barrientos, "Teleoperated robots for live power lines maintenance (robtet)," in *Electricity Distribution. Part 1: Contributions. CIRED. 14th International Conference and Exhibition on (IEE Conf. Publ. No. 438)*, vol. 3, 1997, pp. 31/1 –31/5 vol.3.

[11] "Safety rules (third edition) guidance notes," National Grid UK Electricity Transmission plc, Tech. Rep. NGUK/PL/ETSR/GN Issue 4, 2011.

[12] J. Roberson. Composite polymeric overhead line insulator systems. Online. The University of Manchester. [Online]. Available: http://www.eee.manchester.ac.uk/research/groups/eeps/docs/Robertson.pdf

[13] Guide 5, 2004, guide for visual identification of deterioration and damage on suspension composite insulators. Online. STRI. [Online]. Available: http://ewh.ieee.org/soc/pes/iwg/NCIEvaluation/STRIGuide503.pdf

[14] E. da Silva. Reliability of composite insulators. Online. The University of Manchester. Supergen - AMPerES. [Online]. Available: http://www.supergen-amperes.org/Other%20Information/OHL%20Meeting%2008/Reliability_of_Composite_Insulators_Seminar.pdf

[15] S. Gubanski, A. Dernfalk, J. Andersson, and H. Hillborg, "Diagnostic methods for outdoor polymeric insulators," *Dielectrics and Electrical Insulation, IEEE Transactions on*, vol. 14, no. 5, pp. 1065 –1080, october 2007.

[16] P. Lang, D. Allan, and Y. Zhou, "The investigation of insulation defects in transmission line disc insulators using remote detection techniques," in *Properties and Applications of Dielectric Materials, 1994., Proceedings of the 4th International Conference on*, vol. 2, Jul. 1994, pp. 868 –871 vol.2.

[17] G. Jaensch, H. Hoffmann, and A. Markees, "Locating defects in high voltage transmission lines," in *Transmission Distribution Construction, Operation Live-Line Maintenance Proceedings, 1998. ESMO '98. 1998 IEEE 8th International Conference on*, Apr. 1998, pp. 179 –186.

[18] D. Giorgio, R. Romeo, P. Giovanni, and V. Claudio. (Accessed 2006) Studies for the diagnostics of composite insulators. Online. [Online]. Available: http://www.dke.de/NR/rdonlyres/ 5EE04422-B139-4B38-81B0-7CEE46F95C16/20908/etg_3_13.PDF

[19] "National safety instruction: Guidance notes: Nsi 4: Work on or near high voltage overhead lines," National Grid UK Electricity Transmission plc, Tech. Rep. NGUK/PM/ETSR/NSI/04/GN Issue 4, 2011.

[20] P. Moore, I. Portugues, and I. Glover, "Radiometric location of partial discharge sources on energized high-voltage plant," *Power Delivery, IEEE Transactions on*, vol. 20, no. 3, pp. 2264 – 2272, Jul. 2005.

[21] *MPEG-1: Coding of Moving Pictures and Associated Audio for Ditigal Storage Media at up to about 1.5 Mbit/s*, Online, ISO Std. ISO/IEC JTC1/SC29WG11, 1996. [Online]. Available: http://mpeg.chiariglione.org/ standards/mpeg-1/mpeg-1.htm

[22] *Information Technology - Digital Compression and Coding of Continuous Tone Still Images - Requirements and Guildlines*, Online, ISO/IEC; ITU Std. ISO/IEC IS 10 918-1; ITU-T Recommendation T.81, 1992. [Online]. Available: http://www.jpeg.org/jpeg/index.html

[23] E. Van Atta and E. L. White, "Radio interference from line insulators," *American Institute of Electrical Engineers, Transactions of the*, vol. 49, no. 1, pp. 1 –5, Jan. 1930.

[24] G. E. Adams, "Radio interference from high-voltage transmission lines as influenced by the line design," *Power Apparatus and Systems, Part III. Transactions of the American Institute of Electrical Engineers*, vol. 77, no. 3, pp. 54 –62, Apr. 1958.

[25] R. Mather and B. Bailey, "Radio interference from high-voltage lines," *Power Apparatus and Systems, IEEE Transactions on*, vol. 82, no. 68, pp. 775 –782, Oct. 1963.

[26] W. Pakala and V. Chartier, "Radio noise measurements on overhead power lines from 2.4 to 800 kv," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-90, no. 3, pp. 1155 –1165, May. 1971.

[27] J. Langton and E. Bradshaw, "Radio interference from discharges on high-voltage line insulators," *Electrical Engineers, Journal of the Institution of*, vol. 75, no. 455, pp. 643 –652, Nov. 1934.

[28] J. Reichman and J. R. Leslie, "Radio interference studies on extra-high-voltage lines," *Power Apparatus and Systems, Part III. Transactions of the American Institute of Electrical Engineers*, vol. 80, no. 3, pp. 261 –266, Apr. 1961.

[29] V. Chartier, D. Blair, R. Stearns, and D. Lamb, "Effect of bundle orientation on transmission line audible and radio noise," *Power Delivery, IEEE Transactions on*, vol. 9, no. 3, pp. 1538 –1544, Jul. 1994.

[30] F. Warburton, T.-W. Liao, and N. Hoglund, "Power line radiations and interference above 15 mhz," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-88, no. 10, pp. 1492 –1501, Oct. 1969.

[31] M. Moreau and C. Gary, "Predetermination of the radio-interference level of high voltage transmission lines i - predetermination of the excitation function," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-91, no. 1, pp. 284 –291, Jan. 1972.

[32] P. Maruvada and N. Trinh, "A basis for setting limits to radio interference from high voltage transmission lines," *Power Apparatus and Systems, IEEE Transactions on*, vol. 94, no. 5, pp. 1714 – 1724, Sept. 1975.

[33] V. Chartier, S. Lowder, S. Rodick, and M. Vogt, "Radio and audible noise performance of t2 conductors," *Power Delivery, IEEE Transactions on*, vol. 11, no. 3, pp. 1464 –1474, Jul. 1996.

[34] J. Silva and R. Olsen, "Use of global positioning system (gps) receivers under power-line conductors," *Power Delivery, IEEE Transactions on*, vol. 17, no. 4, pp. 938 – 944, Oct. 2002.

[35] S. Nayak and M. Thomas, "Computation of emi fields generated due to corona on high voltage over head power transmission lines," in *Electromagnetic Interference and Compatibility, 2001/02. Proceedings of the International Conference on*, 2002, pp. 15 – 19.

[36] D. Huang and J. Ruan, "Discussion on electromagnetic environment of 1000 kv ac overhead transmission lines in china," in *Power Engineering Conference, 2007. IPEC 2007. International*, Dec. 2007, pp. 752 –757.

[37] D. Huang, J. Ruan, and F. Huo, "Study on the electromagnetic environment of 1000 kv ac double-circuit transmission lines in china," in *Power Systems Conference and Exposition, 2009. PSCE '09. IEEE/PES*, Mar. 2009, pp. 1 –7.

[38] M. Hikita, T. Kato, H. Yamashita, and H. Okubo, "Phase-resolved partial discharge measurements and their electromagnetic spectrum in sf6 gas," in *Partial Discharge, 1993., International Conference on*, Sep. 1993, pp. 121 –122.

[39] M. Hikita, H. Yamashita, T. Kato, N. Hayakawa, T. Ueda, and H. Okubo, "Electromagnetic spectrum caused by partial discharge in air under ac and dc voltage application," in *Properties and Applications of Dielectric Materials, 1994., Proceedings of the 4th International Conference on*, vol. 2, Jul. 1994, pp. 570 –573 vol.2.

[40] M. Hikita, H. Yamashita, T. Hoshino, T. Kato, N. Hayakawa, T. Ueda, and H. Okubo, "Electromagnetic noise spectrum caused by partial discharge in air

at high voltage substations," *Power Delivery, IEEE Transactions on*, vol. 13, no. 2, pp. 434 –439, Apr. 1998.

[41] T. Babnik, R. Aggarwal, P. Moore, and Z. Wang, "Radio frequency measurement of different discharges," in *Power Tech Conference Proceedings, 2003 IEEE Bologna*, vol. 3, Jun. 2003, p. 5 pp. Vol.3.

[42] S. Xiao, P. Moore, M. Judd, and I. Portugues, "An investigation into electromagnetic radiation due to partial discharges in high voltage equipment," in *Power Engineering Society General Meeting, 2007. IEEE*, June 2007, pp. 1–7.

[43] M. Kezunovic, C. N. Georghiades, and A. Shapoury, "Power sytem monitoring using wireless substations and system-wide communications part i," Texas A&M University, Texas A&M University Department of Electrical Engineering College Station, TX 77845-3128, Tech. Rep. PSERC 02-46, November 2002.

[44] Z. ze Liang, E. Li, M. Tan, G.-P. Liu, and D. Rees, "Control of inspection robots for power transmission lines based on remote video over internet," in *Computational Intelligence for Modelling, Control and Automation, 2005 and International Conference on Intelligent Agents, Web Technologies and Internet Commerce, International Conference on*, vol. 1, Nov. 2005, pp. 877–882.

[45] Q. Shan, I. Glover, P. Moore, I. Portugues, R. Watson, and R. Rutherford, "Performance of zigbee in electricity supply substations," in *Wireless Communications, Networking and Mobile Computing, 2007. WiCom 2007. International Conference on*, Sept. 2007, pp. 3871–3874.

[46] Q. Shan, I. Glover, P. Moore, I. Portugues, R. Watson, R. Rutherford, R. Atkinson, and S. Bhatti, "Laboratory assessment of wlan performance degradation in the presence of impulsive noise," in *Wireless Communications and Mobile Computing Conference, 2008. IWCMC '08. International*, Aug. 2008, pp. 859–863.

[47] IEEE, "Ieee standard for information technology - telecommunications and information exchange between systems - local and metropolitan area networks

- specific requirements. - part 15.1: Wireless medium access control (mac) and physical layer (phy) specifications for wireless personal area networks (wpans)," *IEEE Std 802.15.1-2005 (Revision of IEEE Std 802.15.1-2002)*, pp. 1 –580, 2005.

[48] ——, "Ieee draft standard for information technology - telecommunications and information exchange between systems - local and metropolitan area networks - specific requirements - part 15.4: Wireless medium access control (mac) and physical layer (phy) specifications for low rate wireless personal area networks (wpans)," *IEEE P802.15.4REVi/D04, September, 2010*, pp. 1 –334, 17 2011.

[49] Z. Yu, R. Zeng, M. Li, R. Li, L. Liu, B. Zhang, Z. Zhang, and H. Zhang, "Tests on electromagnetic environment of ultra hvdc transmission lines in high altitude region," in *Electromagnetic Compatibility (APEMC), 2010 Asia-Pacific Symposium on*, Apr. 2010, pp. 1676 –1680.

[50] Z. Yu, R. Zeng, M. Li, R. Li, L. Liu, D. Yang, Z. Zhang, B. Zhang, and F. Tian, "Radio interference of ultra hvdc transmission lines in high altitude region," in *Electromagnetic Compatibility (APEMC), 2010 Asia-Pacific Symposium on*, Apr. 2010, pp. 1672 –1675.

[51] Q. Zhou, C. Sun, L. Liu, W. Sima, and W. An, "Electromagnetic environment of the ehv transmission line and its effect," in *Electrical Insulating Materials, 2001. (ISEIM 2001). Proceedings of 2001 International Symposium on*, 2001, pp. 229 –232.

[52] B. Bailey, "Test line experience with hvdc overhead transmission," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-89, no. 7, pp. 1625 –1634, Sep. 1970.

[53] F. Bologna, N. Mahatho, and D. Hoch, "Infra-red and ultra-violet imaging techniques applied to the inspection of outdoor transmission voltage insulators," in *Africon Conference in Africa, 2002. IEEE AFRICON. 6th*, vol. 2, oct. 2002, pp. 593 – 598 vol.2.

[54] R. A. Fernandes and R. L. Sieron, "Transmission line sensor apparatus operable with near zero current line conditions." U.S. Patent 4 808 917, 1989.

[55] R. A. Fernandes, J. E. Burbank-III, W. R. Smith, and R. L. Sieron, "Apparatus for measuring the temperature and other parameters of an electric power conductor." European Patent 84 302 500, 1984.

[56] R. A. Fernandes, "Electrical power line parameter measurement apparatus and systems, including compact, line-mounted modules." European Patent 87 309 870, 1987.

[57] K. Lau, J. Yee, R. Mayeda, and M. Riccomini, J.and Ilyin, "ault sensor device with radio transceiver," U.S. Patent 5 565 783, 1996.

[58] S. Hirose and S. Aoki, "Development of the insulator washer robot," in *Robotics and Automation, 1995. Proceedings., 1995 IEEE International Conference on*, vol. 2, may 1995, pp. 1783 –1789 vol.2.

[59] S. Murakami, K. Yanoyasuo, K. Higashijima, H. Wakizako, K. Takaoka, T. Irie, J. Goto, and T. Hasegawa, "Application of live line work robots for distribution work: Kyushu electric's challenges for fully-automated robotic system," *Advanced Robotics*, vol. 15, no. 3, pp. 339–344, 2001. [Online]. Available: http://www.ingentaconnect.com/content/vsp/arb/2001/00000015/00000003/art00012?crawler=true

[60] S.-y. Jiang, Y. Hu, Y. Wang, H. Jaio, and L. Ren, *Intelligent Robotics and Applications*, ser. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2008, vol. 5314, ch. Development of Hanging-Arm Inspection Robot for High-Voltage Transmission Line, pp. 1089–1098.

[61] S. Montambault and N. Pouliot, "The hq linerover: contributing to innovation in transmission line maintenance," in *Transmission and Distribution Construction, Operation and Live-Line Maintenance, 2003. 2003 IEEE ESMO. 2003 IEEE 10th International Conference on*, april 2003, pp. 33 – 40.

[62] ——, "About the future of power line robotics," in *Applied Robotics for the Power Industry (CARPI), 2010 1st International Conference on*, oct. 2010, pp. 1 –6.

[63] *BS 5049-1:1994 Radio interference characteristics of overhead power lines and high voltage equipment - Part 1: Description of phenomena*, British Standards Institute Std. ISBN: 0-580-22 822-3, 1994.

[64] P. Moore, I. Portugues, and I. Glover, "Partial discharge investigation of a power transformer using wireless wideband radio-frequency measurements," *Power Delivery, IEEE Transactions on*, vol. 21, no. 1, pp. 528 – 530, Jan. 2006.

[65] *Wireless Telegraphy Act 2006*, Online, Her Majesty's Government Std., 2006. [Online]. Available: http://www.legislation.gov.uk/ukpga/2006/36/ introduction

[66] (Accessed June 2008) The uk frequency allocation. Online. Roke Manor Research Ltd. [Online]. Available: http://www.roke.co.uk/download/ datasheets/uk-frequency-allocations.pdf

[67] *Interface Requirements 2030: Licence Exempt Short Range Devices*, Online, Ofcom: The Office of Communications Std., 2006.

[68] F. Liptrot, *IEE Power and Engineering Series 32: High Voltage Engineering and Testing*, 2nd ed., ser. IEE Power and Engineering Series. London: The Institution of Electrical Engineers, 2001, ch. Overhead Lines, pp. pp. 167–211.

[69] *MCHPFSUSB v1.3 USB Framework*, Microchip Inc, 2007, accessed January 2008. [Online]. Available: http://www.microchip.com/stellent/idcplg?IdcService=SS_GET_ PAGE&nodeId=2124&param=en532204&page=wwwFullSpeedUSB

[70] *Pseudo-Random Number Generation Routine for the MAX765x Microprocessor*, Maxim IC, Dec. 2007. [Online]. Available: http: //pdfserv.maxim-ic.com/en/an/AN1743.pdf

[71] Errata notes cc2500. Online. Texas Instruments. Austin, Texas, US. [Online]. Available: http://www.ti.com/litv/pdf/swrz002d

[72] L. Chappell, *Wireshark Network Analysis*, 1st ed. Protocol Analysis Institute, LLC, 5339 Prospect Road, 343 San Jose, CA 95129, USA: Protocol Analysis Institute, LLC, 2010.

[73] *PRF 2581 - TCP Congestion Control*, Online, The Internet Engineering Task Force Std., 1999. [Online]. Available: http://tools.ietf.org/html/rfc2581

[74] T. Zhao and M. Comber, "Calculation of electric field and potential distribution along nonceramic insulators considering the effects of conductors and transmission towers," *Power Delivery, IEEE Transactions on*, vol. 15, no. 1, pp. 313 –318, jan 2000.

[75] W. Que and S. Sebo, "Electric field and potential distributions along dry and clean non-ceramic insulators," in *Electrical Insulation Conference and Electrical Manufacturing Coil Winding Conference, 2001. Proceedings*, 2001, pp. 437 –440.

[76] B. Zhang, S. Han, J. He, R. Zeng, and P. Zhu, "Numerical analysis of electric-field distribution around composite insulator and head of transmission tower," *Power Delivery, IEEE Transactions on*, vol. 21, no. 2, pp. 959 – 965, april 2006.

[77] D. Huang, J. Ruan, Y. Chen, F. Huo, S. Yu, and S. Liu, "Calculation and measurement of potential and electric field distribution along 1000 kv ac transmission line composite insulator," in *Electrical Machines and Systems, 2008. ICEMS 2008. International Conference on*, oct. 2008, pp. 428 –433.

[78] W. Que, S. Sebo, and R. Hill, "Practical cases of electric field distribution along dry and clean nonceramic insulators of high-voltage power lines," *Power Delivery, IEEE Transactions on*, vol. 22, no. 2, pp. 1070 –1078, april 2007.

[79] R. Zeng, Y. Zhang, W. Chen, and B. Zhang, "Measurement of electric field distribution along composite insulators by integrated optical electric field sensor," *Dielectrics and Electrical Insulation, IEEE Transactions on*, vol. 15, no. 1, pp. 302 –310, february 2008.

[80] Uk cantennas. Online. Private Individual. [Online]. Available: http://www.tropic.org.uk/~edward/homepage/cantenna/index.html

[81] *RFC 1350 - TFTP Protocol (Revision 2)*, Online, The Internet Engineering Task Force Std., 1992. [Online]. Available: http://tools.ietf.org/html/rfc1350

[82] O. Ribeiro. (2006, June) 374931 - ffserver - single jpeg output stream - debian bug reprt logs. Online: Debian Bug Report Log. Accessed: August 2009, Available:http://bugs.debian.org/cgi-bin/bugreport.cgi?bug=374931. [Online]. Available: http://bugs.debian.org/cgi-bin/bugreport.cgi?bug=374931

Appendix A

Data Types Extracted

The data collated during preliminary high voltage investigations is too vast to included in print. The processed error data can be found on the attached DVD under the './data/<test>/Excel/' directory; the test voltage used is specified as part of the filename. Error reports can be found in the './data/<test>/Error Report/' directory; the meaning of the different file extensions can be found below.

Ext.	Description	Type
err	Core error file, contains raw error data accumulated when processing raw transmission data.	Text
errlst	Documents channel, total number of bit errors in that channel and the average number of bit errors when averaged over the number of error packets	CSV & XLS
errbepp	Error Bits Per Packet, i.e. number of packets with 1 bit error, number with 2 bit errors etc. - documented per channel - for a total of 48 entries per channel - 0 bit errors are not included as there should be none in this analysis	CSV
errmp	Error - number of missing PRNs between local and remote values. When the number of missing numbers is small, then the likely explanation is missed packets; otherwise the likelihood is data corruption.	Text
errbp	Number of errors in each bit position of packets when remote instances are the same.	CSV

Table A.1: Error Report File Extensions

This page is intentionally blank.

Appendix B

Video File Listing

B.1   Videos

A video showing the electrical discharge from an energised line to the developed robot's Faraday cage can be found on the DVD in the ./video/live/ directory.

Videos of the energised video streaming experiments can be found on the DVD in the ./video/streaming directory. Six videos are available:

- At shed 36 un-energised
- At shed 36 energised to 231 kV
- At shed 54 un-energised
- At shed 54 energised to 231 kV
- At shed 63 un-energised
- At shed 63 energised to 231 kV

B.2   Mechanical Animations

In order to better visualise the mechanisms described in this report, a number of animations have been produced. The animations can be found on the attached DVD in the ./video/animation/ directory; the files are described below.

Name	Animation Description
Three Part Door	This indicates how Fernades et al[56, 55, 54] 'jaw' mechanism would operate when applied to a larger device. This animation however does not depict the actual mechanism that will perform the opening and closing.
Motorised Pivot Points	This animates the initially proposed insulator string locomotion mechanism. Two sets of three motorised mounting points are used to support and to move the robot from shed to shed. This animation only depicts one from each set for improve clarity. One set of mounting points retract and moves along the length of the robot and subsequently latches onto the next shed. The second set retract and moves along with the body of the robot to the same position as the first.
Mechanical Pivot Points	This refines the previous concept. The previous solution places the weight of the robot on three or six motorised pivot points; the motors would need to provide enough torque to support this weight. A motorised ring contains three equally placed notches into which pivots can 'fall' when correctly aligned. Use of two sets of these pivot rings mounted above one another will allow two sets of pivots to be used independently. If these rings are then able to move along the body's length, a shuffling motion is then possible.
Camera Motion and Tracking	This animation shows the motion of the carriage around a shed of the insulator. The wheel depicted outside the main ring and the overhanging part of the carriage will not be possible in reality due to the shell of the robot, located near the outer edge of the main ring. A prism mirror will be attached to the end of a linear actuator and will be able to pass over the surface of the shed.
Prism Mirror Rotation and Sensor Usage	This animation shows a minor problem with the use of a prism mirror in conjunction with the use of a fixed position camera. As the prism mirror moves towards the centre of the shed, the lens starts to 'see' a finite area around the surface of the prism. This then means that the portion of the sensor used to capture surface data is reduced for central locations compared to edge locations.

Table B.1: Animation Details

Appendix C

Webcam Video Listing

The application `FFmpeg` has many input and output options and is therefore very versatile. While it attempts to make a guess as to what the parameters should be, it does not always get it correct. This was the case for the two webcams that were tested. The table below details the camera, the configuration and a description of the resultant video file. The video files can be found on the enclosed DVD in the ./video/ffmpeg directory.

| ID | Filename | Camera | Input | | | | | Output | | | |
			Format	Size	Codec	Pixel Format	Rate	Format	Codec	Description	
1	20090728-1.mpeg	Logitech	rawvideo	cif	N/A	N/A	15fps	mpeg	N/A	'Static' video	like
2	20090728-2.mpeg	Creative Labs	rawvideo	cif	N/A	N/A	15fps	mpeg	N/A	'Static' video	like
3	20090728-1.swf	Logitech	rawvideo	cif	N/A	N/A	N/A	swf	N/A	'Static' video	like
4	20090729-1.swf	Logitech	rawvideo	qvga	rawvideo	N/A	N/A	swf	N/A	'Static' video	like
5	20090729-2.swf	Logitech	rawvideo	vga	N/A	N/A	30fps	swf	N/A	'Static' video	like
6	20090729-3.swf	Creative Labs	rawvideo	vga	N/A	N/A	N/A	swf	N/A	'Static' video	like

Table C.1: Camera Configurations Part 1

ID	Filename	Camera	Input					Output		Description
			Format	Size	Codec	Pixel Format	Rate	Format	Codec	
7	20090729-4.swf	Logitech	rawvideo	vga	N/A	rgb8	N/A	swf	N/A	'Static' like video
8	20090729-5.swf	Logitech	rawvideo	vga	N/A	rgb24	N/A	swf	N/A	'Static' like video
9	20090729-6.swf	Logitech	rawvideo	vga	N/A	yuyv422	N/A	swf	N/A	'Static' like video
10	20090729-7.swf	Logitech	rawvideo	vga	N/A	bgr24	N/A	swf	N/A	'Static' like video
11	20090729-8.swf	Logitech	rawvideo	vga	N/A	pal8	N/A	swf	N/A	'Static' like video
12	20090729-9.swf	Logitech	rawvideo	vga	N/A	pal8	1fps	swf	N/A	'Static' like video
13	20090729-10.swf	Logitech	mjpeg	N/A	N/A	N/A	N/A	swf	N/A	Blury Video, Dropping Frames, same for 20090729-11.swf, 20090729-12.swf, 20090729-13.swf

Table C.2: Camera Configurations Part 2

| ID | Filename | Camera | Input | | | Output | | | Description |
			Size	Codec	Format	Size	Codec		
14	20090729-14.swf	Logitech	vga	N/A	swf	vga	N/A		Blury Video
15	20090729-15.swf	Creative Labs	vga	N/A	swf	vga	N/A		Out of Focus
16	20090729-16.swf	Creative Labs	cif	N/A	swf	cif	N/A		Out of Focus
17	20090729-17.swf	Creative Labs	N/A	N/A	swf	N/A	N/A		Out of Focus
18	20090729-1.mpg	Creative Labs	N/A	N/A	mpeg	N/A	N/A		Out of Focus, but smooth video
19	20090729-1.avi	Creative Labs	N/A	copy	mpjpeg	N/A	N/A		No Video
20	20090729-18.swf	Creative Labs	N/A	copy	swf	N/A	N/A		No Images
21	20090729-1.jpg	Creative Labs	N/A	N/A	mjpeg	N/A	N/A		+ 20090729-2.jpg, 20090729-4.jpg are 'washed' out

Table C.3: Camera Configurations with MJPEG Input Format Part 1

ID	Filename	Camera	Input			Output		Description
			Size	Codec	Format	Size	Codec	
22	20090729-3.jpg	Logitech	N/A	N/A	mjpeg	N/A	N/A	Over Exposed
23	20090729-19.swf	Creative Labs	N/A	N/A	swf	N/A	N/A	Change focus while capturing, does get into focus
24	20090729-20.swf	Logitech	N/A	N/A	swf	N/A	N/A	Change focus while capturing, does NOT get into focus
25	20090729-21.swf	Logitech	N/A	N/A	swf	vga	N/A	Second focus test
26	20090729-2.avi	Logitech	N/A	copy	avi	N/A	N/A	Blurry video and frame error
27	20090729-3.avi	Logitech	N/A	copy	avi	N/A	N/A	Blurry video and frame error
28	20090729-4.avi	Creative Labs	Size	N/A	avi	N/A	N/A	Blurry video and frame error
29	20090729-5.avi	Creative Labs	N/A	copy	avi	N/A	N/A	Blurry video and frame error
30	20090729-3.mpeg	Creative Labs	N/A	copy	N/A	N/A	mpeg1video	Incorrect Codec

Table C.4: Camera Configurations with MJPEG Input Format Part 2

This page is intentionally blank.

Appendix D

Engineering Drawings

Full size engineering drawings can be found on the enclosed DVD in Portable Document Format (PDF) and are also presented here for quick reference.
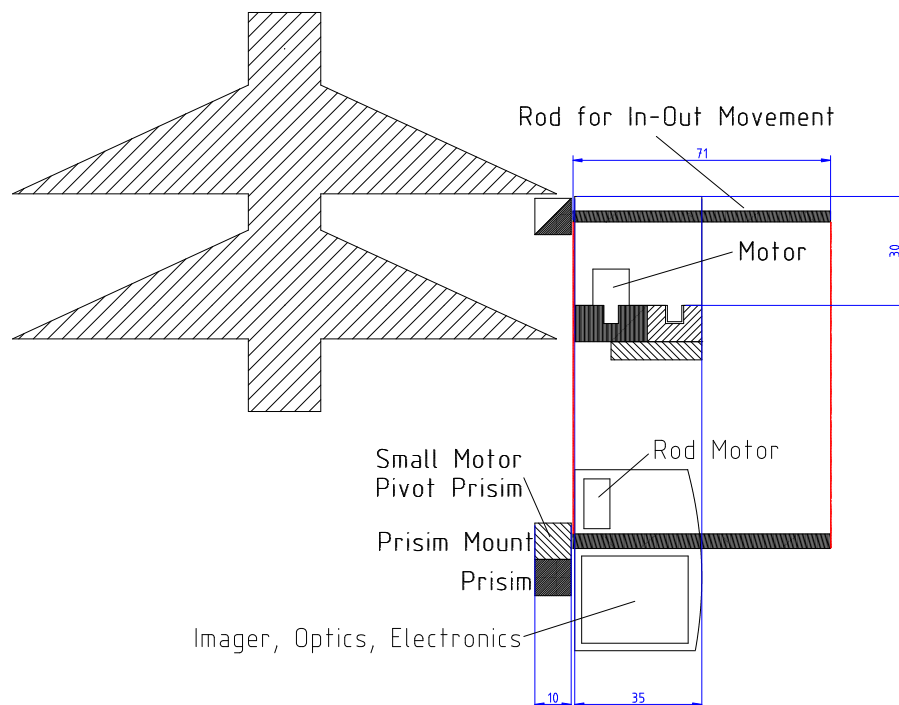
Figure D.1: Engineering drawing of initially proposed carriage
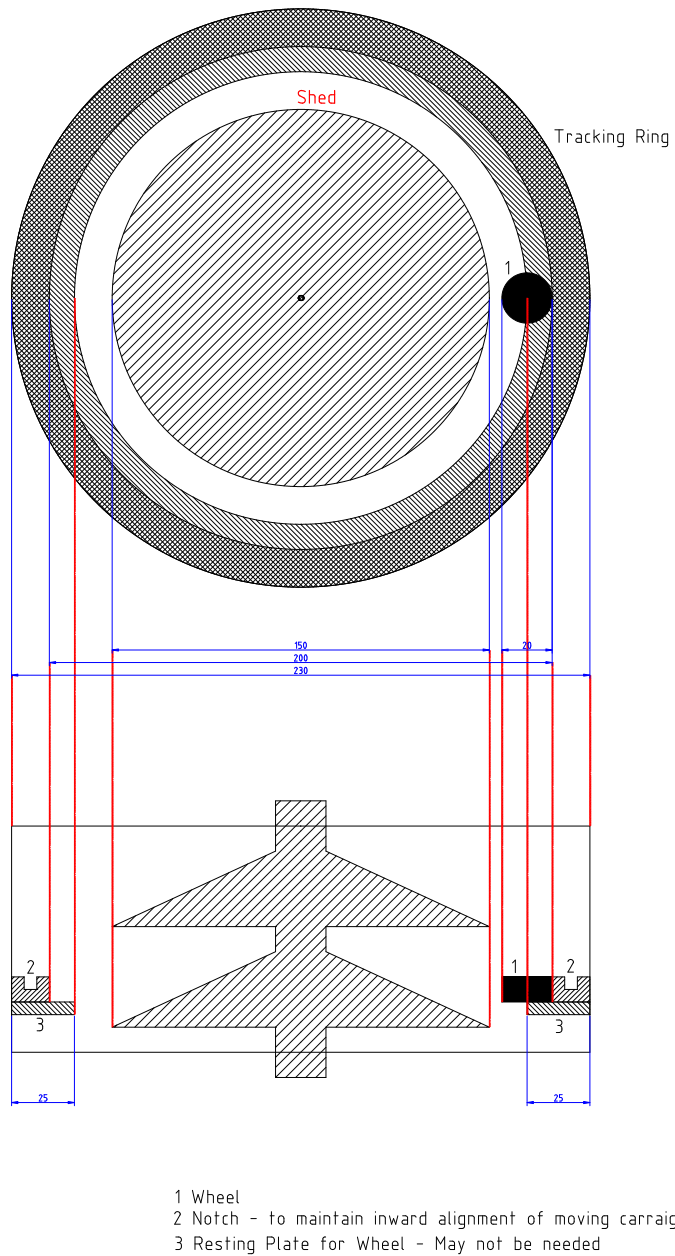
Figure D.2: Engineering drawing of proposed circular tracking mechanism

Appendix E

The Composite Cross Arm Project

Chapter 12 was a case study on a potential use of the technology developed within this research. This appendix details work undertaken to develop a monitoring system for the Composite Cross Arm (CCA) to assess its mechanical properties in real world conditions.

E.1   The Author's Involvement

When asked to be involved with the continued development of the CCA, the product had developed to a stage where a prototype has been electrically tested in the High Voltage Laboratory at The University of Manchester. The next stage was to arrange for field trials of a full scale product to be undertaken. Two such trials were scheduled with the cooperation and assistance of Scottish and Southern Energy (SSE).

The first of these trials was to take place in the vicinity of The Lecht Ski Centre in the Cairngorms, on a line which was due to be removed within a few years at the time of writing. This was to be an unenergised trial where the mechanical properties of the CCA could be verified under real-world harsh-environmental conditions. Other aspects such as ice and snow loading in comparison to a traditional steel cross arm were also to be assessed.

The second trial was initially due to take place at Peterhead on the site of a electricity substation in a coastal location. This was to be an energised trial to verify the CCA under conditions as close to real-world as possible, in addition to experiencing coastal/salt spray. Although energised, only a low current will flow, so it is still somewhat different from an active line.

Both these trials required heavy monitoring and its associated data collection and storage. Due to the future applicability of the Ph.D.'s research within deployments of CCA technology and the wider realm of monitoring, an opportunity arose to be involved in the ongoing testing of the CCA product within the above mentioned trials. The requirement was to design and develop a complete remote monitoring system for the two trials. This included the selection of sensing and monitoring hardware, the design and implementation of monitoring software and installation of the completed system. The need for custom hardware solutions in certain areas was also envisaged; for this there was a requirement to design, construct and program microcontroller based circuits.

E.2 Monitoring Requirements

For ease of implementation, National Instruments was selected as the provider of the core monitoring hardware. Specifically their Compact Real-time Input Output (cRIO) platform was chosen due to its rugged construction, the range of available input/output modules and its claimed operating facilities/abilities.

A number of modules were ordered for this product based on the initially intended monitoring objectives; however as the project proceeded requirements changed and monitoring plans were altered. Some of these changes were due to a combination of cost and time issues. For example, it was initially the project's intention to monitor/record sound throughout the trial. This was on the basis that a previous trial had discovered that composite insulators, in one particular scenario, generated audible noise during windy conditions. This trial was the only known event of this phenomena and the project team was made aware of it verbally. Therefore, exactly what was to be measured was unknown. Given that a high quality microphone and associated pre-amplifier, which was certified and was compatible with the Integrated Electronic PiezoElectric (IEPE) interface, cost approximately £2500,

it was decided to pre-plan mounting locations on the monitored towers, but neither purchase them nor develop the associated capture and processing software.

E.2.1 The Lecht Trial - Mechanical

The trial utilised a transmission line route that was in the process of being removed. The plan involved maintaining six transmission towers in close proximity to The Lecht Ski Centre until the end of the line's removal schedule. The two end towers were to act as 'tie-back' points from which the remaining span would be supported; these towers would retain their steel cross arms. The four centre towers would have composite cross arms fitted to the centre phase on the side of the tower closest to The Lecht Road (the A939). The opposite centre phase cross arm would be retained for comparative purposes. Only the centre two cross arms were to be monitored, providing both redundancy (should one system fail) and a method by which results could be verified.

As this trials' primary goal was mechanical assessment, the majority of data acquisition requirements were sought from the project's industrial partners. Their primary target was to acquire loading data from the various cross arm members in different planes and then post process the results to calculate the loading on the cross arm as a whole. As is natural with projects of this nature, these requirements changed over time.

Initially the request was to implement two full Wheatstone strain gauge bridges in each of the horizontal members to measure the horizontal and vertical loading. Load pins would then be used to measure the tension in the diagonal members, from which the compression of the horizontal members could be calculated. A load cell would be placed at the nose of the cross arm arrangement to assess the load the line is placing on the structure. Finally, two accelerometers would be placed near the nose; one on the nose connection joint to measure vibrations and one on a lower linkage to assess the inclination of the wire bundle.

It had already been decided that the load cell and the load pins were to operate using a 4-20 mA current loop. Combined with the four Wheatstone bridges, each installation would require seven current measurement channels. Based on available modules for the cRIO, current inputs (of the required type) can be added in blocks

of eight, therefore a single module was sufficient for the project's monitoring needs.

During later discussions with the project's industrial partner, it was stated that it would be prefered if the tension/compression of the horizontal members were measured with strain gauges in addition to the previously specified load pins. This meant there was a requirement for nine current measurement channels; it then followed that an additional current measurement module was required.

However due to the modular cost, it was unjustifiable to purchase an additional module for the addition of one sensor. The University of Manchester team therefore asked the project's collaborators to decide between the use of load pins or additional strain gauges. The decision was made to use additional strain gauges. As the time to the scheduled installation data approached, this would soon become the only option as the load pins were quoted to have an eight week lead-time.

E.3   Sensors

In addition to the above sensor systems, the following components were required:

- Weather Station - to correlate environmental data with mechanical performance
- Video Cameras - to record images of physical mechanical phenomena
- A Power Scavenging System - to provide power for the system throughout the trial period
- A Wireless Data System - to relay data to location for permanent storage and retrieval.

The majority of sensors were to be interfaced to National Instruments modules located within the cRIO base hardware. This hardware solution as a whole was then programmed through the company's LabView software solution. These areas are explained in further detail in the following subsections.

E.3.1   Strain Gauge Bridges

A Wheatstone bridge is normally excited by either a 5 or 10 V source. A differential voltage is then produced dependant on the strain seen by the strain gauges. If the

measurement device is close to the Wheatstone bridge arrangement it is sufficient to directly measure the differential voltage. However if the measurement system is to be many meters away, then cable resistance can start to effect the measurement result.

To ease installation duties, it was this project's intention to place the measurement equipment as close to the tower base as operationally feasible. This would then mean cable distances to the Wheatstone bridge would be in the ten's of meters. Although sense wires could be employed to ascertain the actual excitation voltage, it was decided to employ a 4-20 mA current loop system to relay the signal from the CCA to the measurement station (at the tower's base).

As the output for a Wheatstone bridge is a differential voltage, a converter was required to allow it to operate within a 4-20 mA loop system. This was located in the form of a Strain Gauge Amplifier from Applied Measurements Limited. This is a very versatile box which can amplify a strain gauge input and supply an output in various formats, including the required 4-20 mA. This unit also had the ability to set the sensitivity level and apply an offset to allow for maximum measurement scope.

At this point, it had been decided to implement three bridges per horizontal member. Through the use of four strain gauges per bridge, rather than employing 'dummy' resistors, temperature compensation could be achieved.

Horizontal and vertical bridge arrangements were tested on similar material and produced reasonable results. However the horizontal bridges could not be located in final design as initially proposed and therefore problems ensured. A compromise was reached which theoretically should yield similar results, although during testing on a sample profile it was seen that vertical movement could also be detected by the horizontal Wheatstone bridge.

The larger problem related to the recently requested compression/tension measuring Wheatstone bridge. This was requested late in the manufacturing cycle and therefore could not be pre-tested. This meant the first attempt was on a member that was to be installed in the field. As this required mechanical knowledge of the horizontal members of the CCA, The University of Manchester had to let EPL Composite Solutions lead in making the decision as to where this measurement bridge was to go.

This resulted in one failed attempt, where the compression/tension bridge also measured horizontal and vertical displacement. It was then stated that the ends of the profile would experience the least horizontal and vertical displacement and would thus be the ideal position to measure tension/compression; which was assured to be evenly distributed throughout the profile. However due to the high resolution/sensitivity requested, this solution too seemed far from ideal; additionally showing some deviation from static when horizontal and vertical loads were applied.

However, based on EPL Composite Solutions' insistence that the set-up should work as required, the situation was accepted on the basis that EPL Composite Solutions would be solely responsible for the placement and mechanical bonding of the strain gauges to the horizontal members to be deployed into the field.

The Strain Gauge Amplifiers were wired for 10 V excitation and configured to produce a 4-20 mA output signal. The bridge sensitivity was set according to the mechanical designers' needs and the offset was configured to be approximately 12 mA. This would then allow for an equal measurement swing for mechanical movement in either direction.

E.3.2   Accelerometer

As noted there was a need for two accelerometers per installation. One was to measure vibrations in the cross arm itself and the second to allow the calculation of the angle of the conductor bundle.

Although a number of commercial accelerometer systems are available pre-built, they come at a price; quotes received ranging from a few hundred pounds to over one thousand pounds. Various interface types are available, from analogue voltage, 4-20 mA, SPI, I^2C, RS232 and RS485. The most appropriate would probably have been 4-20 mA, as the signalling levels would not be affected by cable length. However a three axis accelerometer would require three signal lines and therefore a requirement for an additional six current sensing inputs on the cRIO.

While the use of SPI and I^2C protocols are possible, they are generally only suitable for short distance (usually intra-board) communication between integrated circuits. The problem arises due to their inherent use of a master driven clock

line to latch data onto and back-from the communication line. Over a long distance the clock and the data signal may become separated in time and therefore incorrect data may be read. For completeness it should be noted that the cRIO can support such protocols through the use of a digital IO module and specific LabView programming.

It was thought that the use of R485 would provide a better solution; the RS485 specification includes the use of asynchronous differential transmission, thereby negating the risk of a drifting clock and minimising the issues caused by long cable lengths. The RS485 protocol was designed for use over distances and its use has previously been proved in the field.

While pre-built units were available with many output types, the accelerometer chips available from silicon manufactures like Analog Devices were only available in SPI and/or I²C formats. Given the cost of the pre-built units, it was decided that it would be more cost effective to build custom accelerometer units. This device would utilise a Micro Electro-Mechanical System (MEMS) accelerometer, a microcontroller to convert between serial formats and a differential serial driver to create the appropriate bus signal levels for the RS485 protocol.

The accelerometer hardware also includes space for the inclusion of a 3-axis gyroscope, thereby making the unit a Six Degrees of Freedom (6DOF) Inertial Measurement Unit (IMU). Initial firmware development commenced for the use of both the accelerometer and gyroscope, however due to certain issues this was not used for the initial deployment. Issues included the speed of the RS485 bus, transmitter-receiver synchronisation and data encoding.

A native data module is available for the cRIO which supports both RS422 and RS485, which also contains a master oscillator & frequency divider system which enables all common serial communications system speeds to be accurately generated. This accuracy is not true of the microcontroller, which obtains its clock reference from the main oscillator. The main oscillator was selected to be 20 MHz to allow for a reasonable processing speed of 5 Mega Instructions Per Second (MIPS). A bus speed of 115,200 kBaud was arrived at, which has an error of approximately 0.08%. It should be noted that due to the protocol being asynchronous, the clock used to maintain timing is reset for every byte received - therefore it is conceivable that a bus speed with a high error on the microcontroller side could have been

employed.

The next issue related to Data Terminal Equipment (DTE)<->Data Communication Equipment (DCE) synchronisation. Although the protocol itself is asynchronous, meaning the bytes do not have to flow at a constant rate, the receiving system needs to know what it is receiving and its position within the transmitted frame. The simplest way to achieve this is to transmit a header then a fixed number of packets. Therefore on initialisation the receiving system waits to see the predefined header, after which it should be able to track its position in the data stream. If the system was ever to get 'lost' it would simply wait for a new header.

The final issue related to the encoding of the data: ASCII or raw binary. Initially the former was attempted, supplying both the accelerometer and gyroscope data converted and scaled in blocks of 100 readings. Minimum and maximum values for each axis were also tracked and transmitted after 10 frames (1000 readings) as formatted strings. The accelerometer was set to provide data at 1000 Hz and the gyroscope when available. It was then seen that the formated text strings could not be transmitted in the 'down-time' between the availability of new accelerometer data.

In hindsight it should have been forethought that transmitting pre-formated text strings would saturate the bus. The alternative would be to transmit the raw binary output from the sensor and allow the conversion and scaling to be done on the cRIO. This was then implemented with the reduction in accelerometer data speed to 100 Hz and the removal of the gyroscope.

After implementation, the accelerometer unit was interfaced to the LabView monitoring program, which is discussed in more detail in Section E.3.4. After appropriate binary to decimal conversion and scaling, the produced data stream appeared sensible. The MEMS accelerometer used within the unit was tri-axial and was referenced to gravity. The axis facing gravity was shown to produce an output approximately equal to 1 (or -1).

One problem was discovered however during the LabView interfacing phase; the unit seemed to be producing too much data - i.e. the data rate on the RS485 line seems to exceed the programmed 100 Hz. Although the PCB hardware was designed to allow the use of interrupts (from the accelerometer to the microcontroller), this facility was not used in current firmware. The line was polled instead,

with the data being transfered from the accelerometer when data was available, i.e. when the line being polled was at the correct logic level. Under this scenario the MEMS accelerometer is the device controlling the data rate. Hence, this observation suggests that there is a fault with the interrupt generation/clearing on the accelerometer IC. Due to time constraints, this matter was not investigated further.

Once a working unit prototype had been built, four units specifically for field deployment were fabricated. Each unit was identically constructed and programmed with the same firmware, however it was discovered that only one of the four units functioned flawlessly, another showed intermittent operation (usually after a power cycle), while the last two ceased to function at all.

This issue was initially thought to be related to failed solder joints and hence the entire board was checked and all joints retouched. This however did not solve the intermittent nature of operation. During investigation it was realised that the bulkhead connector system used for data and power provision may be at fault. This was on the basis that 'stressing' the connector would sometimes cause the device to operate. As the installation date was fast approaching, further debugging duties were passed to the School's Electronics Workshop.

Early analysis suggested that the problem may be due to the power sequencing requirements of the 'LF' J-Series PIC Microcontroller chosen to power the accelerometer. The J-Series uses a 2.5 V core, while the (low power) LF variant omits the on-board 2.5 V regulator found on the 'F' variant. This meant the a separate 2.5 V regulator had to be implemented on the PCB in addition to the 3.3 V regular used to power the microcontroller peripherals and accelerometer IC. These were left un-sequenced and allowed to 'power-up' unmonitored. Sequencing the power supplies and forcing a microcontroller reset seemed to resolve this issue. The unexpected data rate was traced to a oversight in the configuration of the microcontroller's baud rate generator.

The IMU PCB was redesigned to use the F variant containing the on-board regulator. At the same time it was decided to re-enforce the mounting points for the accelerometer and RS485 driver and include strain-relief points for the off-board connector.

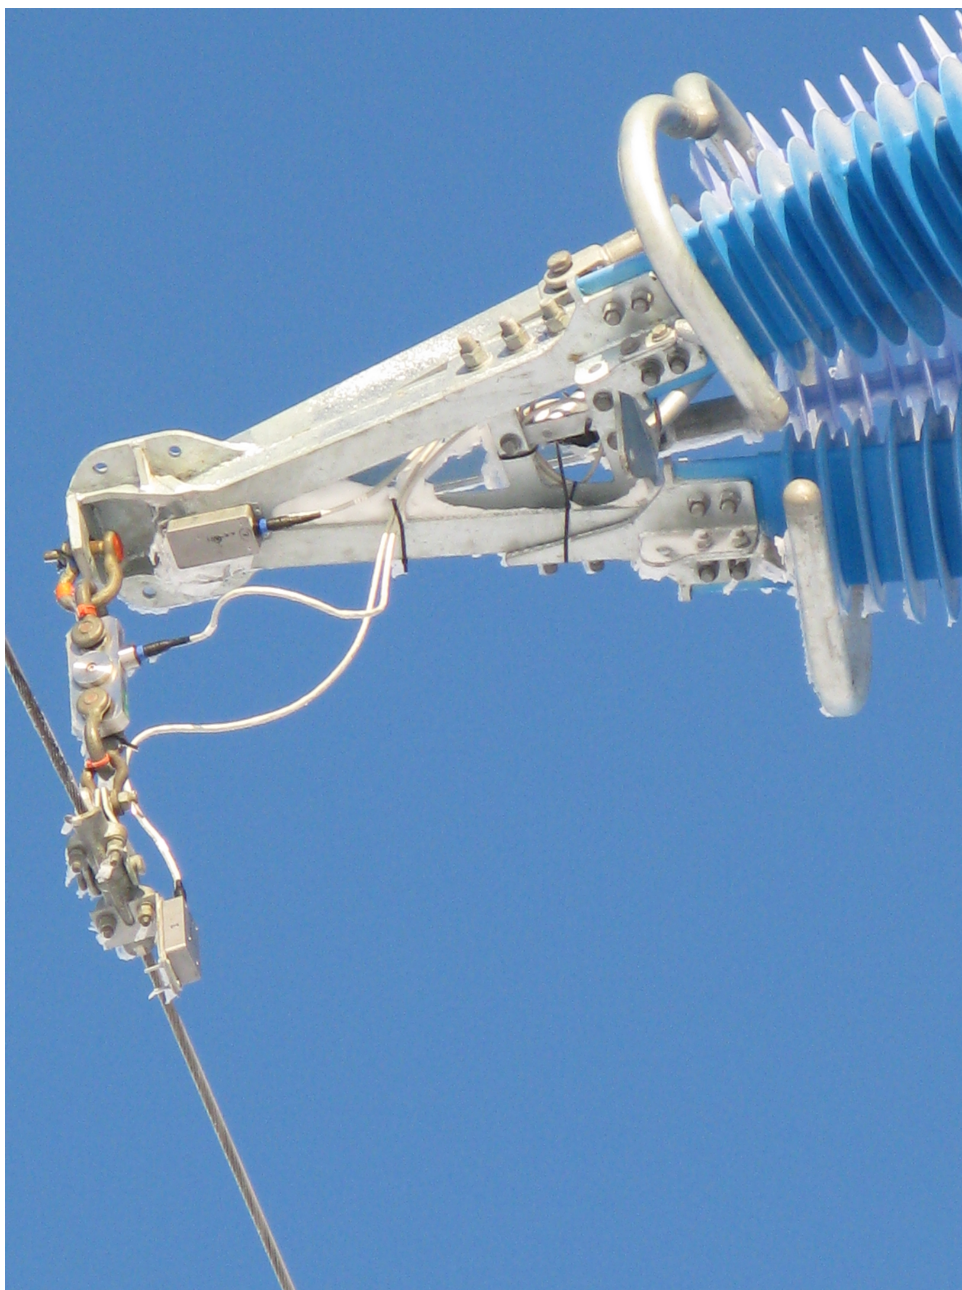An instrumented cross arm can be seen in Figure E.1[1].

Figure E.1: Instrumented Composite Cross Arm Nose

---

[1]Courtesy Arago Technology

E.3.3   Weather Station

Another requirement was to record weather data for correlation to data received from the strain gauges, load cells and accelerometers. This was to allow investigators to see if certain weather conditions produce certain mechanical phenomena on a repeating basis, or if the mechanical results seen are due to another cause. The following parameters were required

- Temperature
- Humidity
- Wind Speed
- Wind Direction
- Rain Fall
- Precipitation

It was initially suggested that these values can be obtained from discrete sensors. While this may be appropriate in some situations, there are some concerns about use of this methodology here. The typical way to achieve this would be through the use of sensors which output an analogue voltage signal proportional to the quantity being measured. These sensors have to be then combined into one sensor platform[2] for deployment onto the tower. For measured variables to be useful, they would need to be captured at near the same elevation as the CCA itself, as such the steel cross arm opposite the CCA is the ideal monitoring location.

In this discrete sensor configuration, each sensor would require a separate signal cable running to the cRIO near the base of the tower, while the cRIO would require a analogue voltage measurement module. Such measurement modules had already been purchased, so procurement was not an issue. Such analogue voltage output sensors would suffer a similar problem as non-amplified strain gauge bridges - that of signal degradation due to cable length and cable resistance.

Unmodified, these sensors would usually be intended for use in close proximity to the measurement system. With the visualised set-up, cable runs would be between 20-30 meters, over which a significant potential difference may be developed dependant on cable type. This would then potentially lead to reduced measurement resolution and accuracy.

---

[2]Sensor platform meaning a single mounting box housing all required sensors.

An alternative is to procure a pre-built weather station capable of measuring the required variables and uses a signalling protocol designed for long distance communication. As the accelerometers were designed to use the RS485 serial protocol, weather stations which used similar protocols were investigated. The selected device was the MetPakII from Gill Instruments. This is a versatile device which is capable of measuring the following parameters:

- Temperature
- Humidity
- Wind Speed
- Wind Direction
- Barometric Pressure
- Dew Point

It can communicate using RS232, RS422 and SDI; with the use of the provided converter cable the device can also be used via a USB port using serial emulation, additionally drawing power from the same port. It has a maximum update rate of 1 Hz and transmits data as clear text, thereby allowing a simple terminal application (e.g. HyperTerminal) to view the streaming data. A polled mode of operation is also available. Device numbering allows up to 26 devices to be used on one system. The use of clear text for data transmission also simplified the requirements for the LabView application, as discussed in the next section.

The one issue with this particular weather station was that it did not measure rainfall or precipitation. While they would have been a welcome addition to the dataset, they were not absolutely required for the non-energised trial.

E.3.4 LabView Software

LabView is National Instruments programming system/language for driving the majority of the company's products. National Instruments products are versatile measurement devices, which can be customised into Virtual Instruments (VIs) through the creation of LabView Applications which harness the required facilities of the product being used.

For the CCA project, the cRIO platform was chosen due to its robust field-orientated and industrial product billing. It consists of two main parts, the controller and the chassis. This chassis also contains a Field Programmable Gate Array (FPGA) which acts as the interface between the controller and the modules pluged into the chassis. There are two modes in which the cRIO can be programmed: Scan-Mode or FPGA-Mode.

In Scan-Mode the maximum sampling rate is 1000 Hz regardless of the capabilities of the installed module; while in FPGA-Mode the module dictates the maximum capture capabilities[3]. Some modules only function through the FPGA, for example the RS422/RS485 module which has precise timing requirements for the capture of serial data.

In reality regardless of the mode selected the FPGA is actually used. However in Scan-Mode the use of a set of pre-defined routines to create VI's allows for a simplified FPGA bit-file. In FPGA-Mode, there is scope to create more intricate VI's and hence the compilation of a bit-file is a more complex and time-consuming operation. The two modes can use combined, in which case LabView integrates the Scan-Mode and FPGA-Mode routines into a single bit-file.

Due to the quantity of work required to be completed in the short time to the installation date, it was not possible for the author to undertake the development of the LabView Virtual Instrument (VI) personally. As an alternative two undergraduate students were tasked to undertake this development work under the supervision of the author. For this reason this section does not specify in detail the development sequence or encountered problems and their resolutions. Rather it specifies the operation of the system implemented under the guidance of the author.

The implemented VI consists of numerous data collection operations (loops) executing in parallel. The major collection algorithm is the one sampling data from the analogue current channels, those that are measuring the loading on the CCA via the six strain gauge bridges and the load cell. The basic requirement was for the system to capture at 100 Hz and store it in an unspecified format for later processing.

It was felt that it would be beneficial to capture data in one hour 'blocks', with

---

[3]This is assuming there are no programmatic blocks to the operation in the VI.

the view to ease processing and storage requirements. Further more, if the power source were to fail and the data file corrupted, one hour of data at the most would be lost.[4] In addition, it was felt that capturing at a higher rate would be beneficial during 'periods of interest'. This was defined as 'if the average of the 1000 previous samples deviated by more than 10% of the long term average'. This long term average was updated on an hourly basis and was also recorded to file. This latter stage was to assist in recovery from power failure; in such an event the previous long term average can be loaded rather than have to be recalculated.

At the default sample rate (100 Hz), the real-time average is completely refreshed every 10 seconds; at the increased rate (1000 Hz), it is refreshed every second. The comparison of the real-time average to a long term value allows sudden changes to be detected and subsequently captured in high resolution. The sampling rate remains at 1000 Hz until the real-time average falls to within 10% of the long term average. A fixed comparison average was initially considered, however this could have potentially lead the real-time average to permanently drift beyond the 10% limit. A long term average was then implemented therefore allowed 'creep' in the CCA to be accounted for. The 10% deviation range and the increased sampling rate can be easily adjusted should the need arise.

E.3.5   Video Cameras

Another feature that was requested by the CCA design team were cameras to visually record photographs and video footage of the CCA configuration under investigation. This was to allow the numerical loading data to be related to visual images. Another important aspect of this image capture was to allow visual inspection of ice and snow loading on a CCA in real-world conditions. This was initially limited to examining the line attached to the CCA and a complete overview of the CCA itself. This was to be implemented on both monitored towers to give some redundancy.

Later, at the request of a design member whose special interest is the development and performance of insulation materials, an additional two cameras were specified for mounting onto one of the monitored towers. The purpose of these cameras were to provide a 'side-by-side' comparison of a composite and steel cross arm under

---

[4]This however does not take into account total failure of the storage medium.

simultaneous and 'near' identical conditions[5].

There are multiple ways in which video and photographic data could be captured in analogue and digital domains. Due to the environment that the cameras were to operate, cameras intended for the CCTV security industry were thought to be ideal. In this domain cameras which used either analogue composite video signals or digitally compressed video over an IP network (IP Webcam) were located.

Many manufactures produce IP Webcams, however due to the use of a Axis Communications webcam in the HV laboratory, it was decided that this manufacturer should be used for the cameras on this project. The camera in use in the HV laboratory has a resolution of 640x480 pixels; to improve useablity it was decided to locate a camera with a megapixel resolution. It also discovered that a number of cameras were available which could accommodate a Secure Digital High Capacity (SDHC) memory card for local storage. With this feature, it would be possible to locally record video and still images as a backup method, in the event that streaming capture mechanism should fail. The camera chosen was the Axis P1344.

It is not possible to directly use the camera in an outdoor environment without a protective housing. The protective housing selected was the Verso from VideoTech. This product has a heated perspex window and can operate from a 12 V d.c. source. A heater and fan for the housing is available but requires a minimum 24 V a.c. supply, which was not available on this system as will be revealed in the next section. The heater should ensure that no condensation/fogging occurs on the window, thereby reducing visibility. A sun shield is also provided which should provide additional protection from harsh environmental conditions.

The two additional cameras that were specified were to be of the same type, but in one case use a different lens to allow for better image quality at greater focus distances. The purpose of these additional cameras changed during development of the system with the final positions being as follows. The first camera was to be on tower one and focused on a glass insulator string located on the adjacent steel cross-arm; this was in order for ice accretion to be viewed and assessed. The second camera was to be located on adjacent tower directed towards the first viewing both

---

[5]'Near identical' because each cross arm will be adjacent and not occupy an identical space. Due to the main tower body, there is a possibility that each arm may experience slightly different environmental conditions.

the steel and composite cross-arms and along the line connecting the two towers.

The P1344 has the ability to start recording based on an external trigger signal; this is ideal to allow the cRIO to trigger the capture of video based on the values it samples from connected sensors. The design team also wished for period pictures/video to be captured. While the camera allows timed events to be configured, during investigation it was discovered that these events appeared to be 'one-shot' rather that periodic as was needed. A solution was to implement an external circuit (PIC Microcontroller based) to invoke a periodic trigger signal onto the camera's trigger input. As the camera only as a single trigger input, the PIC Microcontroller would also be required to accept the trigger from the cRIO and pass it onto the camera as required.

E.3.6   Power Scavenging

It was calculated that the monitoring system required (at most) 150 W to operate all the various monitoring systems. As with most transmission towers there are no low-voltage power sources in close proximity. Even if this trial was being conducted on an energised line, it would not be possible to source power from the line itself as it poses too much of a safety risk.

There were a number of options discussed by the design team. These are presented below including reasons they could not be chosen.

- Use of the line to transmit the required power - the loss on the line would be too great for the energy required.

- Wrap a cable around the line to transmit the required power - too labour intensive.

- Run a power cable at ground level from The Ski Centre - would need to be at least two feet below the surface and due to the site being a Site of Special Scientific Interest (SSSI), the surface could not be penetrated.

The possible solutions were either an in-house design or the commissioning of an independent power generation system. In order to minimise cost and development time the latter option was chosen. A local company, 12Voltz Renewable

Generation, was chosen to design and provide the required system. The initial design was to consist of both solar and wind generation. The final system only consists of wind generation.

E.4 Wireless Data System

As previously noted, the mechanical trial site is located in the Scottish Highlands. The prime objective was simplified data retrieval. Ideally team members wanted all captured data to be transfered back to The University of Manchester for onward distribution and processing. However due to the site's remote nature common communication methodologies were either not available or possible.

While the major population centres in the United Kingdom have good-to-excellent fixed and mobile broadband coverage, the same is not true here. A local business, whose premises' was hired as a operational based during installation, had only an Integrated Services Digital Network (ISDN) connection of only limited performance and marginal stability. Mobile operator presence was known[6] and an initial site survey revealed that a poor General Packet Radio Service (GPRS) connection on the Vodafone network was available.

It was clear that neither remote retrieval nor remote control/(re)configuration was going to be a reality. The team's best option was automated local retrieval to a centralised location, for later physical retrieval by a team member. A local business agreed to allow the team to house some equipment for this purpose. While the team knew that complete data retrieval was not going to be possible, a heartbeat signal and/or daily sensor averages/maximums/minimums along with battery status information would allow the system's performance and operational status to be remotely monitored. This would then simplify the planning of maintenance visits.

As previously noted, installation of the CCA was to occur on four adjacent transmission towers, each approximately 200 m apart. Only the centre two towers were to be monitored. The centralised collection point was then some 400 m away from the closest monitored tower. As the area has been declared as a SSSI, it meant that a wired communication solution could not be implemented, as the site's sur-

---

[6]Based on results from the Ofcom Site-Finder database.

face could not be penetrated for the installation of cables. A wireless solution was therefore needed.

The author assisted in the creation of a wireless network strategy based on WAPs operating in bridging mode. The bridging capabilities were tested by assessing the effect of power-cycling one end to simulate power loss and to establish if communication automatically resumes. It was initially thought that the use of high gain antennas would expedite the bridging of the 400 m expanse directly, but early assessment revealed that the link quality (hence data rates achievable) were only marginal; this was also exasperated by potential line of site issues due to the terrain.

This then meant that another bridging point was required. The availability of a disused SSE emergency shelter/hut, located between the first monitored tower and the collection point, with direct line-of-sight to both, facilitated this requirement. This location also needed its own power scavenging solution, but would also allow the implementation of a secondary backup station based on the low current BeagleBoard platform (or a variant) if required. A seventh camera was also located here to provide an overview of the installation site.

External communication of status information was initially going to be achieved via a mobile broadband modem attached to the storage computer, with data being uploaded to a database hosted on-line. The wireless network architecture was subsequently adapted to include a 3G Broadband Router. This would then allow all devices attached to the network to communicate with the 'wider-world'. This was of particular interest in the case of the network cameras, which could be configured to email and upload regular image captures to a server; another indication of wider-system operation and current environment.

E.5 Data Recover Mechanism

As previously mentioned, this field test of the CCA was to be undertaken in a remote area in the Scottish Highlands; as such communication infrastructure is poor and physical access can be problematic from late autumn to early spring. Two possible onward communication mechanisms were noted in the previous section, with neither considered suitable for data recovery back to the University of

Manchester. The previous section established that a local business was to be used as the data collection/recovery point. This section discusses the mechanisms implemented to create data redundancy & security and to keep researchers at The University informed of the state of the system.

The measurement data is initially stored locally on each of the two cRIO's with video/photo files stored on memory cards installed in the cameras. This data is then transfered to the storage location on a periodic basis via FTP. In order to secure and protect files, downloads were made to an external USB based drive employing Redundant Array of Inexpensive Disks (RAID) technology. Therefore, if one drive should fail, the data should be safely protected on the second drive. A second standard USB drive would be used to physically transfer data logs back to The University of Manchester during maintenance visits. A BeagleBoard with large local storage and running a similar FTP transfer algorithm in 'The Hut' was considered as as another backup solution. The storage of the captured data in multiple locations thereby provides multiple layers of redundancy.

While the wireless network implemented high gain panel antennas to improve link performance, it was known that there may be eventualities where the bridged links may fail, for example if the antennas are covered in snow. In the Scottish Highlands such a situation may persist for significant periods of time. With only limited storage on the cRIO and with cameras limited to 32 Gb per device, storage may quickly be exhausted.

With potentially long durations in which the wireless link cannot be re-established, an alternative recovery mechanism was required. To combat this situation each monitored tower implemented an additional WAP operating in client mode. This was to allow a team member localised access to all devices operating on the tower in question via a laptop. Although the data being recorded would be of little use to another outside the project team, WPA2-PSK with a long security key was used to protect the connection.

Once collected it was intended for the data to be brought back to The University of Manchester for distribution to design partners and analysis.

E.5.1 FTP Access Mechanism

It was necessary to retrieve data from two cRIOs and seven network cameras. As each device included an FTP server, data and images could be collected using industry standard methods. Typically FTP is used in a manual fashion with the user 'driving' a FTP client. This case FTP actions needed to be performed automatically and periodically; a scripted solution was required. A FTP solution was located and appropriately coded to connect to each device in turn and downloaded available files based on the following principles.

The cRIO does not have an automated file cleaning process, files need to be either deleted via appropriate FTP commands or by formating the storage location. Based on the limited storage available in comparison to the daily data generation, files would have to be regularly and automatically deleted. To facilitate this the script designed to firstly request a list of the previous day's data files. Files would then be downloaded, with a successful download triggering a delete instruction to be issued against that file. The total number of files available for the day in question was recorded as part of the FTP log file.

A 32 GB[7] flash memory card was installed into each of the network cameras for local storage purposes. The cameras had the ability to automatically manage their available storage capacity by automatically deleting old files when a particular threshold was met. Based on this on-board management capability, a simple one-way synchronise function was implemented for camera downloads. This mechanism downloads all files that are present on camera that are not present on the downloading machine - files are not checked in the reverse direction.

The scripting language allowed external applications to be executed; one such application was 'blat!', a command line based email client recommended by the creators of the FTP Client. This mechanism was used to send a status email regarding the FTP process to the team. Windows Scheduler is used to run the script on a daily basis.

There is however a potential problem with cRIO file transfers; in particular if a file transfer fails it will not be scheduled for transfer again the next day. In order to resolve this, a Java program was created to scan the FTP transfer log and extract

---

[7]This is the biggest card available under the SDHC specifications.

the number of files expected for a particular day. This is then compared to the number of file located on local hard drive. If there are more files expected than available on the local disk, a second FTP script was to be executed to retrieve the missing files.

A second Java program was created to parse each data file and create hourly averages for each sensor. These averages were then to be sent to the team via email, again using 'Blat!'. Windows Scheduler was again to be used to execute the utility on a daily basis.

E.5.2   Device Data Storage Implementation

The application running on the cRIO logs data to four files: one for the analogue samples and one each for the serial devices (two accelerometers and one weather station). As previously noted, new files are created every hour to aid the management of generated data. Testing indicated that about ~110 MB over four files is generated per hour, equating to ~2.64 GB/day. Internal cRIO storage is specified as 4 GB, but after formating and loading the firmware/operating system & user application this is restricted to little over 2.5 GB. This then suggested that internal storage will not be sufficient.

The cRIO allows storage to be expanded though the use of USB based flash drives and external hard drives; hence it was decided that an external 32 GB should be used. It was subsequently revealed that only certain drives are recognised by the cRIO and no definitive list of functional devices exists. A drive that was found to operate with the cRIO was subsequently found to offer poor simultaneous read/write performance. Under this scenario the drive's FTP transfer rates could not be assessed as the FTP software repeatedly timed-out attempting to retrieve the directory listing.

When the read operation was assessed independently the FTP transfer rate was little more than 100 kb/s. It should also be noted that this is through a wired (Ethernet) connection as opposed a wireless link. While this is theoretically sufficient to transfer the expected 2.64 GB of data within 8 hours, the monitoring process would need to be halted. This was not possible with the implemented solution neither would it be practical; the system would effectively loose 1/3 of

its monitoring capability. The conclusion was that this particular drive model implemented a poor quality USB memory controller.

A Solid State Drive (SSD), noted for their fast read/write speeds, was examined via a SATA to USB converter. In this case speed will be limited by the USB interface at 480 Mb/s against SATAI at 1.5 Gb/s. This solution not only improved transfer rates, it allowed the FTP process to occur with the cRIO was actively logging data.

One caveat however was the random failures of the drive, whereby the cRIO would loose it's reference to the drive. It was initially thought that this was linked to either simultaneous FTP download & upload activity or poor FTP protocol implementation, however these theories were eliminated through lab investigations using multiple FTP utilities.

It was then discovered that the drive appeared to be powering down at random times. At this point it suggested that a powered USB hub be installed in-line between the cRIO and the SSD in an attempt to power the later directly; due to time constraints this advise was not heeded. During a later visit by a National Instruments Engineer, monitoring the current draw of the cRIO and SSD combination during power-up revealed that the cRIO was hitting some internal limit. This suggests that the cRIO does not fully implement the USB specification.

Adding an independently powered USB hub in-line resolved the power-down and current-limit problems. However, due to the hub's crude power system implementation, the SSD is not power-cycled when the cRIO is rebooted. This therefore potentially leaves the SSD in a random state with possible unfinished write operations or open file references pending. This then causes the cRIO to fail to recognise the drive on reboot. The simplistic and implemented solution was to power-cycle both the cRIO and USB hub together if ever the former is required.

E.6   Continuing Work

The Composite Cross Arm project is an ongoing research and development activity which will stretch beyond the completion of this Ph.D. In order to draw focus back to this work's primary research goal, approximately six months was committed to

the development needs of the CCA.

In this time the outline structure of the cRIO LabView application has been designed and its creation and testing in code has been supervised. Accelerometer units have been designed, programmed and tested and a basic wire-routing plan with regards to tower installation was put in place.

A wireless bridging system combined with a file transfer mechanism for the localised collection of captured sensor data was additionally developed. Wider communication strategies were investigated via a 3G Router to provide a remote indication of system status.

As with any project, initial work will be refined and superseded as a project continues and the CCA is no different. The first iteration of the cRIO application is being re-written to improve and rectify mistakes and omissions. Cabling and connector options are being upgraded to better cope with the harsh conditions in the Scottish Highlands. Regardless of the future changes, a basis has been provided from which to expand and improve upon as required.

This page is intentionally blank.

Appendix F

Publications

Veerappan, C.A.; Green, P.R.; Rowland, S.M.; ,'High Data Rate Communication in High Voltage Environments,' *Student Poster Competition, 2010 IEEE PES Transmission and Distribution Conference and Exposition*, Apr. 2010

Veerappan, C.A.; Green, P.R.; Rowland, S.M.; , 'Visual live-line condition monitoring of composite insulators,' *Applied Robotics for the Power Industry (CARPI), 2010 1st International Conference* on , vol., no., pp.1-6, 5-7 Oct. 2010, doi: 10.1109/CARPI.2010.5624439, URL: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5624439&isnumber=5624409

Zachariades, C.; Rowland, S.M.; Cotton, I.; Green, P.R.; Veerappan, C.A.; Chambers, D.; 'A Trial Installation of High Voltage Composite Cross-Arms,' *High Voltage Engineering, 2011 17th International Symposium* on, 22-26 Aug. 2011