Monitoring and analysis of data captured from industrial equipment for diagnostic purposes

LICKES, Andrew Daniel

Monitoring and Analysis of Data Captured from Industrial Equipment for Diagnostic Purposes

Andrew Daniel Lickes

A thesis submitted in partial fulfilment of the requirements of
Sheffield Hallam University
for the degree of Master of Philosophy

September 2017

Collaborating Organisation: Baldwin & Francis Ltd

Abstract

Preventative maintenance is a key consideration for equipment manufacturers. In industries such as mining, oil & gas and rail, the site operators prioritise minimising disruption caused by downtime which leads to significant losses in terms of revenue and damage to reputation. Ensuring equipment is healthy is therefore a major concern for all stakeholders to ensure continued reliable operation of a site.

Current methods of delivering such diagnostics typically form part of a SCADA suite which takes a site-wide approach to displaying information. While correct use of these suites has shown to bring benefits to both customer and supplier, it can be argued that they are not suitable for customers with small scale implementations, especially given their expensive nature and difficulty to configure. SCADA systems are also prone to being ineffective when troubleshooting low-level equipment related problems, especially in scenarios where automated equipment is deployed in remote sites for extended periods. In such cases, service engineers report scenarios where alarm logs were too full to effectively analyse and draw meaningful conclusions from.

The aim of this project is to design and develop a solution that delivers enhanced product support through diagnostics and preventative maintenance. A graphical interface will be developed for engineers that displays technical information about equipment, aiding the troubleshooting process and reducing time spent on-site. A report-based graphical interface will cater for non-technical customers, providing an intuitive display of how healthy the system is. The report will also provide advice to aid with decision making when equipment is not being used optimally or is approaching the end of its life.

Equipment manufacturers can take advantage of the solution to improve relationships with existing customers, appeal to new customers and generate a new revenue stream through enhanced support programmes. The solution is designed to recognise that continued site operation is of paramount importance for a customer. Production environments impose stringent safety restrictions which restricts invasive access to equipment when it is operating. A non-invasive wireless access mechanism is therefore a key requirement to ensure that downtime is minimised.

Acknowledgements

I would like to thank Sheffield Hallam University supervisors Alan Holloway and John Holding for their continued input and guidance throughout this project. Their assistance with supplying hardware enabled the collection of detailed data which was used to inform the decision making process throughout the project. Being able to draw from their knowledge enabled me to gain new insight into methods and techniques that I would have not have originally considered.

I would also like to thank Baldwin and Francis for the opportunity to work in a challenging and exciting environment. In particular would like to give thanks to Gareth Andrews who supervised the project. As a team member his input, support and management throughout the project ensured goals were consistently met. Also from the company I give thanks to Paul Marshall, whose technical expertise contributed to innovative methods, particularly related to hardware interfacing and low-level development.

Finally I would like to give thanks to all parties involved in the 'Knowledge Transfer Partnership' (KTP) programme, particularly Trevor Gregory from Innovate UK. I am proud to have taken part in this successful KTP and have found the opportunity to serve as a link between academia and industry to be truly unique. I cannot emphasise enough how effective KTP programmes are for opening new opportunities and bringing benefits to all stakeholders.

TABLE OF CONTENTS

I Glossary

General Terms

ATEX – European safety standards

Customer – End user of industrial equipment

IECEx – Global safety standards

IS – Intrinsically Safe

Technical Terms

API – Application Programming Interface

ASCII – American Standard Code for Information Interchange

CAN – Controller Area Network

CMSIS – Cortex Microcontroller Software Interface Standard (Standardised drivers and libraries for an ARM based microcontroller)

CRC – Cyclic Redundancy Check

DFSS – Distributed Feeder Soft Start (Bespoke equipment for the rail industry)

DHCP – Dynamic Host Configuration Protocol

EMI – Electromagnetic Interference

FLP – Flameproof

FTP – File Transfer Protocol

GPIO – General Purpose Input Output

GUI – Graphical User Interface

IDE – Integrated Development Environment

IP – Internet Protocol

ISR – Interrupt Service Routine

LRC – Longitudinal Redundancy Check

MCU – Microcontroller Unit

NTP – Network Time Protocol

RDP – Remote Desktop Protocol

RF – Radio Frequency

RTC – Real Time Clock

RTU – Remote Terminal Unit

RTOS – Real Time Operating System

RTS – Request To Send

RX – Receive

SCADA – Supervisory Control and Data Acquisition

SQL – Scripted Query Language

TCP – Transmission Control Protocol

TTL – Transistor-Transistor Logic

TX – Transmit

USART – Universal Synchronous/Asynchronous Receiver/Transmitter

VLAN – Virtual Local Area Network

VPN – Virtual Private Network

VSD – Variable Speed Drive (Bespoke equipment for the mining industry)

WAP – Wireless Access Point

XPS – Protection System (Bespoke equipment for the mining industry)

II List of Figures

III List of Tables

1 Introduction

This project will begin by investigating the core concepts of diagnostics and the purposes they serve. Existing solutions that deliver diagnostic functionality will be investigated to inform the direction this project will take.

An industrial scenario of a coal mine will be used as an example production environment, where a customer is defined as the user of the equipment. Equipment used throughout a mine will be examined to determine where diagnostics can be applied to improve site efficiency.

Implementation will begin by developing a robust method to capture data from equipment that uses the Modbus RTU protocol to communicate. This protocol is used extensively by equipment in a range of industrial sectors including oil & gas, rail and mining. Intensive data processing will be carried out to analyse this data and store it long-term in an intelligently structured database. Data will then be manipulated for presentation on a platform-independent graphical user interface. The interface will display relevant information for both technical engineers and non-technical customers.

An Agile development methodology will be utilised to manage the project, ensuring that stakeholders have input throughout the project and that any changes are accommodated.

2 Literature Review

2.1 Background

Automation, protection and continued reliability of equipment used in industrial environments are growing considerations for organisations that operate within these sectors. Equipment manufacturers aim to provide customers with robust and reliable products, often providing warranties to incentivise and give customers the confidence that their product will meet their demands.

Manufacturers therefore need to ensure their equipment is used within their defined parameters to prevent damage and/or failure. Such scenarios would reflect negatively on the reputation of the manufacturer and incur costs where work is required as described under the equipment warranty. Goulao, Paulo, Salvado, Martins and Granjeia (2011) discusses that in cases of dispute, where equipment is thought to have been misused, it is often difficult to prove what caused the failure, meaning the manufacturer is responsible for carrying out costly repairs.

As documented in the study by Filev, Chinnam, Tseng and Baruah (2010), industrial equipment does not traditionally provide the capabilities for built-in self-diagnostics. Manufacturers therefore realise the potential benefits to be gained from a dedicated mechanism capable of obtaining diagnostic information from their equipment. Evidence that demonstrates equipment misuse can be used to avoid loss incurred from work being carried out for free under warranty. The diagnostic information can also be used by service engineers and customers alike to gain a better understanding of how healthy the equipment is and whether any action is required to ensure continued operation.

2.2 Equipment Failure

Any type of complex equipment, be it industrial or consumer, is prone to failure and ultimately has a limited lifespan. The consequences of failure are far reaching, especially in industrial scenarios. Downtime results in significant loss to the customer, both monetary and reputational. A study by Vegunta and Milanovic (2011) states that monetary costs include factors such as the loss of raw material that could have been obtained during normal operation, increased labour costs to repair the fault and lost opportunities. Eti, Ogaji, and Probert (2006) analyse downtime that directly relates to equipment failure, stating that this can amount to between 2-16% of annual revenue. Reputational damage arises when the time lost due to failure leads to missing an order deadline.

Failure in a safety critical application can have disastrous consequences, resulting in loss of life. An early example of this is the Flixborough incident in June 1974 which led to 28 deaths and the destruction of the plant (Hendershot, 2009). Given such consequences, the need to reduce the likelihood of such failures has been a natural concern since the dawn of engineering.

Initial methods of risk reduction relied upon reactive techniques such as learning from mistakes, but as industries evolved so did their dependence on hazardous materials and thus increased the risk of tragedy in the event of failure (Smith, 2017). Failure mode analysis techniques were thus developed as a mechanism to reduce the likelihood of failure. Such techniques do however incur costs related to increased design and development time. Methods of mitigating failure, such as having standby systems, also increase costs due to the need for secondary hardware. Given that a project budget is limited, there is a need to find a balance between the costs of failure reduction with the value of enhanced performance.

2.3 Preventative Maintenance

Preventative techniques, such as condition based maintenance (CBM) and time based maintenance (TBM), aim to take a proactive approach to reducing disruption. TBM programmes establish set timeframes when maintenance is performed and does not take into account the actual health of equipment. In contrast, CBM is performed only when the health of the equipment deteriorates and therefore negates the need to perform unnecessary maintenance when the equipment is healthy, thus reducing costs. CBM is therefore more effective than a traditional reactive approach (Jardine, Lin, & Banjevic, 2006). Effective implementation of CBM relies on a number of factors, chiefly that the data collecting approach is implemented in a way that data captured is accurate and contains appropriate information to draw conclusions from. Eti et al. (2006) discuss that when good preventative maintenance programmes are observed a customer may see up to a 30% increase in profits related to the increased efficiency, decrease in repair costs and reduction in downtime that results in lost sales revenue.

Prognostics and diagnostics are fundamental to a CBM programme. Prognostics seek to determine the likelihood of a fault occurring using methods such as trend analysis. A study by Engel et al. (2000) argues that this type of analysis does not provide enough information to arrive at a decision for the correct course of action. The study gives evidence to prove that the prognostics process only becomes accurate as equipment health deteriorates, owed to the fact that problems become more pronounced during this stage and are therefore easier to identify.

Diagnostics aim to identify when equipment is not performing as expected and offer troubleshooting advice that relates to the natures of the problem. Failures can also be the result of multiple causes, making up a 'reaction chain'. Diagnostics and troubleshooting techniques work through this chain to identify the most probable root cause (Bloch & Geitner, 1999). Common causes include poor design, poor installation/assembly, defects in materials used, wear and tear and equipment misuse.

In the context of harsh environments, a study by Sheng, Li, Qin, Guo and Zhang (2011) documents that the failure of a mechanical component accounts for 60% of breakdowns. This is related to the continued varied stresses the equipment is exposed to throughout its operating life. The study goes on to discuss that diagnostic information from these mechanical faults can be obtained by analysing data from normal operation compared with data captured when stationary. Pattern recognition methods can subsequently be employed to identify abnormal data to predict future faults.

2.4 Diagnostics

A modern example of effective diagnostics can be seen in the automotive industry where, by law, all vehicles have been fitted with an onboard diagnostics (OBD) system since 1996 in the US and 2001 in Europe (Táutiva, Roewer, & Furlan, 2013). As documented in BS ISO 15031-4 (2014) the system is capable of providing live readings from a car such as engine RPM, coolant temperature, speed and fuel status. Alongside this, the system is used to illuminate malfunction indicators on the dashboard, as shown in Figure 1.

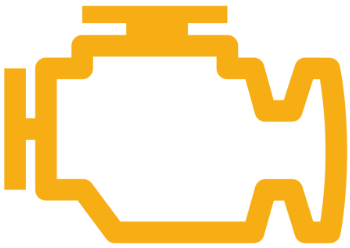

Figure 1 – Check engine light found in modern vehicles

During a service, a mechanic can access the OBD to perform diagnostic tests to assess whether further work is necessary to ensure the continued reliability of the car. This type of maintenance not only helps the mechanics work more efficiently but also benefits the customer by reducing the risk of vehicle breakdown in future.

The 'Enhanced vehicle onboard diagnostic system and method' (2016) patent discusses that despite modern advancements there are still a number of weak areas that could be improved upon to deliver enhanced functionality, suggesting that more proactive monitoring and storage of historical data could be used to gain a better insight into deterioration and overall health of individual components in the vehicle. This data could drive an improved preventative maintenance programme to help the user avoid disruption when using the vehicle. Another area of improvement is identification of usage patterns to gain an insight into vehicle performance and further understanding of the environmental impact being made. A fundamental constraint of such functionality is that it must be easy to use, given that the majority of users are unfamiliar with the technical concepts in their vehicle.

Many parallels from the OBD system can be drawn with diagnostics solutions employed in industrial sectors such as oil & gas, mining and rail. This study focuses on delivering diagnostics for mining equipment as an example use case, but the same rationale applies to equipment found in oil & gas (e.g. ring main units) and rail (e.g. principle supply points).

2.5 Industrial Equipment

The mining industry relies on heavy duty machinery to carry out different roles within the mine. These machines often require large amounts of power, and protection of this power is vital for ensuring safety. A shearer (shown in Figure 2), for example, cuts the coal from the coalface using cutting drums, requiring 3300V and consuming around 231.88A of current.

Protection systems fulfill their role by monitoring the power that is supplied while these machines operate. The unit is configured to cut the power in failure scenarios such as earth leakage, over current and over temperature. This prevents disastrous consequences, such as an explosion, which could ultimately lead to loss of life. Data provided by these systems include the current and voltage, both ideal candidates for historical logging for providing diagnostic information.

Figure 2 – A shearer used for cutting coal from a coalface

Once the coal is cut it is transported via a conveyor belt (Figure 3) that is attached to a motor driven by a VSD (Variable Speed Drive), requiring between 1100-3330V to operate. These VSDs may have built in intelligence, providing data to a master control system about their running state, such as stack temperature and phase current. These are another example of data that can be monitored to provide diagnostic information to the user.

Figure 3 – Conveyor belt for transporting coal out of the mine

2.6 Intrinsic Safety

The concept of Intrinsic Safety (IS) applies in hazardous environments including coal mines, where the surrounding atmosphere may contain a mixture of flammable gasses. A spark generated by electrical equipment could ignite this gas and cause an explosion. IS became a requirement following two incidents in 1912 where explosions were linked to sparks caused by low power open-wire signalling circuits (A. L. Bartels, 1983).

In Europe, the ATEX directives are part of EU law and define the requirements for equipment used in hazardous environments (Council Directive, 1999). Equipment within a mine would fall under the ATEX device group of either M1 or M2, with a Zone of 0 for environments where "explosive gas/air mixture is continuously present for long periods" or Zone 1 where "an explosive gas/air mixture is likely to occur in normal operation."

IECEx is a similar but voluntary scheme aimed worldwide. Pomme and Sijrier (2010) discuss that due to the global reach, the IECEx does not form part of a legal system in the same way ATEX does within the EU. The paper does however discuss that countries can make use of the worldwide 'United Nations Economic Commission' framework, which endorses the use of the IECEx system, as part of their own national legislation.

2.7 Equipment Communication Protocols

Equipment can use a variety of protocols for communication, such as CAN bus, Modbus or Profibus. These different protocols define the type of network topology supported, how nodes are addressed on the network, how data is transmitted and interpreted, security mechanisms and the types of functionality that are supported.

CAN bus for example is commonly used in automotive applications, such as the previously discussed OBD for vehicles. CAN uses a message arbitration mechanism that ensures messages of the highest priority will always be serviced first, and the limited message size prevents a single node from hogging the bus (Chin-Long Wey, Chung-Hsien Hsu, Kun-Chun Chang, & Ping-Chang Jui, Dec 1, 2013).

Modbus, developed by Modicon, is a much older protocol, dating back to 1979 (Huitsing, Chandia, Papa, & Shenoi, 2008). One of the main advantages of Modbus is its simplicity to both implement and use. It is because of this simplicity and proven track record that, as discussed by Hendrix (2001), Modbus is seen as a 'de-facto' protocol for industrial equipment. As such it is supported by a wide range of equipment types, from protection systems to motor management units. The flexibility of the protocol is demonstrated in studies by Dao-gang Peng, Hao Zhang, Li Yang and Hui Li (2008) and Sen Xu, Haipeng Pan, Jia Ren and Jie Su (2008) that utilise the protocol as part of complex embedded and Linux-based systems

The protocol is master/slave orientated, where a master device (e.g. a centralised control system) initiates all communication. The master sends a poll containing a request of a defined function type to a specified slave, the slave then responds with the data requested.

Different varieties of Modbus exist, the most widely used are: Modbus ASCII, Modbus RTU and Modbus TCP/IP. Both ASCII and RTU use serial communication to transmit messages. In ASCII the messages are transmitted using ASCII characters, whereas RTU simply uses bytes. Both types implement error checking, ASCII using a longitudinal redundancy check (LRC) whereas RTU uses a cyclic redundancy check (CRC). ASCII encoding results in more bytes required to form a message than compared to RTU, meaning RTU is more efficient and thus can achieve better throughput speeds. An example of this for a read holding register request for 3 registers starting from register 40108 is given on Modbus ASCII vs Modbus RTU (2017) as follows:

ASCII message: 3A 3131 3033 3030 3642 3030 3033 3745 0D 0A (17 bytes)

RTU message: 11 03 00 6B 00 03 76 87 (8 bytes)

Another difference between these serial varieties is that ASCII defines delimiting characters as ':' for the start of a message and the 'CR LF' characters for the end of a message. Since RTU uses a series of bytes it is not possible to rely on a unique byte sequence to delimit messages, since a valid message may contain this byte sequence within the payload. RTU therefore relies on a silent time of at least 3.5 character times to delimit messages (as shown in Figure 4)

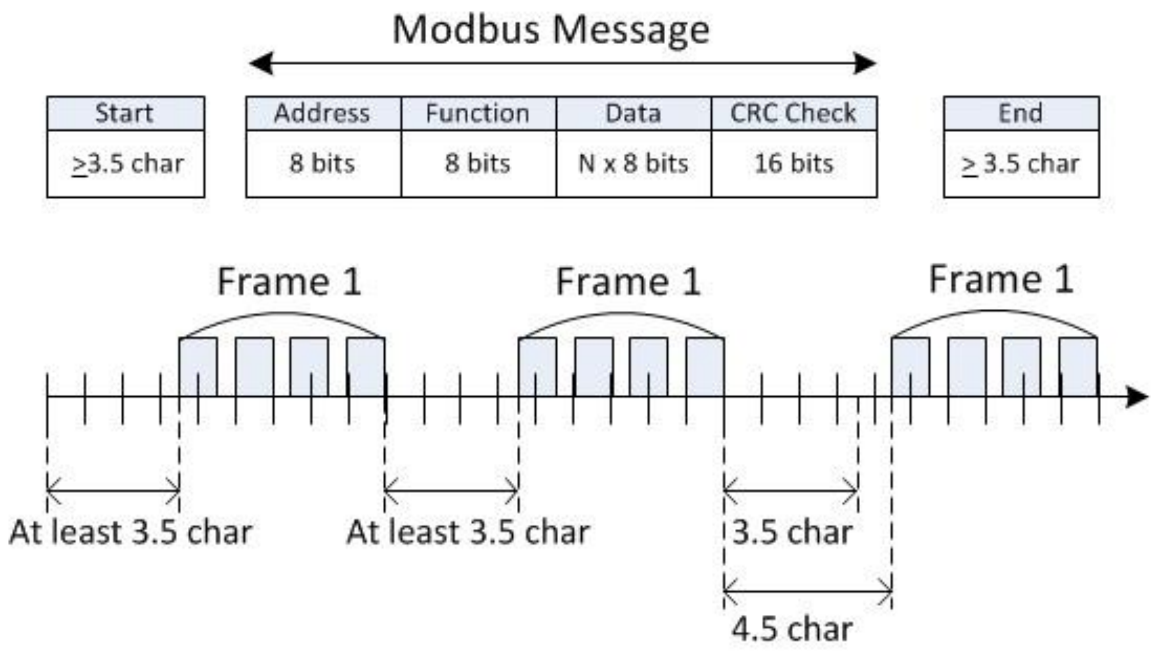

Figure 4 – Silent time used to delimit messages in the Modbus RTU protocol

The TCP/IP variety serves as a wrapper for RTU messages to enable transmissions over Ethernet. The major advantage of this is that the connection-orientated nature of TCP guarantees delivery of all messages, retransmitting in the event of an error. Depending on the implementation it can also allow for more convenient use, given the ubiquitous availability of Ethernet compatible equipment. Gonzalez and Papa (2007) discuss how this ubiquitous availability also brings about new risks related to security, attributed to the possibility of interconnecting these systems to corporate IT networks. The study goes on to demonstrate how such a network can be monitored to capture the Modbus data, achieved using a network scanner, transaction validator and network mapper. The outcome of this study discusses that this functionality could be used to form the foundation of a diagnostics solution that would be similar to the goal of this project. It does however also prove that appropriate security mechanisms need to be in place when a system is attached to a corporate network to prevent unauthorised traffic snooping.

2.8 Existing Solutions

A large section of commercial diagnostic solutions form part of SCADA (supervisory control and data acquisition) suites. These are fully featured suites that take a ‘system-wide’ approach to monitoring, providing functionality such as displaying asset mimics, site connectivity mimics, alarm management, trend viewing and data archiving. These suites are often expensive, for example the cost of the 'Ignition Pro' solution amounts to US$14,995 (Ignition, 2016). Babovic and Velagic (2009) discuss that additional equipment costs for these solutions may also be expensive, giving example costs of US$28,000 for a three branch system. Such high costs would be unjustifiable for smaller manufacturers, who may instead opt to use alternatives such as bespoke solutions that are tailored to their own hardware.

Off the shelf SCADA solutions take a universal approach to asset management. The main benefit of this is that the suite can be applied across a broad range of industries, offering vast customisability options. The downside of this however is that this level of flexibility, as discussed by Stefanczak (2013), results in integration and maintenance complexity that is undesirable for a customer, as such the integration process is often carried out by integration specialists. This would also lead to difficulties related to software support, since each industry and application is unique and it is unlikely that a developer would have the level of understanding required to address an individual customer related issue.

Existing standalone protocol monitoring applications, such as 'Serial Port Monitor' as shown in Figure 5, are aimed at serving engineering needs as opposed to customer needs. As such the user interface is presented from a technical standpoint as opposed to being easy for a novice user to operate.

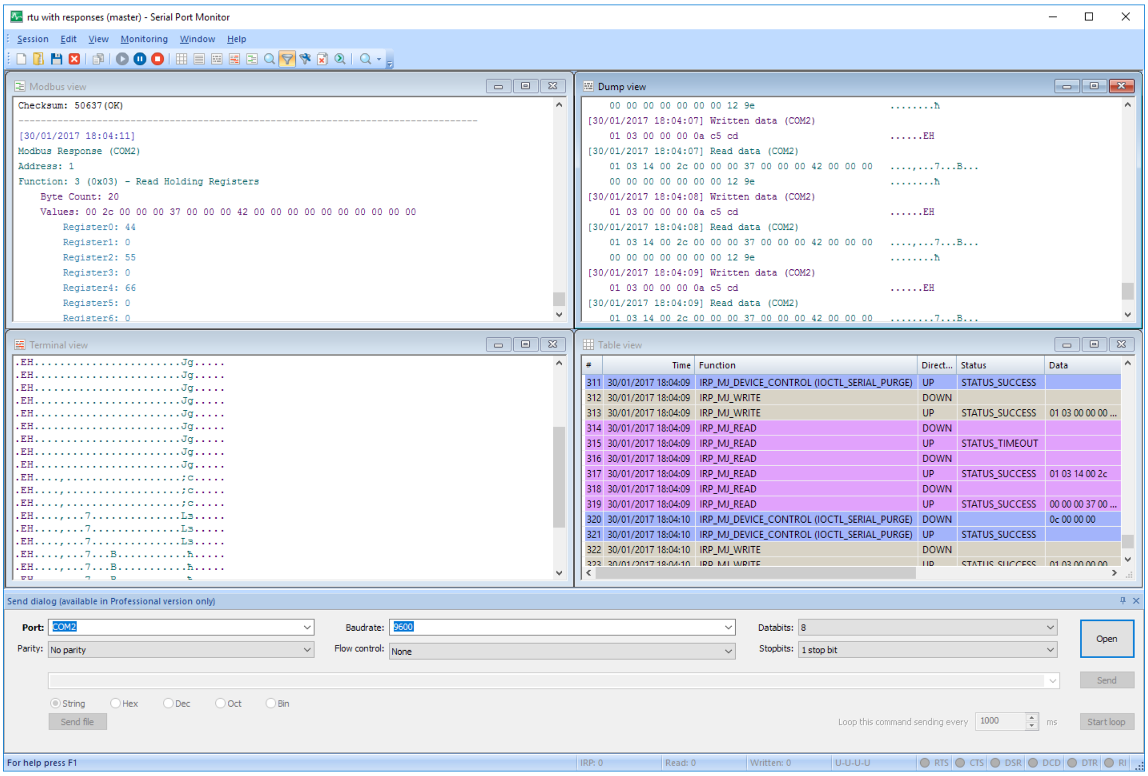

Figure 5 – Serial port monitor application

These applications also do not provide an intuitive mapping method from register assignments to real-world parameters, making troubleshooting difficult. They also do not provide functionality for historical logging into a database and do not enable a 'high level' view of mapped equipment in a graphical user interface.

USB hardware, such as the RS-422/485 ComProbe from Teledyne LeCroy, enables interfacing with an RS485 bus and can potentially be used to monitor Modbus RTU. The drawback with this is the cost - as of 2017 the hardware alone costs US$1,395. The software used with this hardware is aimed at experienced technical users and revolves around displaying protocol specific data. It is therefore not suitable for use by customers that require intuitive diagnostic information about their equipment. Furthermore this software is purchased separately and is similarly expensive, adding another US$600 to the total cost.

2.9 Remote Monitoring

Remote monitoring adds further opportunity for manufacturers to reduce costs, where off-site access to data can be analysed by technical specialists to identify and address faults within a system (Lampert, 2015). Security of such systems would need to be addressed from the start of the project, ensuring that proposed hardware and transmission mediums are developed and hardened in a way that ensures they cannot become compromised, such as in the high profile ‘Stuxnet’ incident.

The Stuxnet worm was developed with the intention of targeting industrial controllers from a specific manufacturer. Unlike traditional worms which serve to compromise information, the Stuxnet worm was able to cause physical damage to equipment. This was achieved in the form of a ‘man-in-the-middle’ type attack where communication between physical I/O and application code was intercepted and manipulated in such a way to cause physical damage (Langner, 2011). While Stuxnet did not exploit remote access directly, it demonstrates the potential for what could occur in the event such a system was to become compromised.

3 Methodology

3.1 Development Phases

This project will be split into four distinct phases. Phase one will involve researching existing solutions and will examine the types of equipment they monitor, the types of data provided and how diagnostics information is presented. This research will help determine the most efficient route to achieve the functionality required. This phase will also examine the breadth of industrial equipment types, enabling familiarisation of the industries that use them and gaining an understanding of the different environments the system may be deployed into.

The second phase will use the information gathered from the research to evaluate and justifiably select appropriate equipment that is capable of delivering the diagnostics solution. This will include market research of existing hardware, liaising with suppliers to arrange evaluation of selected products and testing of these units in production-like environments.

The third phase covers the design, development and testing of the software that the system comprises of.

The final phase deals with non-invasive delivery of the content. Market research and evaluation testing will be performed to select appropriate Wi-Fi equipment that is capable of reliably providing this functionality.

Management of the project will take advantage of the Agile methodology, utilising the scrum framework to separate tasks into manageable development ‘sprints’ that have a duration of up to two weeks. Each sprint will end with a stakeholder meeting to discuss progress and accommodate any changes to requirements that may arise.

3.2 Scope

The project scope defines key requirements that the prototype must satisfy. Backward compatibility with existing products and support for newly developed products are important to ensure that both new and existing customers will be catered for by the solution.

The need to ensure that operation of customer sites is not disrupted by the solution is to be tackled by using a non-invasive wireless access mechanism. In environments where intrinsically safe equipment is a requirement this must involve the use of approved equipment only.

The diagnostics and troubleshooting content will be delivered using an intuitive graphical interface aimed at both service engineer and customer users. A two-step secure access model will also be in place to restrict access to content depending on the type of user viewing the content.

3.3 Solution Role

The end goal is to deliver enhancements for all stakeholders: equipment manufacturers, service engineers and customers. Manufacturers can expect to gain an increase in ROI with the generation of a new revenue stream by offering enhanced product support to customers. Different levels of service could be made available enabling flexibility to suit a particular customer needs, for example

- Basic level: First year health check and access to basic “System Overview” of the diagnostics system.

- Middle level: Yearly health check for X years and access to detailed system operating parameters.

- Full support: Regular health checks determined by the reliability of the system, access to detailed system operating parameters and ability to view operating trends of selected parameters.

The manufacturer can also use the system to ensure their equipment is being used within the defined parameters. In scenarios where equipment has failed due to customer misuse the data provided will serve as evidence that a warranty would be invalidated and the customer is responsible for replacing or repairing the equipment.

3.4 Engineer Use Case

Service engineers will see an increase in productivity when using the system. Their workload while diagnosing problems on-site will be reduced by providing an easy to use, engineer focussed graphical interface to aid with fault finding. Equipped with a tablet PC, the engineer can attend a customer site and interrogate the equipment to see the live running status and assess performance.

More detailed analytics will enable the engineer to ‘drill down’ on an asset by asset basis to gain more detailed technical information. Historical view of usage patterns will enable the possibility of identifying potential causes for concern. Crucially, this can all be achieved without having to disconnect the power, minimising downtime which is of paramount importance to the customer.

In the event of equipment failure, the service engineer can use the system to determine whether equipment has been used as intended over the course of its operating life. This data serves as an audit trail that can be used as evidence to prove if equipment has been misused.

The solution will also offer the potential for remote assistance via the Internet, provided that customer infrastructure is capable of providing robust security (such as a firewall and VPN). This will enable engineers to help support customers in a much more responsive way than is currently possible, given that customers are often based in remote locations worldwide.

3.5 Customer Use Case

A customer will be presented with a less technical interface than what an engineer would see, taking the form of a system health report. This intuitive interface will display asset health using a globally recognised traffic light system where green means healthy, amber means warning and red means action required.

Information will be delivered to help the customer maximise efficiency by ensuring equipment is used within the agreed operating parameters. The system will also provide notifications to help the customer prepare for scenarios where equipment is approaching the end its life.

The customer could access the system via their own Wi-Fi enabled tablet or via a PC at the surface which has network connectivity to the equipment. The customer will have the option to print out the information to form a report which can be used when discussing future decisions.

3.6 Production Environment

This project will use the mining industry as an example production environment. A coal mine has unique challenges due to the harsh nature of the environment, as such the solution will be:

- Subject to shock
- Subject to exposure to dirt, dust and hazardous gasses
- Subject to loss of power in failure scenarios

In countries such as India, temperature and humidity will also be a key consideration. Equipment must be capable of operating in ambient temperatures up to 60°C and humidity ranges of 50% to 90%, non-condensing with a maximum wet-bulb of 28°C.

Inside the mine there are areas where methane gas may be present. As such, equipment used in these zones must either be flameproof (FLP), meaning they able to withstand an explosive event and cool any vented gases to below the flashpoint, or intrinsically safe (IS), where equipment is designed such that is incapable of generating an incendive spark. ATEX and IECEx safety standards govern the stringent requirements that equipment must conform to in order to receive certification that it can be used in hazardous zones.

This prototype will be deployed into a flameproof container, such as the one as shown in Figure 6. Such a unit may have an attached IS chamber (coloured blue on the left) where only IS approved equipment can be used. The equipment enables three methods of implementation, either:

- Deploy all hardware into the FLP container
- Deploy all hardware into the IS chamber
- Deploy some hardware into the FLP container and some into the IS chamber.

Figure 6 – Flameproof container used in the mining industry

If all equipment is to be placed inside the IS chamber then it is a requirement that only IS approved hardware is selected. This may prove difficult for niche components whose limited use means there is no existing approved equipment on the market. In this case it would be necessary to undertake the approval process as part of the project. This is a complex, lengthy and expensive process that falls outside the scope of this project. This involved nature of approving equipment also means that products in the market that are certified have an inflated price tag. A zone 1 approved tablet PC, the ‘Rough Pro’ from supplier ‘Astron’ for example, has a price of £685. In comparison a non-approved unit with the same specification costs under £100.

The advantage of deploying in the FLP means that cheaper, non-approved equipment can be used. One consideration is that it would still be preferable to select industrial grade equipment as opposed to consumer equipment due to the harshness of the environment.

The main issue with deploying inside the FLP is that a Wi-Fi signal would be attenuated by the thick steel walls. This could result in poor wireless performance that would affect the core access mechanism of the solution. In contrast the IS chamber has much thinner walls and would be much better suited to housing an antenna. Deploying the antenna inside the chamber also enables the option for external mounting via access glands on the side of the chamber. For this reason the option of taking advantage of both IS and non-IS equipment will be explored.

4 Hardware Selection

In the production environment explored, the equipment uses the Modbus RTU protocol. This brings up distinct challenges, such as interfacing with the transmission media and message delimitation, which are not present on a Modbus TCP implementation. Therefore it is important that a hardware platform that is capable of reliably monitoring and capturing this data is selected.

4.1 Communication Monitoring Hardware

Monitoring of low-level Modbus RTU communication can be achieved using a variety of platforms. In order to determine the most suitable for this task, each platform was evaluated to determine the strengths and weaknesses to formulate a justified selection.

4.1.1 Field Programmable Gate Array (FPGA)

An FPGA consists of programmable logic blocks which gives the advantage of making them extremely customisable and thus can be programmed to suit any application. A study by Tian & Benkrid (2010) shows that with effective programming, that takes advantage of parallelism and limits conditional branching, an FPGA vastly outperforms a high end general purpose CPU. The study also discusses how energy efficient an FPGA is compared to a PC, in this case the FPGA performed 336x more energy efficient.

An Altera Cyclone IV FPGA, shown in Figure 7, was evaluated using the ‘Quartus II 15.0 Web Edition’ development environment. This free evaluation software imposes functionality limitations such as single thread compilation and restricted access to advanced utilities. Nevertheless, it was possible to gain an understanding of the scale of work that would be involved to achieve an implementation that would meet the project needs. The ‘blank slate’ nature requires a substantial amount of work to implement even basic functionality, such as flashing an LED. The evaluation board does however provide multiple timers that offer a high degree of accuracy which would be ideal for message delimitation of Modbus RTU messages.

Figure 7 – Altera Cyclone IV FPGA

The main drawback is the complexity involved in implementation. The amount of time spent in development would not be enough to justify the gains in speed and customisability. It can also be argued that the speed of the FPGA is better served in process intensive applications, such as real-time video processing, meaning it is overkill for this project.

Furthermore, given that the end solution would be responsible for multiple complex tasks it would be necessary to make use of an RTOS to manage task scheduling. Given the complexity involved in implementing even basic tasks it makes it unlikely that the project would have enough time/man power to implement the functionality for this project.

4.1.2 Micro Controller (MCU)

An MCBSTM32F400 development board, shown in Figure 8, was selected for evaluation of the microcontroller platform. The board is equipped with the STM32F407IG microcontroller and has additional components, including external memory for increased storage capacity. The board also provides interfaces for a number of peripherals, such as the USART and SD card, which are required for this project.

Figure 8 – MCBSTM32F400 development board

Initial proof-of-concept development was carried out using the evaluation version of the Keil 'uVision' IDE. This was suitable for testing basic functionality but was limited to a maximum of 32k of code, meaning that complex functionality could not be tested.

To overcome this, a freeware alternative called 'CooCox' IDE along with the Arm GCC tool chain was used, enabling unrestricted access to the board. Development involved the use of standard ST peripheral drivers along with the 'freemodbus' stack, ‘freeRTOS’ and other freeware middleware components for implementing core functionality such as networking and file system control. Interoperability of these components presented a number of compatibility issues that required unanticipated time and effort to rectify. This dependency on freeware components resulted in a lack of confidence around robustness of the code. These components are community developed open-source libraries and give no guarantee of operation and no comeback in scenarios where problems occur. This is not acceptable in an industrial environment. Given the observed issues it is unlikely that the independently developed components would work together in a single application without spending an unknowable amount of time troubleshooting code written by others.

If an MCU is to be selected it is therefore required to use trusted drivers and middleware such as the unrestricted version of Keil's development kit. This kit enables access to standardised CMSIS drivers, access to middleware libraries and fully featured debugging utilities. This license will require an initial investment but would give royalty free access to libraries that are guaranteed to work. Support would also be available should this be required.

The main benefit of a microcontroller is their capabilities for very accurate timing, meaning that real-time monitoring of communications is achievable. Using an RTOS means that tasks can be prioritised to ensure that critical operations, such as monitoring, will always occur. This guarantee of message capture is fundamental to the project.

The downside is that implementation of end-user functionality will be more difficult than compared to a PC based system. Functionality such as data interpretation and presentation is more suited to developing with a high level language, such as C# or Java. Attempting to do so in low-level C would be unwise and increase the amount of development time required. Having the microcontroller responsible for this would also add additional strain to the device which has limited resources.

4.1.3 Embedded PC

The B&R APC 2100 shown in Figure 9 was selected for evaluation of a PC based system.

Figure 9 – B&R APC 2100

Using a PC provides access to an operating system such as Windows, which in-turn enables access to high level programming languages and simplified integration of services. Data storage services, such as SQL, allow structured data storage and enable the processes required to produce the troubleshooting information for display to the end user. Other services, such as NTP, can be implemented with ease on the PC.

Third-party applications, such as 'mdbus', prove that monitoring Modbus communication from Windows is possible. Figure 10 shows the application monitoring under controlled conditions using simulation software to act as both master and slave.

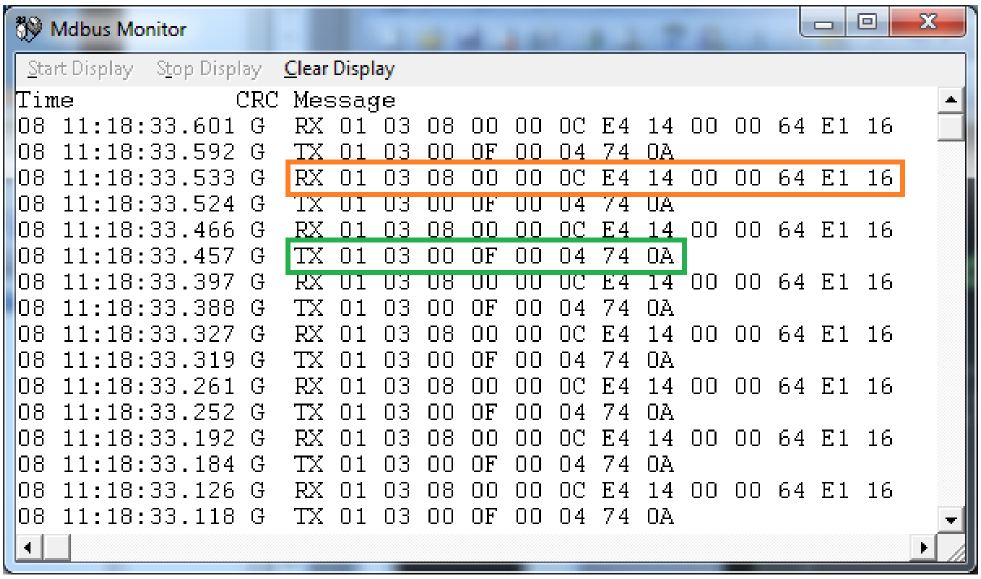

Figure 10 – Mdbus running in monitor mode with correctly parsed data

In practice, however, the reliability has been shown to vary between different hardware. Figure 11 shows how, with a shorter delay between polls, the application incorrectly interprets a poll from a master and a response from a slave as a single message. Since Modbus relies on a silent timer to delimit messages (as discussed in Section 2.7) there is no reliable way to programmatically determine the end of one message and the start of the next.

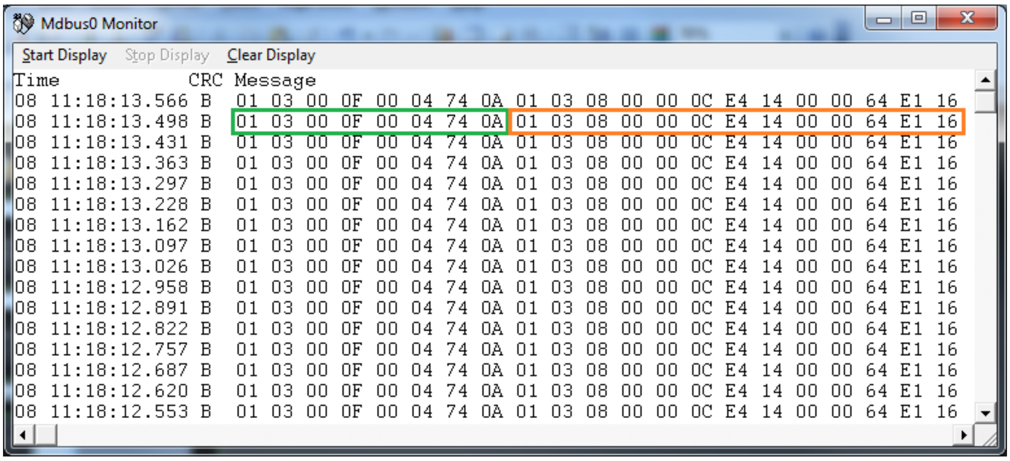

Figure 11 – Mdbus running in monitor mode with incorrectly parsed data

Using an embedded PC for data monitoring therefore leaves the potential for messages to be missed or misinterpreted, this would undermine the validity of information which is vital for the project.

This problem was exacerbated when connecting the PC to real world equipment. Figure 12 shows the extent to which TX/RX messages were being interpreted as a single message, rendering the data unusable.

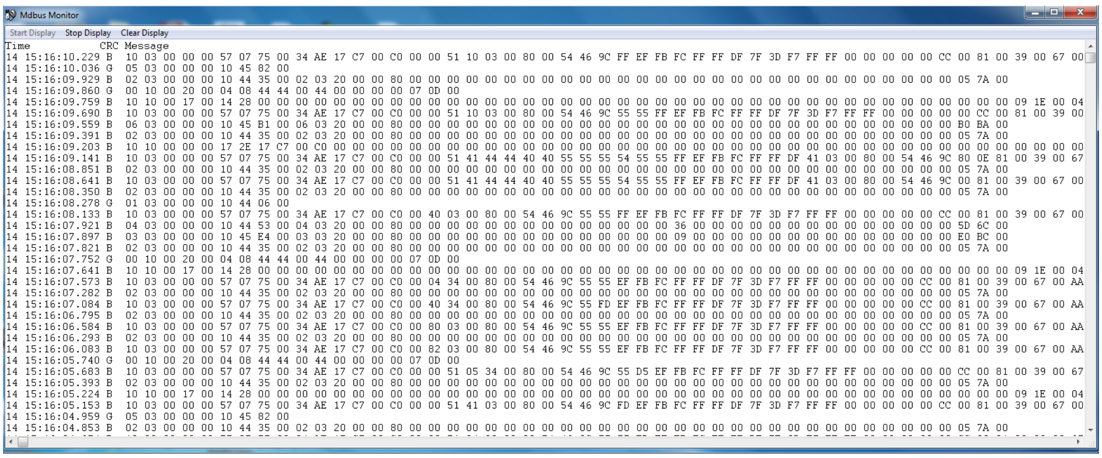

Figure 12 – Mdbus monitoring industrial equipment with incorrectly parsed data

Another downside of using a PC is in scenarios where a catastrophic event occurs. In these situations it is vital that data logging occurs as quickly as possible once power is restored as this data could provide valuable insight into a start-up problem. An embedded PC would need time to boot the operating system and launch system applications before user-applications run. On Windows this can take around 30seconds even with solid-state disks as stated by Jonathan Parkyn (2015). In contrast, a dedicated embedded platform could begin logging within milliseconds of being powered on. Furthermore, a PC generates much more heat than an embedded platform. Observations from the APC showed that when idle the unit ran at around 36°C. In contrast the MCU generates such little heat that it does not require a heat sink.

The major benefit that a PC would bring is the capability to enable heavy data processing. This core project requirement is suited to being developed on a high-level PC based solution. This processed data can be subsequently presented in a way that allows the remote end user device to act merely as a thin client, meaning no additional software is required. This brings the advantage that there is no risk of mismatching software versions between the client device and the embedded PC.

4.1.4 Summary of Findings

A breakdown of pros and cons for each platform is shown in Table 1.

Table 1 – Pros and cons of hardware for monitoring data communication

Platform	Pros	Cons
Field Programmable Gate Array (FPGA)	• Extensive low-level flexibility • Very fast operation • CAN/Ethernet stacks available to ease development • Generates very little heat	• A steep learning curve means increased development time • Speed of operation is overkill for this project • Fewer online support resources
Micro Controller (MCU)	• Small form factor • Ideal for developing low-level C code to interface with other hardware • Enables the option of using a real-time operating system for time critical operations • CAN/Ethernet/Modbus libraries available to ease development • Instant operation on power up • Generates very little heat	• Difficult to implement high-level applications • Limited processing capabilities • Application libraries and hardware drivers require expensive licences
Embedded PC	• Ideal for developing high-level code (C#/Java) suitable for end user applications • Enables the use of a variety of intensive end user services (e.g. SQL) • CAN/Ethernet/Modbus libraries available to ease development	• Larger form factor • Less robust • Runs a standard operating system (unsuitable for real-time applications) • Generates more heat than dedicated solutions • Longer boot times

4.1.5 Selection Justification

The outcome from the platform evaluation was the decision to make use of both an embedded PC and an ARM-based microcontroller. Since both platforms have distinct strengths that could be applied in this project, the logical decision is to take the best of both worlds and develop a method of transferring data between the two.

A fundamental aspect of the project involves monitoring Modbus communication - this has been shown to produce inaccurate results when using a non-deterministic, PC-based system. A microcontroller, with guaranteed real-time processing, is suited to the task of data acquisition and will give confidence that the data can be trusted.

Another vital aspect of the project involves heavy data processing and the potential for simplified deployment of third-party services such as SQL. These kinds of services are unsuited for deployment on a microcontroller, and attempting to do so would add months to development time. There is also the risk that the extra load from such services could compromise the vital data acquisition process. An embedded PC is therefore the obvious choice to meet this need.

4.1.6 Keil Licence Justification

Satisfying the goal of the project requires the use of peripherals on the microcontroller. It is vital that the low-level drivers for these are robust and fit for purpose.

Keil are a trusted organisation owned by ARM, the industry standard manufacturer of processors for microcontrollers. The software they provide is built with these devices in mind, and as such they will be committed to ensuring they provide robust and fully functional standardised libraries. Taking advantage of these libraries will result in a reduction of development time required. In the unlikely event of a problem with this software, they can be contacted directly to fix the underlying issue.

Since the library is standardised, it also gives the advantage that the same code can run on a different microcontroller should the need arise in future versions of the product. Furthermore, the same licence can be used for multiple products without incurring any additional costs.

An alternative to this would be to develop bespoke drivers and middleware. This would however take a considerable amount of time and could be seen as ‘reinventing the wheel’. With the Keil stacks, teams of specialised developers work together to create the fully functional, robust stacks. It is unrealistic to assume that a single developer will be capable of creating stacks that are of the same professional standards that will be suitable for use in an end product. Furthermore, the scope of this project does not include the creation/implementation of such stacks.

Community developed libraries are another option, but these present significant risks and negative factors that make them unsuitable for use in a professional end product. One issue is that interoperability issues exist when integrating the multiple different libraries, drivers and stacks that have been developed by different communities. This would lead to increased development time to rectify problems. The end product could also be vulnerable to bugs and flaws in code that was written by a community who do not provide support/fixes. It would therefore be the responsibility of an in-house developer to fix such bugs, this could be considerably difficult if the code is poorly documented. Crucially, there is absolutely no guarantee of continued reliable operation and no comeback in scenarios where problems are encountered.

Implemented low-level code using community libraries would also be device specific. This limits flexibility and means future developments would incur increased development time should a change of hardware occur.

4.2 Data Processing Hardware

With the communication monitoring hardware selected it was apparent that a more capable PC-based platform will still be required to handle the data processing and content delivery. Market research (detailed in Appendix I) was performed to identify potential candidates for use in the prototype, these units will then be evaluated by carrying out tests that aim to assess whether they will be suitable for use in the production environment. Three units from three different manufactures were selected with the aim of covering all suitable product types within the embedded PC marketplace.

The production environment is an enclosed flame-proof container located underground where ambient temperatures will be more extreme than those observed above ground. Consumer products would be unsuitable as they rely on fan-based cooling methods which would lead to an increased risk of a spark caused by metal-to-metal contact. Furthermore, the conditions within a mine would lead to build up of dirt on the fan blades which, overtime, would lead to reduced cooling performance and eventual failure of either the fan or the component that relies on it.

Industrial grade ‘fan-less’ devices are designed to be capable of operation within more extreme ambient temperature ranges and also make use of specially designed heat sinks that prevent the build-up of dust and dirt.

4.2.1 Minimum Requirements

In order for the PC to fulfil the requirements the unit must have the following specification as a minimum:

- Intel Atom 1.5 GHz
- Dual Core
- 4GB RAM
- 128GB Storage
- Fan-less
- Dual Gigabit Ethernet

4.2.2 B&R APC2100 Industrial PC

The B&R APC2100 that was evaluated for communication monitoring was tested further to assess suitability for data processing applications. The unit cost £1,314 as of July 2015.

4.2.2.1 Specification

- Intel Atom E3845 1.91 GHz
- Quad Core
- 4GB SD RAM
- 128GB CFast storage
- Fan-less
- Dual Gigabit Ethernet
- POWERLINK
- CAN
- RS232
- 1x USB 2.0
- 1x USB 3.0

4.2.2.2 Discussion

The B&R APC2100 is a powerful unit that includes all the interfaces required for this project. Other B&R products, such as the 'Power Panel', have been used reliably in the production environment for many years. This gives added confidence that the APC will be of similar robust quality and will meet the requirement for continuous reliability. The addition of a CAN interface may prove useful in future developments if CAN is to be supported. The major drawback is the high cost of the unit. Table 2 shows a breakdown of the pros and cons.

Table 2 – Pros and cons of the B&R APC2100

Pros	Cons
• Fully finished product • All required components packaged together in single unit • Previous positive experience with B&R gives confidence that the unit will work as intended • Support for all interfaces required by this project • CAN interface • DIN rail mountable	• Expensive

4.2.3 Amplicon Impact-D

The next unit selected was the Amplicon Impact-D shown in Figure 13. This unit cost £690 as of July 2015.

Figure 13 – Amplicon Impact-D

4.2.3.1 Specification

- Intel Atom E3845 1.91 GHz
- Quad Core
- 2GB DDR3L
- 32GB Industrial SSD storage
- Fan-less
- Dual Gigabit Ethernet
- 4x USB3.0
- 4x DB9 COM ports (COM1/3 & 4 supports RS232 - COM2 supports RS232/422/485)

4.2.3.2 Discussion

The Impact-D uses the same CPU as the B&R APC and is half the cost. The addition of SSD capability will also result in faster boot times than the CFAST card utilised by the APC. This is an important consideration as long boot times would lead to delayed access to the diagnostics system. In scenarios where power failure occurs, a service engineer would rely on the diagnostics system for accessing information to determine what caused the failure. It is therefore vital that system allows for timely access to reduce the amount of downtime.

First impressions of the aesthetics of the unit suggest that it is in the early development stages. This raises some doubt as to whether the unit will be suitable for use in a production environment. Table 3 shows a breakdown of the pros and cons for this unit.

Table 3 – Pros and cons of the Amplicon Impact-D

Pros	Cons
• Cheaper alternative to APC 2100 • All required components packaged together in single unit • RS485 interface • Support for all interfaces required by this project • DIN rail mountable • Generous number of interfaces provides more flexibility with connectivity	• 'Less finished' feel of product suggests it is an early development model.

4.2.4 Adlink PC104

The final unit selected was the Adlink PC104 shown in Figure 14. As of July 2015 this unit cost around £667, depending on the additional components selected.

Figure 14 – Adlink PC104

4.2.4.1 Specification

- Intel Atom E3845 1.91 GHz
- Quad Core
- 4GB DDR3L
- Storage not included
- Fan-less
- Dual Gigabit Ethernet
- 1x USB 3.0
- 2x USB 2.0
- 2x SATA 3Gb/s
- 1x mSATA
- 8x GPIO

4.2.4.2 Discussion

The Adlink PC104 is a single board PC that utilises the same CPU as both the Impact-D and APC2100. A major benefit is that the unit can be certified to operate in extreme temperature ranges up to 85°C, which is an important consideration for this project.

The PC104's modular design enables greater flexibility when determining which additional components to use on a project-by-project basis. An example of this is in scenarios where an existing PSU located in the flameproof enclosure can be used to supply power, meaning a dedicated PSU would not be required, resulting in cost savings. Pros and cons of the PC104 are shown in Table 4. A full breakdown of costs is detailed as follows:

- Individual Components
  - Board only: £299
  - 4GB DDR3L RAM: £34
  - 128GB mSATA SSD: £97
  - Cableset: £56
  - Extended temperature testing: £46
  - PSU: £135
- Complete System (Board and components): ~£667

Table 4 – Pros and cons of the Adlink PC104

Pros	Cons
• 'Extreme rugged' variant ensures operation in temperatures between -40°C to +85°C • Small form factor • Support for all interfaces required by this project • Cost savings in scenarios where additional components are not required	• Not supplied as a complete system - additional components purchased separately (RAM, SSD, Power Supply) • Board is not encased and would therefore require additional design time

4.2.5 Suitability Testing

Contact was made with suppliers for the units identified. A one-month evaluation of each unit was then arranged to allow for suitability testing. Table 5 shows the scope of these tests.

Table 5 – Scope of suitability testing for the PC platform

In Scope	Out of Scope

• Temperature testing
  ○ Idle
  ○ Under max load
  ○ Under max load in an enclosed space with no ventilation
• Raised ambient temperature (PC104 Only)
  ○ Under max load
	• Humidity testing
	• Data throughput testing

4.2.6 Temperature Testing

Load testing was performed using the 'IntelBurnTest' application configured to the maximum stress level. Temperature readings were taken using the 'Speedfan' application which displays readings from sensors located on the main board and CPU core. Intel documentation states that the Intel Atom is capable of operating with a maximum core temperature of 99°C. Temperatures were closely monitored to ensure this was not exceeded.

All tests were performed continuously over a period of 8 hours. Units were placed inside a sealed box for the enclosed space testing with no ventilation.

An important consideration is that both B&R APC2100 and Amplicon Impact D are fully encased units that include a power supply. The Adlink PC104, however, is not and was tested using a standard consumer ATX power supply that uses fan-based cooling. This 300W supply is designed for use with consumer PCs and is capable of delivering 12V, 5V and 3V - as such it generates a lot of heat on its own. The PC104 can however be powered from 5V only, in production environments this can be provided by existing power supplies located inside the FLP. It was therefore decided not to place the ATX supply in the box for the enclosed space test.

4.2.7 Results

Table 6 shows the results observed when testing using the B&R APC, Table 7 shows results for the Impact-D and Table 8 shows the PC104 results. A comparison of results is graphed in Figure 15.

Table 6 – Results for B&R temperature testing

Test	**Maximum Sensor 1 Temperature**	**Maximum Sensor 2 Temperature**
Idle, open ventilation	34°C	36°C
Max stress, open ventilation	47°C	49°C
Max stress, enclosed space, no ventilation	63°C	63°C

Table 7 – Results for Amplicon Impact-D temperature testing

Test	**Maximum Sensor 1 Temperature**	**Maximum Sensor 2 Temperature**
Idle, open ventilation	55°C	56°C
Max stress, open ventilation	67°C	71°C
Max stress, enclosed space, no ventilation	93°C	98°C

Table 8 – Results for Adlink PC104 temperature testing

Test	**Maximum Sensor 1 Temperature**	**Maximum Sensor 2 Temperature**
Idle, open ventilation	41°C	44°C
Max stress, open ventilation	48°C	55°C
Max stress, enclosed space, no ventilation	55°C	64°C

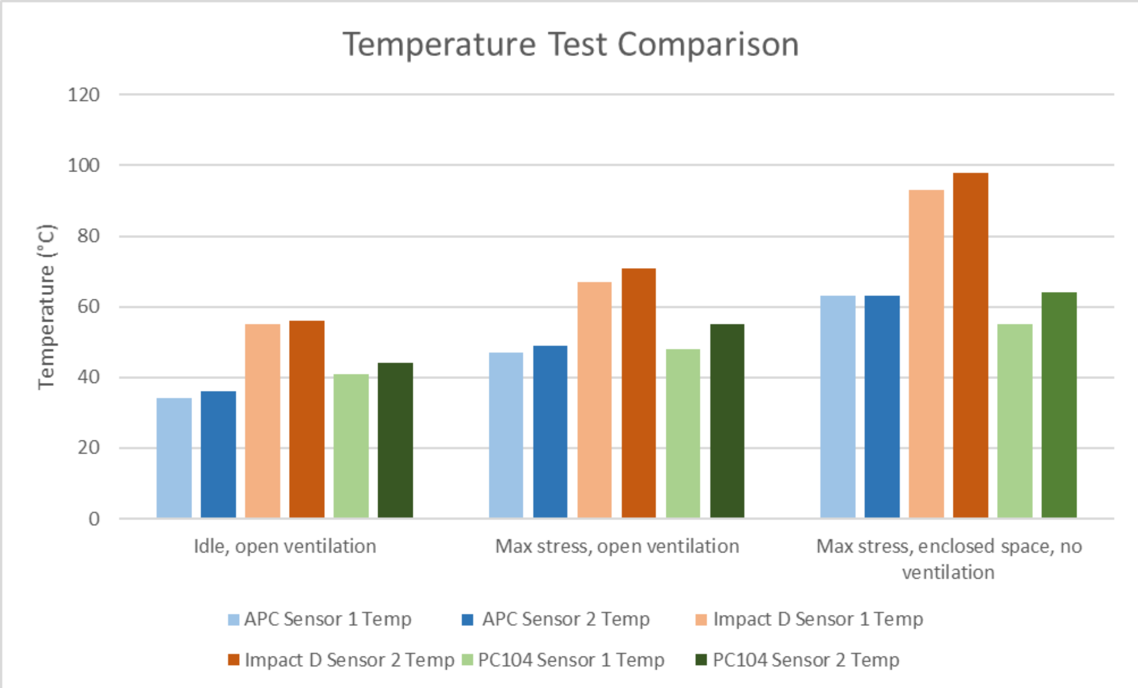

Figure 15 – Comparison of results from temperature testing

4.2.7.1 Discussion

Temperature testing has shown that the B&R APC2100 and Adlink PC104 show similar readings when idle and under full load. The Impact-D however showed significantly higher readings, to such extremes that the enclosed space test had to be abandoned early due to the high levels being observed. These findings were discussed with Amplicon who were initially surprised, but later confirmed they were able to recreate the observations and claimed to have fixed the problem. This reaffirms original first impressions that this unit is in the early prototype stages and not fit for a production environment.

It is important to note that, as stated in Section 4.2.6, the PC104 board was tested without a power supply located in the enclosed space. The heat produced by the consumer grade 300W ATX supply would give an unfair disadvantage to the PC104 because it is not intended to be used with a small PC104 board.

4.2.8 Raised Ambient Temperature Testing

Further testing was performed on the PC104 board to give added confidence that the unit is capable of operating in environments where the stated operating temperature was being reached. This test was not performed on the Amplicon Impact-D since it had already failed the enclosed space test. As for the B&R APC, there is enough confidence that this unit will meet project needs from positive experience with other B&R equipment as previously discussed.

The evaluation PC104 board provided by the supplier did not have the 'Extreme Rugged' certification that guaranteed operation in temperatures up to 85°C, instead the maximum was listed as 60°C. For this reason, it was decided to test with an ambient temperature of 55°C. To reach this temperature the board was placed inside an oven. Network connectivity was achieved by feeding an Ethernet cable through an access gland, meaning that temperatures could be monitored using remote desktop. The unit was tested under maximum stress using IntelBurnTest for one hour. Table 9 and Figure 15 show the results.

Table 9 – Results from raised ambient temperature testing with the Adlink PC104

Time	**Oven Temperature**	**Sensor 1 Temperature**	**Sensor 2 Temperature**
14:40	55°C	56°C	64°C
14:50	58°C	61°C	64°C
15:00	65°C	71°C	74°C
15:10	65°C	81°C	84°C
15:20	58°C	83°C	84°C
15:30	55°C	83°C	84°C
15:40	55°C	84°C	82°C

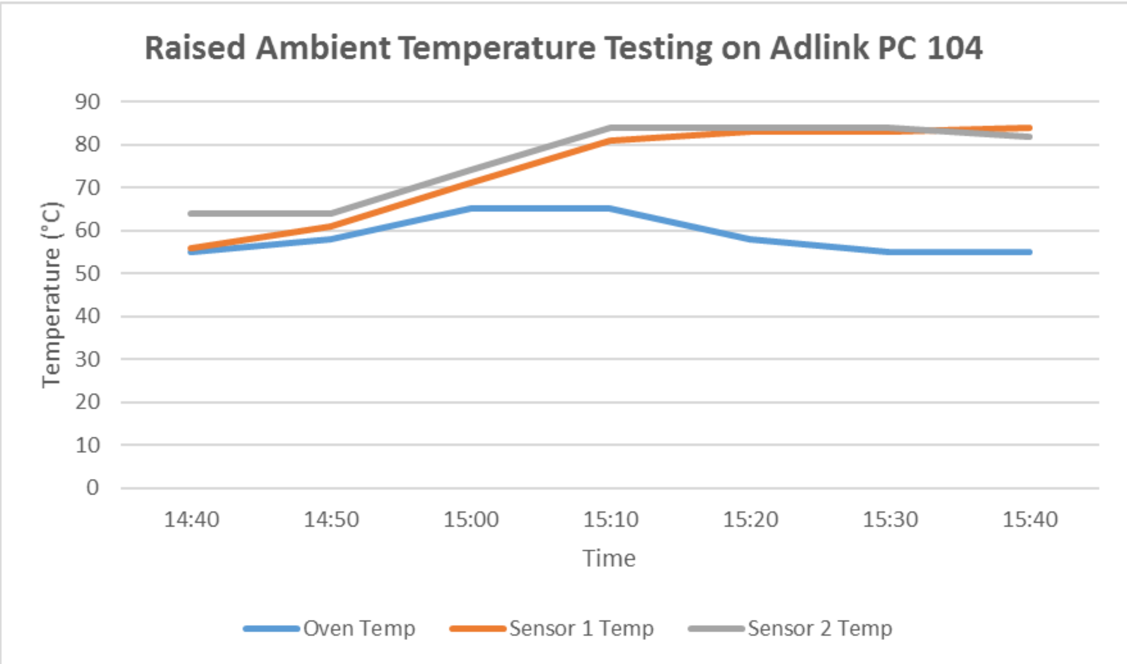

Figure 16 – Graphed results from raised ambient temperature tests with the PC104

This test showed that the PC104 is capable of operating without interruption in environments with an ambient temperature of over 55°C. Temperature readings levelled off 30 minutes into the test recording a maximum of 84°C. Although this test was performed on a board that had not undergone extreme temperature conformance testing by the supplier, it does give confidence that the unit will continue to operate when approaching, and even exceeding, the stated operating temperatures.

4.2.9 Conclusion

Temperature testing under maximum load showed that the Amplicon Impact-D was unsuitable due to the high readings when running in an enclosed environment. Intel documents that the maximum operating temperature of the Atom CPU is 99°C, and observations from the Impact-D reached 98°C before abandoning the test. These findings were discussed with Amplicon who have since claimed to have fixed an issue that was causing this problem. This has however led to doubt about the maturity of the product. Future products provided from Amplicon would need to be re-tested to ensure they meet the requirements where there is a high ambient temperature.

The B&R APC and PC104 performed consistently well in enclosed environment testing with maximum stress, where temperatures did not exceed 64°C. Further testing on the PC104 with raised ambient temperatures also showed promising observations. Here, operating temperatures levelled out at 84°C after 30 minutes and continued to stay the same for the remaining 30 minutes of the hour-long test.

In terms of cost, the PC104 board at £667 is significantly cheaper than the APC2100 at £1314 despite their similar specifications. The difference of £647, coupled with the fact that the PC104 can be certified for use in extreme ambient temperatures (up to 85°C) makes the PC104 the most suitable option for this prototype.

4.3 PC104 Heat Sink

Selection of the PC104 was made based on testing that proves the unit will continue to function in environments with high ambient temperatures.

Testing in Section 4.2.7 was performed with a unit that uses an E3815 Intel Atom CPU. Results from these tests gave promising results, showing the unit operating at a maximum of ~84°C in an ambient temperature of 60°C. The decision was made to purchase a board with a more powerful CPU along with added ‘Extreme Rugged’ temperature testing, which is documented as ensuring the unit is capable of operating in ambient temperatures up to 85°C. This would be more than enough to meet the requirements for this project as the ambient temperature within a flameproof enclosure would not exceed 60°C.

The purchased board contains a more powerful E3845 Intel Atom CPU which results in more processing capabilities, but also leads to an increased heat generation. Testing performed on this unit found that operating temperatures were much higher than the evaluation unit. Discussion with the supplier found that the heat spreader on the board would require an additional heat sink in order to see similar results as the evaluation board.

4.3.1 Thermal Resistance

It is important that the selected heat sink has a thermal resistance that ensures it is capable of cooling the CPU sufficiently. Finding an appropriate thermal rating can be achieved using parameters given in the thermal design guide for the CPU. Intel (2016) defines these as follows:

- Thermal design power (TDP) = 10W
- Thermal resistance from junction-to-heat sink = 0.3°C/W
- Temperature drop between CPU junction and heat sink surface (ΔT1) = 3°C
- Maximum junction temperature for CPU = 110°C

4.3.1.1 Example Calculation

The following example runs through the calculations performed to discover the thermal rating for a heat sink with the following target characteristics:

- Surrounding ambient temperature: 60°C
- Target operating junction temperature: 75°C

The heat sink surface temperature was calculated as follows:

$$Heat\ sink\ surface\ temperature = Junction\ temperature - \Delta T1$$
$$= 75 - 3$$
$$= 72°C$$

The temperature drop between heat sink surface and ambient (ΔT2) was calculated as follows:

$$\Delta T2 = heat\ sink\ surface\ temperature - ambient\ temperature$$
$$= (72 - 60)°C$$
$$= 12°C$$

The thermal resistance of the heat sink could then be found as follows:

$$Thermal\ resistance = \frac{\Delta T2}{TDP}$$
$$= \frac{12}{10}$$
$$= 1.2°C/W$$

These findings result in a heat sink that would be unsuitable for this application due to its large size. As discussed previously, the Intel Atom is capable of operation up to 99°C, thus the calculations were repeated adjusting the target junction temperature to determine a trade-off between operating temperature and heat sink size.

Further analysis was also carried out by reducing the ambient temperature to determine the type of heat sink required in order to achieve these targets. Table 10 documents these results.

Table 10 – Thermal rating of heat sinks for a range of target parameters

Ambient Temp (°C)	**Target Junction Temp (°C)**	**Heat Sink Surface Temp (°C)**	**Thermal Resistance of Heat Sink (°C/W)**
60	75	72	1.2
60	85	82	2.2
60	88	85	2.5
60	90	87	2.7
55	75	72	1.7
55	85	82	2.7
55	88	85	3.0
55	90	87	3.2

4.3.2 Market Research

From these results an appropriate heat sink can be found, Table 11 contains suggestions for a range of heat sinks with different thermal ratings.

Table 11 – Suggested heat sinks for a range of thermal ratings

Thermal Rating (°C/W)	**Dimensions**	**Cost**	**Profile**
1.4	100 x 96 x 40mm	£14.77	
1.9	75 x 88 x 25mm	£6.43	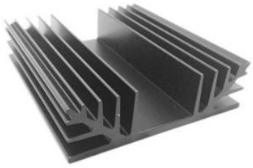
2.5	50 x 88 x 25mm	£4.99	
2.5	37.5 x 88 x 35mm	£2.82	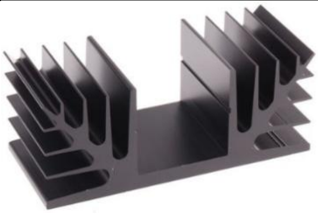
2.5	37.5 x 37.5 x 33mm	£5.42	

2.6	41.91 x 25.4 x 63.5mm	£0.99 Price each (In a Pack of 10)	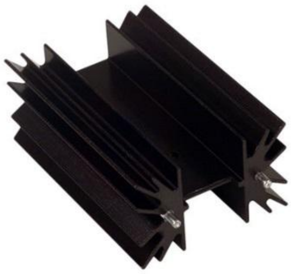
3.1	100 x 50 x 28mm	£5.68	

4.3.3 Heat Sink Selection

Findings from the calculations show that the large heat sink, with a rating of 1.4°C/W, should be more than capable of cooling the unit and will therefore be selected for testing. In production environments where a lower ambient temperature is expected, the smaller finned sink with a thermal rating of 2.5°C/W may serve as a middle ground between size and cooling efficiency.

4.3.4 Temperature Testing

Initial tests were performed by placing the unit in an enclosed box with no ventilation. Maximum stress was applied to the CPU using the 'Intel Burn Test' application which was configured to carry out 10 tests. The time taken to complete each test was recorded, with shorter times indicating better performance. Temperature readings were taken using 'Speed Fan' and CPU speed was monitored using 'CPU-Z'. Raised ambient temperature tests were performed by placing the board in an oven and using the same software.

The less powerful evaluation board that uses the E3815 CPU was tested first with no heat sink attached to serve as a baseline. Tabulated results are shown in Table 12 and graphed in Figure 17.

Table 12 – Temperature tests using the E3815 with the standard heat spreader

Test	CPU speed Min	CPU speed Max	Temp 1 Sensor (°C)	Max Core Temp (°C)	Time (s)	Speed (Gflops)
1	1466MHz	1466MHz	64	66	1172	1.35
2	1466MHz	1466MHz	64	69	1152	1.38
3	1466MHz	1466MHz	64	69	1151	1.38
4	1466MHz	1466MHz	64	69	1154	1.37
5	1466MHz	1466MHz	64	69	1167	1.36
6	1466MHz	1466MHz	64	69	1168	1.36
7	1466MHz	1466MHz	64	69	1177	1.35
8	1466MHz	1466MHz	64	70	1170	1.36
9	1466MHz	1466MHz	64	70	1166	1.36
10	1466MHz	1466MHz	64	70	1166	1.36

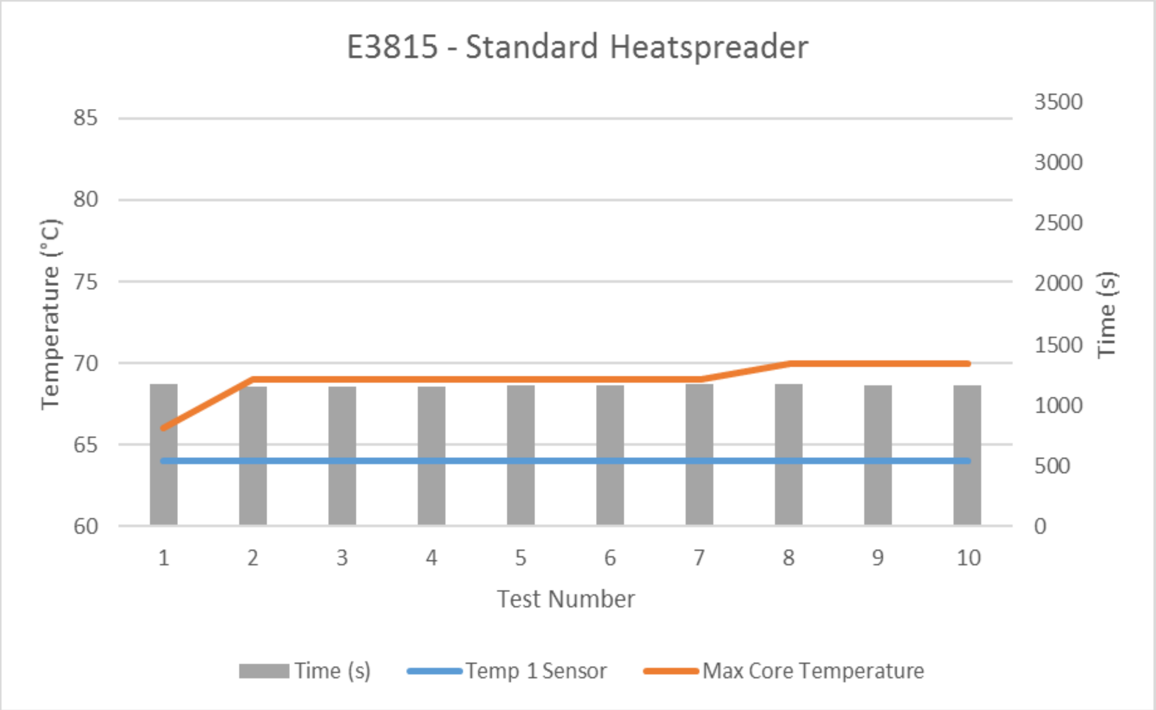

Figure 17 – Graphed temperature test results using the E3815 with standard heat spreader

The more powerful purchased board that uses the E3845 CPU was then tested with a copper block attached to the heat spreader. The block has a poor thermal rating and thus demonstrates how the unit performs in a worst case scenario. Tabulated results are shown in Table 13 and graphed in Figure 18.

Table 13 – Temperature tests using the E3845 with a copper block heat sink

Test	CPU speed Min	CPU speed Max	Temp 1 Sensor (°C)	Max Core Temp (°C)	Time (s)	Speed (Gflops)
1	1915MHz	1915MHz	74	76	813	7.1
2	1915MHz	1915MHz	76	79	826	6.9
3	1915MHz	1915MHz	78	81	815	7
4	1915MHz	1915MHz	79	84	814	7
5	499MHz	1915MHz	73	74	3383	1.7
6	666MHz	666MHz	72	74	2266	2.5
7	666MHz	666MHz	72	74	2292	2.5
8	666MHz	666MHz	72	74	2288	2.5
9	666MHz	666MHz	72	74	2270	2.5
10	666MHz	666MHz	72	74	2282	2.5

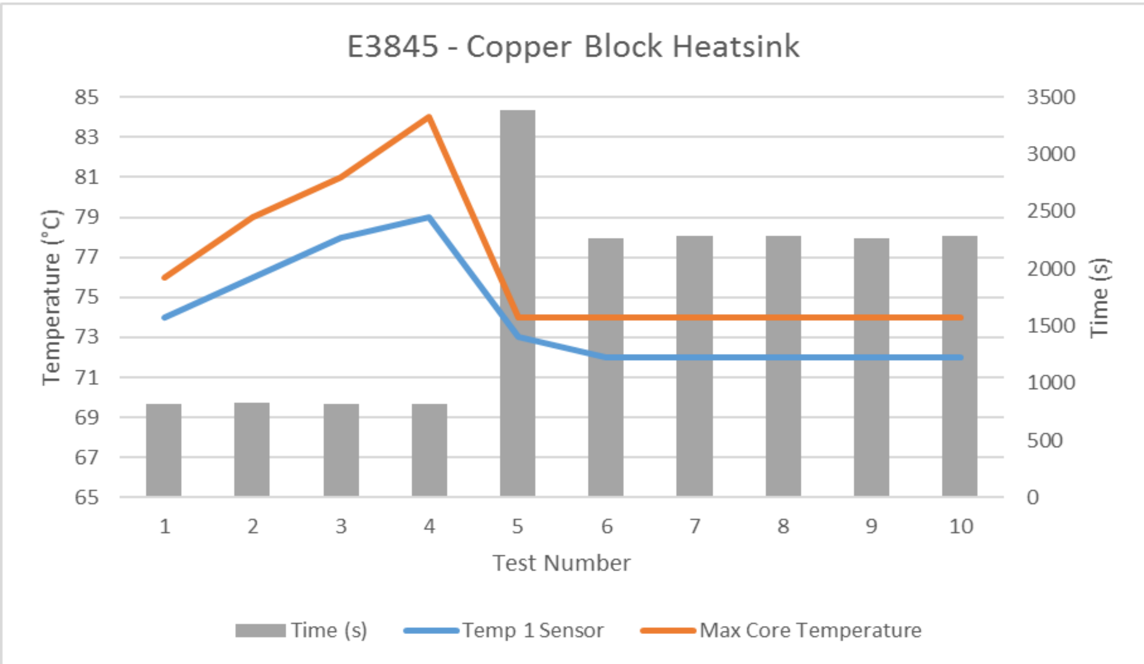

Figure 18 – Graphed temperature test results using the E3845 with standard heat spreader

The test was then repeated by attaching the large heat sink to the heat spreader. Tabulated results are shown in Table 14 and graphed in Figure 19.

Table 14 – Temperature tests using the E3845 with a large heat sink rated at 1.4°C/W

Test	CPU speed Min	CPU speed Max	Temp 1 Sensor (°C)	Max Core Temp (°C)	Time (s)	Speed (Gflops)
1	1915MHz	1915MHz	44	47	811	7.04
2	1915MHz	1915MHz	44	53	810	7.04
3	1915MHz	1915MHz	54	55	810	7.04
4	1915MHz	1915MHz	54	57	809	7.05
5	1915MHz	1915MHz	54	58	810	7.04
6	1915MHz	1915MHz	54	59	810	7.04
7	1915MHz	1915MHz	54	59	820	6.95
8	1915MHz	1915MHz	54	61	823	7.03
9	1915MHz	1915MHz	54	61	817	6.9
10	1915MHz	1915MHz	54	61	810	7.04

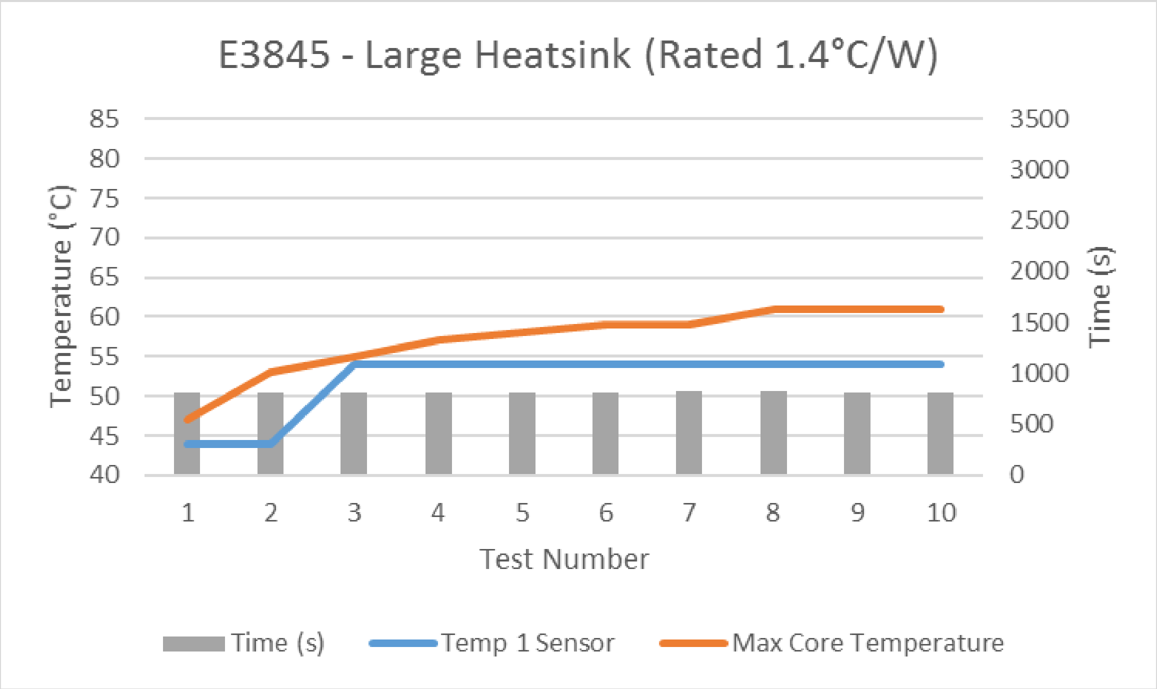

Figure 19 – Graph showing temperature test results using the E3845 with a large heat sink rated at 1.4°C/W

A final test was carried out to determine whether the heat sink will perform effectively in a production like environment with an ambient temperature of 60°C. This was achieved by placing the unit inside the same oven used in Section 4.2.8. Tabulated results are shown in Table 14 and graphed in Figure 20.

Table 15 – Temperature tests using the E3845 with a large heat sink rated at 1.4°C/W in a target ambient temperature of 60°C

Test	**CPU speed Min**	**CPU speed Max**	**Max Ambient (°C)**	**Temp 1 Sensor (°C)**	**Max Core Temp (°C)**	**Time (s)**	**Speed (Gflops)**
1	1915MHz	1915MHz	70	74	76	489	7.13
2	1915MHz	1915MHz	65	74	77	488	7.13
3	1915MHz	1915MHz	60	74	78	488	7.14
4	1915MHz	1915MHz	63	74	81	489	7.13
5	1915MHz	1915MHz	63	74	81	489	7.13
6	1915MHz	1915MHz	65	74	82	488	7.15
7	1915MHz	1915MHz	65	74	83	488	7.14
8	1915MHz	1915MHz	65	74	83	488	7.15
9	1915MHz	1915MHz	65	74	83	488	7.14
10	1915MHz	1915MHz	65	74	83	488	7.14

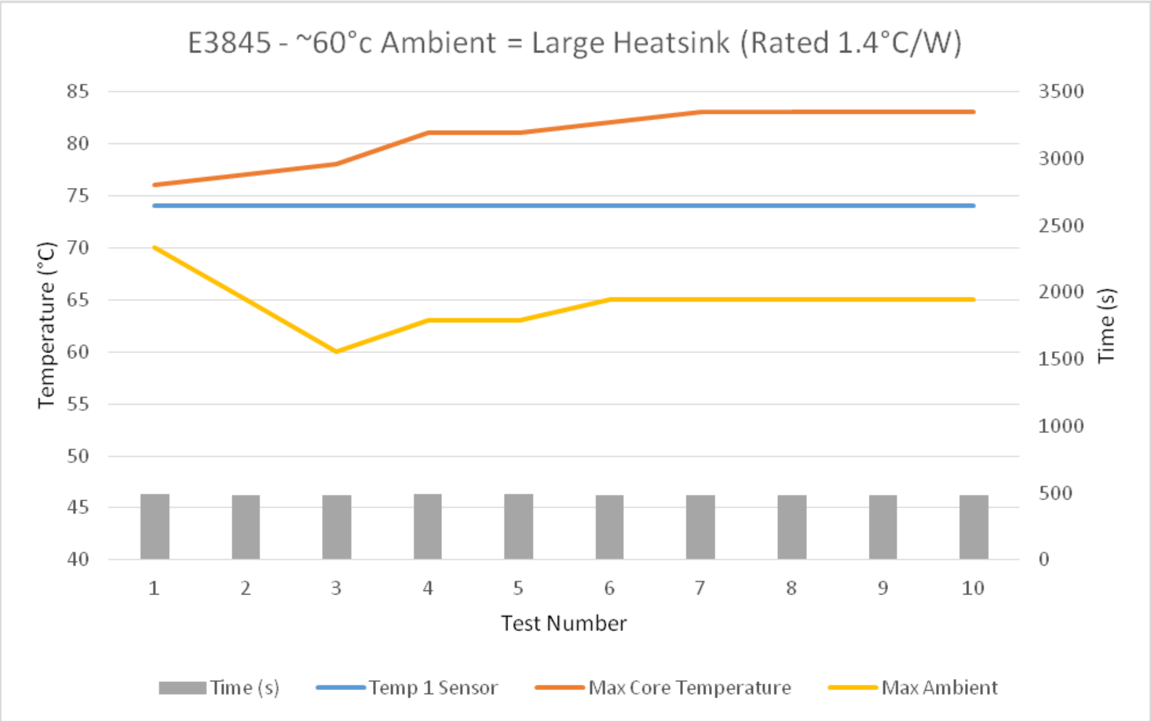

Figure 20 – Graphed temperature test results of the E3845 with a heat sink rated 1.4°C/W in 60°C

4.3.4.1 Discussion

Tests performed on the lower specification E3815 evaluation board with the supplied heat spreader showed temperatures levelling out at 70°C under maximum stress. Each test took an average of 20 minutes to complete.

The same tests on the higher specification E3845 purchased board using a copper block heat sink attached to the heat spreader showed temperatures much higher than those observed with the evaluation board. The first few tests did however complete much faster, as expected, taking an average of 14 minutes which reflects the higher performance offered by this CPU. Once the maximum core temperature exceeded 85°C, however, the CPU speed was throttled, initially to 499MHz and subsequently levelling off at 666MHz. This resulted in temperatures being stabilised to 74°C but also resulted in tests taking much longer to complete, averaging at 38 minutes. This increased time was due to the performance loss caused by throttling and shows that the lower specification E3815 board actually performs faster under these conditions.

Results from testing with the selected heat sink on the E3845 showed a significant improvement, where CPU throttling did not occur and temperatures levelled out at 61°C in the enclosed box. This meant that the CPU was able to continue operating at its maximum clock speed for all tests, resulting in completion times averaging 14 minutes.

Following these findings, it was decided to determine whether the unit would see similar results in a real-world scenario where an ambient temperature of 60°C would be present. The unit was placed inside an oven with a target ambient set to 60°C. The benchmarking application was configured to complete tests within half the time of the initial tests to enable more regular readings to be recorded. Maintaining an ambient of 60°C proved difficult due to the oven being unreliable, after initial fluctuations this eventually settled above the target at 65°C. CPU core temperatures slowly rose over the course of the tests, but levelled out at 83°C over the course of all 10 tests, which took just over an hour. These results fall very close to the throttling temperature of 85°C, but the 5°C increase of ambient temperature should be factored in to this. Furthermore, if throttling was disabled the CPU would not automatically turn off unless a temperature of 95°C was reached.

The original calculations showed that for an ambient temperature of 60°C a junction temperature of 75°C could be reached with a heat sink of 1.2°C/W. The next step up from this was a junction temperature of 85°C which was calculated as requiring a heat sink rated at 2.2°C/W. Since the heat sink used was rated between these two at 1.4°C/W there is confidence that the calculations closely align with reality.

4.3.5 Heat Sink Conclusion

Calculations were performed based on parameters given in the CPU technical documentation to find the required thermal rating of the heat sink. An ambient temperature of 60°c was chosen to represent the worst case scenario of a production environment.

In order to ensure that the unit remains as cool as possible it is important that a heat sink with the lowest thermal rating is selected. The main issue with this however is that lower thermal ratings result in an increase in physical size of the heat sink. A selection was made to use a heat sink with a thermal rating of 1.4°C/W which, based on calculations, should be more than capable of cooling the CPU.

Temperature testing was performed to determine whether these calculations align with the reality of a real-world scenario where ambient temperatures reach up to 60°C. The results showed temperatures that correlate with the results from the calculations, thus giving confidence that the heat sink selected is suitable.

5 Implementation

With the platform equipment selected the project could move onto the next phase which focuses on designing, developing and testing the software that delivers the prototype diagnostics system.

As discussed by Jardine et al. (2006) a condition based maintenance programme consists of three key stages: data acquisition, data processing and maintenance related decision making (referred to as business logic throughout this section). Implementation of the solution will tackle each of these stages individually, with data acquisition being handled by the low-level microcontroller and data processing and application of business logic taking place on the high-level PC platform.

5.1 Low level Solution – Data Acquisition

5.1.1 Requirements

The low level solution is developed for the STM32F4 microcontroller platform that was selected as discussed in Section 4.1. Its main task is to acquire the raw Modbus RTU data that is transmitted by equipment during operation. Core to this is the requirement for the solution to not play an active role in a conversation between equipment. Modbus RTU is a single master protocol using a shared medium. If the MCU attempts to communicate on the same medium a collision would occur when the two separate devices attempt to communicate over this shared medium at the same time. This would result in data corruption and messages from the master to the slave would be lost. It is therefore important that the system passively monitors the data and does not interfere as shown in Figure 21.

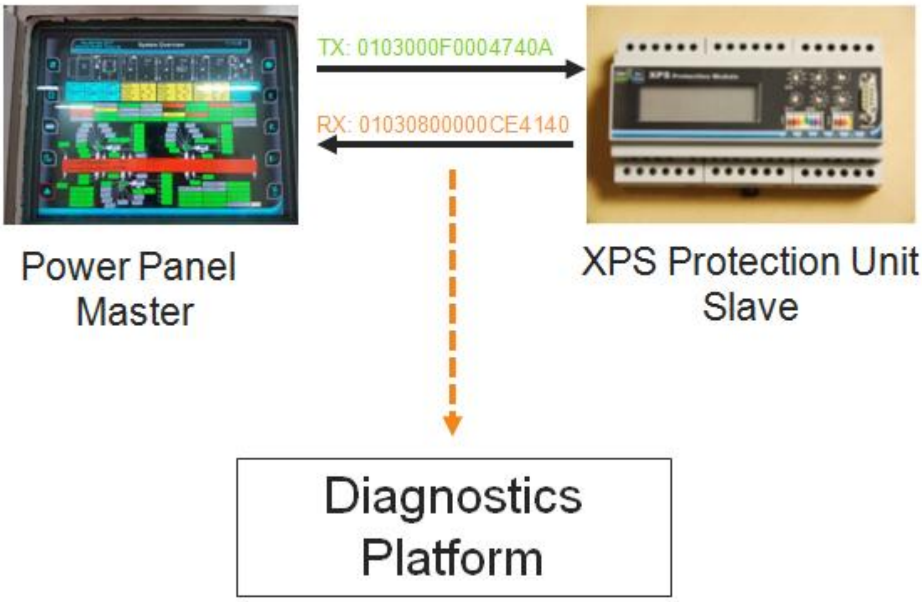

Figure 21 – Passive communication monitoring between a Modbus master and slave device

Aside from monitoring, the solution must also be responsible for other tasks to provide required functionality. The raw data must be streamed to the high level solution which is responsible for analysing the data and presentation to the end user. Another task is required to store the raw data on a local storage medium on the MCU, such as an SD card.

The nature of the MCU platform means that only one task can be performed at a time. It is therefore necessary to have a mechanism for scheduling tasks to ensure that each has an opportunity to run and perform its duty. A real time operating system (RTOS) provides this functionality and once implemented will manage the running order of tasks. The RTOS also enables prioritisation and pre-emption of tasks which is a vital concern for this solution since the data acquisition task must always take precedence over any other tasks to prevent data loss.

5.1.2 RS485 Communication Interface

The selected evaluation board provides an RS232 to interface with the onboard USART. For industrial applications however, an RS485 interface is generally used for communication as it offers the following advantages:

- differential transmission providing rejection of common mode noise
- greatly-increased distance over RS232 or TTL (transistor-transistor logic) – up to 1km
- can support ‘multi-drop’ operation, allowing simultaneous connection of many nodes

An RS485 transceiver uses a pair of differential lines, known as ‘A’ and ‘B’. A single transceiver can work in either transmit or receive mode, but never both simultaneously (known as ‘half duplex’ operation). To transmit data, the hardware takes control of the bus and physically drives the lines to the required state. During reception, it ‘releases’ the bus and can receive data by looking at the differential voltage present on the lines. ‘Full duplex’ operation is possible using 2 transceivers, but is not commonly used.

The diagnostics system is required to interface directly with ‘half duplex’ RS485 on equipment. A bespoke RS485 interface was purposely designed and developed for the project. To achieve this, 3 signals are required from the evaluation board – ‘TX’ (transmit), ‘RX’ (receive) and ‘RTS’ (request to send). An additional design challenge was present in that the interface signals on the board are 3.3V logic levels, but the RS485 transceiver used 5V logic levels. Figure 22 shows the design of the circuit complete with interconnections between components

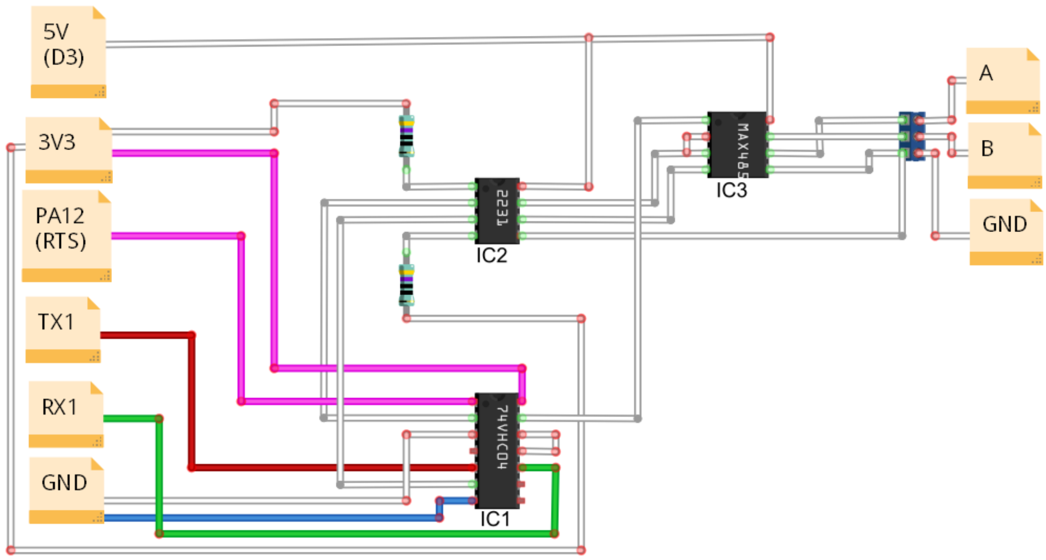

Figure 22 – Circuit design of an RS485 interface for the MCU

IC1 is a hex logic inverter designed for operation on 3.3V or 5V supplies. It is used here for signal inversion and 5V to 3.3V interfacing of the 'receive' signal from IC3.

IC2 is a dual-channel, high-speed logic opto-coupler designed for use on a 5V supply. It is used here for 3.3V to 5V interfacing of the 'transmit' and 'RTS' signals to IC3. The LEDs and output stages in both channels of IC2 provide the required 3.3V to 5V level shift.

IC3 is an integrated RS485 transceiver from Maxim (MAX485) which converts the single-ended 'transmit' and 'receive' signals into a differential bus ('A' and 'B'). It requires a 5V supply to meet the RS485 specification.

Figure 23 shows a prototype of the circuit implemented on Veroboard.

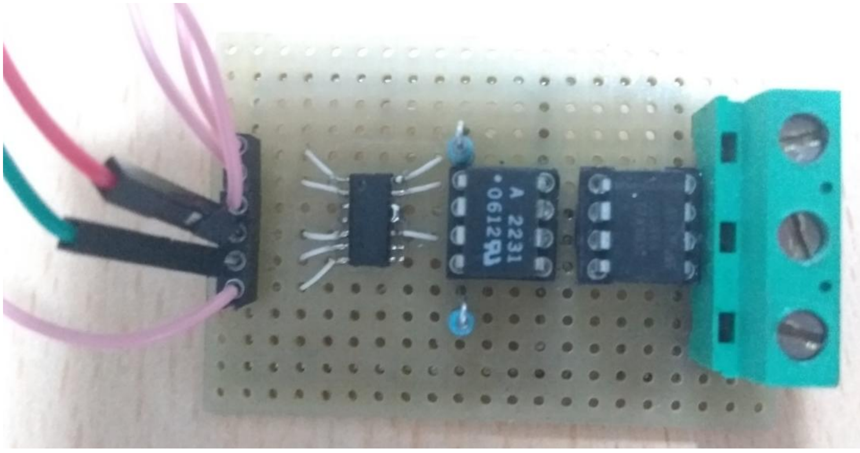

Figure 23 – Prototype of RS485 interface for the MCU using Veroboard

5.1.3 Theory of Operation – RS485 Interface

In receive mode, the MCU holds the RTS line into IC1 pin 1 low (logic 0 or 0V). IC1 pin 2 will thus be high, and no current will flow through the LED element in IC2, causing its output (pin 7) to be held low. IC3 pin 2 '/RE' (receive enable) and pin 3 'DE' (drive enable) will both be low, which holds IC3 in a 'receive' condition. In this condition, the state of IC3 pin 1 'RO' (receiver output) depends on the differential voltage between IC3 pin 6 ('B') and IC3 pin 7 ('A'). If A > B by 200mV, 'RO' will be high (5V), and if A < B by 200mV, 'RO' will be low (0V).

If another active transceiver drives the RS485 bus, it can thus send a binary logic signal, either by driving A to 5V and B to 0V ('RO' high) or A to 0V and B to 5V

('RO' low). The 'RO' signal is taken back to IC1 and inverted twice by 2 series-connected stages. The first stage converts the 5V to 3.3V for compatibility with the evaluation board, and the second stage restores correct polarity. The level-shifted signal goes into the evaluation board and passes to the serial USART device for data reception.

In transmit mode, the MCU takes the RTS line into IC1 pin 1 high (logic 1 or 3.3V). IC1 pin 2 will go low and the LED element in IC2 will switch on, causing its output (pin 7) to go high. This puts IC3 into a 'transmit' or 'drive' condition. Now, pins 6 and 7 are actively driving the RS485 bus. If IC3 pin 4 'DI' (driver input) is high, 'A' will be at 5V and 'B' will be at 0V; taking pin 4 low will cause pins 6 and 7 to change over. Figure 24 shows an oscilloscope trace during a TX event. The blue trace is probing RTS and yellow is probing the TX line, both at the 3.3V stage of the circuit.

When the MCU transmits data, its USART device sends a TX signal of either high (passive) or low (active) logic state. This is inverted by another stage in IC1 (pins 5/6) and used to control the second channel of IC2. When TX is high, IC2 outputs 5V on pin 6, taking IC3 pin 4 high; when TX is low, IC2 outputs 0V and IC3 pin 4 is low. This is transferred to the RS485 bus on pins 6 and 7 of IC3, producing the required signalling action on the bus. At the end of the transmission, the MCU returns the RTS signal low and the MAX485 releases the bus, allowing further reception to take place.

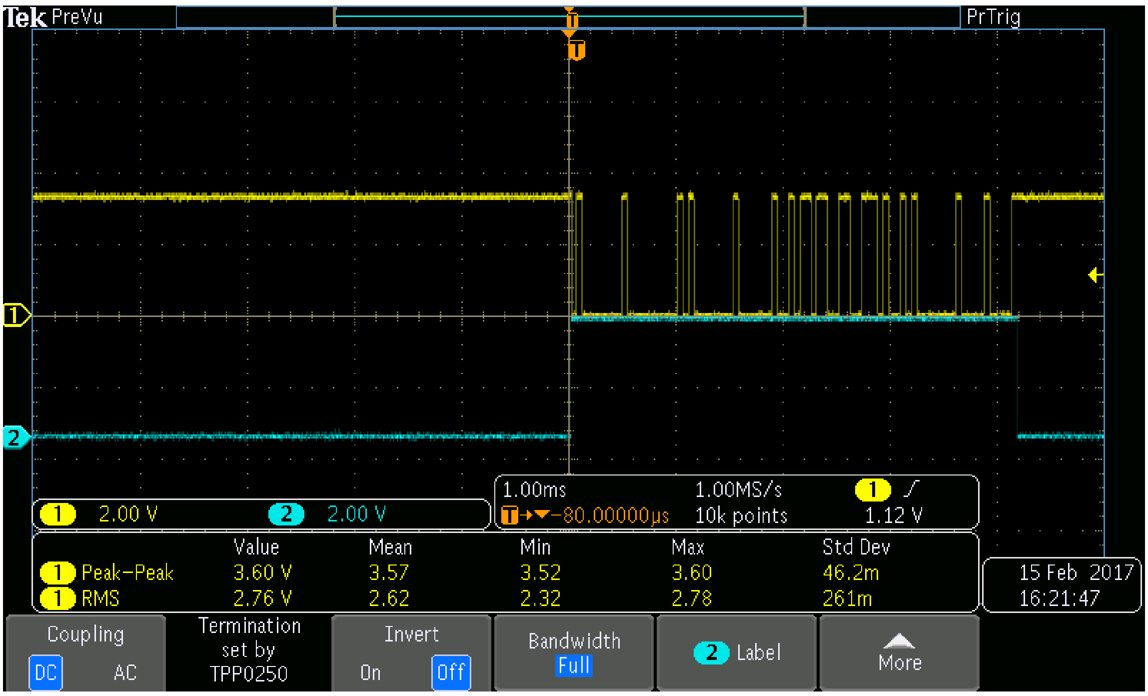

Figure 24 – Oscilloscope trace during a transmit event using the RS485 interface

5.1.4 Data Acquisition Solution

With a suitable communication interface implemented, work could begin on developing software for acquiring data. A single solution was developed using the C programming language, taking advantage of the Keil middleware to provide support for peripherals and services. The solution consists of 7 separate threads and uses the Keil RTOS to manage these tasks. Table 16 details the purpose and priority of each task. Interrupts are also used as shown in Table 17.

Table 16 – Threads used in the data acquisition solution

Purpose	C file	Priority
Data acquisition	mbThread.c	High
Raw data storage	fsThread.c	Above Normal
Raw data streaming	netThread.c [netSendThread function]	Normal
Network service handlers (FTP/TCP)	netThread.c [netThread function]	Normal
Network time synchronisation	ntpThread.c	Normal
Heartbeat	heartbeatThread.c	Normal
Signal management	sigThread.c	Normal

Table 17 – Interrupts used in the data acquisition solution

Purpose	Peripheral	Initialisation function	Handler function	Preempt Priority	Sub Priority
Modbus silent timer	TIM2	Timer_Init [main.c]	TIM2_IRQHandler [mbThread.c]	0	0
Modbus data receiver	USART1	Init_Mb_Thread [mbThread.c]	myUSART1_callback [mbThread.c]	1	1

Inter-task communication is achieved using a feature of the RTOS known as signals. These signals are used throughout the application whenever one thread needs to communicate with another. They are defined and enumerated within a header file as follows:

- SIGNAL_MODBUS_THREAD
- SIGNAL_NTP_FIRST_RUN
- SIGNAL_NTP_THREAD

- SIGNAL_FS_THREAD
- SIGNAL_NET_SEND_THREAD
- SIGNAL_HEARTBEAT_THREAD

The idle running state of the solution prioritises the netThread task which runs a middleware API function for handling all network services. Once this task completes, the signal thread is allowed to run to determine if any new data has been received that needs processing. In debugging this can be seen in the event viewer as shown in Figure 25, where blue sections indicate thread operation for a duration of 50ms.

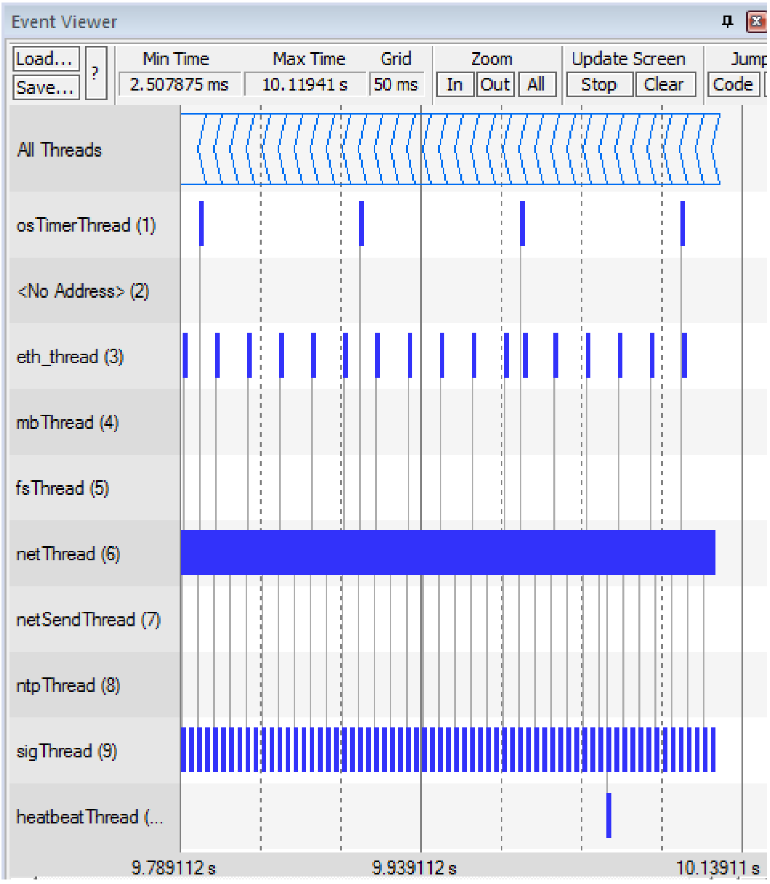

Figure 25 – Event viewer debug utility showing expected idle operation of the data acquisition solution

Whenever Modbus data is received via the USART the associated ISR runs and delegates to a call-back function. This is defined by the middleware and minimises user code within the ISR. The call-back sends a signal to the Modbus thread to run as shown in Figure 26.

Ensuring that user code in an ISR is minimised is vital for time critical embedded systems where determinism is core to achieving the desired operation. This is because whenever an interrupt occurs, any task that is presently running will be stopped in order to service the interrupt. Execution only returns to the main task once the ISR routine has finished (Nuggehalli Ramachandra & Kannur, 2008).

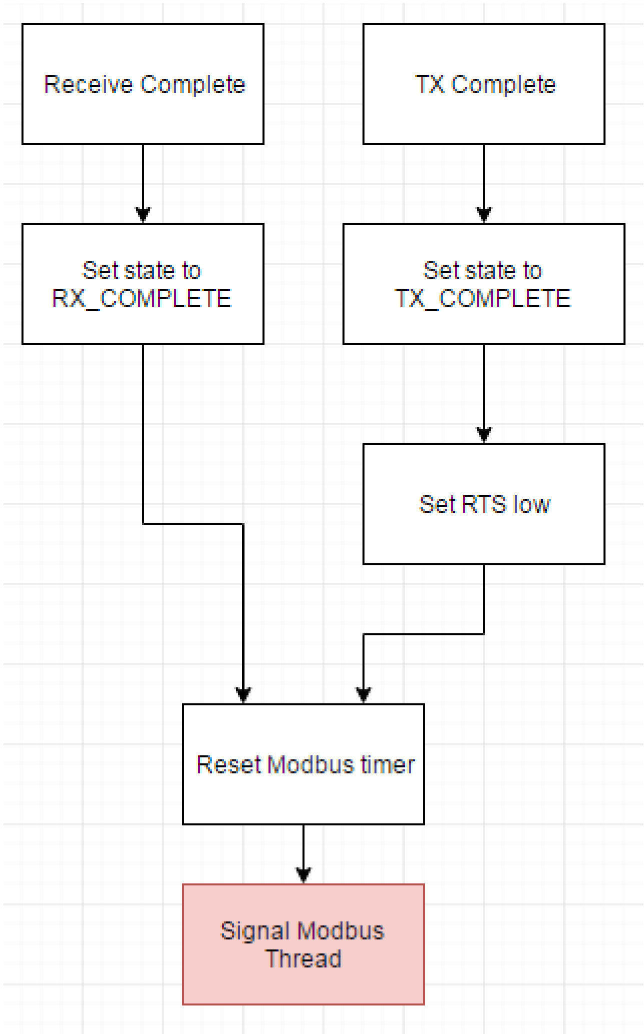

Figure 26 – Program flow of the USART call-back function

The ISR pre-empts any presently running task and runs immediately, ensuring that no data is lost. The Modbus thread then stores the new data in a circular buffer (code documented in Appendix II) using the 'writeIntoBuffer' function. Program flow of the thread is shown in Figure 27.

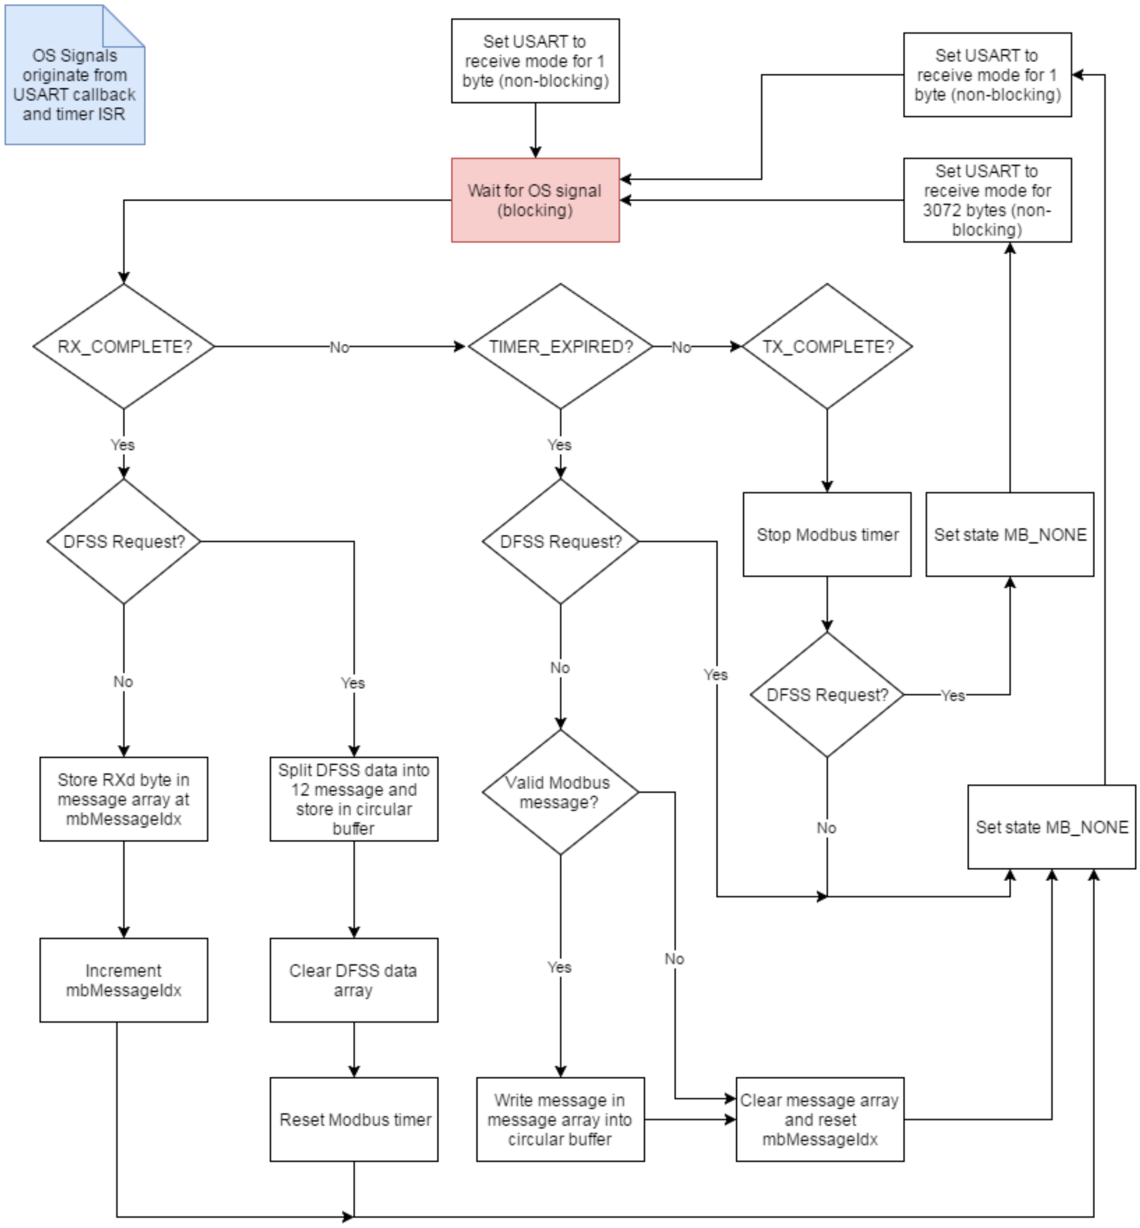

Figure 27 – Program flow of the Modbus thread

Whenever new data is available in the buffer, the signal thread sends a signal to the file system thread to run. This thread performs formatting on the raw message, adding key information that relates to the message properties and the status of the capture process. Delimitation of this information is achieved using the pipe character as shown in Figure 28. Once this formatting is performed, the message is written to a text file on a new line.

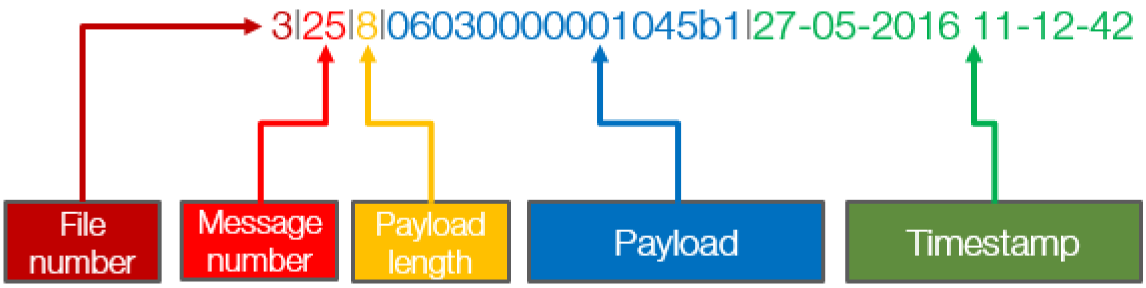

Figure 28 – Raw Modbus data with additional status information added

Each line of the raw text files contains a single formatted message as shown in Figure 29. The default configuration is to log 1000 messages per file. Each file is stored in a directory structure that identifies the year, month and day when it was captured. To prevent clogging up a single 'day' directory with many files, subdirectories are used. Each time a subdirectory reaches 100 files a new one is created. An example of this structure for data captured on the 25th August 2016 is as follows: '\2016\08\25\1\data.txt'

```
1|513|8|0103000000104406|25-08-2016 12-33-03
1|514|37|0103200000800000000000000000000000000000000000000000000000000000000000727a|25-08-2016 12-33-03
1|515|8|0203000000104435|25-08-2016 12-33-03
1|516|37|020320000080000000000000000000000000000000000000000000000000000000000057a|25-08-2016 12-33-03
1|517|8|03030000001045e4|25-08-2016 12-33-03
1|518|37|03032000008000000000000000000000000000000000000000000000000000000000029ba|25-08-2016 12-33-03
1|519|8|0403000000104453|25-08-2016 12-33-03
1|520|37|0403200000800000000000000000000000000000000000000000000000000000000000eb7a|25-08-2016 12-33-03
1|521|8|0503000000104582|25-08-2016 12-33-04
1|522|37|0503200000800000000000000000000000000000000000000000000000000000000000c7ba|25-08-2016 12-33-04
1|523|8|06030000001045b1|25-08-2016 12-33-04
1|524|37|0603200000800000000000000000000000000000000000000000000000000000000000b0ba|25-08-2016 12-33-04
1|525|8|1003000000570775|25-08-2016 12-33-04
1|526|179|1003ae17c700000000555555555555555555555555ffeffbffffffdf7ffdf7ffff0000000000000000000000000000000000000
0000000000000ca0000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
00000000000009d0000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
0000000000000000000000000bfab|25-08-2016 12-33-04
1|527|8|100300800054469c|25-08-2016 12-33-04
```

Figure 29 – Raw text file showing formatted messages that contain Modbus data

This same message format is used when sending messages via the TCP socket to the high level solution on the PC that will be listening for data. In scenarios where the file system fails (e.g. an SD card is not present) the file system thread will still run since it is responsible for adding the formatting to the message. Functionality has been implemented to ensure that any failures to write to the file system are handled and recovery attempts are made.

Date and time are fundamental to providing accurate information to the end user. The MCU must therefore have functionality for keeping track of the current date and time. The STM32F4 is equipped with a real time clock (RTC) for this purpose and can be configured to make use of a coin-cell battery. This ensures that the time is not lost when main system power is disconnected or when the system restarts.

On start-up, the application automatically attempts to synchronise the RTC with the PC using the network time protocol (NTP). If this fails, the application will wait 1 minute and try again. By default, there will be 5 attempts to sync with the NTP, if all these fail a critical fail flag is set and no more retries will occur. Further attempts can be initiated from the Modbus Monitor application on the PC.

5.1.5 Modbus RTU Silent Timer

As discussed in Section 2.7, the Modbus RTU protocol employs a silent timer to delimit between individual messages. Implementing this on the MCU was achieved using the internal timer 'TIM2'. The timer was configured to use an interval that complies with the Modbus standard of 3.5 character times between messages. For a baud rate of 19200 and a configuration using 8 data bits and no parity (8N1) the silent time was calculated as follows:

Calculate the bit time (t) for the baud rate:

$$t = \frac{1}{baud\ rate}$$

$$t = \frac{1}{19200} = 5.208^{-5}s = 52\mu s$$

Individual character time could then be found as follows:

$$52\mu s \times 10 = 5.208^{-4}s = 0.52ms$$

Finally, the Modbus RTU silent time is calculated as follows:

$$0.52ms \times 3 = 1.56ms$$

To implement this on the embedded application requires a prescaler and a period. These can be found using formulae given in the MCU technical manual as follows:

$$TimerDefaultFrequency = \frac{SystemCoreClock}{2} = \frac{168}{2} = 84MHz$$

$$Prescaler = \left(\frac{TimerDefaultFrequency}{1000000}\right) - 1 = \left(\frac{84000000}{1000000}\right) - 1 = 83$$

The period can be found with the following calculations, involving conversion of the required interval for the Modbus silent timer found above (0.00156s) to Hz:

$$Period = \left(\frac{1000000}{RequiredInterval}\right)$$

$$f = t^{-1}$$

$$0.00156^{-1}s = 641.0256 \text{ Hz}$$

$$Period = \left(\frac{1000000}{641.0256}\right) = 1559$$

This results in a prescaler of 83 and a period of 1559 and can be implemented in code when initialising the timer as follows:

- htim2.Init.Prescaler = 83;
- htim2.Init.Period = 1559;

Validation of this was achieved by placing a test trigger in the ISR for TIM2. This trigger sets a GPIO pin high then low which could be traced via oscilloscope as shown in Figure 30. This shows a frequency of 641Hz which matches the calculations.

Figure 30 – Oscilloscope trace validating TIM2 interrupt frequency

Once tested, the ISR for the timer was replaced with user code to handle the case that a complete message has been received and needs to be written into the buffer ready to move on to the next one. This ISR code is kept as short as possible, using flags that can be referenced from the Modbus thread as shown in Figure 31.

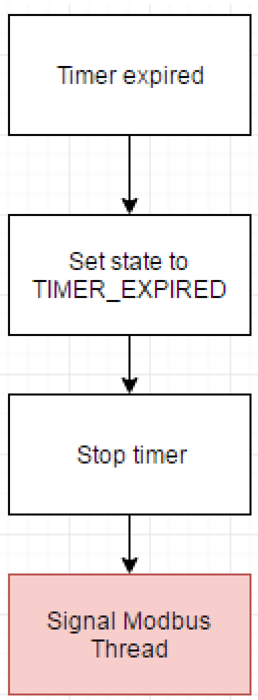

Figure 31 – Program flow of the TIM2 interrupt service routine

5.1.6 Edge Case

Functionality has also been implemented to ensure that edge cases are handled. One such case is scenarios where the storage medium (SD card) becomes full. Attempts to write data to a full SD card without appropriate checks in place would result in new data being lost. Code has been implemented that determines how much free space is available on the SD card each time a new raw file is required. If this is below 10%, a delete function runs which deletes the oldest month's worth of data. This function ensures the card will never get full. The flow of this can be seen in Figure 32.

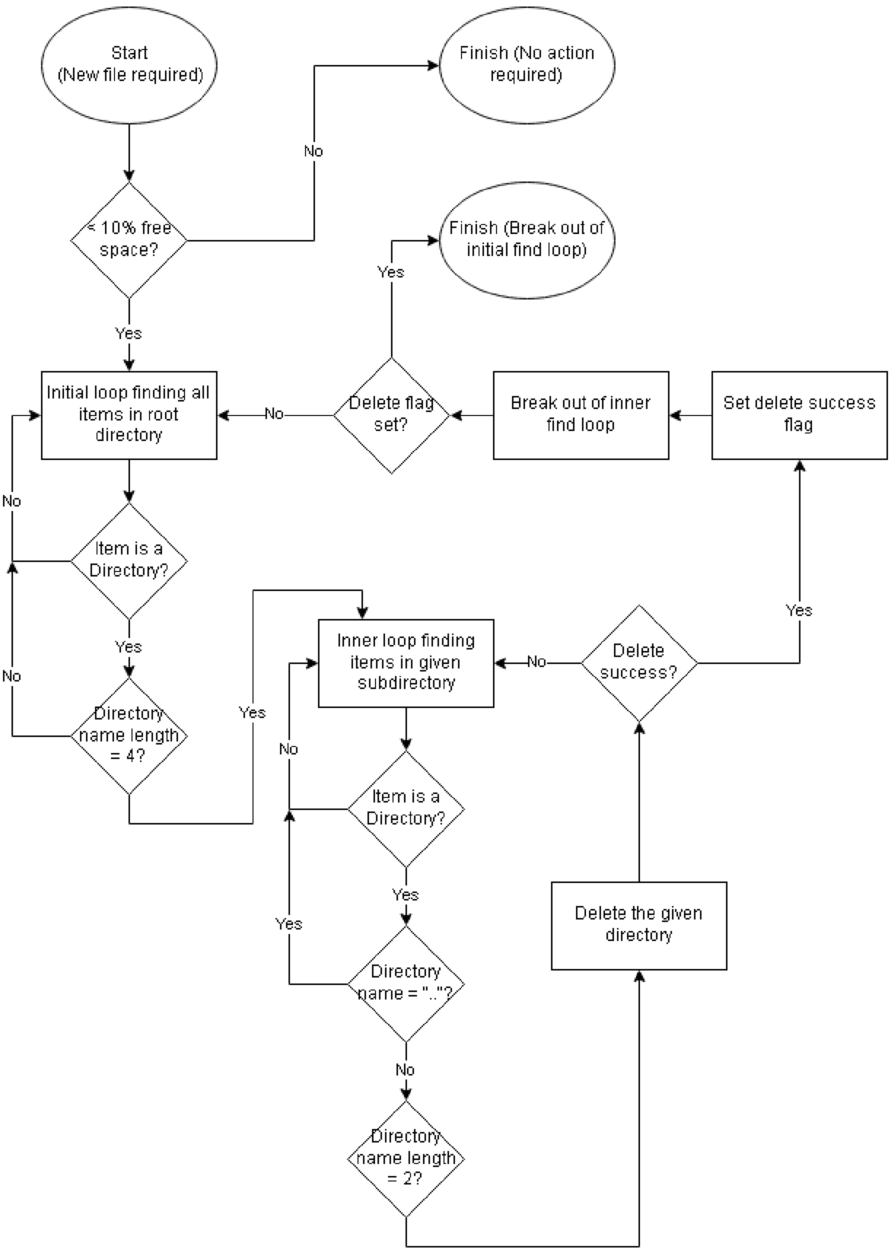

Figure 32 – Program flow of the function that deletes old data when less than 10% space remains on the MCU storage medium

5.1.7 Other Services

A telnet interface has also been implemented to provide information about the operating status. This can be accessed using a telnet client, such as PuTTY, and connecting to the MCU IP address on port 23. The text based interface allows for commands to be sent to view details such as available storage space, thread status and contents of the data buffer. Entering ‘help’ shows detailed information about available commands.

System status is also display using LEDs that are available on the evaluation board. These are defined in the ‘status.h’ file as follows:

- LED_HEARTBEAT 0
- LED_TCP 4
- LED_FILESYSTEM 5
- LED_NTP 6

The heartbeat LED flashes every second to show the system is active. The remaining LEDs are illuminated when they are running as normal. These LEDs are connected via standard GPIO ports and can be reconfigured for different pins if a different PCB is to be used in future.

Other functionality includes an FTP server which can be accessed via any FTP client (such as Filezilla) using the MCU IP address and default FTP port of 21.

Application related variables are configured within the ‘config.h’ file, making use of the Keil Configuration Wizard as in Figure 33.

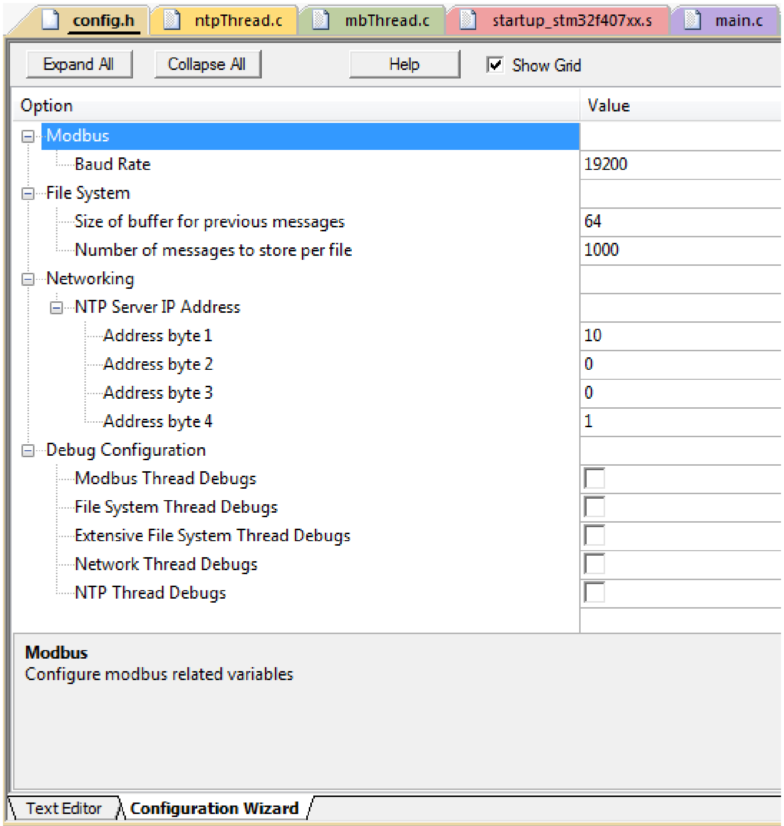

Figure 33 – Configuration interface for variables on the data acquisition solution

The MCU lies at the core of the diagnostics system. Figure 34 shows an overview of the functionality the developed solution provides.

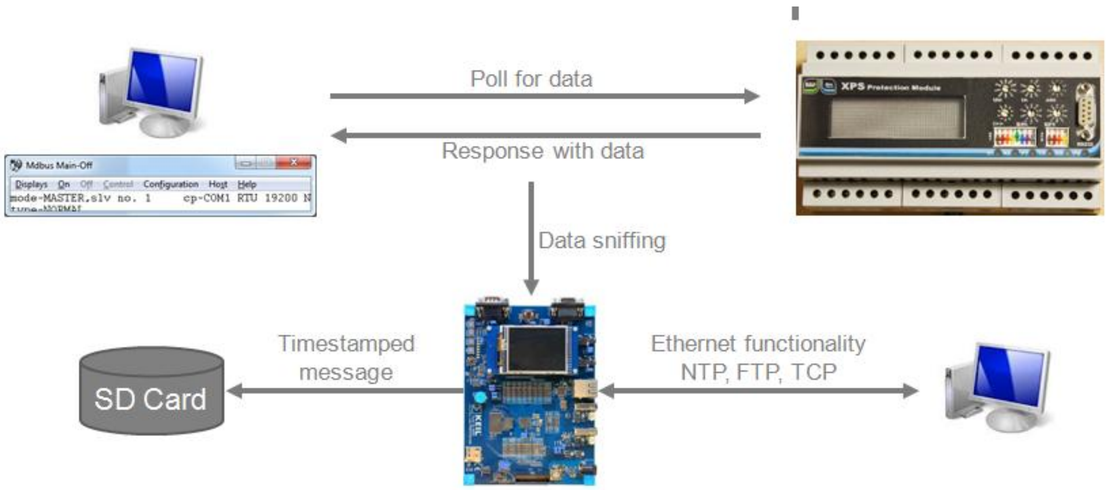

Figure 34 – Overview of functionality provided by the MCU

5.2 High level Solution – Data Processing

With data being captured and TCP streaming enabled, the project could move onto the next phase of processing this data. This task is required to operate on a Windows based PC and uses the embedded PC104 selected in Section 4.2 as the target platform.

5.2.1 Requirements

Key requirements for this development are:

- Receive acquired Modbus data from MCU as it is being captured
- Log individual messages into a file along with a timestamp
- Determine if values need to be logged, and if so, to update the associated table in the database
- Ensure data consistency in scenarios of connection failure
- Continual automated operation without need for manual intervention

This platform will be host for a number of other services and bespoke applications that, alongside the low level solution, will form the remainder of the diagnostics functionality.

5.2.2 Theory of Operation

A single solution was developed using the C# language to handle all data processing functionality. The solution consists of distinct classes to separate code that implements the differing functionality provided. Table 18 documents these classes and the functionality they provide.

Table 18 – Classes used within the data processing application

Class Name	Purpose
BfBusinessRules	Defines rules for driving the customer facing health report
BfDbConnector	Handles connecting to the database
BfDbReader	Read data from the database
BfDbWriter	Write data to the database
BfDfssData	Defines properties to store data from a bespoke equipment type
BfDfssProcessor	Process data from a bespoke equipment type
BfFileHandler	Handles writing raw data to text files
BfLogger	Logs messages, exceptions and debugs
BfMbProcessor	Process a Modbus message and store data in a BfMbTransaction object
BfMbSync	Handles data synchronisation between MCU and PC
BfMbTopology	Defines properties that describe the topology of the system being monitored
BfMbValidator	Validates Modbus messages
BfTcpSocket	Handles the TCP socket connection to the MCU
BfTransaction	Defines properties that are specific to a transaction

Underpinning the solution is the requirement to establish a connection to the MCU and receive the raw Modbus data as it is being captured. This has been achieved by implementing a TCP socket client which connects via Ethernet to the TCP socket server running on the MCU. Figure 35 shows a Wireshark capture of the TCP communication occurring between the MCU and PC.

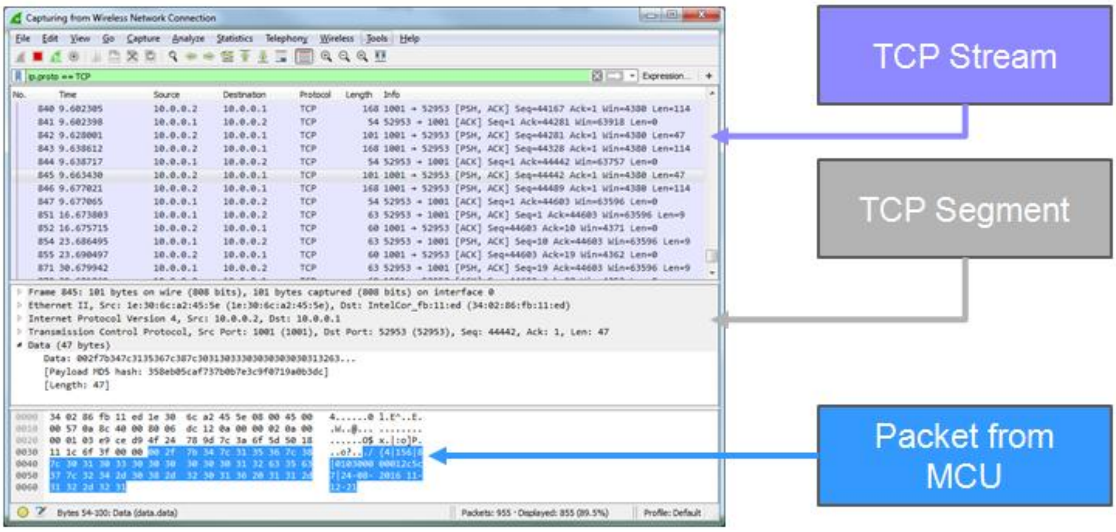

Figure 35 – Wireshark capture of TCP communication between MCU and PC

Separate threads are used to create the socket and receive the data. Effective handling of this connection and threads is vital to ensure that any interruptions are gracefully recovered from, without the need for manual intervention. Program flow of this functionality is shown in flow charts in Figures 36 to 39.

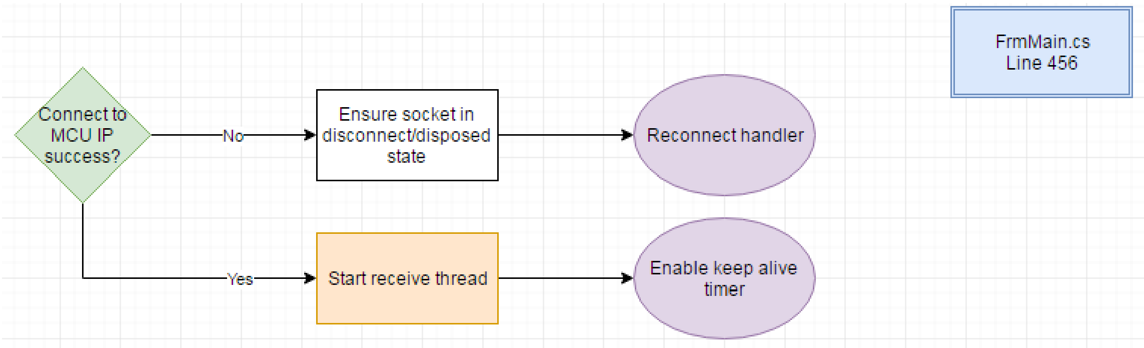

Figure 36 – Program flow of the connect thread in the data processing application

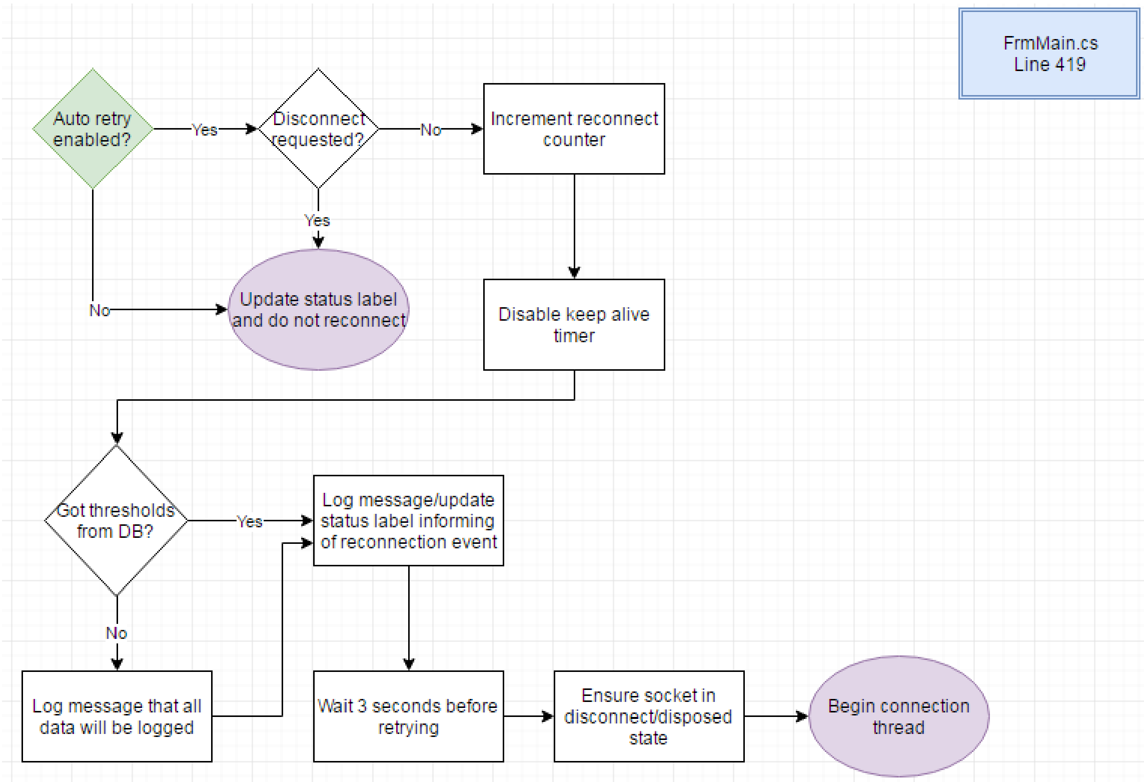

Figure 37 – Program flow of the reconnect handler method in the data processing application

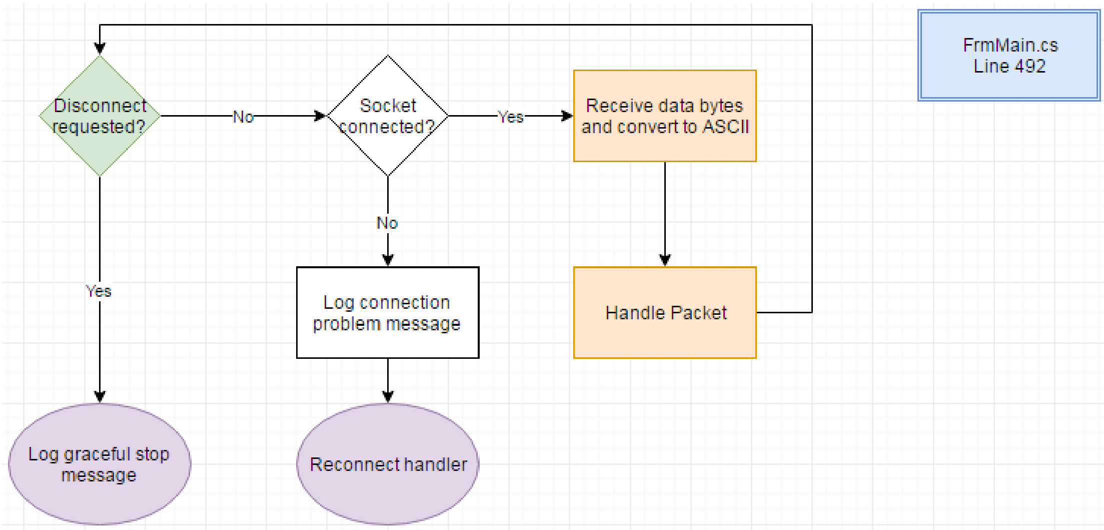

Figure 38 – Program flow of the receive thread in the data processing application

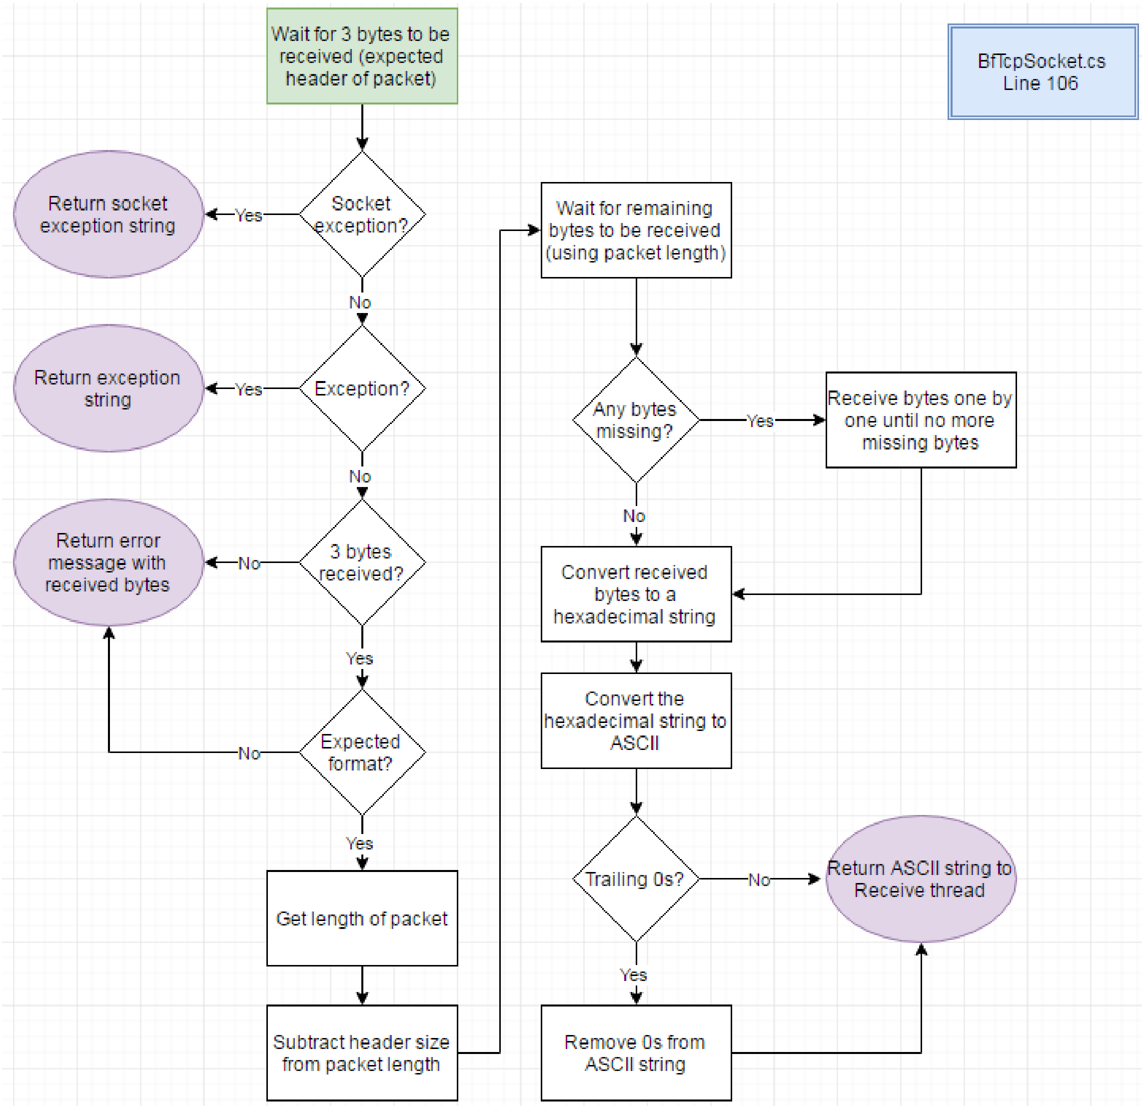

Figure 39 – Program flow of the receive data method in the data processing application

Once a packet has been received, it is split into its component parts and logged into a text file in the same format as on the MCU. This means that both the PC and MCU will have identical copies of the raw data so long as the TCP connection is established. The full flow chart displaying packet handling is shown in Figure 40. A synchronisation mechanism has been developed to handle scenarios where the TCP connection is disconnected, e.g. due to connection problems, and is discussed further in Section 5.2.4.

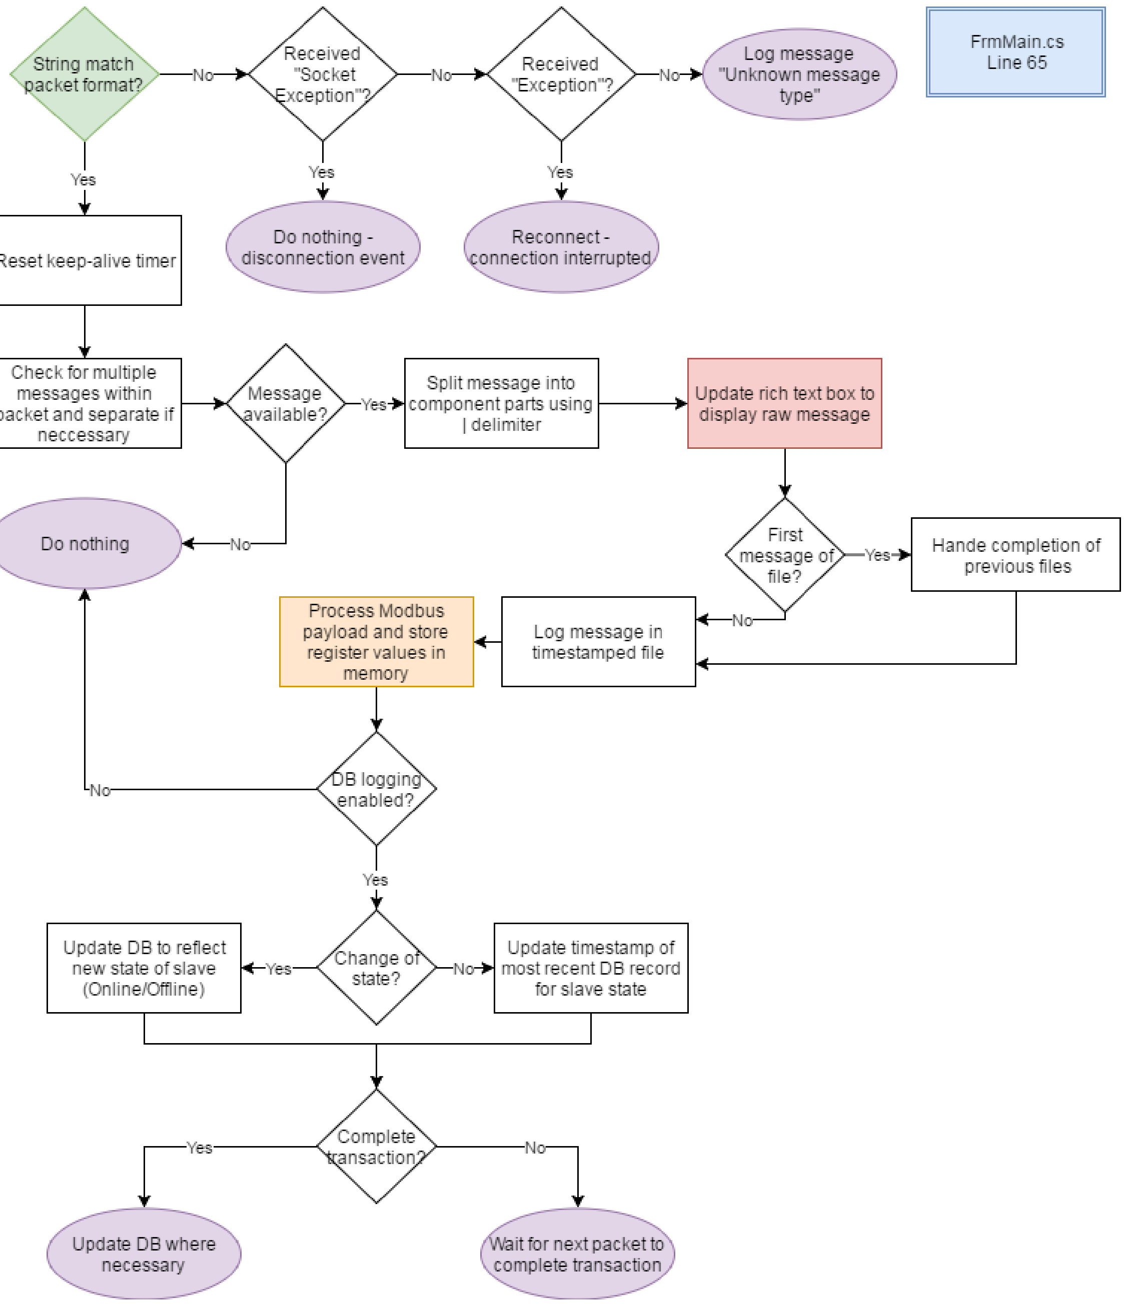

Figure 40 – Program flow of the packet handling method used in the data processing application

After the packet has been logged, the Modbus payload that was split from the packet is sent to a separate processing class. This class performs validation of the message (shown in Figure 41) and decodes the information stored in the data.

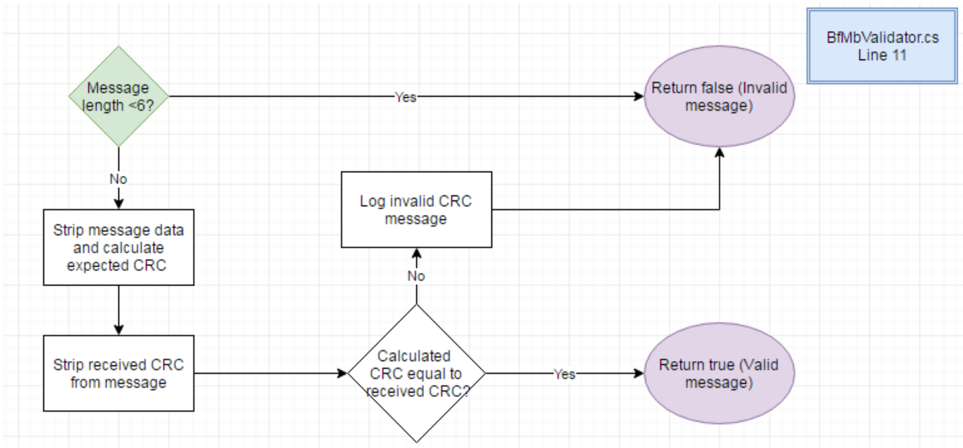

Figure 41 – Program flow of the Modbus validation method used in the data processing application

The first 2 bytes of each Modbus message always define the same parameter. The first byte identifies the slave address and the second identifies the function code. Subsequent bytes depend upon the function code and whether the message is a query or a response. Each query is followed by a response from the addressed slave and can be referred to as a ‘transaction’. Table 19 breaks down a valid transaction between a master and slave device, showing the meaning of each byte.

Table 19 – Breakdown of an example Modbus transaction between a master and slave

<table>
<tr><th colspan="8">Query</th></tr>
<tr><td>01</td><td>03</td><td>00</td><td>00</td><td>00</td><td>02</td><td>44</td><td>06</td></tr>
<tr><td>Slave address: 1</td><td>Function code: 3 [Read holding registers]</td><td colspan="2">Address of first register to read: 40001 [High byte | Low byte]

(Holding registers start at address 40001)</td><td colspan="2">Number of registers required: 2</td><td colspan="2">CRC [High byte | Low byte]</td></tr>
</table>

<table>
<tr><th colspan="9">Response</th></tr>
<tr><td>01</td><td>03</td><td>04</td><td>00</td><td>65</td><td>FF</td><td>FF</td><td>43</td><td>3A</td></tr>
<tr><td>Slave address: 1</td><td>Function code: 3 [Read holding registers]</td><td>Number of data bytes to follow: 4</td><td colspan="2">Value of register 40001: 101</td><td colspan="2">Value of register 40002: 65535</td><td colspan="2">CRC [High byte | Low byte]</td></tr>
</table>

A mechanism has been implemented to ensure transactions are valid, negating any risk of mismatch between query and responses as shown in Figure 42.

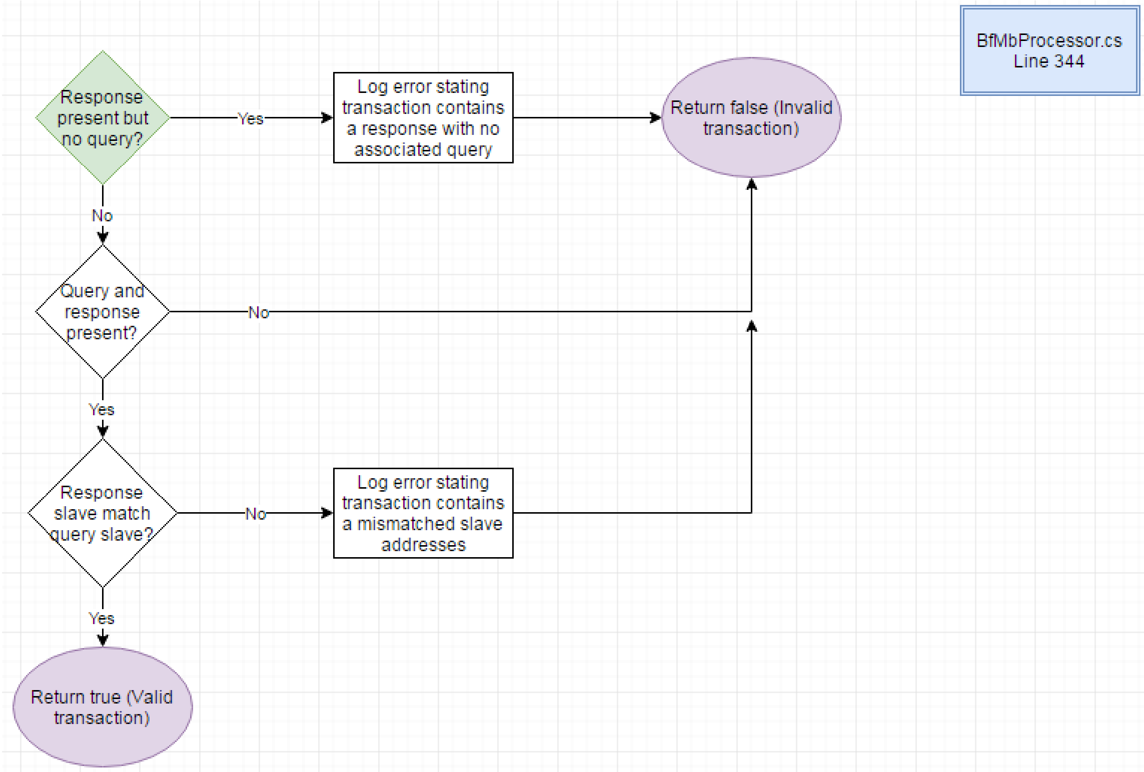

Figure 42 – Program flow of the transaction validation method used in the data processing application

A complete and valid transaction means that data from a slave is ready for processing. For ‘read holding register’ functions, this involves saving the value for each requested holding register in memory. This is achieved using a key value pair dictionary that is defined within the system topology class as shown in Figure 43.

```
Dictionary<Tuple<int, int>, Tuple<int, bool>> register { get; set; }
```

Figure 43 – Definition of the register dictionary used in the data processing application

Both key and value of the dictionary use the ‘Tuple’ class. The key uses two integers to identify the slave address and holding register. The value uses an integer and a Boolean to store the value of the registers for the given slave and to flag whether a database update is required.

Further to this, an optional threshold can be set for a register, meaning that the database is only updated with a new value once the threshold has been exceeded. These thresholds are stored in a similar key value pair dictionary, where the key defines the registers and slave, and the value is the threshold value as shown in Figure 44. Program flow of this functionality is shown in Figure 45.

```
Dictionary<Tuple<int, int>, int> thresholds { get; set; }
```

Figure 44 – Definition of the thresholds dictionary used in the data processing application

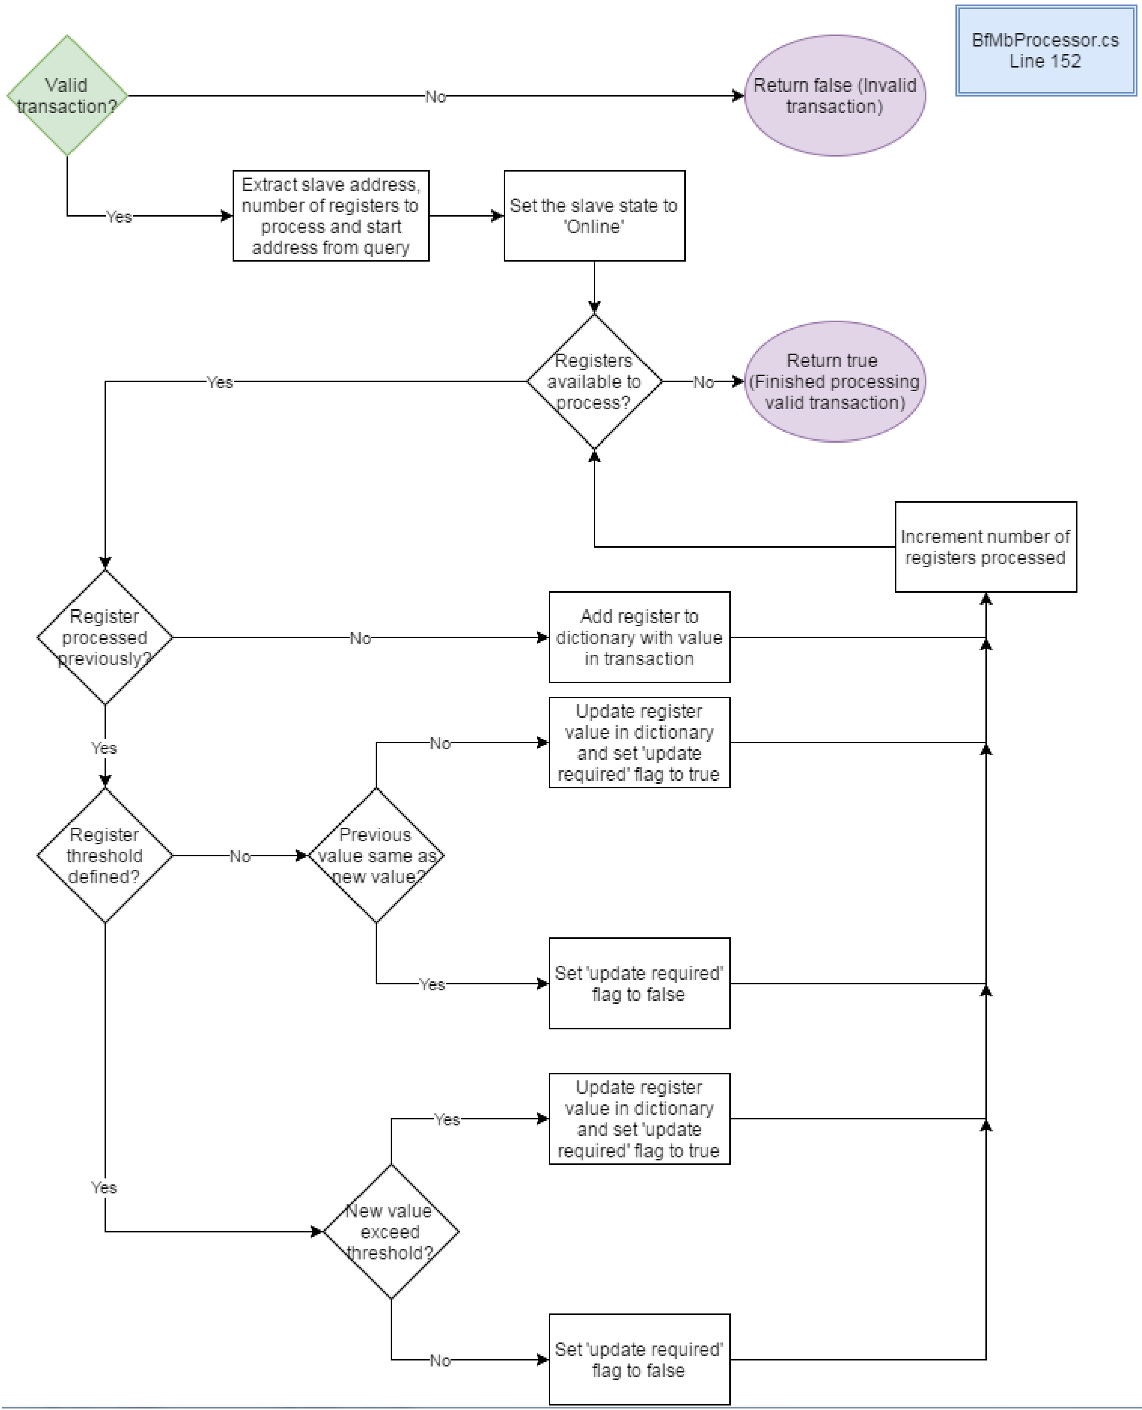

Figure 45 – Program flow of updating the dictionary used to store register values in the data processing application

The finished solution is presented using a ‘Windows Forms’ based interface as shown in Figure 46. This interface enables configuration of system parameters, such as the IP address of the MCU, database IP and credentials and logging depth. By default, the system is configured to automatically connect and reconnect but these options can also be customised. There is also functionality for displaying live data as it is received to validate the system is working as expected. A memory viewer is also present that displays live status of slaves and register values as well as error counters.

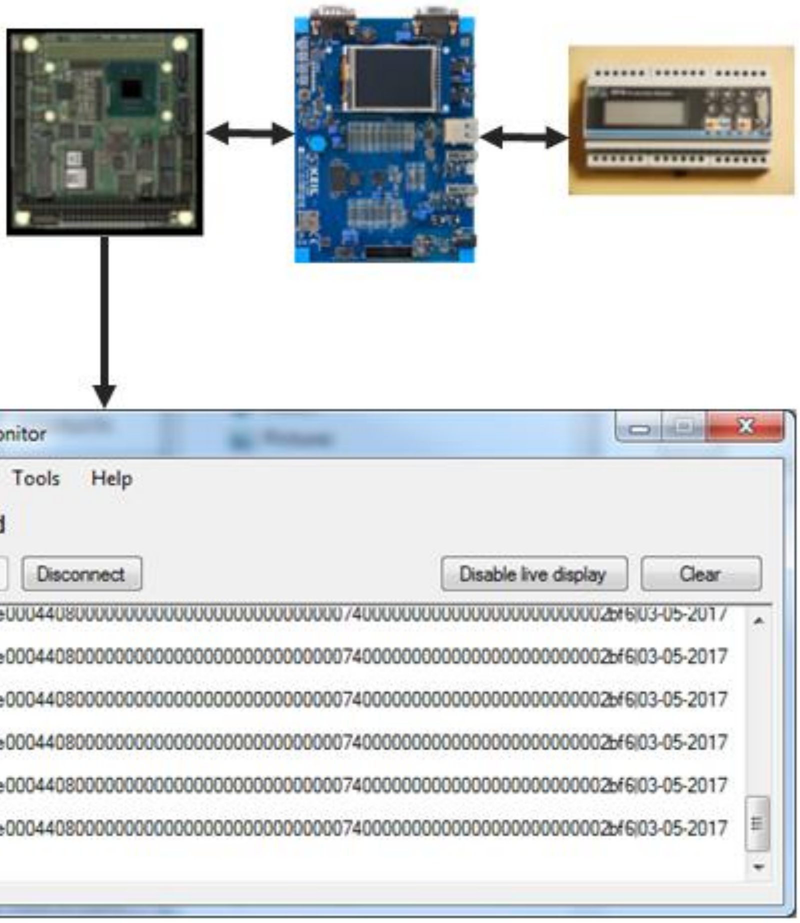

Figure 46 – Data processing application presented using a Windows Forms based interface

5.2.3 Data Synchronisation

In scenarios where the TCP connection between PC and MCU is not established, there will be no data available for the PC to process. The raw data on the PC will be out of sync with data that is available on the MCU and must be processed once the connection is re-established. The developed mechanism ensures that any data that is not present on the PC is downloaded via FTP once the TCP socket connection is re-established.

This is achieved by having a synchronisation log file on the MCU that keeps track of any files that have been written to when the socket was not established. Once the connection is re-established the PC can download this log file, analyse it and download the raw files it lists. The PC also keeps its own version of this synchronisation log file to ensure that future disconnection events do not trigger a complete re-download of files listed.

Processing of synced data uses the same classes as live data, except a distinct object of the ‘BfMbTopology’ type is used. This enables synced data to be processed and written to the database at the same time as new data is being received over the TCP socket.

5.3 Historical Data Storage

Once the data has been processed, it can be written to a database and stored in a structured way such that it can be manipulated and interpreted for display to an end user. Structured storage can be achieved in a variety of ways using widely available third-party platforms. SQL and NoSQL are examples of widely used solutions. SQL solutions are used when the data types being stored are already known and integrity of this data is a major concern. NoSQL on the other hand is a newer technology which lends itself more to data whose type is fluid and subject to future change. Such change would be very difficult to accommodate in an SQL solution because the schema is fixed when the database is created.

For this solution, an SQL database is most suitable since the data lends itself to a relational model. The values recorded from the equipment are of known types and will never change. Furthermore, the requirement for delivering accurate information to the end user means that enforcing data integrity is vital. In SQL databases, this is achieved using primary and foreign keys.

5.3.1 Database Platform Selection

Many third-party SQL solutions are available such as MySQL, Microsoft SQL (MSSQL), OracaleDB, SQLLite and PostgreSQL. Each has their own advantages and disadvantages. MSSQL, for example, is a well-supported platform that offers seamless integration with the Visual Studio development suite, simplifying implementation. A disadvantage is that it requires an expensive licence to take advantage of the fully featured offering.

PostgreSQL is an open source alternative that is widely used by big name internet companies such as Cisco and Apple. Advantages of this platform are that it is free to use and its wide use gives confidence that it the platform is robust. Disadvantages include increased complexity of setting up and more difficult integration with development suites such as Visual Studio.

PostgreSQL has been selected as the database to use for this prototype given its cost effectiveness and proven reliability under demanding heavy traffic conditions with large volumes of data (Gilmore & Treat, 2006). Table 20 lists the operating limits of the platform as per the official documentation. These capabilities far surpass the requirements for this project, which has relatively few tables and will have very low traffic when compared to Internet-based projects that use the platform.

Table 20 – Operating limits of the Postgres database platform

Limit	Value
Maximum Database Size	Unlimited
Maximum Table Size	32 TB
Maximum Row Size	1.6 TB
Maximum Field Size	1 GB
Maximum Rows per Table	Unlimited
Maximum Columns per Table	250 - 1600 depending on column types
Maximum Indexes per Table	Unlimited

5.3.2 Database Schema

The schema contains the tables, relationships and functions that make up the database. This includes primary and foreign keys that are fundamental components of any relational database. Each table has a primary key which uniquely identifies each record in the table. An example of this is the ‘valueid’ column in ‘mbxpsvals’ table. This stores an auto incrementing number which uniquely identifies each value.

A foreign key is used to enforce data integrity and can be applied to multiple columns of a table. An example of this is in the ‘slaveaddr’ column of the ‘mbxpsvals’ table which has a foreign key relation with the ‘slaveaddr’ column of the ‘topology’ table. If an insert to the ‘mbxpsvals’ table is attempted for an XPS slave with an address of ‘3’ and there is no matching XPS slave in the topology table, a foreign key violation will occur and the value will not be logged.

5.3.3 Equipment Mapping Tables

The mapping tables are used to describe how equipment is configured and to link numeric identifiers to human readable textual descriptions. These mappings are standardised for all units of the same type, meaning these tables can be populated when the database is created.

Mapping for Modbus function codes to a textual type is stored in the 'mbfunctions' table. A breakdown of the columns is shown in Table 21.

Table 21 – Breakdown of columns used in the mbfunctions database table

Column Name	Description	Data type	Example data	Notes
funccode	Modbus function code	integer	03	Primary Key
functype	Function type	text	Read Holding Registers	

Data concentrator devices collect data from external inputs and outputs (e.g. switches). Mapping to their associated holding register is stored in the 'dchrmapping' table. Table 22 shows the columns that make up this table.

Table 22 – Breakdown of columns used in the dchrmapping database table

Column Name	Description	Data type	Example data	Notes
inputid	Input identifier	text	ai1	Primary Key ai1 = Analogue input 1
port	Port the input is attached to on the data concentrator	text	a	
holdingreg	Holding register that the input is associated with	integer	40016	
boardnum	Number of the board the input is attached to	integer	7	
statusreg	The register that holds the status of the board	integer	40001	

For protection units, holding registers are mapped to a textual parameter in the 'xpshrmapping' table. A breakdown of columns is shown in Table 23.

Table 23 – Breakdown of columns used in the xpshrmapping database table

Column Name	Description	Data type	Example data	Notes
holdingregid	Holding register identifier	integer	40005	Primary Key
parameter	Textual descriptor for this holding register	text	Voltage	
accesstype	Type of access permitted by this register	text	RO	RO = Read only RW = Readable and Writeable

5.3.4 Customer Mapping Tables

Customer specific mappings are unique to each implementation and will be populated when a system is being commissioned. The 'topology' table stores the types of equipment that make up an implementation as shown in Table 24.

Table 24 – Breakdown of columns used in the topology database table

Column Name	Description	Data type	Example data	Notes
slaveaddress	Associated slave address	integer	1	Primary Key
eqtype	Type of equipment	text	xps	
duty	Designation of equipment that the protection unit is connected to	text	Shearer	A shearer is a large piece of machinery responsible for cutting coal from a coalface

If a data concentrator is part of a system, the customer will have their own switches and sensors attached to it. Analogue inputs are mapped in the 'dcanamapping' table with a column breakdown shown in Table 25. Digital inputs use the 'dcdigmapping' table as shown in Table 26. Mapping for multifunction outputs are stored in the 'dcmfmapping' table as shown in Table 27.

Table 25 – Breakdown of columns used in the dcanamapping database table

Column Name	Description	Data type	Example data	Notes
inputid	Input identifier	text	ai1	Primary Key
slaveaddress	Associated slave address of data concentrator	integer	16	
port	Port the input is attached to on the data concentrator	text	a	
designation	Textual descriptor of input	text	Speed proxistor 1	
function	Textual descriptor of function of input	text	Belt speed sensor	
type	Type of input from the following options: Type A – 0.4-2V, Type B – 4-20mA (100R resistor across 0.4-2V input terminals) Type C – pulser Type D – PT100	integer	1	1 = Type A 2 = Type B 3 = Type C 4 = Type D
zeroscale	Lower bound of the scale that the input operates between	integer	0	
fullscale	Upper bound of the scale that the input operates between	integer	200	
units	Unit of measurement associated with the input	text	Hz	
accuracy	Number of decimal points the readings should be rounded to	integer	2	Default value: 2
threshold	How much a value is allowed to change before it is written into the database e.g. ‘20’	integer	20	Default value: 0

Table 26 – Breakdown of columns used in the dcdigmapping database table

Column Name	Description	Data type	Example data	Notes
inputid	Input identifier	text	di1	Primary Key
slaveaddress	Associated slave address of data concentrator	integer	16	
port	Port the input is attached to on the data concentrator	text	a	
designation	Textual descriptor of input	text	Conveyor remote start	
diodefunction	Textual descriptor for the function that occurs when in diode state	text	start	
description	Textual descriptor for further details	text		
switch	Textual descriptor for the type of switch	text	Green n/o push button	
tristate	Textual descriptor for the third state with 3-way switches	text	remote control	

Table 27 – Breakdown of columns used in the dcmfmapping database table

Column Name	Description	Data type	Example data	Notes
outputid	Output identifier	text	do1	Primary Key
slaveaddress	Associated slave address of data concentrator	integer	16	
port	Port the input is attached to on the data concentrator	text	a	
designation	Textual descriptor of output	text	Conveyor running	
function	Textual descriptor for the function this output provides	text	Closed when belt detected above user-set speed. Drives In-Bye sequence.	

5.3.5 Equipment Values Tables

When the system is live, the data processing application updates the specific values table whenever new data is available. The online/offline status of each asset is stored in the ‘mbslavestatus’ table. A breakdown of columns is shown in Table 28.

Table 28 – Breakdown of columns used in the mbslavestatus database table

Column Name	Description	Data type	Example data	Notes
statusid	Unique identifier for this status entry	serial	1	Primary Key, auto increments
slaveaddress	Associated slave address	integer	2	
online	Whether or not the slave is online e.g. ‘True’	boolean	true	
timestamp	Time this update occurred	timestamp without time zone	2016-09-23 12:10:42	Time from MCU (GMT)

The data concentrator stores the online/offline status of the boards that are attached to the unit in status registers 40001 and 40002. Table 29 shows a breakdown of the columns.

Table 29 – Breakdown of columns used in the dcstatus database table

Column Name	Description	Data type	Example data	Notes
valueid	Unique identifier for this value entry	serial	1	Primary Key, auto increments
slaveaddress	Associated slave address	integer	16	Foreign Key (with topology table)
funccode	Modbus function code	integer	3	Foreign Key (with mbfunctions table)
holdingreg	Holding register associated with this value	integer	40001	
value	Value of the holding register	integer	6087	
timestamp	Time this update occurred	timestamp without time zone	2016-09-23 12:10:42	Time from MCU (GMT)

Readings for different types of equipment have their own tables that follow the same structure but have differing foreign keys to enforce data integrity. Values captured from a data concentrator are stored in the ‘dcanavals’, ‘dcdigvals’ and ‘dcmfvals’ tables.

Tables 30, 31 and 32 show a breakdown of the columns for these. Values from a protection unit are stored in the 'mbxpsvals' table as shown in Table 33.

Table 30 – Breakdown of columns used in the dcanavals database table

Column Name	Description	Data type	Example data	Notes
valueid	Unique identifier for this value entry	serial	1	Primary Key, auto increments
slaveaddress	Associated slave address	integer	16	Foreign Key (with topology table)
funccode	Modbus function code	integer	3	Foreign Key (with mbfunctions table)
holdingreg	Holding register associated with this value	integer	40016	Foreign Key (with dchrmapping table)
value	Value of the holding register	integer	30	
timestamp	Time this update occurred	timestamp without time zone	2016-09-23 12:10:42	Time from MCU (GMT)

Table 31 – Breakdown of columns used in the dcdigvals database table

Column Name	Description	Data type	Example data	Notes
valueid	Unique identifier for this value entry	serial	1	Primary Key, auto increments
slaveaddress	Associated slave address	integer	16	Foreign Key (with topology table)
funccode	Modbus function code	integer	3	Foreign Key (with mbfunctions table)
holdingreg	Holding register associated with this value	integer	40004	Foreign Key (with dchrmapping table)
value	Value of the holding register	integer	21845	
timestamp	Time this update occurred	timestamp without time zone	2016-09-23 12:10:42	Time from MCU (GMT)

Table 32 – Breakdown of columns used in the dcmfvals database table

Column Name	Description	Data type	Example data	Notes
valueid	Unique identifier for this value entry	serial	1	Primary Key, auto increments
slaveaddress	Associated slave address	integer	16	Foreign Key (with topology table)
funccode	Modbus function code	integer	3	Foreign Key (with mbfunctions table)
holdingreg	Holding register associated with this value	integer	40004	Foreign Key (with dchrmapping table)
value	Value of the holding register	integer	21845	
timestamp	Time this update occurred	timestamp without time zone	2016-09-23 12:10:42	Time from MCU (GMT)

Table 33 – Breakdown of columns used in the mbxpsvals database table

Column Name	Description	Data type	Example data	Notes
valueid	Unique identifier for this value entry	serial	1	Primary Key, auto increments
slaveaddress	Associated slave address	integer	16	Foreign Key (with topology table)
funccode	Modbus function code	integer	3	Foreign Key (with mbfunctions table)
holdingreg	Holding register associated with this value	integer	40009	Foreign Key (with xpshrmapping table)
value	Value of the holding register	integer	1730	
timestamp	Time this update occurred	timestamp without time zone	2016-09-23 12:10:42	Time from MCU (GMT)

5.3.6 Example Queries

Queries are used to extract pertinent information from the database. An example of this is displaying XPS values for a single day, e.g. for a trend view graph on the web front end.

```
SELECT     *
FROM       bfdiag.mbxpsvals
WHERE      TIMESTAMP > '2016-09-21 00:00:00'
AND        TIMESTAMP < '2016-09-22 00:00:00';
```

Another example is to display the textual parameter mapping of a value instead of the holding register number, e.g. register 40005 should show the parameter 'Voltage'.

```
SELECT       *
FROM         bfdiag.mbxpsvals
INNER JOIN bfdiag.xpshrmapping
ON           mbxpsvals.holdingreg = xpshrmapping.holdingregid
ORDER BY     holdingreg ASC;
```

This is a more complex query since values are stored in one table and textual mappings of registers are stored in another. This requires the 'INNER JOIN' operation which works on multiple tables with the clause given after the 'ON' operator. Figure 47 shows how this applies to the given statement.

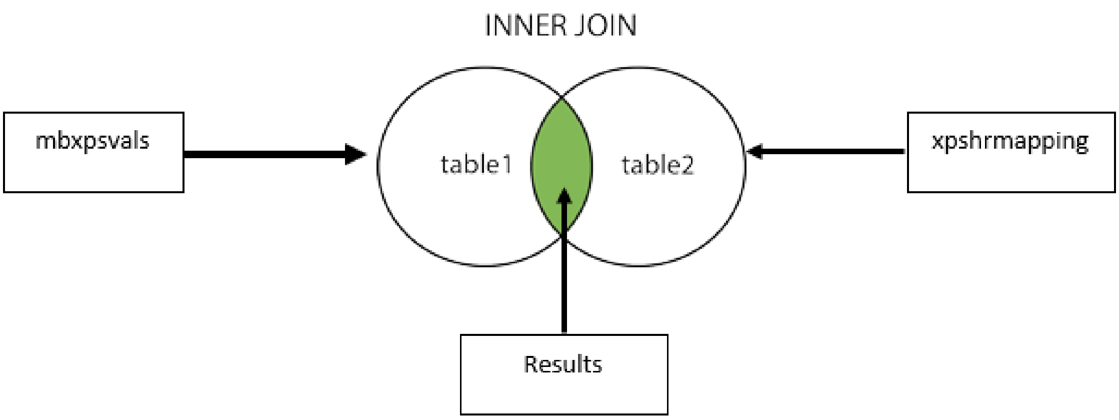

Figure 47 – Example of an inner join from the mbxpsvals and xpshrmapping database tables

5.3.7 Example Data

Figure 48 shows an example of the data captured from a protection unit whose reading for the voltage parameter is decreasing over time.

Edit Data - PostgreSQL 9.5 (localhost:5432) - modbusDB - bfdiag.xpsvals

File Edit View Tools Help

100 rows

	valueid [PK] serial	slaveaddress integer	funccode integer	holdingreg integer	regvalue integer	timestamp timestamp without time zone
35	14501	1	3	40005	1171	2017-04-06 13:52:41
36	14502	1	3	40005	1170	2017-04-06 13:52:45
37	14503	1	3	40005	1172	2017-04-06 13:52:51
38	14504	1	3	40005	1175	2017-04-06 13:52:54
39	14505	1	3	40005	1174	2017-04-06 13:52:58
40	14506	1	3	40005	1173	2017-04-06 13:53:01
41	14507	1	3	40005	1165	2017-04-06 13:53:04
42	14508	1	3	40005	1135	2017-04-06 13:53:07
43	14509	1	3	40005	1134	2017-04-06 13:53:10
44	14510	1	3	40005	1133	2017-04-06 13:53:14
45	14511	1	3	40005	1134	2017-04-06 13:53:17
46	14512	1	3	40005	1132	2017-04-06 13:53:20
47	14513	1	3	40005	1159	2017-04-06 13:53:23
48	14514	1	3	40005	1178	2017-04-06 13:53:27
49	14515	1	3	40005	1176	2017-04-06 13:53:30
50	14516	1	3	40005	1178	2017-04-06 13:53:33
51	14517	1	3	40005	1174	2017-04-06 13:53:36
52	14518	1	3	40005	1159	2017-04-06 13:53:40
53	14519	1	3	40005	1139	2017-04-06 13:53:43
54	14520	1	3	40005	1141	2017-04-06 13:53:46
55	14521	1	3	40005	1142	2017-04-06 13:53:52
56	14522	1	3	40005	1139	2017-04-06 13:53:56
57	14523	1	3	40005	1138	2017-04-06 13:54:02
58	14524	1	3	40005	1137	2017-04-06 13:54:05

100 rows.

Figure 48 – Example data showing a decreasing voltage on a protection unit

5.4 Presentation of Diagnostic Content

The final phase of development work was to implement the graphical user interfaces that the engineer and customer will use to interrogate the equipment. Key requirements include:

- Platform independent
- Dynamic web pages displaying diagnostic content for a service engineer
- Dynamic web pages displaying system health report content for a customer
- Query the database to get the desired results
- Controlled access using a secure log-in
- Large controls for easy navigation using the touch interface of a tablet PC

5.4.1 Service Engineer Interface

The graphical interface for an engineer is a web-based application developed using ASP.NET. The web-based nature means that the content can be accessed from any device that has a web browser, allowing for maximum flexibility. The interface has been designed to be easily navigated using a tablet PC but can just as easily be viewed on a smartphone, laptop, Mac or desktop PC.

The ASP.NET platform provides different patterns for implementing dynamic web pages such as Web Forms and MVC. Web Forms have existed since 2003 and enables rapid application development that is suited to small scale applications (Iliya Nedyalkov, 2013). MVC is a modern pattern which delivers improved separation of concerns into layers of data, business logic and content presentation (Dino Esposito, 2011). This style is recommended for large scale solutions and requires more development time during initial implementation. Given the scale of the solution being developed and time constraints, it was decided to opt to follow the Web Forms pattern.

Figure 49 shows an example of the components that make up a Web Forms page. A master template contains mark-up that defines the visual style used throughout the application. Subpages have their own content and user interface elements defined in an '.aspx' file, these pages inherit the style from the master page. Each subpage also has an '.aspx.cs' file which contains programming code that defines the functionality of controls on the page.

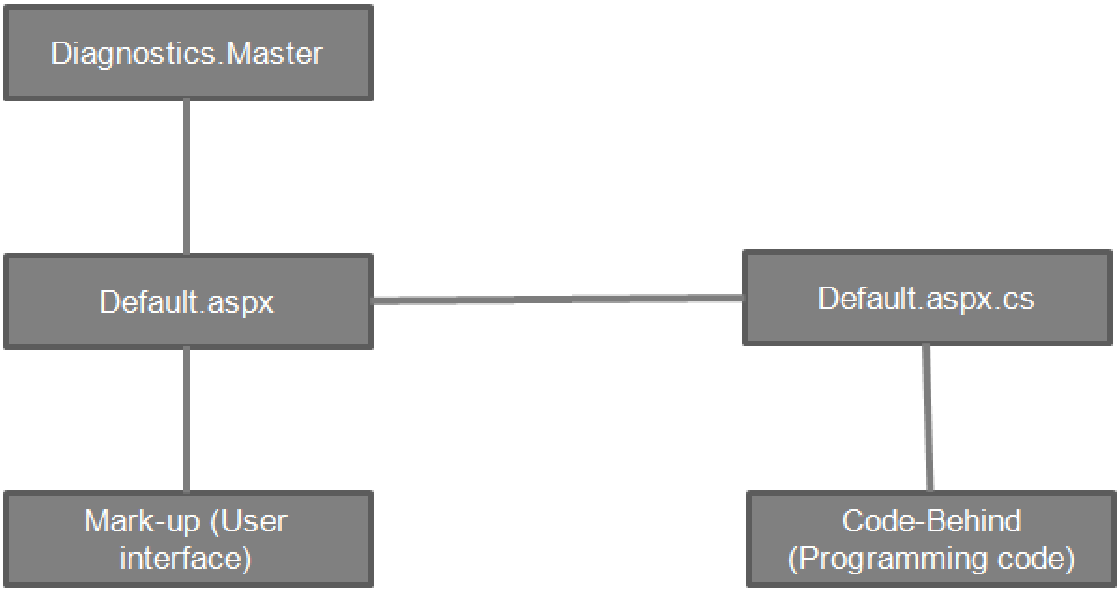

Figure 49 – Components that make up a Web Forms page in ASP.NET

The proof of concept solution focuses on delivering information at a level that an engineer would understand. In contrast, a customer facing interface would need to be simplified with more attention given to the visual aesthetics of the page design. Table 34 shows a breakdown of the pages developed. Figure 50 displays a hierarchical view of how an engineer would navigate around the system.

Table 34 – Breakdown of pages developed for the engineer graphical interface

Page	Description	Purpose
Default.aspx	System Overview	Display a list showing the status of all equipment attached to the system being monitored
About.aspx	About page	Display information related to the organisation, customer and software version
AnalyzeXps.aspx	Register data for selected XPS device	Display a list showing live readings from the selected XPS slave. Registers are displayed as the textual mapping defined in the database and a timestamp is given for the most recent reading
AnalyzeDc.aspx	Register data for selected data concentrator device	Display a list showing the status of I/O boards
AnalyzeVsd.aspx	Register data for selected VSD device	Display a list showing live readings from the selected VSD slave. Registers are displayed as the textual mapping defined in the database and a timestamp is given for the most recent reading
AnalogueValues.aspx	Analogue values for a selected I/O board	Display a list showing the status of analogue inputs from the selected I/O board
DigitalValues.aspx	Digital values for a selected I/O board	Display a list showing the status of digital inputs from the selected I/O board
BitFlag.aspx	Trend view of XPS status registers	Display a graph showing the trend of bit flags over time
Eventlog.aspx	Paginated table showing timestamp events	Displays online/offline state changes of a selected asset
Statushistory.aspx	Paginated table showing timestamp events with customisable parameters	Displays state changes of the status registers on a selected asset
TrendsXps.aspx	Trend view of XPS registers	Display a graph showing the trend of a selected XPS register over time and allow the user to customise the timescale
TrendsVsd.aspx	Trend view of VSD registers	Display a graph showing the trend of a selected VSD register over time and allow the user to customise the timescale
TrendsDigital.aspx	Trend view of digital inputs	Display a graph showing the trend of a selected digital input over time and allow the user to customise the timescale
TrendsAnalogue.aspx	Trend view of analogue inputs	Display a graph showing the trend of a selected analogue input over time and allow the user to customise the timescale
Error.aspx	Error page	Error page which can be shown to users instead of the default ASP.NET error page.

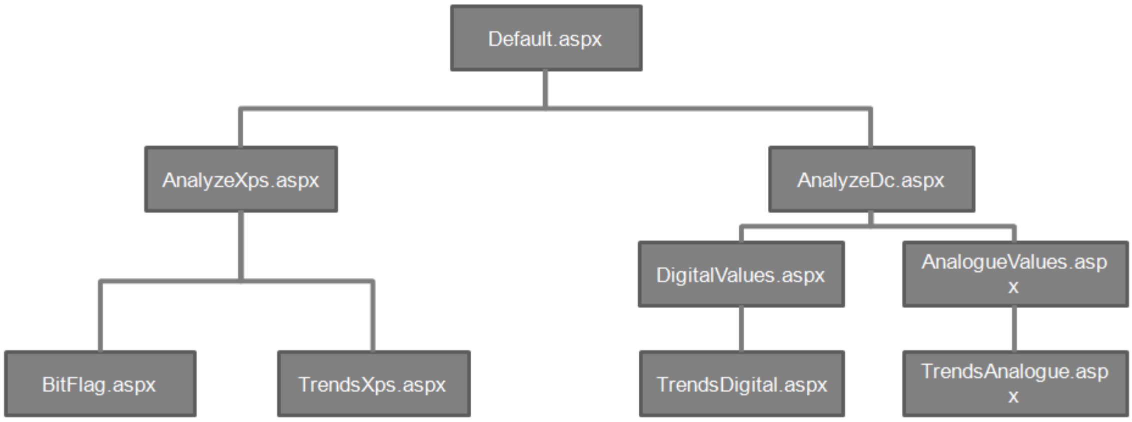

Figure 50 – Example navigation of the engineer graphical interface

5.4.2 System Status

The system status page displays the topology of the equipment at a particular site. Each unit of equipment can be in the states shown in Table 35. Selecting a particular device will display detailed readings.

Table 35 – Equipment states shown on the system status page of the engineer graphical interface

 **{Equipment Mapping}**	**Unknown state** The equipment has not been polled recently and its current state is unknown
 **{Equipment Mapping}**	**Fault state** The equipment has a fault flag active on one of the status registers
 **{Equipment Mapping}**	**Offline state** The equipment is not replying to polls
 **{Equipment Mapping}**	**Online state** The equipment is online and healthy

5.4.3 Detailed Register Readings

Modbus registers are mapped to real world values as defined in the database. For a protection unit, this may include parameters such as voltage and current. For input/output units this includes implementation specific inputs such as temperatures and carbon dioxide detectors. Table 36 shows an example register reading as it is displayed on the service engineer interface. A working example of this can be seen in Figure 51.

Table 36 – Example register reading as displayed on the engineer graphical interface

Voltage: **3300**	**Register detail**
	Displays the name of the real world value mapped to a particular register along with the most recent reading.
Updated: 20/07/2016 15:35:57	The time displayed is when the most recent reading was taken.

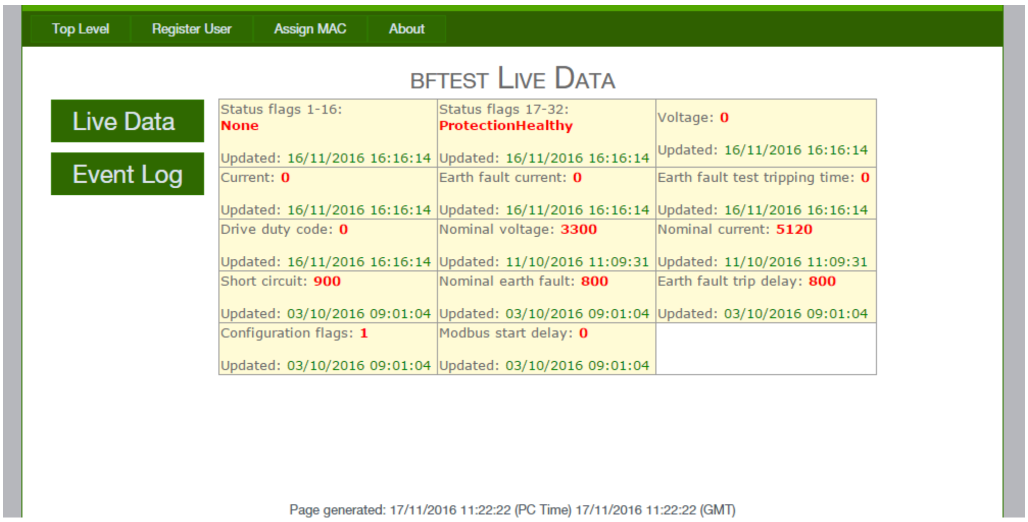

Figure 51 – Working example of register readings from a live system as displayed on the engineer graphical interface

Trend graphs can be drawn from historical data in the database to give a visual representation of system behaviour over time. An example is shown in Figure 52. Changes of status flags can be visualised in a historical log as shown in Figure 53

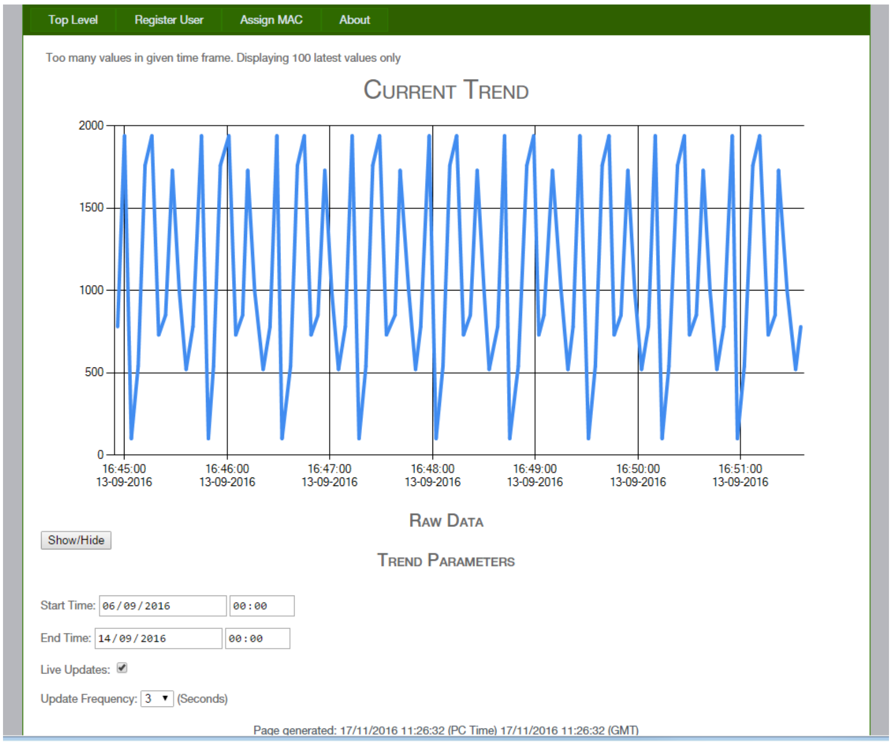

Figure 52 – Working example of a trend view from a live system as displayed on the engineer graphical interface

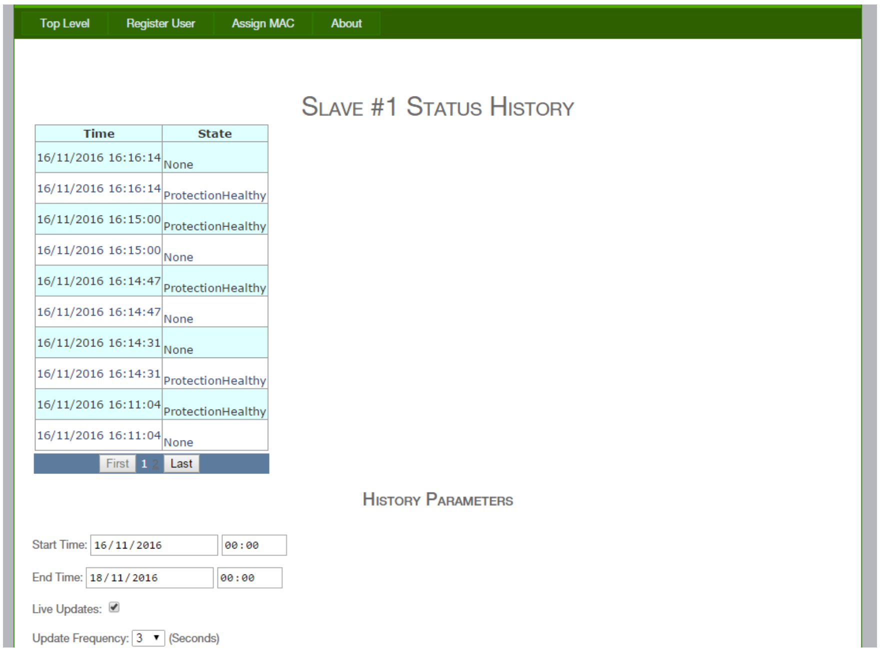

Figure 53 – Working example of status history from a live system as displayed on the engineer graphical interface

5.4.4 Customer Interface

The customer facing interface requires a professional 'look and feel' and thus greater emphasis was placed on the design of the pages. This type of user also adds a number of constraints to the design, such as ensuring that the information delivered is aimed at users who do not have a technical understanding of the components that make up the system. This was tackled by opting for a health report style interface which provides a simplified view of how the system is operating, allowing users to quickly determine where attention is required.

Another constraint is the global customer base, meaning that users may not have a good understanding of the English language. To solve this, a globally recognised traffic light system was used to identify asset status, where green means healthy, amber means warning and red means action is required. This can be customised to display online/offline status of equipment using user-defined colours. Figure 54 shows an example of the design for equipment status over time.

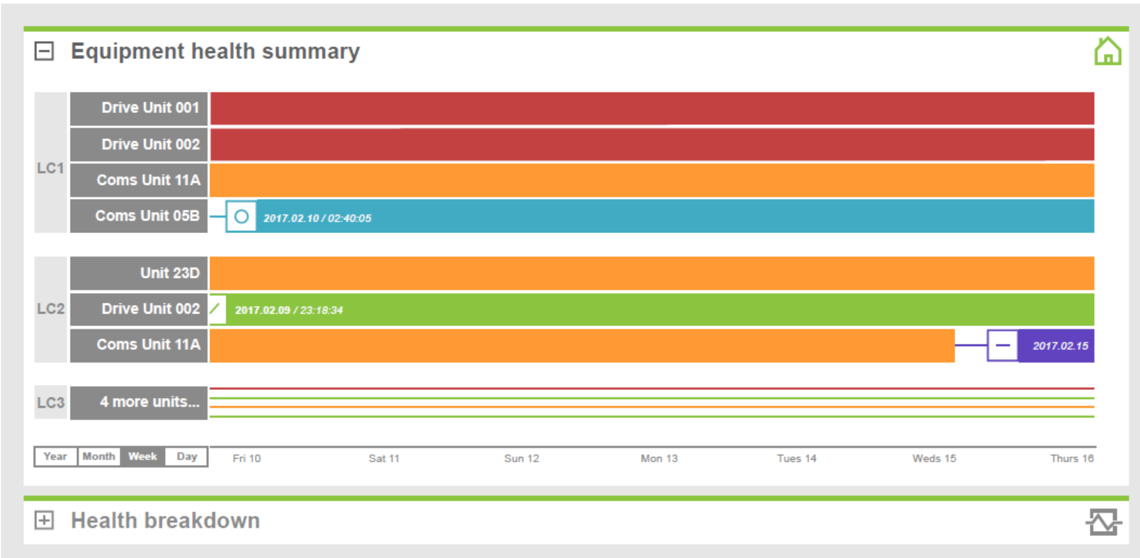

Figure 54 – Mock up of equipment health summary as displayed on the customer graphical interface

Diagnostic information to help the customer understand the meaning of this health status is provided on a status page when a device is selected as shown in Figure 55. This information gives the customer the opportunity to make decisions to improve system performance, for example providing advice for adjusting operating parameters (e.g. voltage) to increase the lifespan of a component. In scenarios where a component is approaching the end of its life, the suggestion can be made for the customer to order a replacement part before failure occurs.

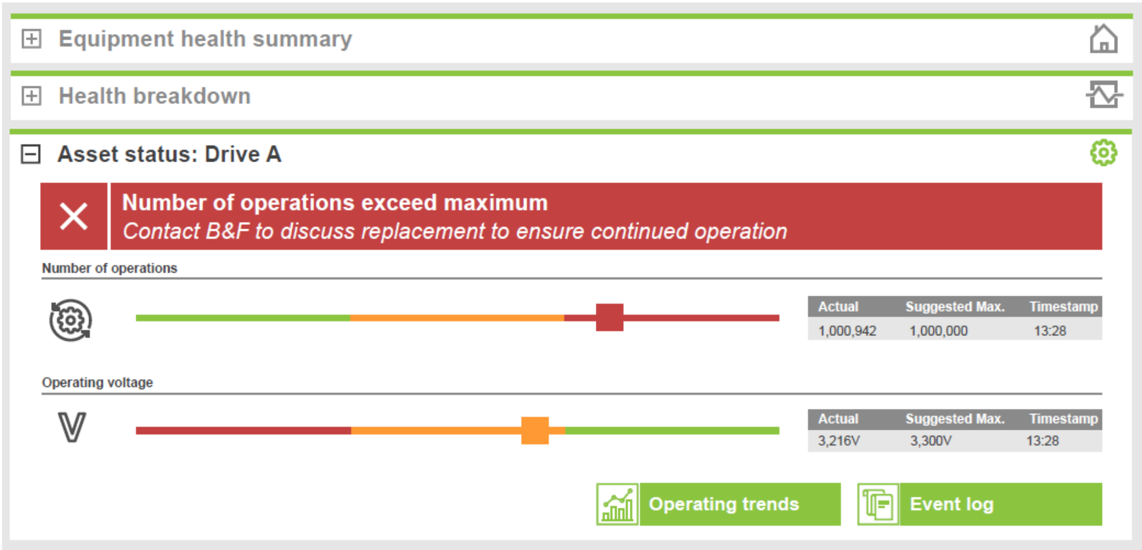

Figure 55 – Mock up of asset status as displayed on the customer graphical interface

Implementation of this requires a layer of business logic to translate the system data into the customer friendly health status identifiers and troubleshooting advice. Two tables were added to the database to support this functionality. The 'healthstatus' table stores the health of all equipment that make up a customer system. Table 37 shows a breakdown of this.

Table 37 – Breakdown of columns used in the healthstatus database table

Column Name	Description	Data type	Example data	Notes
statusid	Unique identifier of this health status	integer	1	Primary Key
slaveaddress	Associated slave address	integer	1	
colour	Health colour to be displayed on the interface	text	Green	Green = Healthy Amber = Warning Red = Action required
statuscode	A code that relates the health status to a cause	integer	2	Foreign key with statuscodeid in 'statuscodes' table
timestamp	Time this update occurred	timestamp without time zone	2016-09-23 12:10:42	Time from MCU (GMT)

Translation of status codes to a human readable text description are stored in the ‘statuscodes’ table. Table 38 shows the breakdown of columns.

Table 38 – Breakdown of columns used in the statuscodes database table

Column Name	Description	Data type	Example data	Notes
statuscodeid	Unique identifier of this status code	integer	1	Primary Key
statuscodedescription	Mapping	text	Irregular usage pattern	

5.4.5 Access Control

A two-step authentication mechanism has been developed to address the need to control access different parts of the system. The first step uses traditional credential based logins implemented using the AspNet.Identity framework. The second step uses hardware identifiers to limit access to particular devices. Having this two-step mechanism adds a further layer of protection in the event that an elevated user has their credentials compromised.

Pages were added to the developed application to handle logins and add new users as shown in Table 39. Figure 56 shows the process followed when a user attempts to login

Table 39 – Pages used for the access control system used in the graphical interfaces

Page	Description
Login.aspx Logout.aspx	Handle system logins
Register.aspx	Allows an admin to add a new user and assign roles
Addmac.aspx	Allows an admin to assign a new MAC address to a user
Denied.aspx	Access denied page which is shown when users attempt to access content they are not permitted to see

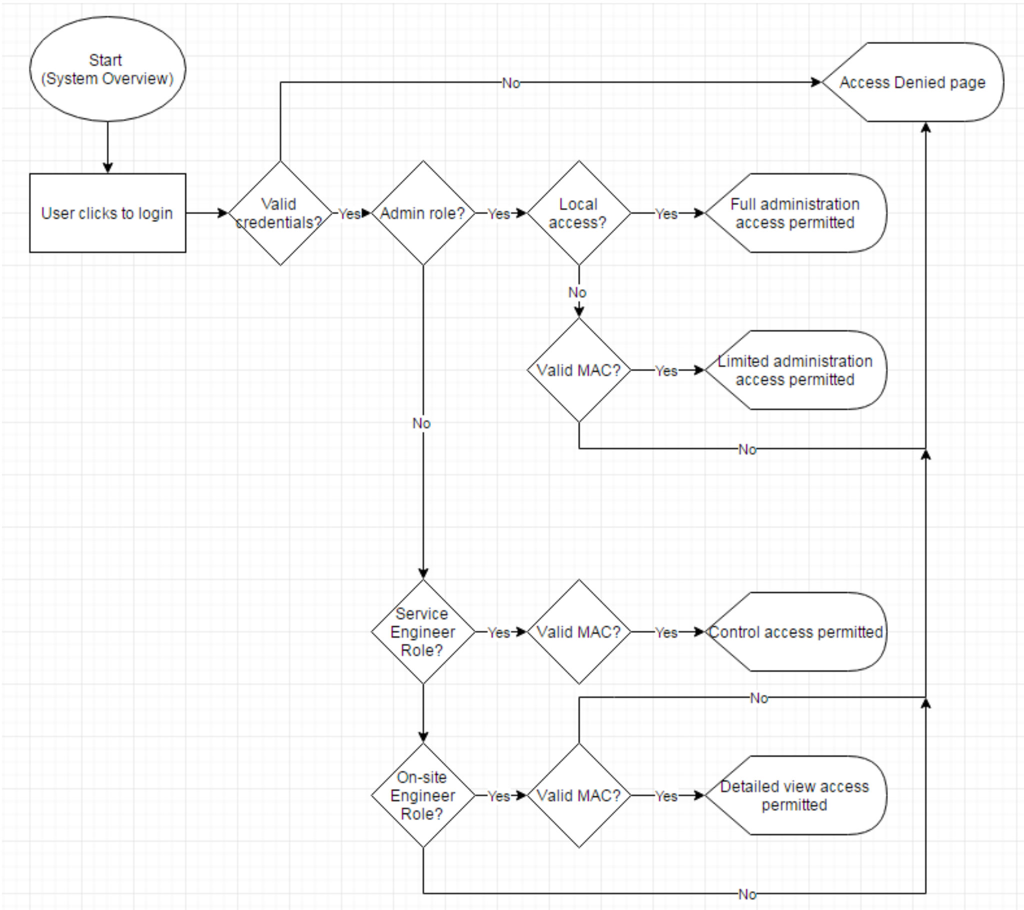

Figure 56 – Program flow of the access control mechanism used in the graphic interfaces

5.5 System Testing

The developed solutions were regularly tested throughout the development process. Tests were carried out with each new revision of software using real world equipment to validate how the system would perform in a production environment. Equipment was monitored continuously overnight to ensure reliability of the results. Test outcomes are documented in Table 40.

Table 40 – System test results

Test	Revision	Test Description	Expected output	Observed output	Notes	Outcome
1	Modbus Monitor PC: 115 Web Interface: 120	Run Modbus Monitor overnight on PC104 monitoring an XPS which changes 5 register values every 3 seconds.	Both applications continue to operation for the duration All XPS value changes are logged to the database Web front-end displays trends from the historical data	Applications ran without failure DB updated as expected Trends displayed very slowly	Slow trends were due to the amount of data in the database and the default timescale set to show an entire day	Modbus Monitor MCU/PC: Pass Web Interface: Partial Pass
2	Modbus Monitor PC: 115 Web Interface: 120	Run Modbus Monitor overnight on PC104 monitoring the real world simulator with 7 XPS's and 1 data concentrator.	As above	Monitor application ran smoothly up until ~200 raw files were created. An exception was then encountered and monitoring stopped.	This was likely due to a similar issue observed during initial tests where having too many files in the same directory caused a delay on the micro controller. Increasing the number of messages from 100 to 1000 stopped this for a single XPS. On the simulator however many more files were created in a single day, meaning the issue was re-encountered. This caused the write delay to reoccur. The delay led to messages being missed, meaning that a scenario may arise where at the end of one file a query for a certain slave may be followed by a response from a different slave at the start of the next file. E.g: 185|1000|8|06030000001045b1|25-08-2016 15-25-01 186|0|37|05032000008000000000000000000000000000000000000000000000000000000000000	Fail

					00000c7ba|25-08-2016 15-25-04 The delay can be fixed by increasing the number of messages per file. A fix also needs to be issued to the monitor to ensure that every response matches the slave address in the query.	
3	Modbus Monitor PC: 123 Modbus Monitor MCU: 122	Run Modbus Monitor with continuous polling at max speed with 1 XPS.	Subfolders should be created within the current date directory. Each subfolder should hold 100 raw text files each.	Monitor application appeared to run smoothly on the PC until 1000 files were created. At this stage the client slowed down indicating a problem.	Attempts to access the MCU via FTP resulted in failure (FTP connected but the directory listing was empty) Moving the SD card to the PC revealed that subfolders 8 onward were corrupt. Allowing Windows to scan and repair this resulted in the folders no longer being present on the SD card, meaning the data has been lost.	Fail
4	Modbus Monitor PC: 126 Web Interface: 120 Modbus Monitor MCU: 127	Re-run test 2	Both applications continue to operate for the duration All XPS value changes are logged to the database Web front-end displays trends from the historical data	Applications ran without failure DB updated as expected Logs contained multiple occurances of “ArgumentOutOfRangeException” Trends displayed very slowly Live status was not shown as expected	The issue observed in test 2 has now been solved. Files are stored in subfolders for each new date, with 100 files per subfolder. Slow trend display still needs to be fixed by adjusting the default time span. The live status not being displayed is due to the time on the MCU following GMT, and the PC is currently on BST. By morning, the time on the MCU was out by 30minutes. (PC: 09:35 BST, MCU: 09:04 GMT). MCU therefore running too fast.	Modbus Monitor PC: Fail Web Interface: Partial Pass Modbus Monitor MCU: Fail
5	Modbus Monitor PC: 130 Web Interface: 129 Modbus Monitor MCU: 127	Re-run test 2	As Above	Applications ran without failure DB updated as expected Logs contained multiple occurances of “Invalid message received” Trends	Invalid message received logs were included in commit 130. Invalid messages include those that do not meet the minimum length requirement of normal Modbus messages and messages that fail to pass a CRC check. This could be due to an issue with the Modbus RTU silent timer on the MCU or maybe an incorrect poll sent from the master.	Modbus Monitor PC: Pass Web Interface: Partial Pass Modbus Monitor MCU: Partial Pass

				displayed very slowly	Appendix III documents the observed output from log files.	
6	Modbus Monitor PC: 130 Web Interface: 129 Modbus Monitor MCU: 127	Re-run test 2	As Above	File system thread on MCU failed DB updated as expected Logs contained multiple occurances of "Invalid message received" Trends displayed very slowly	The MCU stopped logging messages to text files on the SD card. Inserting the SD card into a PC reveals the card was corrupted. Invalid messages in log are as described above.	Modbus Monitor PC: Pass Web Interface: Partial Pass Modbus Monitor MCU: Fail
7	Modbus Monitor PC: 157 Web Interface: 142 Modbus Monitor MCU: 139	Re-run test 2	As Above	File system thread on MCU continued as normal DB updated as expected Trends update in a responsive way	The file system thread no longer crashes; investigation found that the issue could be recreated by physically tapping the evaluation board near the SD card holder, suggesting a hardware problem. Code has been added to ensure the file system is re-initialised if this occurs, preventing future crashes. Trends on the web interface are no longer slow to update.	Modbus Monitor PC: Pass Web Interface: Pass Modbus Monitor MCU: Pass
8	Modbus Monitor MCU: 171	Test the new state machine overnight to ensure correct operation	MCU should log and stream data in the same way it did previously	The MCU logged numerous messages incorrectly, displaying on the PC as bad CRCs, invalid lengths and invalid transactions.	While the new state machine worked as expected for much of the time, there was numerous instances where messages were incorrectly logged. An example of this is the following request for slave 3: The correct transaction should be as follows: 15|222|8|03030000001045e4| 23-02-2017 16-01-01 15|223|37|0303200000800000000000000000000600000000000900000000000000000000000000086da|23-02-2017 16-01-01 One poll however returned the following: 15|866|8|03030000001045e4| 23-02-2017 16-01-31 15|867|36|03200000800000000000000000060000000000009	Modbus Monitor MCU: Fail

					0000000000000000000000000000086da|23-02-2017 16-01-31 The first byte was missing from the start of the message and the PC interpreted this as a bad CRC.	
9	Modbus Monitor MCU: 169	Re-do test 8 with an earlier revision to ensure this issue is not caused by other external factors	As above	The MCU logged all received values correctly	These findings suggest that the state machine is causing the problem.	Modbus Monitor MCU: Pass
10	Modbus Monitor MCU: 177	Re-do test 8 with the stated new revision to determine whether the issue has been fixed	As above	The MCU logged all received values correctly	Problems with the state machine have been fixed.	Modbus Monitor MCU: Pass

These tests demonstrate the nature of issues that arose throughout the development process. Fixes were issued and further testing was performed to ensure correct operation. Extended testing was subsequently performed where the system was monitored continuously for a period of weeks. This final acceptance testing demonstrated that the system was reliable and ready for use in a production environment.

5.5.1 Data sizes

To determine how much space would be consumed by the data, a test was carried out monitoring 6 protection units and 1 data concentrator for a full day. To simulate real world conditions, a protection unit was configured to update five status registers every second. Raw data storage in text files was configured to store 1000 messages per file.

At the end of the 24 hour test, a total of 1553 files had been created across 16 subfolders. Each text file varied in size by between 109kb-112kb. The total size of these files on disk came to 171MB. In contrast, the sizes of the tables in the database were much smaller, as shown in Table 41. This can be attributed to the fact that the raw text files contain every byte of every single poll and response from the attached equipment, whereas the database only contains the register values. Table 42 shows extrapolated sizes for various time frames based on these results

.

Table 41 – Number of values stored in the database for each equipment type over a 24 hour period

Equipment Type	Number of Values	Database Table Size
Protection Unit	126,643	7.6MB
Digital Inputs	50	8.1KB
Analogue Inputs	29,596	1.7MB
Multifunction I/O	86	8.1KB

Table 42 – Extrapolated sizes for set time frames based on the number of values captured in a 24 hour period

Time Frame	Number of Raw Text Files	Size of Raw Text Files	Database Values	Size of Database Tables
Hourly	65	7MB	1,239	0.38MB
Daily	1,553	171MB	29,732	9.31MB
Monthly	48,143	5.3GB	921,692	288.61MB
Yearly	577,716	63.61GB	11,060,304	3.46GB

In terms of disk space, a total of 67GB would be required each year to accommodate these findings. The mechanisms developed to control disk space ensure that there will always be space available for new data, automatically deleting old data on a month by month basis once the storage media is approaching its maximum capacity.

6 Content Delivery

The final phase of the project covers ‘non-invasive’ delivery of content. This requires the use of a wireless access point (WAP) that can enable a tablet PC with Wi-Fi to communicate with the solution.

Market research was performed to identify potential candidates to meet this need as detailed in Appendix IV. Selected units were then evaluated by carrying out tests to assess whether they will be suitable for use in the production environment.

Outcomes of this research found a range of units available with differing specifications, capabilities and connectivity options. ATEX mining approved units were also on the market but were much more expensive at around £2000, in contrast to a cost of around £300 for a non approved unit. Since the production environment can take advantage of the flameproof enclosure this means that a non-approved WAP could be used. The only caveat to this is that to ensure a stable signal would require an antenna located in the IS chamber which would need to be fed from the WAP. Fortunately an approved RF isolator, shown in Figure 57, is available from a supplier called Extronics at a cost of £150. This isolator ensures the output from the WAP to the antenna meets IS requirements.

Figure 57 – Extronics RF Isolator

Two industrial units were selected for evaluation from different manufactures. The option of using a consumer grade WAP was also examined since this would result in a greater cost saving compared to industrial grade units. It should however be noted that consumer units traditionally have a suggested maximum operating temperature of 60°c which is close to a worst case scenario temperature. Industrial grade devices, however, are designed to operate in more extreme ambient temperature ranges.

6.1 Selection Criteria

For the WAP to fulfil project requirements the unit must have the following specification as a minimum:

- Single antenna socket
- Capable of operation in high ambient temperatures
- Access Point (infrastructure) mode
- WPA2 PSK Security
- Under £400

6.2 Antaira APN-210N-T

Figure 58 – Antaira APN-210N-T WAP

First impressions of the Antaira, shown in Figure 58, were positive. The LED display indicators showing signal strength would be useful during commissioning for validating that Wi-Fi is working as expected. The addition of two Ethernet sockets will also give added flexibility when deploying into a wider network. As of October 2016 the cost of this unit was US$359 (approximately £266). A list of pros and cons are detailed in Table 43. Full specification of the unit is detailed in Appendix V.

Table 43 – Pros and cons of the Antaira APN-210N-T WAP

Pros	Cons
• Industrial operating temperatures • DIN rail mountable • Single antenna socket • 2x Ethernet sockets • LED display of Wi-Fi signal strength	• US manufacturer, EU branch based in Poland – increased cost to import.

6.3 B&B AirborneM2M

Figure 59 – B&B Airborne M2M WAP

First impressions of the B&B, shown in Figure 59, were that the compact square form factor will sit physically comfortably in a stack with the PC104 which has a similar square form. The addition of the RS232 ports is however unlikely to be useful within this project. It was also noted that the power connector offered no method of being secured in place. This could result in it becoming loose over time and losing power to the unit. As of October 2016 the cost of the B&B was £362.10.The pros and cons are detailed in Table 44. A full unit specification is listed in Appendix V.

Table 44 – Pros and cons of the B&B AirborneM2M WAP

Pros	Cons
• Industrial operating temperatures • Small form factor would sit well with other components • Single antenna socket • UK based supplier (ADEY Electronics)	• Single Ethernet port • Unnecessary presence of RS232 ports • Poorly designed power connector

6.4 Netgear WAG102

Figure 60 – Netgear WAG102 WAP

First impressions of the Netgear, shown in Figure 60, were that the unit would take up more space than the industrial units and offered difficult power connectivity given that it is a consumer unit. A cable would need to be created to allow connectivity between terminal connections on an industrial power supply to the barrel connector on the unit. Furthermore, the 12V maximum would mean having to use a dedicated supply because the supplies within the FLP output 24V. As of October 2016 the cost of this unit was £35. The pros and cons are detailed in Table 45. A full unit specification is detailed in Appendix V.

Table 45 – Pros and cons of the Netgear WAG102 WAP

Pros	Cons
• Cheap • Readily available from UK based suppliers (e.g. Amazon)	• Consumer product not rated to operate in ambient temperatures over 60°C • Single Ethernet port • Barrel connection for power supply • Requires 12V supply

6.5 Antennas

Three antennas, each with their own strengths and weaknesses, have also been selected. Each antenna will be tested with each access point to assess performance. The most suitable will then be selected for use in the solution.

6.5.1 iAnt212

Figure 61 – iAnt212 Wi-Fi antenna

The iAnt212, shown in Figure 61, is a small antenna that cost £42 as of October 2016. Pros and cons are documented in Table 46. Full specification is detailed in Appendix VI.

Table 46 – Pros and cons of the iAnt212 Wi-Fi antenna

Pros	Cons
• Small • Cheapest option • Less prone to being damaged if mounted externally on the IS chamber	• Very little gain

6.5.2 iAnt215

Figure 62 – iAnt215 Wi-Fi antenna

The iAnt215, shown in Figure 62, is a dome shaped antenna that cost £112 as of October 2016. Pros and cons are documented in Table 47. Full specification and radiation pattern is detailed in Appendix VI.

Table 47 – Pros and cons of the iAnt215 Wi-Fi antenna

Pros	Cons
• Good amount of gain at 2.4Ghz • Form factor enables simplified external mounting on the IS chamber • Highest operating temperature • Dome shape means the antenna is less likely to be damaged when mounted externally	• More expensive than the iAnt212

6.5.3 iAnt216

Figure 63 – iAnt216 Wi-Fi antenna

The iAnt216 is a large antenna that cost £141 as of October 2016. Pros and cons are documented in Table 48. Full specification and radiation pattern is detailed in Appendix VI.

Table 48 – Pros and cons of the iAnt216 Wi-Fi antenna

Pros	**Cons**
• Good amount of gain	• Most of gain is achieved at 5GHz which is not used in this solution • Prone to being damaged if mounted externally on IS chamber • Most expensive

6.6 Connectivity

6.6.1 Physical Connectivity

The WAP and RF isolator will be located within the flameproof enclosure alongside the PC and MCU. This satisfies safety standards as demonstrated in Figure 64 where a brick wall performs the same role as the flameproof enclosure.

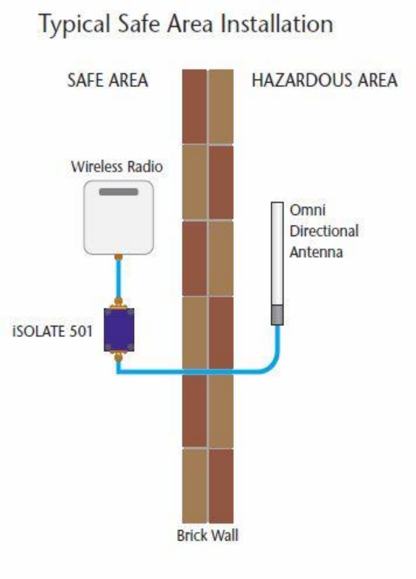

Figure 64 – Safe area installation using a non-approved WAP with the RF isolator

The option of using the antenna mounted inside the IS chamber would bring the added benefit that the antenna is less prone to physical damage when it is on-site. If this is not possible due to space constraints, the antenna can be mounted externally on the IS chamber which will give the benefit of reduced signal attenuation. Both options are shown in Figure 65.

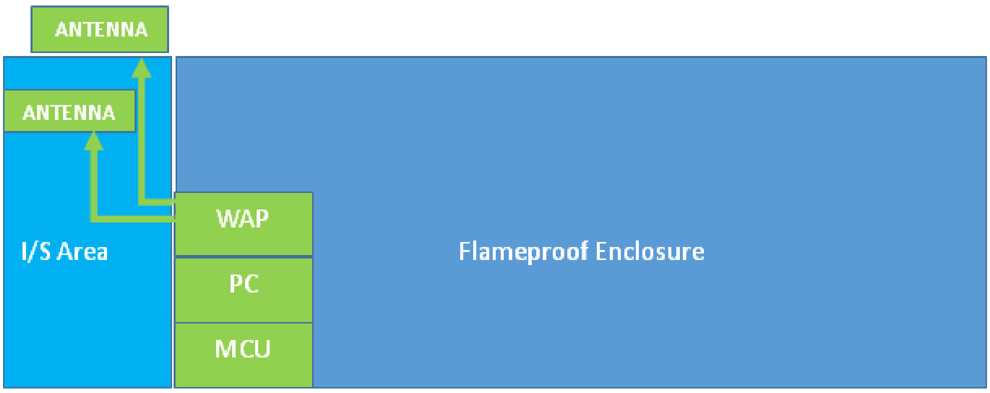

Figure 65 – Deployment options for the physical hardware of the system

6.6.2 Network Connectivity

Two options of network connectivity within the flameproof enclosure have been identified. Figure 66 demonstrates connectivity using a managed switch.

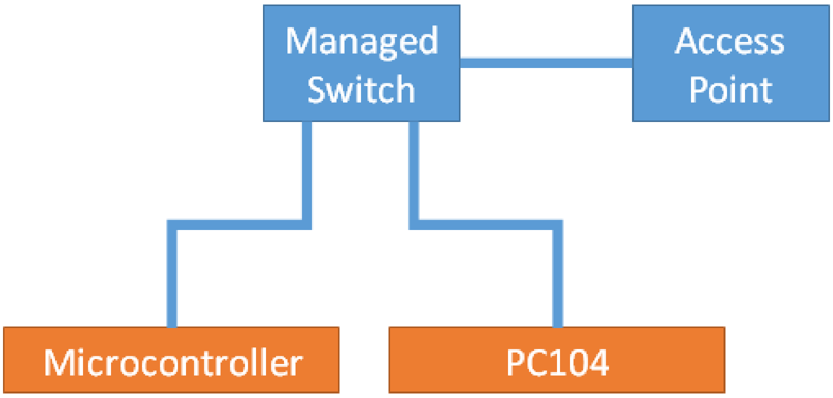

Figure 66 – Network connectivity using a managed switch

This option enables connectivity to other devices on a wider network. A standard installation like this, however, means there is no isolation of the diagnostics equipment. This would be undesirable in a production environment due to security risks. One way to solve this would be to implement a VLAN that the microcontroller and PC104 could be assigned to. Another method would be to have a direct connection from the MCU to the PC104 using the secondary Ethernet port on the PC104 as shown in Figure 67.

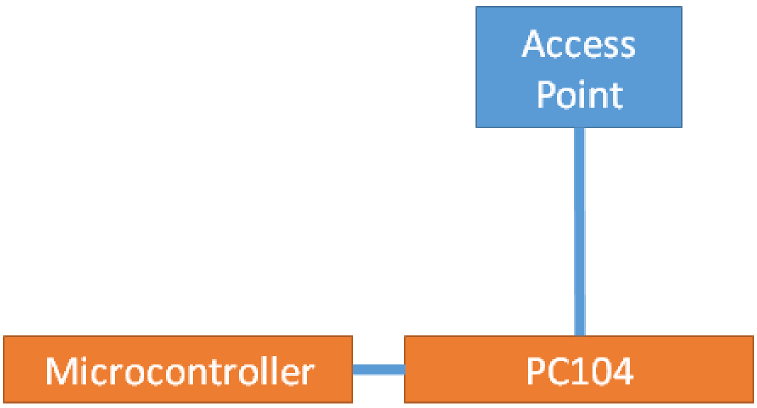

Figure 67 – Network connectivity using a direct connection between PC and MCU

This option ensures the diagnostics content remains isolated on its own network that is only accessible via Wi-Fi. It also ensures that Wi-Fi users cannot access the data acquisition network, thus preventing scenarios of tampering with the sensitive operation of the MCU.

6.7 Wireless Testing

The scope of wireless testing is shown in Table 49.

Table 49 – Scope of wireless testing

In Scope	Out of Scope
• Signal Strength • Interference • Packet Loss • Throughput • Range	• Temperature testing • Humidity testing

6.7.1 Locations/Configurations

Each access point will be tested at the following locations with the given configurations. Each antenna will be used for each test.

- Office – Best case scenario
  - With RF isolator
  - Without RF isolator

- Test Bay – WAP located inside FLP enclosure. Antenna located inside IS chamber. VSD not running
  - Without RF isolator
  - With RF isolator

- Test Bay – WAP located inside FLP enclosure. Antenna mounted inside IS chamber. VSD running at 20% and 80% speed. Production (worst case) scenario.
  - With RF isolator

The dome iAnt215 antenna has a simplified method of mounting which makes it possible to test how the unit would perform if mounted outside the IS chamber. The mounted antenna is shown in Figure 68. A test laptop running Windows 7 located close to the FLP will be used to run all tests.

Testing in a production scenario will be performed with the equipment inside the FLP. Each antenna will be tested with the RF isolator attached and placed within the IS chamber. Readings will then be taken with the VSD powered on at speeds of 20% and 80% which results in increased levels of electromagnetic interference (EMI).

Figure 68 – Flameproof enclosure with iAnt215 mounted externally on the IS chamber

6.7.2 Test Procedure

1. Configure the WAP to act as a DHCP server to offer wireless and Ethernet clients an IP address on the same network. If this is not possible, use static IP addresses on the same subnet.
2. Take a screenshot of Wi-Spy channel activity at the test location
3. Take a reading of the signal strength with the specified configuration using Wi-Spy
4. Ping the PC104 50 times from a laptop connected via Wi-Fi
    a. Command: ping 10.0.0.1 -n 50
5. Run iperf for 1 minute with the PC104 set as the server and the Wi-Fi connected laptop as the client.
    a. Server command: iperf3 -s
    b. Client command: iperf3 -c 10.0.0.1 -t 60
6. Use the best antenna for the production scenario to do a range test with a handheld tablet PC to determine how far the Wi-Fi signal can travel.

6.7.3 Test Location RF Conditions

The least congested Wi-Fi channel will be used in all tests to ensure fairness. A congested channel will impact performance in an unpredictable way due to the variance of traffic on the channel at that moment in time. Given that the production environment is underground it is unlikely that there will be other Wi-Fi enabled equipment present.

In the event that other equipment is present, a site-survey could be performed prior to deploying the WAP. The WAP could then be configured to use the least congested channel.

Figure 69 shows the RF conditions in an office location. Channel 11 is most congested, with Guest and Enterprise APs utilising this channel. Channel 6 is least congested. The test bay location conditions are shown in Figure 70. Channels 6 and 1 were both similarly least congested. As with the office location, channel 11 is most congested.

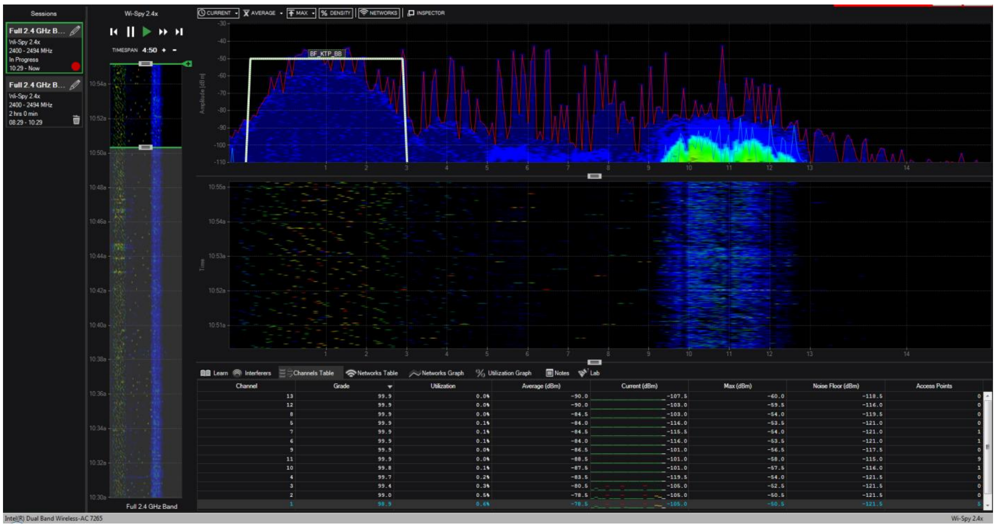

Figure 69 – RF conditions in the office location viewed using Wi-Spy

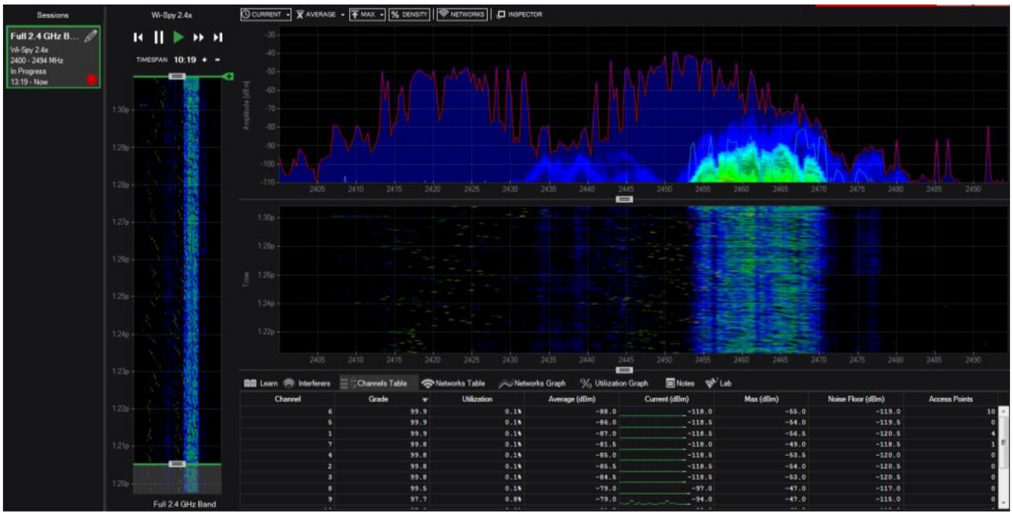

Figure 70 – RF conditions in the test bay location viewed using Wi-Spy

6.8 Wireless Test Results

Tabulated results for each test using the Antaira are shown in Appendix VII, for the B&B in Appendix VIII and for the Netgear in Appendix IX.

6.8.1 Signal Strength Results

Figure 71 shows results from the best case scenario of the office location when using the B&B and Antaira. Figure 72 shows best case results of the Netgear alongside the Antaira. Tests were performed with each antenna both with and without the RF isolator. Worst case results from the B&B vs Antaira are shown in Figure 73. Netgear results alongside the Antaira as shown in Figure 74.

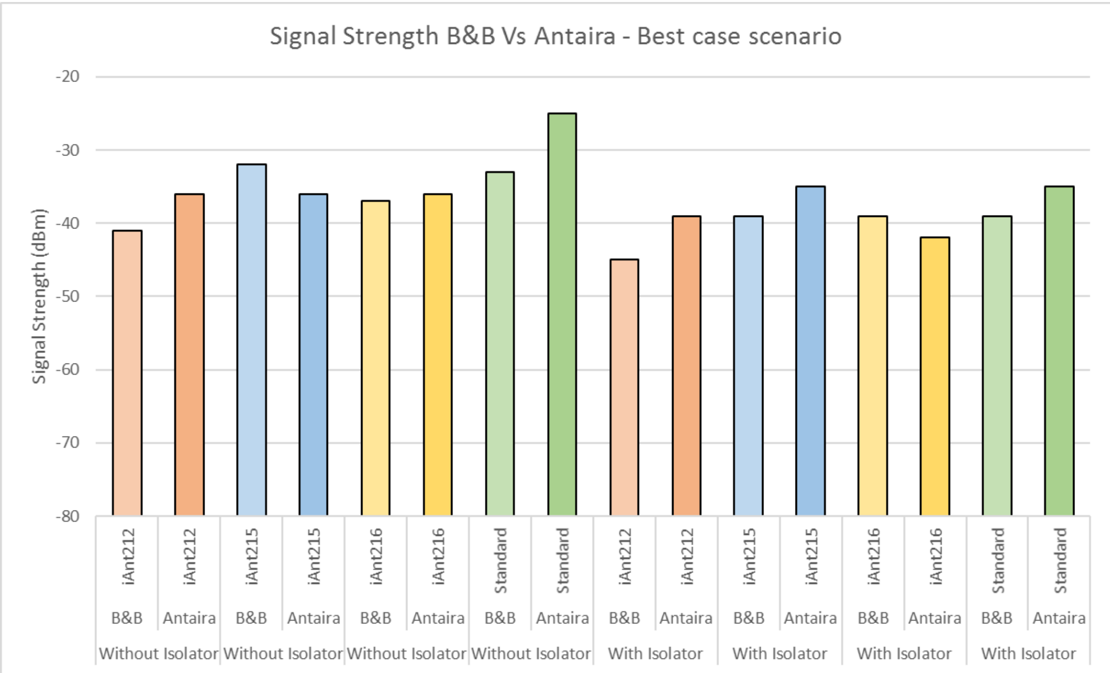

Figure 71 – Graph showing signal strength of the B&B vs. Antaira WAPs in the best case scenario

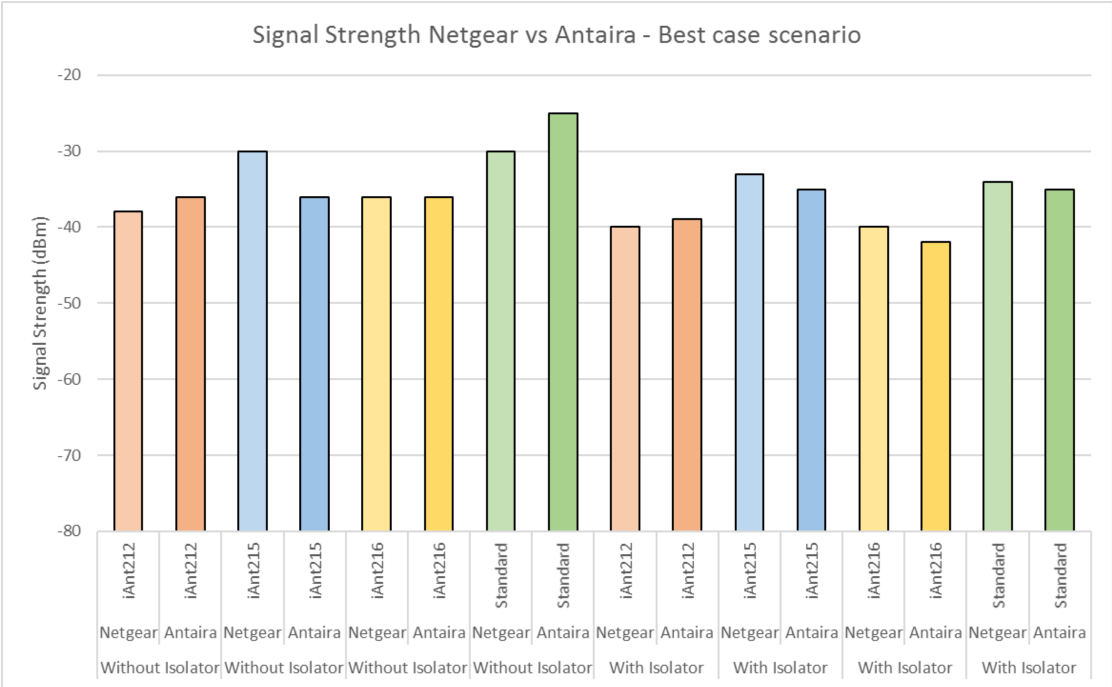

Figure 72– Graph showing signal strength of the Netgear vs. Antaira WAPs in the best case scenario

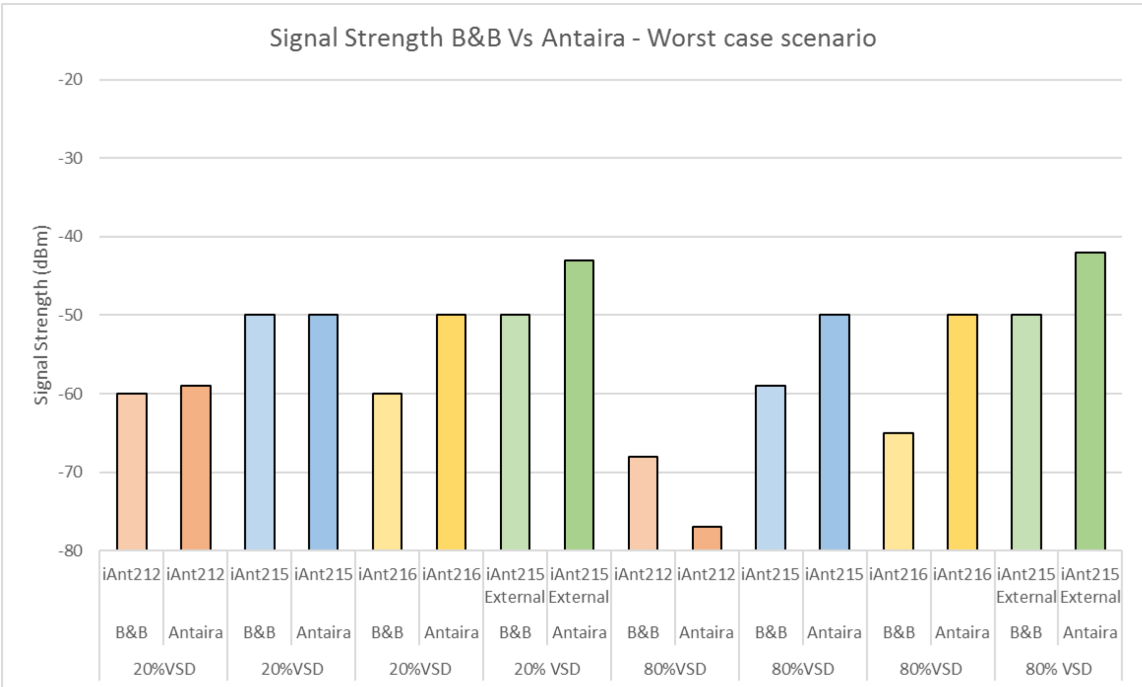

Figure 73 – Graph showing signal strength of the B&B vs. Antaira WAPs in the worst case scenario

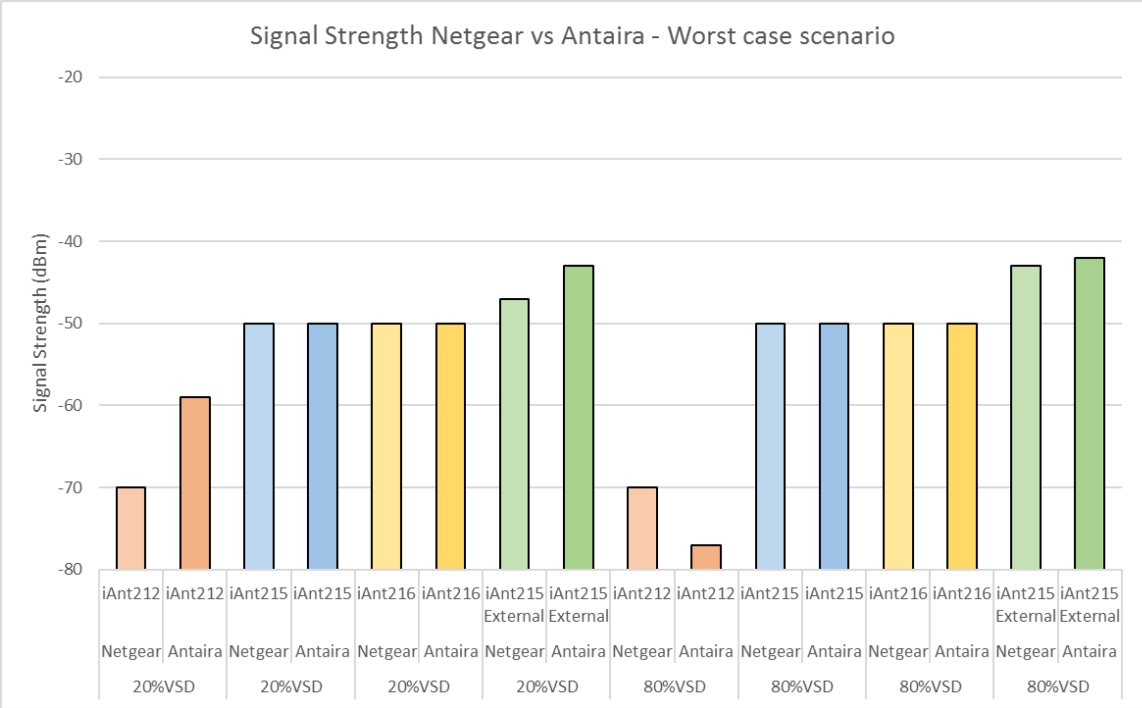

Figure 74 – Graph showing signal strength of the Netgear vs. Antaira WAPs in the worst case scenario

6.8.2 Signal Strength Discussion

Signal measurements taken for each test have consistently shown that the dome iAnt215 performs the best in the given scenarios. This remains true even when the antenna is placed within the IS chamber.

In the office environment signal levels were understandably highest due to the small distance between the WAP and the receiver. These results show consistently good performance regardless of antenna which is unsurprising due to the short and clear path between the test laptop and the WAP.

The introduction of the RF isolator in some cases saw a drop of -10dBm, this can be seen particularly when antennas were placed inside the IS chamber. With this drop in signal also came a drop in performance in terms of increased packet loss, increased ping times and slower data throughput speeds, especially with the small iAnt212.

When the VSD was running, the results show that signal strengths were mostly consistent regardless of what speed the VSD was running at. An exception to this can be seen with the test using the Antaira WAP with the small iAnt212 mounted in the IS chamber. In this test, with the VSD running at 20% speed the signal was -59dBm, whereas with the VSD at 80% speed the signal was reduced to -77dBm. This is a large drop that does not follow the trend of previous tests. Given that the test conditions varied over time (due to other work being carried out within the vicinity) this anomalous result may have been caused by an external factor that was occurring at the time of the test.

Other than the anomalous result the Antaira WAP had the overall best performance, either achieving the same or a slightly better signal than the Netgear and B&B.

In terms of antennas, the results show that the small iAnt212 is not suitable if mounted inside the IS chamber. The dome iAnt215 gave the best results even when placed inside the IS chamber. Furthermore, this antenna enables simplified external mounting, resulting in the signal strength increased by up to 8dBm in a production scenario.

6.8.3 Throughput Results

Throughput results for the best case scenario using the B&B are shown alongside the Antaira in Figure 75. Netgear results are shown in Figure 76. Worst case scenario results for the B&B are shown in Figure 77 and for the Netgear in Figure 78.

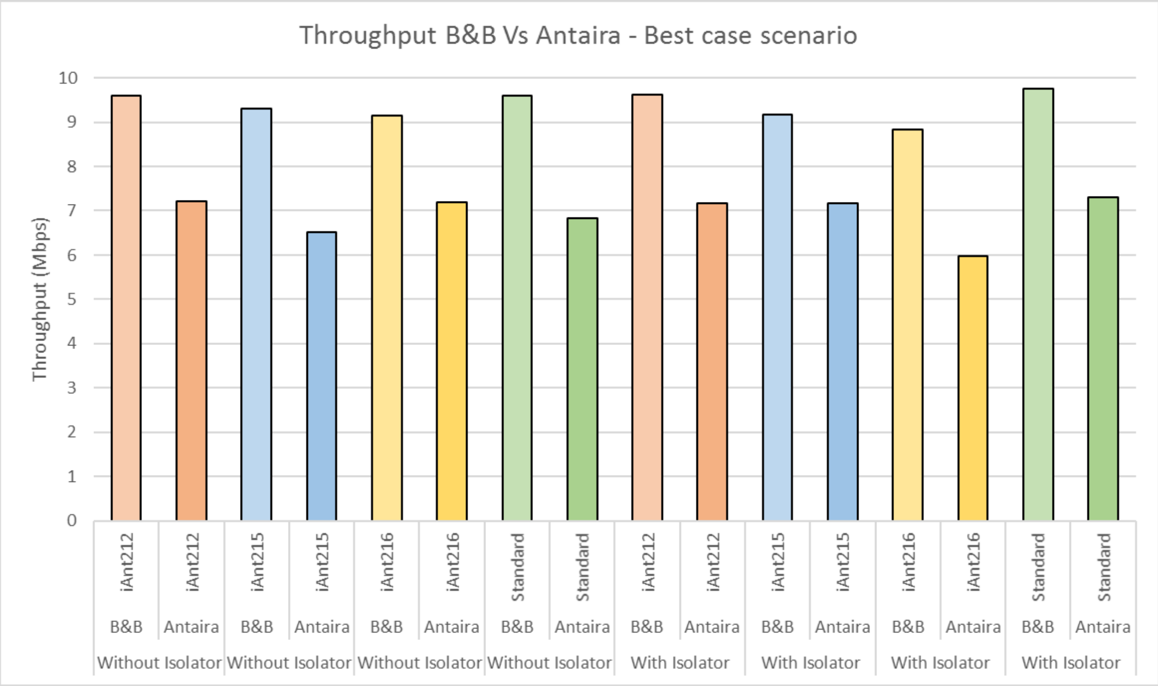

Figure 75 – Graph showing throughput of the B&B vs. Antaira WAPs in the best case scenario

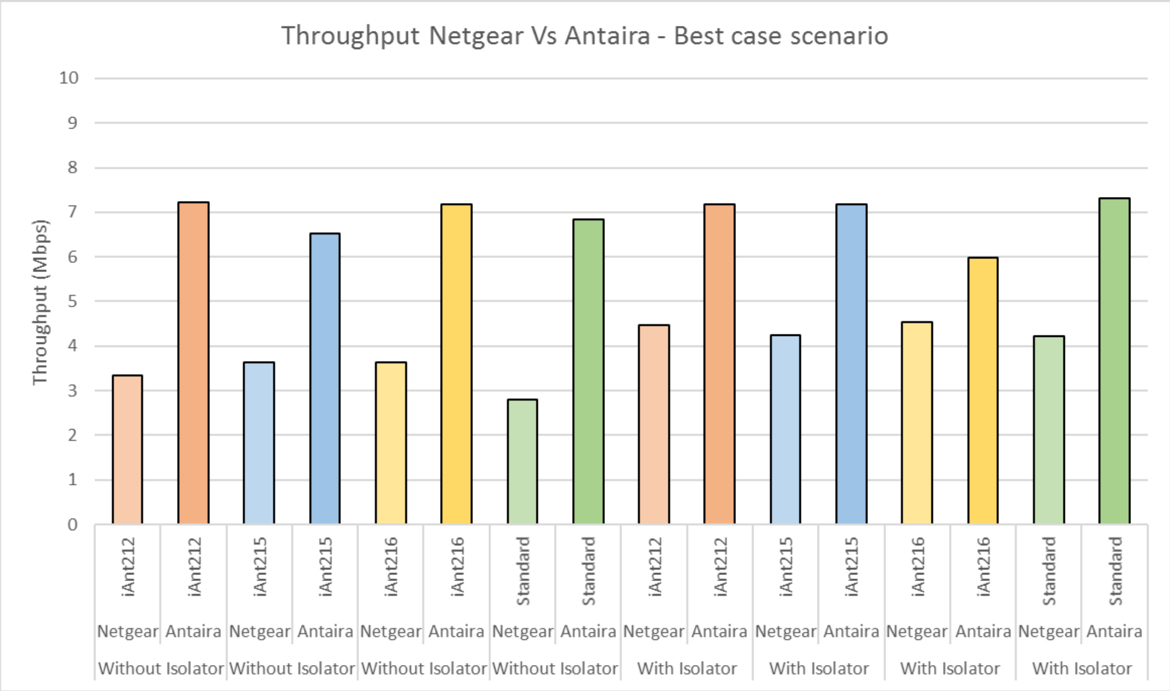

Figure 76 – Graph showing throughput of the Netgear vs. Antaira WAPs in the best case scenario

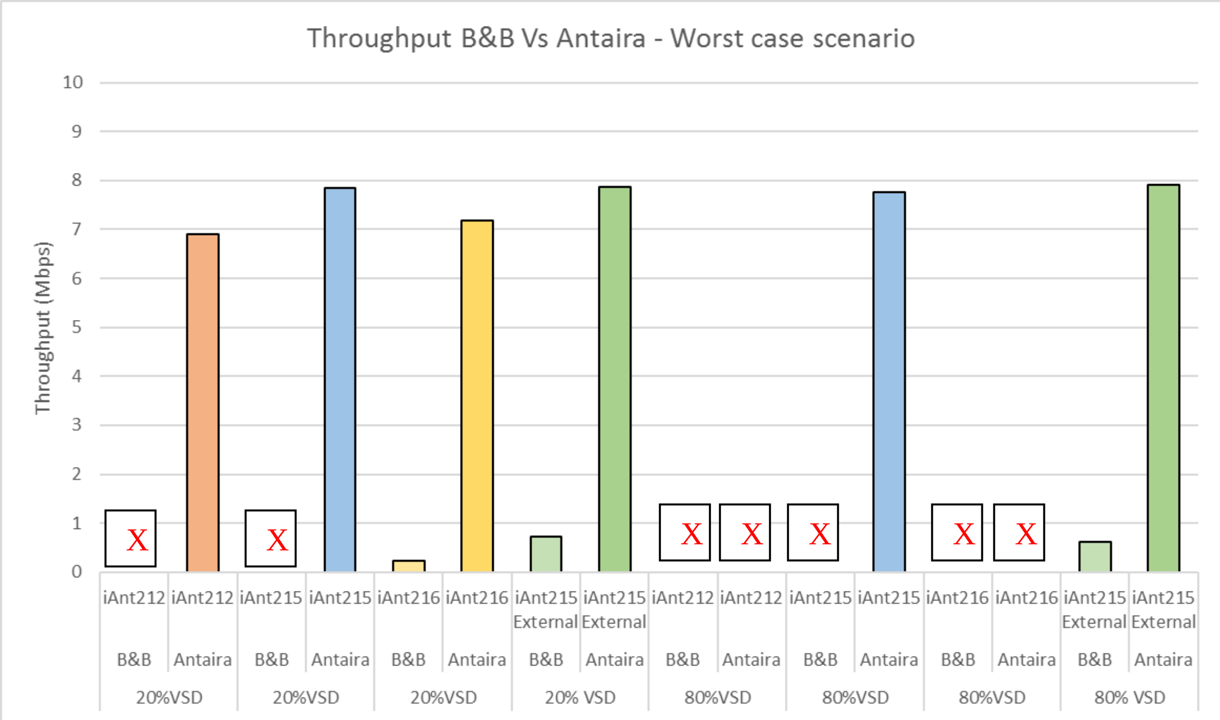

Figure 77 – Graph showing throughput of the B&B vs. Antaira WAPs in the worst case scenario

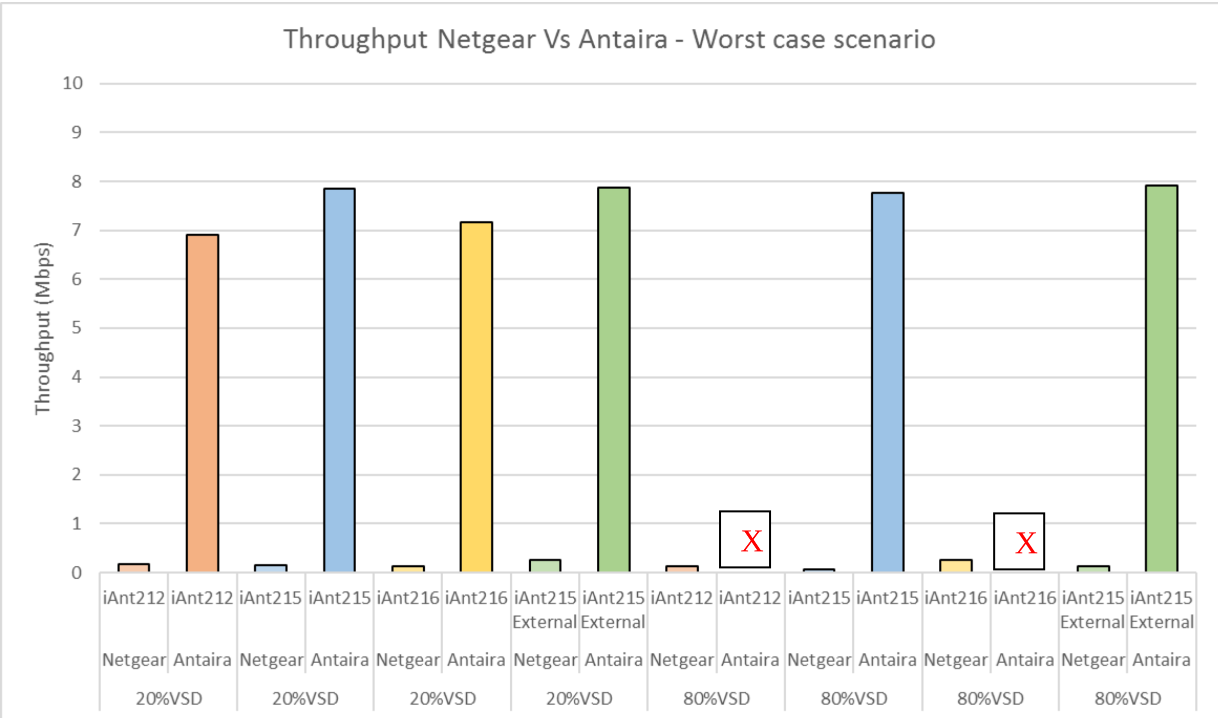

Figure 78 – Graph showing throughput of the Netgear vs. Antaira WAPs in the worst case scenario

6.8.4 Throughput Discussion

The required throughput for this project is minimal, with requests for webpages peaking at 8Kbps. For remote control of the embedded PC104 over Wi-Fi the throughput requirement would increase, with RDP sessions peaking at 25Kbps. To put this into perspecive, the best case scenario saw the WAPs achieving results between 6-9Mbps - more than enough for this project.

During worst case scenario tests, iPerf began reporting instances of 0 throughput for a given interval, an example of this is shown in Figure 79. Observations show that the frequency of these gaps increased with a lower signal strength.

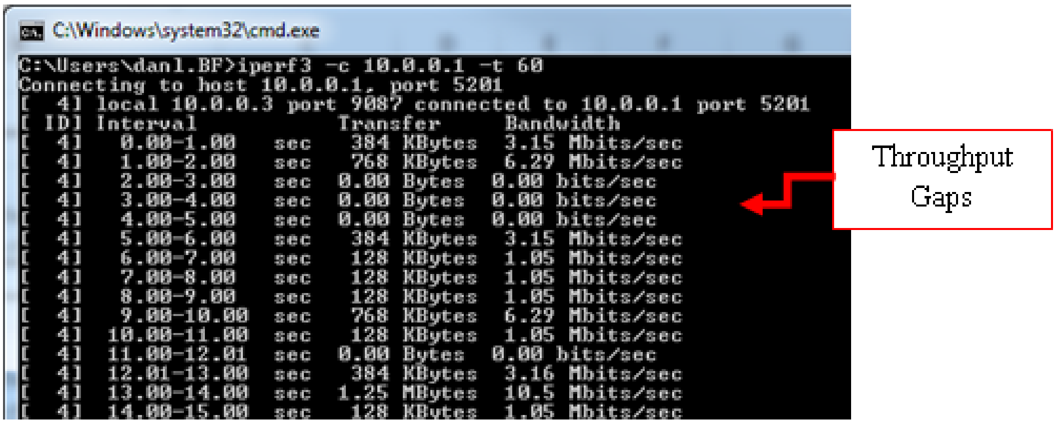

Figure 79 – Gaps observed on iPerf during throughput tests

Production scenario graphs show that throughput for both Netgear and B&B devices was extremely low. Both devices also reported a significant number of gaps on iperf. Further problems were observed when testing with the small iAnt212 and the large iAnt216 antennas, where the wireless connection dropped mid-test. These instances are marked on the graph with an ‘X’ and occurred with both the B&B and Antaira WAPs.

With the dome iAnt215 the Antaira showed robust performance at both 20% and 80% VSD speed - achieving data rates that are similar to those observed in the best-case scenario, there was also minimal occurrences of gaps on iPerf. The Netgear and B&B continued to perform poorly even with the iAnt215 mounted externally.

Tests with and without the isolator show little difference in speed, but occasionally gave results that seem counter intuitive, where tests with the isolator saw slightly more throughput, as observed with the Netgear in the best-case scenario tests.

This can be attributed to the sporadic nature of wireless communications in environments with constantly changing conditions.

6.8.5 Packet Loss Results

Production scenario packet loss results for the Antaira vs. B&B are shown in Figure 80 and for the Netgear in Figure 81. In this instance, graphs that show smaller or no bars indicate better results since high packet loss is undesirable.

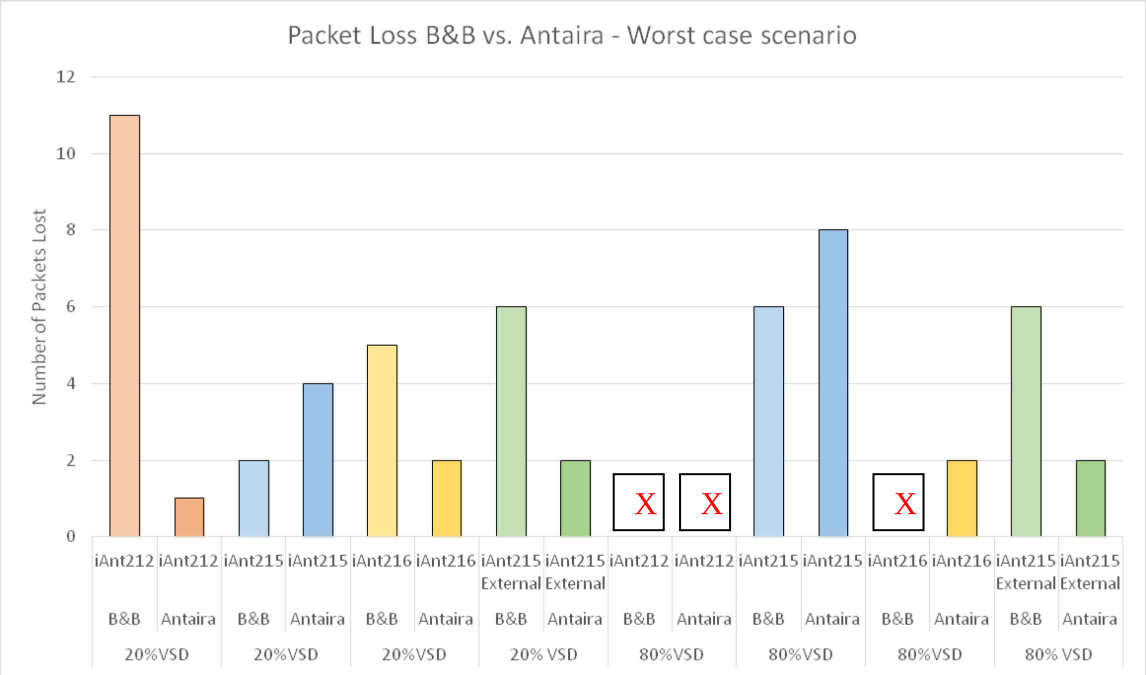

Figure 80 – Graph showing packet loss of the B&B vs. Antaira WAPs in the worst case scenario

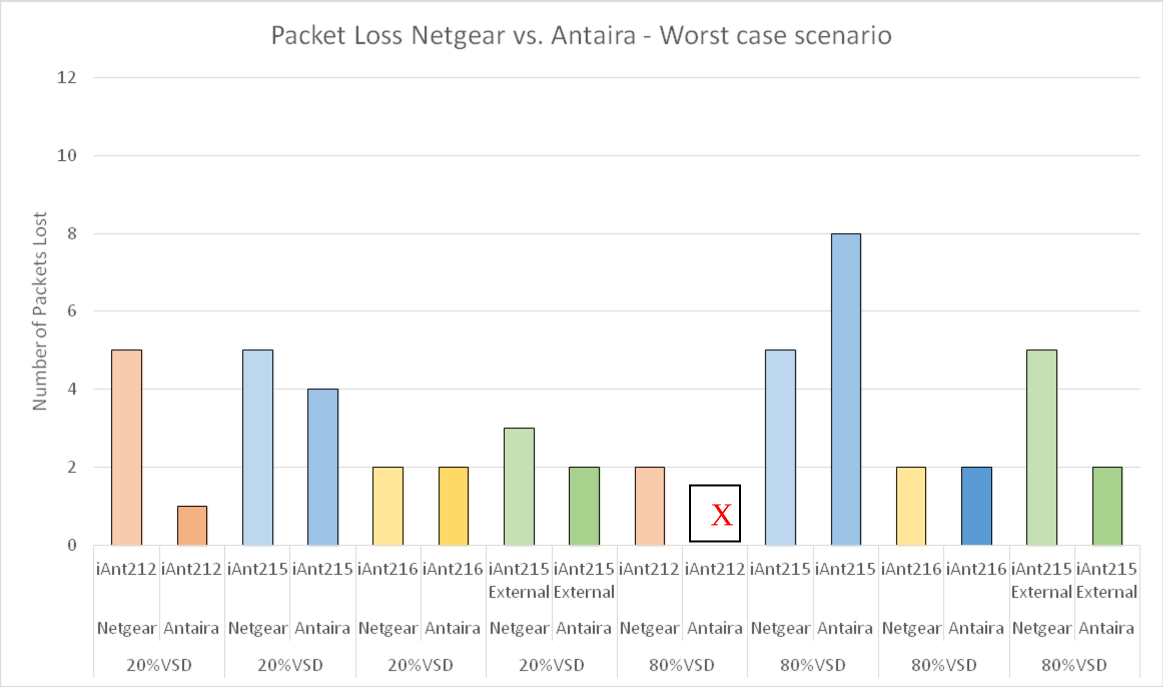

Figure 81– Graph showing packet loss of the Netgear vs. Antaira WAPs in the worst case scenario

6.8.6 Packet Loss Discussion

Each ping test reported the number of packets that were dropped out of 50 that were sent. In the best-case scenario, there was understandably zero packet loss with both the Netgear and Antaira. The B&B however saw a single instance of packet loss; this was unexpected given the short distance between antenna and receiver.

The graphs for the production scenario show instances where tests could not be carried out. This was the same issue observed during throughout testing, only more extreme because a connection to the wireless network was unable to be established at all. The problem only occurred when using the small iAnt212 and large iAnt216 as marked on the graph with an ‘X’.

For tests that completed successfully there was no obvious correlation between number of packets lost and the type of antenna or speed of the VSD. In some cases, the results show more packet loss when the opposite would have been expected, for example with the B&B when the VSD was running at 20% speed, the observed packet loss using the dome iAnt215 mounted externally was greater than with it mounted inside the IS chamber.

One theory is that these sporadic packet losses could be attributed to WLAN rate adaptation. When WLAN rate adaptation is enabled the client will automatically adjust its connection to obtain the highest possible connection speed. If the environment experiences noise to an extent that the transmissions at the current speed can no longer be reliably decoded, it will lower the speed to allow for successful decoding (Acharya, Sharma, Belding, Almeroth, & Papagiannaki, 2010). Whenever these adjustments occur packets may be lost. To get around this, a lower but more robust connection speed could be forcefully set. Since the throughput requirements of this project would not benefit from higher connection speeds this would have no downsides and would help prevent packet loss.

Overall the Antaira had the least amount of loss across the completed tests, giving further confidence that the Antaira should be the AP of choice in a production environment.

6.8.7 Range Testing

A range test was performed with the dome antenna mounted externally on the IS chamber. An Android based tablet with the app ‘Arubua Utilities’ was used to measure signal strength while walking away from the location of the antenna. The route taken is shown in Figures 82 and 83, covering a distance of 45meters through an open environment populated with heavy duty machinery.

Figure 82 – Start of the route taken for Wi-Fi range testing

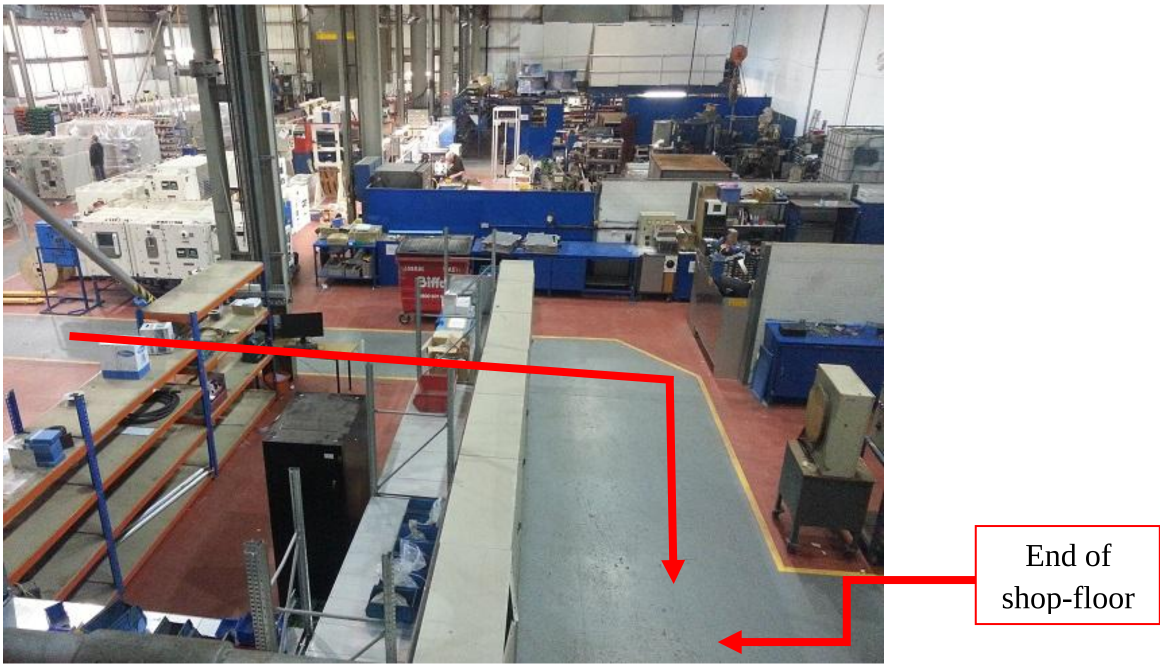

Figure 83 – End of the route taken for Wi-Fi range testing

Figure 84 shows results using the Antaira, Figure 85 shows results for the B&B and Figure 86 shows the Netgear results.

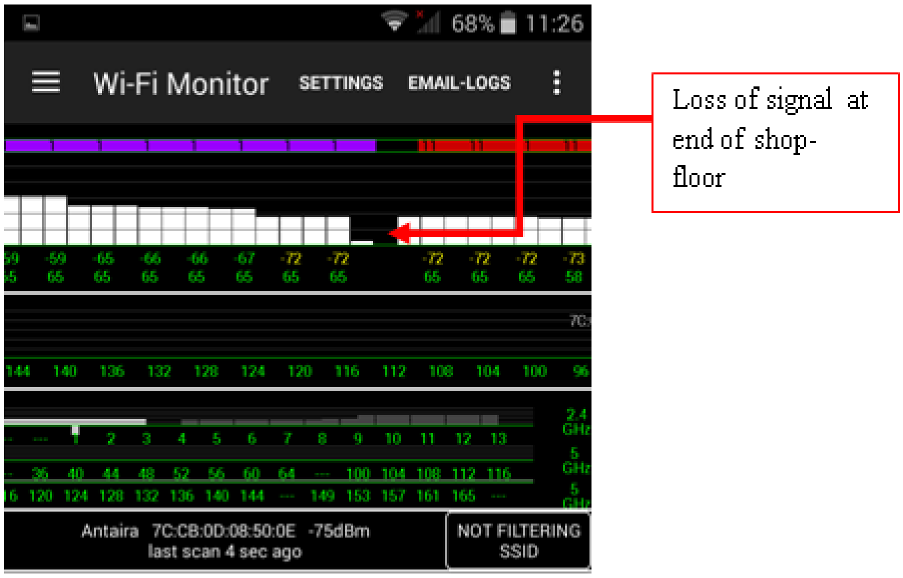

Figure 84 – Range test results using the Antaira WAP

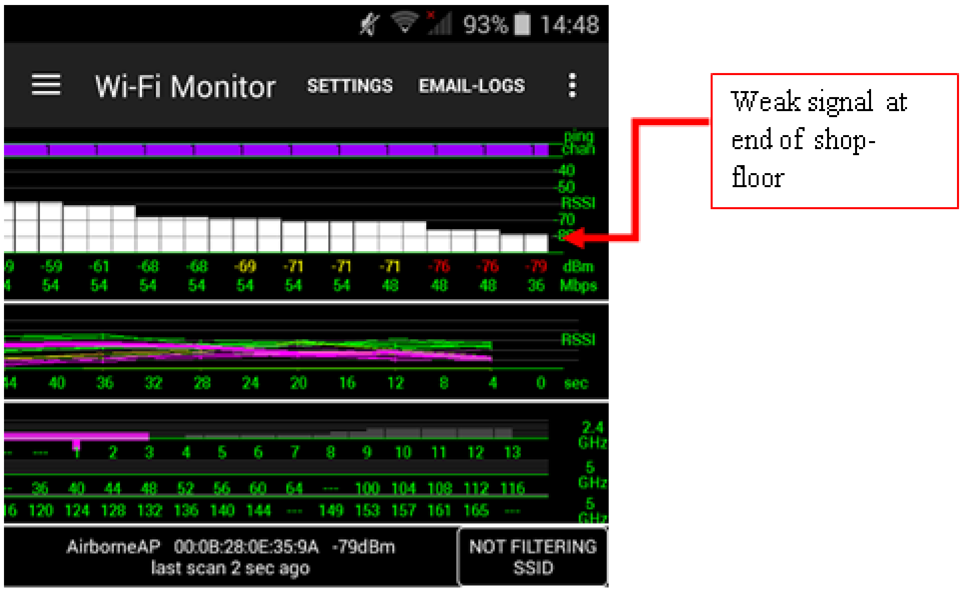

Figure 85 – Range test results using the B&B WAP

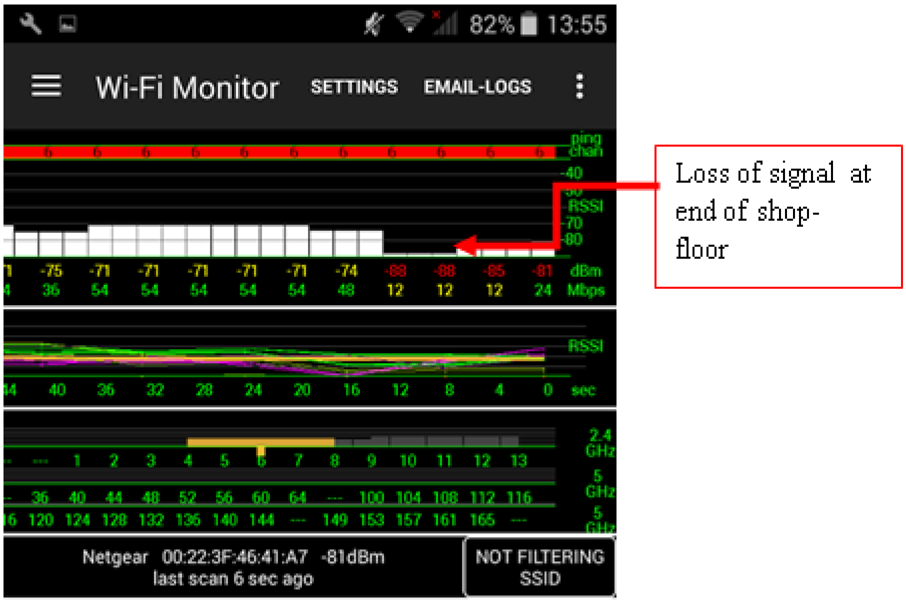

Figure 86 – Range test results using the Netgear WAP

6.8.7.1 Range Discussion

Results from range testing show that the B&B performed better in the given scenario. The distance covered was around 45.72meters but consideration should be given to the fact that the shop-floor was populated with obstacles that would cause attenuation of the signal.

Long range can be seen as feature that is nice to have but not a requirement, given that the production scenario and use cases rely on the user being located in close proximity to the equipment when carrying out diagnostics. Long range is also not feasible in mining scenarios because the walls of the tunnels would certainly block any Wi-Fi signal. These results do however show that in environments that permit, with an externally mounted antenna and a line of sight to the equipment, a user could still access the system up to 45meters away.

6.9 Wireless Access Point Discussion

The square form factor of the B&B would set well with other components of the system, allowing for a single unit to be produced where components are stacked on top of each other. The unit did however have some physical drawbacks - the power connector offered no method of securing it in place, meaning this could potentially become loose over time. Connectivity options were also limited with a single Ethernet port. In contrast, the Antaira has a secure terminal power block which also offers the option for connecting a redundant power supply. Furthermore the unit has two Ethernet ports. These additions demonstrate that the Antaira offers greater flexibility in terms of deployment.

Test results demonstrate that the Antaira gave the best performance, particularly in terms of throughput where the B&B and Netgear units struggled. The importance of a consistent wireless connection in a production environment is vital to this project. Testing has shown that the Antaira is the most suitable for achieving this.

6.10 Antenna Discussion

Externally mounting an antenna on top of the IS chamber would result in the best possible wireless performance. If this was to happen, the iAnt215 would be much less prone to damage due its ruggedised dome form factor than compared to the large iAnt216.

Most importantly, throughput testing has demonstrated that the small iAnt212 and large iAnt216 are unsuitable if mounted inside the IS chamber due to connectivity failures when the VSD was running. The dome iAnt215 saw no such problems even when it was located inside the IS chamber.

Given that the dome iAnt215 antenna is easy to mount and consistently gave the best results it is recommended for use in a production environment.

6.11 Content Delivery Conclusion

Three potential WAPs were identified from market research. The WAP will be responsible for delivering the diagnostic content wirelessly to a tablet PC and must provide a stable and efficient connection. Performance testing was carried out using a range of antennas to identify the best configuration for each WAP.

Production scenario testing was carried out to replicate the conditions of an operating production environment. The results show that the Antaira WAP performed consistently well across all tests. In throughput testing the Antaira, with the dome shaped iAnt215, achieved data rates of around 7Mbps - similar to results observed in a best-case scenario. In the same test the B&B and Netgear both performed poorly, failing to reach data rates above 1Mbps.

Signal measurements have shown that using an externally mounted dome antenna gives the best results. Testing with antennas placed inside the IS chamber results in a signal drop of between 10-20dBm, meaning that the signal will not travel as far from the source than it would if an external antenna was used.

In terms of cost, a consumer WAP like the Netgear would be the cheapest option at around £35. The industrial units are more expensive, with the B&B costing £362.10 and the Antaira costing US$359 (approximately £266). Test results demonstrate that a consumer WAP like the Netgear would be unable to provide the stability and performance that is required for this project given the environment it is to be deployed into. Furthermore, industrial units give additional benefits such as; extended operating temperature ranges, DIN rail mounting and connectivity for industrial power supplies.

Out of the industrial units, the Antaira proved to be the preferred choice. The B&B showed poor results in production testing and has other negative factors such as; a design flaw with the power connector which could result in it becoming loose in transit. The unit also presented undesirable difficulties during the configuration due to a poorly designed web interface.

Out of the antennas, the dome shaped iAnt215 gave the best results and provides a simplified method for externally mounting onto the IS chamber.

7 Critical Reflection

The developed solution has shown that the hardware selected and the developed applications perform their roles reliably with real world equipment in a production scenario over extended periods. The system is therefore at the stage where it can be offered to customers on a trial basis. Outcomes of these trials will give further confidence that the solution is ready to be a marketed product and service.

One area for improvement concerns the amount of space taken up by the raw data, this is due to the fact that polling for data can occur numerous times every second, testing with real world equipment showed 10 polls every second. Space could be conserved by implementing a compression mechanism. Analysis shows that compressing a day’s worth of raw data results in a saving of 95% using ZIP compression or 97.6% using 7z compression. Figure 87 shows the amount of space taken up prior to compression is 159MB and Figure 88 shows how this is reduced to 3.75MB with 7z compression. This high space saving can be attributed to the fact that many polls and responses contain repeated data patterns and therefore, as discussed by Sayood (2012), compresses very well.

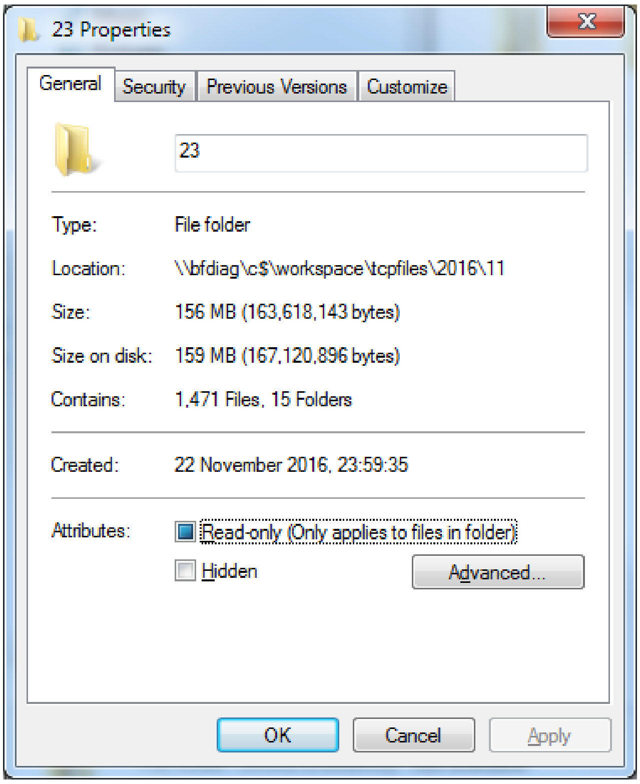

Figure 87 – Disk space used for a day’s worth of raw data when uncompressed

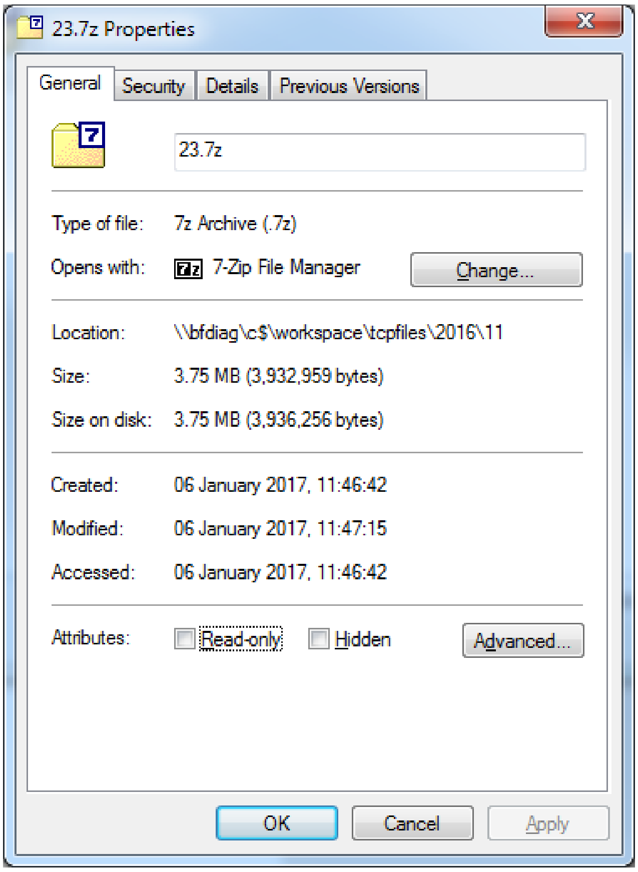

Figure 88 – Disk space used for a day’s worth of raw data when compressed using 7z

Implementation of this functionality on the PC could be done in a trivial method, such as using the Windows task scheduler to run a script that compresses each folder at the end of the month. A more suitable solution however would be to add the functionality as an enhancement to the monitoring solution.

Another potential improvement is related to the cost of the hardware used by the solution. The total cost is around £1500, with the PC104 contributing to the bulk of this. For customers who would find it difficult to justify this cost, an alternative ‘cut down’ service could be provided that consists solely of the MCU platform, reducing the cost to around £840. An embedded web server could be added using the existing middleware to show a basic interface that displays the running status of the equipment. Another alternative would be to explore the use of cheap single board computers in environments that permit. Windows based single board computers such as the ‘Latte Panda’, similar to the Linux based ‘Raspberry Pi’, would be able to run the existing solutions out of the box. As of September 2017 the Latte Panda costs £95 with a modest specification that should meet the requirements for this solution. It is important to note that the use of such boards would be aimed solely towards non-hazardous environments, such as within Principal Supply Points (PSPs) in the rail industry. The harsh environments in oil & gas and mining present constraints, such as high ambient temperatures, which such a board is not certified to operate reliably under.

A future enhancement is to add support for other industrial protocols to the system, such as CAN bus. The selected MCU already provides support for interfacing with CAN bus equipment and the Keil middleware provides drivers that follow the CMSIS standard. This will vastly reduce the development time requirement to add support for the low-level data acquisition stage. High-level data storage support can make use of the existing methods in the Modbus solution, taking advantage of inheritance techniques where necessary to add CAN specific related functionality. Database tables can follow the same structure as Modbus tables, making interfacing with the web-front end a trivial task.

8 Conclusion

Implementing diagnostics for industrial equipment presents a number of unique challenges due to the nature of the equipment being monitored and the hazardous environment it is to be deployed into.

This project has undertaken research and development of a bespoke solution to deliver non-invasive diagnostics of equipment that uses the Modbus RTU protocol for communication. The solution was developed in conjunction with a manufacturer who specialises in mining equipment. It has however been designed to be agnostic to the type of equipment monitored, instead being developed around the fundamental functionality provided by the Modbus RTU protocol. The system is therefore capable of monitoring any Modbus RTU enabled device and perform historical logging without any additional configuration required. Diagnostic information, such as asset status and detailed operating information, can be displayed through trivial configuration of the database that describes the topology of the attached equipment.

Multiple hardware types were evaluated to assess suitability for carrying out different roles required by the solution. A justified selection was made to use a mixture of a dedicated microcontroller for data acquisition alongside an embedded PC for data processing, long term storage and content delivery.

Non-invasive access to viewing content has been achieved by using the Wi-Fi standard. Multiple wireless access points and antenna configurations were tested to ensure robust performance in a production environment where heavy duty machinery results in significant levels of EMI.

Adherence to the stringent ATEX safety standards has been achieved by deploying non-approved equipment inside an approved flame-proof container, ensuring equipment costs are minimised. An intrinsically safe (IS) connection was established to an antenna that was mounted externally on a chamber that houses IS equipment, ensuring a strong and stable wireless signal.

Graphical user interfaces have been developed to provide pertinent information to both engineers and the customers that use the equipment. These interfaces are web-based and are therefore platform agnostic, allowing the content to be viewed using any device with a web browser such as a tablet PC, smartphone, laptop or PC. Security is ensured by employing a two-step authentication mechanism that controls access to different functionality of the solution.

Engineers are presented with detailed technical information about the individual components that make up a system, enabling improved efficiency when troubleshooting onsite. This data also gives manufacturers the ability to ensure their equipment is being operated within the defined limits. In scenarios of equipment misuse the data can be used as evidence if a customer attempts to make a claim for replacement or service work under the warranty.

The customer facing user interface shows an intuitive display of system health alongside troubleshooting advice in the event of an issue, such as highlighting when equipment is approaching the end of its life. This information provides valuable insights and can be used to make informed decisions that ensure downtime is minimised.

Word Count: 26262

9 Bibliography

A. L. Bartels. (1983). Intrinsic safety. *Electronics and Power, 29*(4), 301-304. doi:10.1049/ep.1983.0144

Acharya, P. A. K., Sharma, A., Belding, E. M., Almeroth, K. C., & Papagiannaki, K. (2010). Rate adaptation in congested wireless networks through real-time measurements. *IEEE Transactions on Mobile Computing, 9*(11), 1535-1550. doi:10.1109/TMC.2010.108

Babovic, E., & Velagic, J.Lowering SCADA development and implementation costs using PtP concept. 1-7. doi:10.1109/ICAT.2009.5348454

Bloch, H. P., & Geitner, F. K. (1999). *Machinery failure analysis and troubleshooting* (3. ed. ed.). Houston, Tex: Gulf Publ.

BS ISO 15031-4:2014: Road vehicles. communication between vehicle and external equipment for emissions-related diagnostics. external test equipment (2014). British Standards Institute.

Chin-Long Wey, Chung-Hsien Hsu, Kun-Chun Chang, & Ping-Chang Jui.Enhancement of controller area network (CAN) bus arbitration mechanism. 898.

Minimum requirements for improving the safety and health protection of workers potentially at
risk from explosive atmospheres (1999).

Dao-gang Peng, Hao Zhang, Li Yang, & Hui Li.Design and realization of modbus protocol based on embedded linux system. 275-280. doi:10.1109/ICESS.Symposia.2008.32

Dino Esposito. (2011, Aug 15,). Which ASP.NET is better? *InformationWeek,* , 50.

Engel, S. J., Gilmartin, B. J., Bongort, K., & Hess, A.Prognostics, the real issues involved with predicting life remaining. , *6* 469 vol.6. doi:10.1109/AERO.2000.877920

Enhanced vehicle onboard diagnostic system and method (2016).

Eti, M. C., Ogaji, S. O. T., & Probert, S. D. (2006). Reducing the cost of preventive maintenance (PM) through adopting a proactive reliability-focused culture. *Applied Energy, 83*(11), 1235-1248. doi:10.1016/j.apenergy.2006.01.002

Filev, D. P., Chinnam, R. B., Tseng, F., & Baruah, P. (2010). An industrial strength novelty detection framework for autonomous equipment monitoring and diagnostics. *IEEE Transactions on Industrial Informatics, 6*(4), 767-779. doi:10.1109/TII.2010.2060732

Gilmore, W. J., & Treat, R. H. (2006). Introducing PostgreSQL. *Beginning PHP and PostgreSQL 8* (pp. 573-577). Berkeley, CA: Apress. doi:10.1007/978-1-4302-0136-6_24

Gonzalez, J., & Papa, M. (2007). Passive scanning in modbus networks. *Critical infrastructure protection* (pp. 175-187). Boston, MA: Springer US. doi:10.1007/978-0-387-75462-8_13

Goulao, V., Paulo, R., Salvado, J., Martins, G., & Granjeia, A.A new monitoring and data logger system for industrial cooling equipment applications. 1-3. doi:10.1109/EUROCON.2011.5929329

Hendershot, D. (2009). Remembering flixborough. *Journal of Chemical Health & Safety, 16*(3), 46-47. doi:10.1016/j.jchas.2009.03.006

HENDRIX, J. (2001). Modbus protocol: Diligence in the detail is key to success. *Heating, Piping, Air Conditioning, 73*(5), 32.

Huitsing, P., Chandia, R., Papa, M., & Shenoi, S. (2008). Attack taxonomies for the modbus protocols. *International Journal of Critical Infrastructure Protection, 1*, 37-44. doi:10.1016/j.ijcip.2008.08.003

Ignition. (2016). Ignition SCADA software pricing. Retrieved from https://s3.amazonaws.com/files.inductiveautomation.com/sellsheets/ignition_78/Ignition-HMI-SCADA-Pricing_en.pdf

Iliya Nedyalkov. (2013). Creating modern web applications based on ASP.NET web forms and ASP.NET MVC. *Proceedings of International Conference on Application of Information and Communication Technology and Statistics in Economy and Education (ICAICTSEE),* , 240. Retrieved from https://search.proquest.com/docview/1550836123

Intel. (2016). *Thermal design guide* Intel.

Jardine, A. K. S., Lin, D., & Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing, 20*(7), 1483-1510. doi:10.1016/j.ymssp.2005.09.012

Jonathan Parkyn. (2015, Feb 18,). WHY YOU NEED an SSD & how to install it. *Computer Act!Ve,* , 56.

Lampert, T. (2015). Remote monitoring of surface mining equipment. *Engineering and Mining Journal, 216*(2), 62-65. Retrieved from https://search.proquest.com/docview/1664024387

Langner, R. (2011). Stuxnet: Dissecting a cyberwarfare weapon. *IEEE Security & Privacy, 9*(3), 49-51. doi:10.1109/MSP.2011.67

Modbus ASCII vs modbus RTU. (2017). Retrieved from http://www.simplymodbus.ca/ASCII.htm

Nuggehalli Ramachandra, A. K., & Kannur, A. K.Analysis of CPU utilisation and stack consumption of a multimedia embedded system. 89-94. doi:10.1109/DELTA.2008.38

Pomme, R., & Sijrier, H. J.IECEX certification schemes versus ATEX directives. 1-8.

Sayood, K. (2012). *Introduction to data compression* Morgan Kaufmann.

Sen Xu, Haipeng Pan, Jia Ren, & Jie Su.Design of the modbus communication through serial port in QNX operation system. *, 2* 434-438. doi:10.1109/CCCM.2008.271

Sheng, C., Li, Z., Qin, L., Guo, Z., & Zhang, Y. (2011). Recent progress on mechanical condition monitoring and fault diagnosis. *Procedia Engineering, 15,* 142-146. doi:10.1016/j.proeng.2011.08.029

Smith, D. J. (2017). *Reliability, maintainability and risk* (Ninth edition. ed.). GB: Butterworth Heinemann.

Stefanczak, C. A. (2013, Sep 1,). Custom automation vs. commercial-off-the-shelf, or both? *Control Engineering, 60,* 45.

Táutiva, O. J. D., Roewer, G., & Furlan, P. (2013). *On-board diagnostics: Possible evolutions of the OBDBr-2.* (). doi:10.4271/2013-36-0206

Tian, X., & Benkrid, K. (2010). High-performance quasi-monte carlo financial simulation. *ACM Transactions on Reconfigurable Technology and Systems (TRETS), 3*(4), 1-22. doi:10.1145/1862648.1862656

Vegunta, S. C., & Milanovic, J. V. (2011). Estimation of cost of downtime of industrial process due to voltage sags. *IEEE Transactions on Power Delivery, 26*(2), 576-587. doi:10.1109/TPWRD.2009.2035366

Appendices

Appendix I – Embedded PC Market Research

Manufacturer / Product	CPU	Cores	RAM	Storage	Fan-less?	Networking	Sockets	Notes	Cost
B&R APC 2100	Intel Atom E3845 1.91 GHz	4	4 GB SDRAM	CFast 128GB MLC	✓	Dual GbE Ethernet - RJ45	POWERLINK CAN RS232 (Rx/Tx)	Evaluated. Spec customisable	£1,314.00
Amplicon Impact-D	Intel Atom E3845 1.91 GHz	4	2GB DDR3L	32GB Industrial SSD	✓	Dual GbE Ethernet - RJ45	1 x VGA (DB15) up to 1920 x 1200 @75Hz 1x HDMI up to 1920 x 1200 2x Intel® i210 Gigabit Ethernet 4 x USB3.0 4 x DB9 COM ports (COM1/3 & 4 supports RS232 - COM2 supports RS232/422/485)	Evaluated. Spec customisable	£690.00
ADLINK PC/104 CM1-BT1-E3815-ER	Intel Atom E3815 1.46 GHz	1	4 GB DDR3L	None included (SSD optional)	✓	Dual GbE Ethernet - RJ45	3x USB 2.0 4x RS-232/485 1x SATA 3Gb/s 8x GPIO	Includes temperature testing up to 85'c. Required components: Cableset = +£56 RAM = +£34 SSD = +£97	Board and temperature testing: £340 Including components: £527
ADLINK PC/104 CM3-BT4-E3845	Intel Atom E3845 1.91 GHz	4	4 GB DDR3L	None included (SSD optional)	✓	Dual GbE Ethernet - RJ45	1x USB 3.0 2x USB 2.0 4x RS-232/485 1x SATA 3Gb/s 8x GPIO	Operating temperature up to 85'c. Required components: Cableset = +£56 Extended Temperature Testing = +£46 RAM = +£34 SSD = +£97	Board only: £299 Including components: £532
Mele PCG-01 Compute Stick	Intel Atom Z3735F up to 1.83GHz	4	2 GB DDR3L	32GB eMMC	✓	802.11b/g/n WiFi	Bluetooth 4.0 microSD 1x USB 2.0 1x Micro USB 2.0	Consumer unit only	£55.00
Intel NUC Board NUC5i5MYBE	Intel Core i5-5300U 2.3 GHz	2	None included (DDR3L)	None included (SSD optional)	X	Single GbE Ethernet - RJ45	2x USB 2.0 4x USB 3.0 1x SATA 6Gb/s	Consumer unit only	£338.94

Appendix II – Circular Buffer Function

```
void writeIntoBuffer(struct T_PREVMSGS_ *mptr, char myData[], int numBytes, int numMsgs)
{
    int lastWrittenSample = mptr->tail;
    int i;
    int j;
    // Next available posisition in buffer
    int idx;

    for (i=0; i < numMsgs; ++i)
    {
        mptr->validDataPos = mptr->tail;
        // Increment the current index
        idx = (i + lastWrittenSample);
        for (j=0; j < numBytes; ++j)
        {
            mptr->prevMsgData[idx][j] = myData[j];
            // also keep track of the message length
            mptr->prevMsgLen[idx] = j + 1;
        }
    }
    // Update the current index of our ring buffer.
    mptr->tail += numMsgs;

    // Wrap around to start of buffer if we've reached the end
    if (mptr->tail == mptr->bufferSize)
    {
        mptr->tail = 0;
    }
}
```

Appendix III – Overnight Test #5 Logs

Example 1:
MonitorLog.txt
[27/09/2016 10:23:13] ModbusMonitor: Invalid message received: 040403000000104453

Raw text file on MCU:
68|883|8|03030000001045e4|27-09-2016 09-23-13
68|884|37|03032000008000000000000000000000000000000000000000000000000000
0000000029ba|27-09-2016 09-23-13
68|885|9|040403000000104453|27-09-2016 09-23-13
68|886|8|0503000000104582|27-09-2016 09-23-13
68|887|37|05032000008000000000000000000000000000000000000000000000000000
00000000c7ba|27-09-2016 09-23-13

Expected normal poll for above:
68|867|8|03030000001045e4|27-09-2016 09-23-12
68|868|37|03032000008000000000000000000000000000000000000000000000000000
0000000029ba|27-09-2016 09-23-12
68|869|8|0403000000104453|27-09-2016 09-23-12
68|870|37|04032000008000000000000000000000000000000000000000000000000000
00000000eb7a|27-09-2016 09-23-12
68|871|8|0503000000104582|27-09-2016 09-23-12
68|872|37|05032000008000000000000000000000000000000000000000000000000000
00000000c7ba|27-09-2016 09-23-12

Interestingly there is no response to the poll in the raw text file, suggesting that this may have been a correctly sniffed message and the problem lies with the sender.

Example 2:
MonitorLog.txt
[27/09/2016 11:23:26] ModbusMonitor: Invalid message received: ff

Raw text file on MCU:
134|253|8|06030000001045b1|27-09-2016 10-23-26
134|254|37|0603200000800000000000004410208000004b003e8000000000000000000
0000000008d38|27-09-2016 10-23-26
134|255|1|ff|27-09-2016 10-23-26
134|256|8|1003000000570775|27-09-2016 10-23-26
134|257|179|1003ae17c700000000055555555555555555555555ffeffbfcffffdf7f3df7ffff
0000000000e5000b007b007200cc000b000c0032003a000800000000000000000000000000
00000000000000000000000000000000000000000000000000000000000000

00000000000000000000bf000000000000000000000000000000000000000000
0000000000000000000000000000000000000000000000000000000000000000
000000f49c|27-09-2016 10-23-26

Expected normal poll:
134|278|8|06030000001045b1|27-09-2016 10-23-27
134|279|37|06032000008020000000000441020800004410208000000000000000000
000000000b97b|27-09-2016 10-23-27
134|280|8|1003000000570775|27-09-2016 10-23-27
134|281|179|1003ae17c700000000555555555555555555555555ffeffbfcffffdf7f3df7ffff
0000000000e4000b007b007200cc000b000c0032003a000800000000000000000000000
0000000000000000000000000000000000000000000000000000000000000000
00000000000000000000bd000000000000000000000000000000000000000000
0000000000000000000000000000000000000000000000000000000000000000
0000000b5e2|27-09-2016 10-23-27

Appendix IV – IS Wi-Fi Market Research

Manufacturer/ Product	Type	Wi-Fi Standards	Wi-Fi Bands	Mining Certified?	Industrial Temperatures?	Notes	Cost
P-Ex WLAN	Access Point	IEEE 802.11 a/b/g/n	2.4GHz 5GHz	✓	X	-	£2,171.00
Solexy SWAR	Access Point	IEEE 802.11 b/g/n	2.4GHz 5GHz	✓	X	Has a Modbus gateway	No response for quote
AW5500 industrial IEEE 802.11 a/b/g/n wireless access point	Access Point	IEEE 802.11 a/b/g/n	2.4GHz 5GHz	X	X	Requires 2x Antenna	£301.00
AWK-1131A	Access Point	IEEE 802.11 a/b/g/n	2.4GHz 5GHz	X	X	Requires 2x Antenna	£343.64
AWK-3131-EU-T IEEE802.11a/g/b/n AP/Bridge/AP Client, EU band, -40 to 75°C	Access Point	IEEE 802.11 a/b/g/n	2.4GHz 5GHz	X	X	Requires 2x Antenna	£953.00
AWK-1121-EU-T, Wireless Access point Entry Client 802.11a, 802.11b/g, EU band, -40 to 75°C	Access Point	IEEE 802.11 a/b/g	2.4GHz	X	✓	POE model available (£359)	£319.00
Belden BAT-C Industrial Wireless LAN Client	Client	IEEE 802.11 a/b/g/n	2.4GHz 5GHz	X	X	Uses M12 connectors (would require converters) CLIENT ONLY – unsuitable for use in this project	TBC
Advantech EKI-6332GN	Access Point	IEEE 802.11 b/g/n	2.4GHz 5GHz	X	X	Requires 2x Antenna	£1,000.00
Lantech IWP-1000b/g	Access Point	IEEE 802.11 b/g	2.4GHz	X	✓	Discontinued	N/A
Lantech IWF3310XH Flexible Industrial Access Point	Access Point	IEEE 802.11 a/b/g/n	2.4GHz 5GHz	X	✓	Requires 2x Antenna	TBC
B&B M2M ETHERNET DUAL BAND (2.4 GHZ,5 GHZ) ETHERNET (ABDN-ER-IN5018)	Client	IEEE 802.11 a/b/g/n	2.4GHz 5GHz	X	✓	CLIENT ONLY – unsuitable for use in this project	€ 304.00

B&B M2M Wireless Access Point, 10/100 Ethernet port, 2 RS-232/422/485 ports (APXN-Q5420)	Access Point	IEEE 802.11 a/b/g/n	2.4GHz 5GHz	X	✓	-	€ 338.14
IAP-120/120+	Access Point	IEEE 802.11 b/g	2.4GHz	X	✓	Suppler based in Taipei, rebranding of Antaira APN-210N-T	TBC
Antaira APN-210N	Access Point	IEEE 802.11 a/b/g/n	2.4GHz	X	✓	Suppler based in US	$359.00 (£266)
IE-WLT-AP-BR-CL-ABG-EU	Access Point	IEEE 802.11 a/b/g	2.4GHz	✓	✓	ATEX approved	TBC

Appendix V – Wireless Access Point Equipment Specifications

Antaira APN-210N-T Specification

- Ethernet Port
  - 2 x 10/100BaseT(X) which supports Switch Mode
- Connector
  - RJ-45
- Isolation Protection
  - Built-in 1.5KV Magnetic Isolation
- Protocols
  - ICMP, IP, TCP, UDP, DHCP, BootP, ARP/RARP, DNS, SNMP MIB II,
  - HTTPS, SNMP v1/v2, Trap, Private MIB
- LED
  - PWR 1, PWR 2, DIAG, Ethernet Link/ACT, Signal Level,
  - PWR1 FAIL, PWR2 FAIL
- Antenna
  - 1 x Reverse SMA
- Modulation
  - OFDM: BPSK, QPSK, 16-QAM, 64-QAM
  - DSSS: DBPSK, DQPSK, CCK
- Frequency Range
  - IEEE802.11b/g/n:
  - 2.412GHz ~ 2.462GHz (US & Canada)
  - 2.412GHz ~ 2.472GHz (Europe)
  - 2.412GHz ~ 2.484GHz (Japan)
- Transmit Power
  - 23dBm Max.
- Encryption Security
  - WEP: 64-bit, 128-bit Key Supported
  - WPA / WPA2: 802.11i (WEP and AES Encryption)
  - PSK (256-bit Key Pre-shared Key Supported)
  - 802.1x and RADIUS Supported
  - TKIP Encryption
- Wireless Security
  - SSID Broadcast Disable, 3 x Virtual Access Points, MAC Address Filter
- Power Input
  - 12 ~ 24VDC in 6-Pin Terminal Block, Dual Power
- Power Consumption
  - 6 Watts Max.
- Reverse Polarity Protection
  - Present
- Operating Temperature
  - -35° C to 70° C
- Operating Humidity
  - 5% to 95% (Non-condensing)
- Storage Temperature
  - -40° C to 90° C
- Casing Enclosure

    - IP40 Protection
- Dimensions
    - 46 x 155 x 115 mm (1.81 x 6.10 x 4.53 in.)
- Stability Testing
    - Shock: IEC60068-2-27
    - Free Fall: IEC60068-2-32
    - Vibration: IEC60068-2-6
- Certification
    - FCC
- Warranty
    - 5-Year Warranty

B&B AirborneM2M Specification

- Wireless Technology
    - IEEE 802.11 a/b/g/n, Wi-Fi Compliant
- Wired Interface
    - 2 ports, RS-232/422/485 (RS-232/422 4 wire or RS-485 2 wire)
    - 10/100 Ethernet port (Bridge, Router (NAT3) Modes)
    - Software selectable
- Frequency
    - 2.4~2.4835 GHz (US/Canada/Europe)
    - 2.4~2.497 GHz (Japan)
    - 5.150~5.350 GHz
    - 5.725~5.825 GHz
- Channels (US/Canada):
    - 11 Channels 802.11 b/g
    - 13 Channels 802.11a
- Channels (Europe):
    - 13 Channels 802.11b/g
    - 19 Channels 802.11a
- Channels (France):
    - 4 Channels 802.11b/g
- Channels (Japan):
    - 14 Channels 802.11b
    - 13 Channels 802.11g
    - 23 Channels 802.11a
- Wireless Data Rates
    - 802.11a/g 54, 48, 36, 24, 18, 12, 9, 6 Mbps
    - 802.11b = 11, 5.5,2,1 Mbps
    - 802.11n 65, 58.5, 42, 39, 26, 19.5, 13, 6.5 Mbps
- Network Protocols:
    - TCP/IP, ARP, ICMP, DHCP, DHS, UDAP, TFTP, UDP, PING, HTTP, FTP
- Wireless Security:
    - Open
    - WEP 64 & 128 bit
    - WPA-PSK (TKIP)
    - WPA2-PSK (AES)
    - 802.1x (EAP)

        - WPA-Enterprise
        - WPA2-Enterprise
        - EAP-TLS/MSCHAPv2
        - EAP-TTLS/MSCHAPv2
        - EAP-TTLS (MD5)
        - EAP-PEAPv0/MSCHAPv2
        - LEAP
            - Zero host security footprint
            - Advanced certificate storage and management
- Secure Communications
    - SSH and SSL tunnelling
    - Encrypted configuration
- Transmit Power
    - 802.11b 15 dBm (31.6mW)
    - 802.11g 12.6dBm (18.12mW)
    - 802.11a 17 dBm (50.1mW)
- Input Voltage
    - 5-36VDC +/-5%, 500mA (MAX)
- PoE
    - PoE using a 802.3af Class 1 PSE device
- Operating Temperature
    - -40° to +85°C
- Regulatory Approvals:
    - FCC Part 15.247, Class B Sub C Modular Approval
    - Industry Canada RSS-210
    - CE
    - ETSI EN300-328 v1.7.1
    - ETSI 60950-1
    - Directive 2004/108/EC
    - ETSI EN 55022:2006 + A1:2007 (emissions)
    - ETSI EN 55024:1998 + A1:2001
    - ETSI EN 55024:1998 + A2:2003 (immunity)
    - FCC Part 15 Subpart B:2007
        - Part 15.107(b) (conducted emissions, Class A)
        - Part 15.109(g) (radiated emissions, Class B)
    - Industry Canada ICES-003:2004, Issue 4
    - AS/NZS CISPR 11:2004 (Australia/New Zealand)
    - RoHS and WEEE Compliant

Netgear WAG102 Specification

- Wireless:
    - Standard: IEEE 802.11a up to 108 Mbps; 802.11g up to
- 108 Mbps Wireless Access Point
    - Antennas: Two (2) detachable antennas:
        - 5 dBi, 2.4 GHz and
        - 5 dBi 5 GHz
- Physical Interfaces:
    - LAN Port: One (1) 10/100BASE-T Ethernet (RJ-45) port

with Auto Uplink™ (Auto MDI-X) with IEEE 802.3af Power over Ethernet (PoE) support

- Power adapter: 12 VDC, 1.0A power supply; plug is localized to country of sale.
- Six (6) LEDs: Power, Test, 100 Mbps, Link/ACT, and WLAN 802.11g
- Antenna Connector: Two (2) reverse SMA antenna connectors

- Network Management:
  - Remote configuration and management through Web browser, SNMP
  - SNMP management supports SNMP MIB I, MIB II,
  - and 802.11 MIB
- Advanced Wireless Features:
  - Point-to-point wireless bridge mode
  - Point-to-multipoint wireless bridge mode
  - Repeater mode
  - Adjustable Transmit Power Control (TPC) from 100 mW down to 0 mW
  - RF Management functionality including auto power transmission, channel selection (Future firmware upgrade)
- Typical Maximum Power Output:
  - 802.11b mode, 1 to 11 Mbps + 19 dBm
  - 802.11g mode, 6 to 24 Mbps + 18 dBm
  - 802.11g mode, 36, 48 and 54 Mbps + 17/16/16 dBm

    Note: Maximum Tx varies based on country and/or region selection to ensure local regulatory compliance.
- Receive Sensitivity:
  - 802.11b

    2 Mbps -94 dBm

    5.5 Mbps -92 dBm

    11 Mbps -90 dBm
  - 802.11g

    6 Mbps - 91 dBm

    9 Mbps - 91 dBm

    12 Mbps - 90 dBm

    18 Mbps - 87 dBm

    24 Mbps - 85 dBm

    36 Mbps - 80 dBm

    48 Mbps - 75 dBm

    54 Mbps - 70 dBm
  - 802.11a

    6 Mbps - 92 dBm

    9 Mbps - 91 dBm

    12 Mbps - 90 dBm

    18 Mbps - 86 dBm

    24 Mbps - 84 dBm

    36 Mbps - 79 dBm

48 Mbps - 72 dBm

54 Mbps - 70 dBm

- Power Over Ethernet Required:
  - 4.3 W
- Security:
  - 40/64-, 128-, and 152-bit WEP encryption
  - Block SSID Broadcast
  - VPN pass-through support
  - MAC address filtering with access control lists – up to 256 users
  - 802.1x RADIUS support with EAP TLS, TTLS, PEAP
  - Wi-Fi Protected Access (WPA2)
  - Secure Socket Layer (SSL) remote management login
- Standards, Safety and Electromagnetic Conformance:
  - IEEE 802.11g (2.4 GHz Frequency Band, DSSS Modulation Type)
  - IEEE 802.3af Power over Ethernet (PoE)
  - FCC Part 15 Subpart B and Subpart C
  - CE EN 60950-1 (2001), EN 301 489-17 V1.2.1 (2002), EN 301 489-1 V1.4.1 (2002), EN 300 328-1 (2001-12), EN 300 328-2 (2001-12), R&TTE Directive
  - 99/5/EC, CSIPR22(B), AS/NZS 3548 (B)
  - Plenum Rated
  - IEEE 802.11g (2.4 GHz Frequency Band, DSSS Modulation Type)
  - IEEE 802.11a
  - IEEE 802.3af Power over Ethernet (PoE)
  - FCC Part 15 Subpart B , Subpart C and E
  - C-Tick: AS/NZ 3548 Class B
  - CE EN 60950-1 (2001), EN 301 489-17 V1.2.1 (2002-08), EN 301 489-1 V1.4.1 (2002-08), EN 300328-1 (2001-12), EN 300 328-2 (2003--4),EN301893(5GHz),EN60601, R&TTE Directive 99/5/EC
- Physical Specifications:
  - Dimensions (H x W x D): 32 x 190.5 x 122 mm (1.25 x 7.5 x 4.8 in.)
  - Weight: .62 kg (1.3 lb)
- Environmental Specifications:
  - Operating temperature: 0 to 40° C (32 to 104° F)
  - Operating humidity: 90% maximum relative humidity, noncondensing

Appendix VI – Antenna Specifications

iAnt212 Specification

- Frequency range
  - 1.7 – 2 GHz
  - 2 – 2.7 GHz
  - 2.7 – 3.8 GHz
  - 4.9 – 5.15 GHz
  - 5.15 – 6 GHz
- Gain
  - 1.5 dBi
  - 2 dBi
  - 1.5 dBi
  - 2.5 dBi
  - 2 dBi
- VSWR
  - 2
- Horizontal beam width
  - 360°
- Vertical beam width
  - 360°
- Impedance
  - 50 Ohms
- Polarisation
  - Spherical
- Operating temperature
  - -40°C to +85°C
- Storage temperature
  - -40°C to +85°C
- Radome material
  - ASA plastic
- Ingress protection
  - IP68
- Dimensions
  - 55mm height, 22mm diameter
- Weight
  - 50g
- Connection
  - N-type

iAnt215 Specification

- Frequency range
  - 2400-5470 MHz
- Gain
  - 6 / 7 / 8 dBi
- VSWR
  - 1.8 / 2 / 1.8
- Polarisation
  - Vertical
- Horizontal 3dB beam width
  - 360°
- Vertical 3dB beam width
  - 25°
- Nominal impedance
  - 50 Ω
- Wind speed
  - 160km/h
- Operating temperature
  - -40°C to +80°C
- Composite power max.
  - 75W @ 25°C
- Flammability rating
  - DIN 5510-2, BS 6853, NF F16-101/102, CEN/TS 45545 (2009)
- Radome material/colour
  - UV stabilised polycarbonate/ RAL 7043
- Ingress protection
  - IP68
- Dimensions
  - H50.6 x D86 mm
- Weight
  - 0.3 kg
- RF connection
  - N-type jack (female)
- RoHS
  - Compliant
- Mounting
  - Supplied with stainless steel wall/pole mount bracket kit

2.4GHz horizontal plane radiation pattern

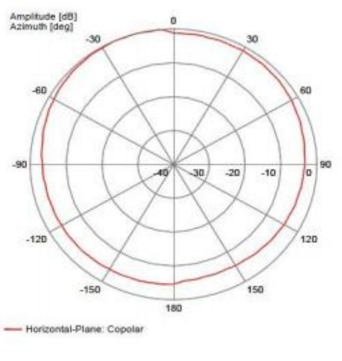

2.4GHz vertical plane radiation pattern

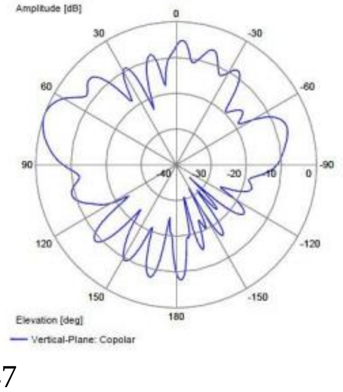

iAnt216 Specification

- Frequency band
  - 2.4-2.5GHz
  - 4.94-5.925GHz
- Gain
  - 5dBi @ 2.4GHz
  - 7dBi @ 5GHz
- VSWR
  - < 2
- Polarisation
  - Vertical
- Horizontal plane beam width
  - 360° 2.4GHz
  - Omni 5GHz
- Vertical plane beam width
  - 25° 2.4Ghz
  - 15° 5GHz
- Composite power max
  - 75W @ 25°C
- Nominal input impedance
  - 50 Ohm
- Wind speed
  - 125mph
- Temperature range
  - -40° C to +70° C
- Radome material
  - White UV-resistant ABS
- Ingress protection
  - IP67
- Dimensions
  - 260 mm x 30 mm
- Weight
  - 300g
- RF connection
  - N-type male

2.45GHz Horizontal Plane Radiation Pattern

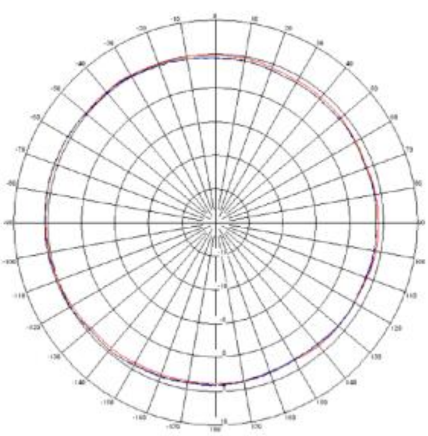

2.45GHz Vertical Plane Radiation Pattern

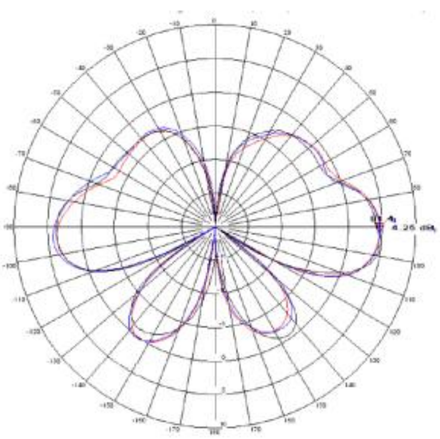

Appendix VII – Antaira WAP Results

Initial tests were performed using a bandwidth of 40MHz to take advantage of faster speeds offered by 802.11n (maximum of 135Mbps with an 800ns Guard Interval). This means the WAP uses frequencies from both channel 1 and channel 6 as shown below.

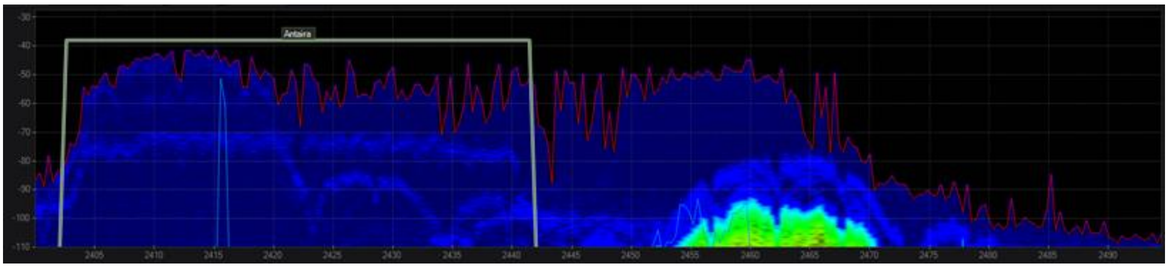

When turning on the VSD it was noted that Wi-Fi channel 6 had become more congested as shown in below.

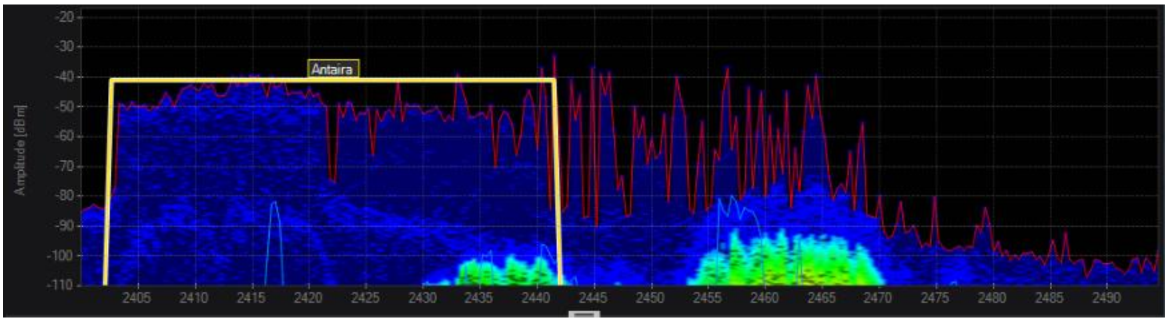

This could be attributed to increased activity from a neighbouring WAP or from interference caused by the VSD. This increased congestion caused problems with stability, it was therefore decided to lower the bandwidth from 40MHz to 20MHz to ensure the WAP would only use the uncongested frequencies on channel 1. This resulted in a stable connection but also a reduced maximum connection speed of 65Mbps.

Best Case Scenario

The following results are from the office location using a bandwidth of 20MHz.

Antenna	Isolator?	Signal Strength	Min Ping	Max Ping	Packet Loss	Throughput	Gaps
Standard	No	-25dBm	1ms	273ms	0	6.83Mbps	0
Standard	Yes	-35dBm	1ms	271ms	0	7.31Mbps	0
Large iAnt216	No	-36dBm	1ms	269ms	0	7.18Mbps	0
Large iAnt216	Yes	-42dBm	1ms	272ms	0	5.98Mbps	0
Dome iAnt215	No	-34dBm	1ms	297ms	0	6.52Mbps	0
Dome iAnt215	Yes	-35dBm	1ms	274ms	0	7.17Mbps	0
Small iAnt212	No	-36dBm	1ms	280ms	0	7.22Mbps	0
Small iAnt212	Yes	-39dBm	1ms	268ms	0	7.17Mbps	0

Test Bay

This section contains results with equipment located inside the FLP without the VSD running. Tests were carried out with channel bonding enabled, resulting in a bandwidth of 40MHz. This doubling enables higher throughput speeds to be achieved. Unless stated otherwise the tests were carried out with the antenna deployed inside the IS chamber.

Antenna	Isolator?	Signal Strength	Min Ping	Max Ping	Packet Loss	Throughput	Gaps
Standard	No	-50dBm	1ms	415ms	3 (6%)	8.62Mbps	2
Standard	Yes	-60dBm	1ms	255ms	2 (4%)	7.5Mbps	5
Large iAnt216	No	-50dBm	1ms	271ms	1 (2%)	12.8Mbps	0
Large iAnt216	Yes	-60dBm	1ms	445ms	1 (2%)	10.8Mbps	4
Dome iAnt215	No	-50dBm	1ms	276ms	1 (2%)	11.1Mbps	0
Dome iAnt215	Yes	-50dBm	1ms	271ms	1 (2%)	11.1Mbps	2
Small iAnt212	No	-50dBm	1ms	274ms	2 (4%)	12.8Mbps	1
Small iAnt212	Yes	-60dBm	1ms	850ms	6 (12%)	1.05Mbps	28
Externally mounted dome iAnt215	Yes	-50dm	1ms	276ms	0	11.2Mbps	3

Worst Case Scenario

This section contains results with equipment located inside the FLP with the VSD running at varied speeds. Due to channel congestion issues channel bonding was disabled, meaning the bandwidth was reduced to 20MHz. Unless stated otherwise the tests were carried out with the antenna deployed inside the IS chamber

Antenna	Isolator?	VSD Speed	Signal Strength	Min Ping	Max Ping	Packet Loss	Throughput	Gaps
Large iAnt216	Yes	20%	-50dBm	1ms	270ms	2 (4%)	7.17Mbps	0
Large iAnt216	Yes	80%	-50dBm	1ms	279ms	2 (4%)	N/A. Wi-Fi kept dropping throughout testing	
Dome iAnt215	Yes	20%	-50dBm	1ms	350ms	4 (8%)	7.85Mbps	1
Dome iAnt215	Yes	80%	-50dBm	1ms	235ms	8 (16%)	7.76Mbps	2
Small iAnt212	Yes	20%	-59dBm	1ms	271ms	1 (2%)	6.9Mbps	0
Small iAnt212	Yes	80%	-77dBm	N/A. Unable to connect to Wi-Fi				
Externally mounted dome iAnt215	Yes	20%	-43dBm	1ms	283ms	2 (4%)	7.87Mbps	0
Externally mounted dome iAnt215	Yes	80%	-42dBm	1ms	283ms	2 (4%)	7.92Mbps	0

Appendix VIII – B&B WAP Results

All tests with the B&B WAP used a bandwidth of 20MHz since there was no support for the 40MHz channel bonding.

Best Case Scenario

Antenna	Isolator?	Signal Strength	Min Ping	Max Ping	Packet Loss	Throughput	Gaps
Standard	No	-33dBm	2ms	121ms	1 (2%)	9.59Mbps	0
Standard	Yes	-39dBm	2ms	124ms	0	9.75Mbps	0
Large iAnt216	No	-37dBm	2ms	116ms	1 (2%)	9.14Mbps	2
Large iAnt216	Yes	-39dBm	1ms	109ms	0	8.84Mbps	2
Dome iAnt215	No	-32dBm	1ms	149ms	0	9.31Mbps	0
Dome iAnt215	Yes	-39dBm	2ms	160ms	0	9.17Mbps	1
Small iAnt212	No	-41dBm	1ms	145ms	1 (2%)	9.59Mbps	0
Small iAnt212	Yes	-45dBm	1ms	142ms	0	9.61Mbps	0

Test Bay

Antenna	Isolator?	Signal Strength	Min Ping	Max Ping	Packet Loss	Throughput	Gaps
Standard	No	-60dBm	2ms	239ms	2 (4%)	6.27Mbps	0
Standard	Yes	-60dBm	2ms	120ms	4 (8%)	3.67Mbps	1
Large iAnt216	No	-50dBm	2ms	138ms	3 (6%)	8.88Mbps	2
Large iAnt216	Yes	-60dBm	2ms	122ms	4 (8%)	4.18Mbps	1
Dome iAnt215	No	-50dBm	2ms	159ms	3 (6%)	6.85Mbps	0
Dome iAnt215	Yes	-60dBm	2ms	148ms	0	5.45Mbps	0
Small iAnt212	No	-60dBm	2ms	122ms	1 (2%)	2.5Mbps	5
Small iAnt212	Yes	-60dBm	2ms	134ms	1 (2%)	6.62Mbps	2
Externally mounted dome iAnt215	Yes	-46dBm	2ms	141ms	2 (4%)	7.17Mbps	0

Worst Case Scenario

Antenna	Isolator?	VSD Speed	Signal Strength	Min Ping	Max Ping	Packet Loss	Throughput	Gaps
Large iAnt216	Yes	20%	-60dBm	2ms	200ms	5 (10%)	253Kbps	45
Large iAnt216	Yes	80%	-65dBm	FAILED TO CONNECT				
Dome iAnt215	Yes	20%	-50dBm	2ms	148ms	2 (4%)	CONNECTION DROPPED MID TEST	
Dome iAnt215	Yes	80%	-50dBm	2ms	148ms	2 (4%)	CONNECTION DROPPED MID TEST	
Small iAnt212	Yes	20%	-59dBm	2ms	138ms	6 (12%)	CONNECTION DROPPED MID TEST	
Small iAnt212	Yes	80%	FAILED TO CONNECT					
Externally mounted dome iAnt215	Yes	20%	-50dBm	1ms	202ms	6 (12%)	734Kbps	29
Externally mounted dome iAnt215	Yes	80%	-50dBm	1ms	170ms	6 (12%)	612Kbps	27

Appendix IX – Netgear WAP Results

As with the B&B WAP, all tests with the Netgear used a bandwidth of 20MHz since there was no support for the 40MHz channel bonding.

Best Case Scenario

Antenna	Isolator?	Signal Strength	Min Ping	Max Ping	Packet Loss	Throughput	Gaps
Standard	No	-30dBm	1ms	243ms	0	2.8Mbps	0
Standard	Yes	-34dBm	1ms	274ms	0	4.23Mbps	0
Large iAnt216	No	-36dBm	1ms	272ms	0	3.64Mbps	0
Large iAnt216	Yes	-40dBm	1ms	288ms	0	4.53Mbps	0
Dome iAnt215	No	-30dBm	1ms	263ms	0	3.64Mbps	0
Dome iAnt215	Yes	-33dBm	1ms	261ms	0	4.24Mbps	0
Small iAnt212	No	-38dBm	1ms	272ms	0	3.34Mbps	1
Small iAnt212	Yes	-40dBm	1ms	252ms	0	4.46Mbps	0

Test Bay

Antenna	Isolator?	Signal Strength	Min Ping	Max Ping	Packet Loss	Throughput	Gaps
Standard	No	-50dBm	1ms	267ms	0	3.93Mbps	0
Standard	Yes	-50dBm	1ms	268ms	0	3.39Mbps	1
Large iAnt216	No	-50dBm	1ms	807ms	0	3.16Mbps	3
Large iAnt216	Yes	-50dBm	1ms	272ms	0	2.8Mbps	2
Dome iAnt215	No	-50dBm	1ms	266ms	0	4.23Mbps	0
Dome iAnt215	Yes	-50dBm	1ms	271ms	0	3.16Mbps	0
Small iAnt212	No	-50dBm	1ms	268ms	0	3.43Mbps	1
Small iAnt212	Yes	-60dBm	1ms	265ms	2 (4%)	2.39Mbps	1
Externally mounted dome iAnt215	Yes	-43dBm	1ms	274ms	0	3.15Mbps	0

Worst Case Scenario

Antenna	Isolator?	VSD Speed	Signal Strength	Min Ping	Max Ping	Packet Loss	Throughput	Gaps
Large iAnt216	Yes	20%	-50dBm	1ms	200ms	5 (10%)	253Kbps	45
Large iAnt216	Yes	80%	-50dBm	1ms	264ms	2 (4%)	256Kbps	50
Dome iAnt215	Yes	20%	-50dBm	1ms	289ms	5 (10%)	157Kbps	51
Dome iAnt215	Yes	80%	-50dBm	1ms	270ms	5 (10%)	69.9Kbps	49
Small iAnt212	Yes	20%	-70dBm	1ms	277ms	5 (10%)	175Kbps	45
Small iAnt212	Yes	80%	-60dBm	1ms	262ms	2 (4%)	122Kbps	50
Externally mounted dome iAnt215	Yes	20%	-47dBm	1ms	267ms	3 (6%)	262Kbps	46
Externally mounted dome iAnt215	Yes	80%	-43dBm	1ms	272ms	5 (10%)	140Kbps	50